

Dynamically Eliciting Unobservable Information*

Christopher P. Chambers[†]

Nicolas S. Lambert[‡]

February 6, 2014

Abstract

We answer the following question: At $t = 1$, an expert has (probabilistic) information about a random outcome X . In addition, the expert will obtain further information about X as time passes, up to some time $t = T + 1$ at which X will be publicly revealed. (How) Can a protocol be devised that induces the expert, as a strict best response, to reveal at the outset his prior assessment of both X and the information flows he anticipates and, subsequently, what information he privately receives? (The protocol can provide the expert with payoffs that depend only on the realization of X , as well as any decisions he may take.) We show that this can be done with the following sort of protocol: At the penultimate time $t = T$, the expert chooses a payoff function from a menu of such functions, where the menu available to him was chosen by him at time $t = T - 1$ from a menu of such menus, and so forth. We show that any protocol that affirmatively answers our question can be approximated by a protocol of the form described. We show how these results can be extended from discrete time to continuous time problems of this sort.

*We are particularly grateful to David Kreps for his invaluable help and advice. We also thank Matthew Jackson and seminar participants at Arizona State University, Microsoft Research, New York University, the Stanford Graduate School of Business, the Toulouse Network on Information Technologies, and the University of Rochester for helpful comments and suggestions. Lambert acknowledges financial support from the National Science Foundation under grant No. CCF-1101209.

[†]Department of Economics, University of California, San Diego; cpchambers@ucsd.edu.

[‡]Graduate School of Business, Stanford University; nlambert@stanford.edu.

1 Introduction

The classic literature on proper scoring rules answers the following question: An (expected-value maximizing) expert has a probability assessment concerning the outcome of a random variable. Can one devise a payoff function, or scoring rule, that rewards the expert based on a report he makes about his prior and the eventually realized outcome, which induces him to report his true prior as a strict best response?

In this paper, we answer the following generalization of the classic question. The random variable will be realized at some future time $T + 1$, and between now ($t = 1$) and time $T + 1$, the expert will receive information that may change his probability assessment. *Knowing nothing about the nature of this information—what it will tell the expert; when it will be received by the expert—can we design a protocol that induces the expert, as a **strict best response**, to reveal at the outset his prior and the structure of information he anticipates receiving, and then to provide truthful updates about information he receives as he receives it?* This protocol must provide the expert with a payoff at time $T + 1$ that depends only on the observable outcome of the random variable and the announcements of the expert prior to time $T + 1$.

When can such information be relevant? To set ideas and give the reader a concrete example to think about, let us consider the following story. On this Sunday, Alice, an expected value maximizer, wants to plan a ski trip for the following Saturday. She must decide whether to book the hotel. The price is \$200 a night, but it is the last day she can benefit from such a price. On Monday, the price will raise to \$250 a night, and will remain fixed at that price until the end of the week. Alice values a day of snow at \$500, and a day of no snow at \$0. (Staying home also gives her a value of \$0.) To help her with her decision, she consults Bob, a weather specialist. If Alice's last chance to book the hotel was this Sunday, then the only information relevant for her decision—booking or not booking—would be Bob's current estimate of the likelihood of snow. Let us suppose Bob gathers all information he has at the current time and, given this information, asserts the likelihood of snow to be .5. With this information alone, Alice's best decision is to book the hotel now to receive an expected payoff of \$50/day.

Now, if Alice has the possibility to book the hotel later in the week—at the increased price of \$250 a night—then the information Bob might learn during the week matters to Alice. If Bob expects to learn nothing new during the week, Alice is still better off booking the hotel on Sunday. But suppose Bob tells Alice that, on Wednesday, he will be able to study again atmospheric conditions. At that time, he will be able to revise his estimate to either .2 or .8. At the current time, he believes that either of these revisions will occur with .5 probability. In that case, Alice optimizes by postponing her decision, and she books the hotel on Wednesday only if Bob revises his assessment upwards, which allows her to raise her expected payoff by an additional \$25/day. In this example, as in most dynamic decision problems, it matters to know both the uncertainty of a future outcome and how uncertainty will unravel over time.

Many real-world decisions have a similar dynamic structure. A standard example is classical futures trading: A decision must be made through time as to how much of a commodity future

to buy at a given price, whose payoff obviously depends on the value of the commodity at the time the future matures. Other examples include revenue management (e.g., how much to price an airline ticket or a webpage advertisement at a given time), production and inventory planning (e.g., what is the optimal capacity production under uncertain market conditions), energy markets (e.g., how should a utility manage its fossil/nuclear electricity production so as to complete its weather-driven wind/hydro electricity productions¹), insurance markets (e.g., what is the value of a medical insurance plan) or irreversible investments (e.g., should an investor take shares in a risky project). This sort of decision problem is extensively studied in the literature on *real options* (see, for example, Dixit and Pindyck (1994)).

If the individual who makes the decision is informed as to the relevant probabilities, then her problem reduces to solving a dynamic optimization problem. However, in many situations, the individual who makes the decision is only partially informed and may consult a better-informed individual—an expert. The classical scoring rule literature constructs payoff functions that can be used to motivate the expert to reveal all the information he knows that can be relevant in solving a *static* decision problem. In this paper, we want to motivate the expert to reveal all the information he knows that can be relevant in solving a *dynamic* decision problem. It includes (probabilistic) information about the future observable outcomes, (probabilistic) information about future private information that the expert anticipates to receive over time, and what the expert privately observes.

In the Alice and Bob example, Bob’s ex-interim information consists of signals—temperatures, precipitation, and so forth. It is assumed that Alice cannot see Bob’s signals. In general, she may not even know the signals’ structure. However, in this example, the relevance of the signals to Alice is reflected solely in the probability estimate of snow it induces. So, all signal realizations become equivalent to the probability of snow they induce; such a probability can be interpreted as a “reduced” or “equivalent” signal. Hence, the information of interest to Alice is summarized by the probability tree induced by the signals. On Sunday, Bob holds probabilities about the probabilities he will hold on Wednesday regarding the outcome snow/no snow (his equivalent signal). Alice’s optimal Sunday decision then depends solely on this probability tree, and her optimal Wednesday decision depends solely on the branch of the tree that has been reached. This observation will be shown to hold true generally.

We show how to motivate the expert to reveal all relevant information with a carefully designed protocol of the following sort. At time T , the expert must choose a security (a function from publicly observable outcomes to payoffs) from a set or menu of such securities, where the menu available to him at time T was chosen by him at time $T - 1$ from a menu of menus. The menu of menus available to him at time $T - 1$ was chosen from a menu of menus of menus available to him at time $T - 2$, and so forth. We show that any protocol that induces the expert to reveal all information of interest can be approximated by the protocols just described, and we show to generalize the protocol to continuous time environments.

Our paper investigates the most difficult possible situation, whereby the individual observes no

¹To be cost-efficient, fossil and nuclear electricity facilities require smooth variation of their production rates.

information prior to the realization of the outcome. It is easy to see, therefore, that our results apply in an intermediate case in which the individual can observe information along the way. The protocols we construct continue to induce a truthful strict best-response even in this environment, though the essential uniqueness is obviously lost. Thus, our mechanism can be used to elicit the probabilities associated with temporal lotteries in the sense of Kreps and Porteus (1978), for example.

Our work is connected to recent work in decision theory, in particular the work of Dillenberger et al. (2012) and Lu (2013). Dillenberger et al. (2012) provide conditions on preference over menus of securities (which pay off as a function of the realization of the random variable) which are equivalent to an individual behaving as if he had a probability over probabilities over these outcomes; and at an interim stage after receiving their signal, chooses optimally. Dillenberger et al. (2012) perform a classical decision theoretic exercise: They place conditions *across* decision problems (that is, on a preference relation over menus) which are behaviorally equivalent to the model in question. Lu (2013) studies a stochastic choice version of the model. In his setup, one observes a probability distribution over each menu—these could correspond, for example, to empirical frequencies of choices from an individual or from society. Lu investigates conditions under which these probability distributions are consistent with the distribution of the conditional choices that would be made upon observing the realization of a random private signal.

In both cases, the model is identified, i.e., one can completely recover all underlying probability measures if one can observe the individual’s entire preference relation (in Dillenberger et al., 2012) or the stochastic choice function (in Lu, 2013). In our Alice and Bob example, this would involve Alice offering to Bob an infinite collection of choices. Such an experiment is a conceptual impossibility: A feasible experiment must bound the number of choices. (Another issue that appears with the multiplicity of choices is the complementarity effect, see for example the discussion in Wold (1952), Allais (1953), or Savage (1954, pp. 29–30).)

In the case of risk neutrality and classical subjective expected utility as axiomatized by Savage (1954) or Anscombe and Aumann (1963), we can elicit an entire preference via a (*strictly*) *proper scoring rule*, which is a single menu of securities, the choice of which completely reveals the individual’s subjective probability. Another early idea to eliciting a preference, first suggested by Allais (1953) and W. Allen Wallis (see Savage, 1954) is to ask the individual to report a choice from a rich enough class of sets, enough that we would be able to uniquely identify preference, and then randomly select one of these sets, paying the individual her announced choice from this set. As this involves only a single choice, this procedure can be experimentally operationalized. In fact, strictly proper scoring rules can be interpreted as an application of Allais’ idea (Savage, 1971). Our paper leverages risk-neutrality and Allais’ idea to construct scoring rules in the dynamic setting in which intermediate information occurs which is unobservable except to the expert.²

The paper is organized as follows: In Section 2, to give the reader an intuition of our results and provide a concrete example of the protocols aforementioned, we analyze our Alice-and-Bob story.

²Risk-neutrality is not binding, as we will observe, our protocols of elicitation extend directly to experts with more general risk preferences.

Section 3 lays down the formal details of the model. The main result, with multiple, discrete time periods is described in full detail in Section 4. We extend the existence results to a continuous-time framework in Section 5. Section 6 provides a brief literature review. Section 7 concludes. The Appendix includes several additional results and the proofs that are excluded from the main body of the paper.

2 A simple example: the two-stage quadratic scoring rule

In this section we explore the simple example of Alice and Bob presented in the introduction. Let X denote the random variable that indicates the event of snow, $X = 1$ (snow) or $X = 0$ (no snow). There are three time periods. From the last period to the first, the following events take place:

- At $t = 3$, X is publicly realized.
- At $t = 2$, Bob receives signals, that he and only he can observe, and from which he infers a posterior assessment of snow p .
- At $t = 1$, given all information received by that time, Bob forms a distribution about the future signals he expects to get at time $t = 2$. The relevant information is summarized by a distribution function F (here, the prior) over the posterior p to be observed (or inferred) at time $t = 2$. Note that the expected value of p under F gives Bob's assessment of the probability of snow at time $t = 1$, which is, in that sense, redundant information.

Alice wants to motivate Bob (an expected-value maximizer) to provide, first, the prior F at time $t = 1$, and then the posterior p at time $t = 2$. She wants to motivate Bob with strict incentives: Bob should respond truthfully at each time and it must be his only best response. We observe that the objects F and p are the only relevant information useful in solving a dynamic decision problem that spans over these three time periods. The decision problem Alice confronts is outside of the model; we simply suppose that there *is* some decision problem for which F and p are relevant to her (such as to whether and how to plan her holiday trip). We impose no particular restriction on F and p .

Our construction uses two types of instruments, *securities* and *menus of securities*. A *security* is an outcome-contingent payoff, represented as a map from outcomes $\{0, 1\}$ to real numbers, or equivalently an element of \mathbb{R}^2 . The simplest instance is the 0-1 security that pays off 1 if it snows, and nothing if it does not. A *menu of securities* is a collection of securities, with the following interpretation: An individual who is offered a menu of securities at time $t = 1$ is required to choose exactly one security from the menu at time $t = 2$.

To grasp the intuition of our result, we begin with a simplified environment in which Alice knows that Bob's prior can be either F_1 or F_2 . For example, F_1 may correspond to a case in which Bob expects to revise his initial probability assessment substantially, while under F_2 Bob remains relatively less informed. If we can design two menus M_1 and M_2 such that the F_1 -type would rather have menu M_1 than menu M_2 , and vice-versa for the F_2 -type, then offering the choice of M_1 and M_2 to Bob at time $t = 1$ induces Bob to reveal his prior.

Let us consider the menu $M_1 = \{S_1, S'_1\}$, with $S_1 = 0$ (security that yields zero payoff) and $S'_1(x) = \alpha - x$ for some fixed constant α . The menu can be interpreted as an option to sell (at time 2) the 0-1 security displayed above for a price α . Bob, with a prior F , derives an expected payoff (at time 1) from the option which is equal to

$$\begin{aligned} \int_{-\infty}^{+\infty} \max(0, \alpha - p) dF(p) &= \int_{-\infty}^{\alpha} (\alpha - p) dF(p), \\ &= F(\alpha)\alpha - \int_{-\infty}^{\alpha} p dF(p), \\ &= \int_0^{\alpha} F(p) dp. \end{aligned}$$

If $F_1 \neq F_2$ then, for some α , $\int_0^{\alpha} F_1(p) dp \neq \int_0^{\alpha} F_2(p) dp$, for example $\int_0^{\alpha} F_1(p) dp > \int_0^{\alpha} F_2(p) dp$. Thus, a Bob with prior F_1 values the option (menu M_1) more than a Bob with prior F_2 . Now let us fix a constant

$$\beta = \frac{1}{2} \left(\int_0^{\alpha} F_1(p) dp + \int_0^{\alpha} F_2(p) dp \right),$$

and fix a degenerate menu $M_2 = \{\beta, \beta\}$ which yields the constant payoff β . Then if Bob has prior F_1 , he will prefer to be offered the option (menu M_1), but if he has prior F_2 , he will prefer to be offered the fixed payoff β (menu M_2).

This shows we can distinguish between any two given prior distributions for Bob with a single menu choice. In addition, it is easy to verify that, if Bob can only have two commonly known posteriors, p_1 and p_2 , then by offering the menu $\{0, \alpha - x\}$ with well-chosen α , we induce Bob to reveal his posterior as a strict best response.

We now return to the general case with no restrictions imposed on the prior or posterior. To distinguish between all prior (and posterior) distributions, we can introduce randomization, in the spirit of Allais (1953) and Becker et al. (1964). In this particular example, we will randomly select two menus of two securities by choosing the parameters α and β at random from the uniform distribution on $[0, 1]$ and $[-1, 1]$ respectively. (Allowing negative values for β ensures that with positive probability, Bob will get a non-trivial menu at time 2, which motivates him to give his posterior.)

There is a natural extension of the Allais idea to the dynamic setting. We suppose Alice never informs Bob which menu was selected, until after all uncertainty about the random variable is resolved. At the first stage, Alice picks the best menu in Bob's interest depending on his announcement. Bob does not know which choice of menu he is really facing—all he needs to know is that Alice will make the best choice for him given his report. At the second stage, Alice chooses the optimal security from the selected menu, given Bob's second stage report \hat{p} . We have chosen the lottery distribution so that, for any pair of prior distributions F and $\hat{F} \neq F$, if F is the true distribution and \hat{F} is the reported distribution, there is a sufficiently rich class of menus distinguishing between F and \hat{F} , in the sense of the preceding paragraph; and similarly for true p and report \hat{p} .

In summary, the elicitation protocol works as follows:

- Alice draws two random numbers α and β independently and uniformly from $[0, 1]$ and $[-1, 1]$ respectively. She then forms the menus $M_1 = \{\alpha - x, 0\}$ and $M_2 = \{\beta, \beta\}$.
- At time $t = 1$, Alice asks Bob to reveal his distribution over posteriors. Based on his report \hat{F} , Alice chooses a menu M^* of M_1 and M_2 which maximizes Bob's expected value, assuming truthful reports, ties broken in favor of M_2 for example.
- At time $t = 2$, Alice asks Bob to reveal his posterior. Based on his report \hat{p} , Alice gives Bob a security S^* from M^* which maximizes Bob's expected value at that time, again assuming a truthful report, ties broken in favor of the constant security, for example.
- At time $t = 3$, the outcome is observed and the security pays off.

For any given choice of two menus identified via α, β , Bob's final payoff is:

$$\Pi(\hat{F}, \hat{p}, x; \alpha, \beta) = \begin{cases} \beta & \text{if } \beta \geq \int \max(\alpha - p, 0) d\hat{F}(p), \\ 0 & \text{if } \beta < \int \max(\alpha - p, 0) d\hat{F}(p) \text{ and } 0 \geq \alpha - \hat{p}, \\ \alpha - x & \text{otherwise.} \end{cases}$$

As described, the dynamic counterpart of the Allais mechanism relies on Alice's ability to commit to making the best choices on Bob's behalf *without* Bob having the ability to verify Alice's actions. But from Bob's viewpoint, the randomized protocol reduces to a deterministic payoff rule that simply corresponds to the average payoff of the protocol over α and β :

$$\Pi(\hat{F}, \hat{p}, x) = \int_0^1 \int_{-1}^1 \Pi(\hat{F}, \hat{p}, x; \alpha, \beta) d\beta d\alpha.$$

Hence, the expected payoff to Bob is a deterministic mechanism. We can calculate an explicit formula for this mechanism, which we refer to as the *two-stage quadratic scoring rule*.

$$\Pi(\hat{F}, \hat{p}, x) = \frac{1}{2} - \frac{1}{2} \int_0^1 \left(\int_0^\alpha \hat{F} \right)^2 d\alpha + \int_{\hat{p}}^1 \left(1 + \int_0^\alpha \hat{F} \right) (\alpha - x) d\alpha.$$

It is straightforward to establish that Π has the right incentive properties at every stage. Suppose Bob's true prior at time 1 is F and his true belief at time 2 is given by p . We first claim that that it is a strict best response for Bob to report truthfully at time 2, no matter his previous declarations. His expected payoff from announcing \hat{p} given his information at time 2 is:

$$\frac{1}{2} - \frac{1}{2} \int_0^1 \left(\int_0^\alpha \hat{F} \right)^2 d\alpha + \int_{\hat{p}}^1 \left(1 + \int_0^\alpha \hat{F} \right) (\alpha - p) d\alpha.$$

If Bob moves his report from p to a different \hat{p} , his expected earnings decrease by

$$\int_p^{\hat{p}} \left(1 + \int_0^\alpha \hat{F}\right) (\alpha - p) d\alpha > 0.$$

We can also see that it is a best response for Bob to report truthfully at time 1. The expected payment, conditional on Bob's information time 1, is

$$\frac{1}{2} - \frac{1}{2} \int_0^1 \left(\int_0^\alpha \hat{F}\right)^2 d\alpha + \int_0^1 \int_p^1 \left(1 + \int_0^\alpha \hat{F}\right) (\alpha - p) d\alpha dF(p),$$

which, after algebraic manipulations, reduces to

$$\frac{1}{2} + \frac{1}{2} \int_0^1 \left(\int_0^\alpha F\right) d\alpha + \frac{1}{2} \int_0^1 \left(\int_0^\alpha F\right)^2 d\alpha - \frac{1}{2} \int_0^1 \left(\int_0^\alpha \hat{F} - \int_0^\alpha F\right)^2 d\alpha.$$

Observe that the last term only affects the incentives, and if $\hat{F} \neq F$, the right-continuity of distribution functions implies that $\int_0^\alpha \hat{F} \neq \int_0^\alpha F$ for an open interval of α 's. So that by reporting \hat{F} instead of F , Bob loses on expectation

$$\frac{1}{2} \int_0^1 \left(\int_0^\alpha \hat{F} - \int_0^\alpha F\right)^2 > 0.$$

Thus we get that the reward scheme has the desired incentive property.

In Appendix A we present a more general mechanism that applies to the case of a categorical random variable, i.e., a variable that takes value in a finite set $\{1, \dots, n\}$. A randomization device generates two menus, in which the securities are drawn independently from any full support distribution. The incentive properties rely crucially on the fact that a distribution over \mathbb{R}^n is uniquely determined by the distributions of the linear combinations of its components, a result known as the Cramér-Wold theorem. The Cramér-Wold theorem generally does not hold for larger outcome spaces, but we will show that the desired incentive properties then be achieved by including more securities and more menus to choose from. Eventually, with multiple time periods, we will be able to recover the entire information structure via a hierarchy of menus.

Finally, it is worth observing that it is not possible to motivate Bob via two successive static mechanisms (i.e., such as one scoring rule to elicit the distribution over the future posterior, and one scoring rule to elicit the realized posterior). One could think that, since there are reward schemes that motivate Bob to reveal a probability assessment of any random variable that is subsequently observed, one could use the following two-stage procedure. First, apply such a reward scheme at time 1 to motivate Bob to tell the true prior, assuming that he will reveal the true posterior, and second, use another such reward scheme at time 2, with strong-enough incentives so to induce Bob to reveal the true posterior. In Appendix B, we show that such a construction is impossible.

3 Model

In the main model of this paper, there are $T + 1$ time periods. An individual has interest in information about an outcome that materializes at $t = T + 1$. The outcomes take value in a compact metrizable space \mathcal{X} (for example, a compact subset of Euclidean space, or a finite set). In a later section we extend our main result to Polish spaces (for example, \mathbb{R}^n) and to continuous time (allowing to track a continuous signal or to have private random times of information arrivals).

At time $t = 1$, an expert holds information about the distribution of the outcome. At every intermediate time $t = 2, \dots, T$, the expert possibly receives additional private information that affects his assessment of the outcome distribution. The expert's information structure captures the information that will be learned over time, and the uncertainty associated with it. It includes what the expert knows, at $t = 1$, about the outcome and its distribution. It also includes what the expert knows about the information he anticipates to receive at every future time $t = 2, \dots, T$. In particular, it includes information about when and how the uncertainty on the outcome gets resolved over time.

The standard way to represent dynamic information in uncertain environments in a general fashion is through filtered probability spaces.³ Therefore, we define the expert's information structure as follows:⁴

Definition 1. An *information structure* is a tuple $(\Omega, \mathbb{F}, \mathbb{P}, X)$ in which:

- Ω is a set of states of the world, endowed with a separable metrizable topology. The state of the world captures every aspect of the world that is relevant to the expert in the current situation. Common examples include finite sets, finite-dimensional spaces, and function spaces.⁵
- $\mathbb{F} = \{\mathcal{F}_1, \dots, \mathcal{F}_T\}$ is a filtration. Every \mathcal{F}_t includes all information (the events) that are known to be either true or false by the expert at time t , with $\mathcal{F}_t \subseteq \mathcal{F}_{t+1}$.
- \mathbb{P} is a full prior, described below.
- $X : \Omega \mapsto \mathcal{X}$ is a random variable that connects the (internal) state of the world to the (observable) outcome of interest.

³This includes as a special case information represented as time-dependent information partitions and information represented via a sequence of signals.

⁴As a simpler alternative, we could consider an expert who gets a signal at every time $t = 1, \dots, T$. The signal at time t is given by a random variable Y_t , taking values in some set \mathcal{Y}_t , for example \mathbb{R} . Each signal can carry information on the distribution of the random outcome, which itself is specified by a random variable X . The information structure, that the expert would be asked to communicate at $t = 1$, would then consist in the space of signals $\mathcal{Y}_1, \dots, \mathcal{Y}_T$, along with a joint prior distribution over (Y_1, \dots, Y_T, X) and, conditional on every signal history $h_t = (y_1, \dots, y_t)$ up to time t , a posterior distribution on (Y_{t+1}, \dots, Y_T, X) given h_t . In addition, at every time $t = 1, \dots, T$, the expert would be asked to provide the observed signal realization. It is a simple matter to adapt our mechanisms to this setting, which is a special case of the current setup. It is worth mentioning that the problem of eliciting the distributions over fixed signal spaces (e.g., real-valued or finite-valued signals) and the subsequent signal realizations is not much simpler than the general setup we consider in this paper. In that sense, the generality of our setup does not impose an additional layer of complexity.

⁵In this paper, every set which is given a particular topology is tacitly endowed with its Borel σ -algebra of events. In many cases of interest, Ω is finite or finite dimensional, and the set of events is as usual.

The information structure is known to the expert, and, in general, the expert only. This means the individual consulting the expert generally does not know the state space Ω , and she does not observe the state of the world. However she eventually observes the *outcome* of the random variable X , and this—together with the expert’s reports—is the only data point she can use in the elicitation procedure.

Given a state space Ω , we let $\Delta(\Omega)$ be the set of all probability measures on Ω . It is endowed with the weak-* topology. The filtration only carries information about what the expert will know for sure in every state of the world. It is the full prior that provides information about the uncertainty over states and events:

Definition 2. Given a state space Ω and a filtration \mathbb{F} , a full prior is a process $P : (t, \omega) \mapsto P_t^\omega$ with values in $\Delta(\Omega)$, such that for all t , and all events E , $\omega \mapsto P_t^\omega(E)$ is \mathcal{F}_t -measurable.

The full prior provides, at every time, the updated distribution over states, given the information available to the expert at that time. At time $t = 1$, in state ω , the expert forms a posterior over states as a function of his information, given by $P_1^\omega \in \Delta(\Omega)$. Then, at every subsequent time $t > 1$, after observing information in \mathcal{F}_t , the expert revises his posterior to P_t^ω . A full prior contains information at every time, as opposed to time $t = 1$ only. This is useful as we must be able to condition on future events in a consistent manner.⁶

It is generally not possible to motivate the expert to reveal *all* of his information about Ω —for instance the expert may have information that does not even concern the future observable outcome. But as the individual’s interest is focused on the future outcome, disclosure of the expert’s information is valuable to her only to the extent that it impacts her beliefs on the random outcome.

In our Alice and Bob example, we assumed Alice only cared to learn the probability tree that can be inferred from Bob’s dynamic information. In the sequel we will continue to assume that, at every time, the individual only wants to learn all probability trees that can be inferred from the expert’s private information. That is, at every time, the individual only cares to learn about (i) an assessment of the outcome distribution, and (ii) an assessment of all the distributions over future posterior distributions (which themselves may concern the outcome, or may concern other future posterior distributions, and so forth). Formally, at every time t_0 , the information of interest includes:

- the best estimate of the outcome distribution, which we refer to as posterior of order 1,
- the distributions on the posterior outcomes, that will be known at a given future time $t_1 > t_0$ (a posterior of order 2, represented as a one-level probability tree),

⁶It is well-known that a prior does not always generate consistent conditional probabilities, and even when it does, these are generally not unique. Embedding such information in a full prior process avoids the need to put overly restrictive assumptions on the state space, and the problem induced by the multiplicity of conditional probabilities. For general conditions that ensure systematic existence and coherence of regular conditional probabilities, see Berti and Rigo (1996).

- the distributions on the posterior that will be known at some $t_1 > t_0$, over the posteriors on outcomes, that will be known at some $t_2 > t_1$ (posterior of order 3, represented as a two-level probability tree),
- and so forth, for all distributions of order up to T .

We let $\Delta^1(\mathcal{X}) = \Delta(\mathcal{X})$ be the set of distributions over outcomes (posteriors or distributions of order 1), and define $\Delta^k(\mathcal{X})$ recursively by $\Delta^k(\mathcal{X}) = \Delta(\Delta^{k-1}(\mathcal{X}))$ (posteriors or distributions of order k).

To link information structures and posteriors of higher orders (probability trees), we define the induced belief tree.

Definition 3. Given an information structure $(\Omega, \mathbb{F}, P, X)$, an *induced belief of order 1* is the map $\varphi : \Delta(\Omega) \mapsto \Delta(\mathcal{X})$ defined as

$$\varphi(Q) = Q(X).^7$$

The *induced belief tree of order $j + 1$ with intermediate times $t_1 < \dots < t_j$* is the map noted $\varphi_{t_1, \dots, t_j} : \Delta(\Omega) \mapsto \Delta^{j+1}(\mathcal{X})$ and defined recursively as

$$\varphi_{t_1, \dots, t_j}(Q) = Q(\varphi_{t_2, \dots, t_j}(P_{t_1}))$$

where P_{t_1} is the random variable of the process P sampled at time t_1 .

That is, the induced belief tree of order 1 corresponds to the posterior over outcomes that can be inferred from the expert's information structure and the expert's posterior over states at a given time; the induced belief tree of order 2 with intermediate time t_1 corresponds to the distribution over what the posterior over outcomes will be at the future time t_1 , and so forth.

Because we infer distributions over future posteriors, we must endow each space of distribution with a σ -algebra of events. We endow each $\Delta^k(\mathcal{X})$ with the weak-* topology and the usual Borel σ -algebra.⁸ We will discuss the reasons that motivate this choice in the next section. It will be fully transparent in the actual elicitation procedure.⁹

Lemma 1. *All induced belief trees are well-defined and measurable.*

Proof. See Appendix D. ■

⁷For a probability measure P and random variable Z taking values in a set \mathcal{Z} , we denote by $P(Z)$ the probability law of Z —i.e., the probability measures over the values taken by Z —and by $\Delta(\mathcal{Z})$ the set of all probability laws of Z . We extend the notation to random vectors and stochastic processes.

⁸The weak-* topology refers to the weakest topology for which for any continuous function, integration with respect to that function is a continuous linear functional.

⁹In the Alice and Bob example, we argued that Alice only cares about the implied probability tree as it is the only information needed to solve any sort of dynamic decision problems. The argument continues to hold in our model when the set of states of the world is finite. For the general case however, the argument does not follow through directly. The issue comes from that the choice of the σ -algebra determines which events are assigned a probability. It is not innocuous because, when the expert's information is rich enough, we may miss useful information if we do not include sufficiently many events. However we show in Appendix C that, under some mild conditions on the kind of decision problems of interest, there is no loss of valuable information.

At a given time t , a posterior of order $k \leq T - t$ can be computed by taking the expectation over the posteriors of higher order. So the only relevant piece of information to the individual is captured by a single posterior, the (only) posterior of order $T - t + 1$ (which gives the “probability tree” relevant at time t).

We can thus restrict attention to protocols in which the expert communicates this information directly. We refer to these as *direct protocols*. All protocols considered in this section and the next are assumed to be direct protocols. Considering such protocols simplify the model, however it also possible (and in certain cases more practical) to extend the model to the more general class of protocols that ask the expert to report the information actually observed (which allows, for example, a direct report of a signal, instead of the whole probability tree). We consider this approach in Section 5.

A direct protocol asks the expert to reveal a probability tree (a posterior of order $T - t + 1$) at every time $t \leq T$. The expert is then rewarded at time $T + 1$ as a function of his reports and the final outcome that obtains. The rewards are specified by a *payoff rule* Π , where $\Pi(q_1, \dots, q_T, x)$ gives the final payoff to the expert who has reported a posterior $q_t \in \Delta^{T-t+1}(\mathcal{X})$ at every time t . We require that the payoff rule be jointly measurable in (q_1, \dots, q_T, x) and that it takes values in some bounded interval $[\underline{L}, \overline{U}]$. Boundedness and joint measurability ensure that payoffs are well-defined random variables whose expectations exist and are finite, independently of the expert’s information.

The expert chooses what to communicate as a function of what he knows. For an expert with information structure $(\Omega^*, \mathbb{F}^*, \mathbb{P}^*, X^*)$, a *strategy* is a tuple $f = (f_1, \dots, f_T)$, in which $f_t : \Omega \mapsto \Delta^{T-t+1}(\mathcal{X})$ maps a state to a distribution of order $T - t + 1$. We require each f_t to be measurable with respect to \mathcal{F}_t^* : The expert can only condition his report on his knowledge at the time of the report.

We are interested in risk-neutral experts, though our results can be generalized.¹⁰ The *time- t value* at state ω^* is the expected payment given the expert’s information at time t . For an expert with information structure $(\Omega^*, \mathbb{F}^*, \mathbb{P}^*, X^*)$ who communicates according to strategy f , it is explicitly written

$$V_t^{\omega^*}(f) = \int_{\Omega} \Pi(f_1(\omega'), \dots, f_T(\omega'), X^*(\omega')) d\mathbb{P}_t^{*, \omega^*}(\omega').$$

It is easily verified that the value process V , that records the expected payoff at every time in every state, is a well-defined stochastic process (Kechris, 1995, Theorem 17.24).

An expert who participates in a given protocol wants to maximize his value at every instant.

Definition 4. Given a direct protocol, a strategy f is *optimal at time t and state ω^** for an expert with information structure $(\Omega^*, \mathbb{F}^*, \mathbb{P}^*, X^*)$ when, for all strategies f', f'' such that $f'_\tau = f''_\tau$ when $\tau < t$ and $f'_\tau = f_\tau$ when $\tau \geq t$, $V_t^{\omega^*}(f') \geq V_t^{\omega^*}(f'')$.

Optimality of a strategy f at a time t requires that, no matter the reports made up to time $t - 1$, the expert is always best off following strategy f from time t onwards.

¹⁰It is also well-known that one can avoid the risk-neutrality assumption by paying experts with securities which pay in lotteries over a good and bad outcome, as in Anscombe and Aumann (1963); see Allen (1987).

Definition 5. A direct protocol is *strategyproof* when, for all information structures $(\Omega^*, \mathbb{F}^*, \mathbf{P}^*, X^*)$:

- There exists a strategy that is optimal at every time and in every state.
- If a strategy f is optimal at time t and state ω^* , then $f_t(\omega^*) = \varphi_{t+1, \dots, T}^*(\mathbf{P}_t^{\omega^*})$, with $\varphi_{t+1, \dots, T}^*$ the induced belief tree associated with the information structure $(\Omega^*, \mathbb{F}^*, \mathbf{P}^*, X^*)$.

Our concept of strategyproofness is robust to deviation from the equilibrium paths: An expert who misreports up to some time still wants to report true distributions at all future times.

4 Main results

4.1 A class of strategyproof protocols

Three types of instruments will be used: *securities*, *menus of securities*, and *menus of (sub-)menus*. A security is a bounded continuous map $S : \mathcal{X} \rightarrow \mathbb{R}$. Menus of securities are collections of securities, and menus of menus are collections of other menus. To distinguish between the different types of menus, we call *menu of order 1* a collection of securities, and *menu of order k* a collection of menus of order $k - 1$. The space of securities is given the usual sup-norm topology.

Menus have the following interpretation: A person who owns, today (at time t), a menu of order 1 must choose without delay to get one (and only one) security from the menu. A person who owns, today (at time t), a menu of order k must choose without delay, among the sub-menus of order $k - 1$ that are included in the menu of order k , exactly one sub-menu to own tomorrow (at time $t + 1$). Thus, after $k - 1$ selections of menus/sub-menus at times $t, \dots, t + k - 2$, the person will eventually choose to own one security from the last menu available to her. When a person makes a selection among the elements of a menu, we say that he or she *exercises the menu*. A menu is *finite* when it contains a finite number of securities or when it contains a finite number of sub-menus, themselves being (recursively) finite. We will denote by \mathcal{M}_t the collection of finite menus of order $T - t + 1$, and by $\mathcal{M}_t^{[a, b]}$ the collection of finite menus of order $T - t + 1$ whose securities takes values in $[a, b]$.

An informal description of our elicitation protocols is as follows. As a preliminary step, the individual who administers the protocol draws a finite menu of order T at random. She keeps the menu hidden from the expert. At every $t = 1, \dots, T$, the expert declares a distribution q_t over the set $\Delta^{T-t}(\mathcal{X})$ of posteriors that he expects to obtain at time $t + 1$ (so, essentially, the expert declares a probability tree), or, if $t = T$, his posterior over the set of outcomes. At every time, the administrator exercises the menu privately and optimally on behalf of the expert, according to the expert's reported information and under the assumption the expert communicates truthfully. At $t = T$, after the expert makes the last report, the individual pays off the expert with the security optimally chosen from the last menu selected.

The protocols we describe use a randomization device on menus. We will show below a simple procedure that generates an appropriate random draw, however to introduce the class of mechanisms in its generality, and use an arbitrary randomization device, we must define the sets of menus to which a probability is attached—the σ -algebra. Noting that a menu is a subset from the space of

sub-menus or the space of securities, it is natural to equip every set \mathcal{M}_t with the Hausdorff metric topology, and then use the usual Borel σ -algebra.

The Hausdorff metric is a standard way to measure distances between sets. If d is a metric on \mathcal{X} ,¹¹ the Hausdorff metric on every \mathcal{M}_t is defined recursively by

$$d(M^1, M^2) = \max \left\{ \max_{S^1 \in M^1} \min_{S^2 \in M^2} d(S^1, S^2), \max_{S^2 \in M^2} \min_{S^1 \in M^1} d(S^1, S^2) \right\}$$

for $M^1, M^2 \in \mathcal{M}_T$,

$$d(M^1, M^2) = \max \left\{ \max_{m^1 \in M^1} \min_{m^2 \in M^2} d(m^1, m^2), \max_{m^2 \in M^2} \min_{m^1 \in M^1} d(m^1, m^2) \right\}$$

for $M^1, M^2 \in \mathcal{M}_t, t < T$.

For any time $t = 1, \dots, T-1$, any finite menu M_t of order $(T-t+1)$ and posterior $p_t \in \Delta^{T-t+1}(\mathcal{X})$ of the same order, we denote by $\pi_t(M_t, p_t)$ the value of menu M_t at time t to a risk-neutral person who possesses information of order $T-t+1$ summarized by p_t . Recursively, we have

$$\pi_T(M_T, p_T) = \max_{S \in M_T} \int_{\mathcal{X}} S(x) dp_T(x),$$

$$\pi_t(M_t, p_t) = \max_{m_{t+1} \in M_t} \int_{\Delta^{T-t-1}} \pi_{t+1}(m_{t+1}, q) dp_t(q), \quad \text{if } t < T.$$

We now formally detail the protocol. Let ξ be a probability measure on $\mathcal{M}_1^{[\underline{L}, \bar{U}]}$. The T -period randomized protocol initiated by ξ proceeds as follow:

- In a preliminary step, a menu M_1^* is drawn at random according to ξ .
- Then, at every time $t = 1, \dots, T-1$, the expert communicates a distribution q_t of order $(T-t+1)$. The individual who administers the protocol selects a sub-menu $M_{t+1}^* \in M_t^*$ that maximizes the expert's time- t value given the expert's declaration:

$$M_{t+1}^* \in \arg \max_{m_{t+1} \in M_t^*} \int \pi_{t+1}(m_{t+1}, p) dq_t(p).$$

If there is more than one sub-menu that is optimal for the expert, the individual selects a sub-menu uniformly at random among all the optimal sub-menus.¹²

- At time $t = T$, the expert reports a posterior over outcomes p_T and is offered a security S^*

¹¹Because menus are finite sets at every level, the σ -algebra of events does not depend on the particular metric on the space of securities, as long as it generates the same topology (see, for example, Aliprantis and Border, 2006, Theorem 3.91).

¹²We will show that selecting a sub-menu uniformly at random guarantees the measurability of the payments. Alternatively, the individual could provide the expert with an equal fraction of all optimal sub-menus, or the individual could choose an optimal sub-menu according to a measurable selection. In the proof of Lemma 3 we show that such a measurable selection always exists.

from the last menu selected that maximizes his expected payment:

$$S^* \in \arg \max_{S \in M_T^*} \int S(x) dq_T(x).$$

If more than one security is optimal for the expert, the individual selects S^* uniformly at random among all the optimal securities.

Given a menu M_1 drawn in the preliminary step, the payoff to the expert, denoted $\Pi^1(q_1, \dots, q_T, x; M_1)$ can be expressed recursively by

$$\Pi^T(q_T, x; M_T) = \frac{1}{|\mathcal{K}|} \sum_{S \in \mathcal{K}} S(x), \quad \text{with} \quad \mathcal{K} = \arg \max_{S \in M_T} \int S(x) dq_T(x),$$

and if $t < T$,

$$\Pi^t(q_t, \dots, q_T, x; M_t) = \frac{1}{|\mathcal{K}|} \sum_{M_{t+1} \in \mathcal{K}} \Pi^{t+1}(q_{t+1}, \dots, q_T, x; M_{t+1}),$$

$$\text{with} \quad \mathcal{K} = \arg \max_{m_{t+1} \in M_t} \int \pi_{t+1}(m_{t+1}, p) dq_t(p).$$

Our first main result asserts that the protocol just described is strategyproof, provided that the probability measure ξ has full support over the finite menus of order T . We recall that, in the current context, a probability distribution over finite menus of order T has full support when for every finite menu M of order T , the probability of drawing a menu M_1^* at most ϵ -close to M is positive (with respect to the Hausdorff metric).

It is also worth noting that, although in principle ξ can be any full-support probability measure, generating an appropriate random draw of menus can be done in a simple and concrete fashion. For example, let F_S be a probability distribution over securities whose values take outcomes in $[\underline{L}, \bar{U}]$ with full support, and F_N a probability distribution over $\{1, 2, \dots\}$ with full support.¹³ We draw a menu of order $T - 1$ at random according to a simple recursive procedure: To draw a menu of order $k \geq 2$ at random, we draw a random number (according to F_N) of menus of order $k - 1$ independently at random; to draw a menu of order 1 at random, we draw a random number (according to F_N) of securities at random (according to F_S) independently. This simple procedure generates a probability measure ξ with full support.

Before we move to the main theorem, we verify that the payoffs are indeed proper random variables, which ensures the protocol is well-defined and allows the computation of expected payoffs. To do so, we rely on the measurable selection theorem, which requires that objective functions be continuous. We therefore begin by proving continuity of the value functions and the step-ahead

¹³When \mathcal{X} is infinite, an example of such a distribution results by taking a countable dense subset of these securities (which exists by (Aliprantis and Border, 2006, Lemma 3.99)), ordering them x_1, x_2, \dots , and giving probability 2^{-k} to x_k . If \mathcal{X} is finite, the space of securities can be identified with $[\underline{L}, \bar{U}]^n$ and one can use the uniform distribution, for example.

value (objective functions). The proofs of the two lemmas can be found in Appendix D.

Lemma 2. *For every t , the time- t value function $\pi_t(M_t, p_t)$, for finite menu M_t associated with posterior p_t , is jointly continuous, and the step-ahead value, $\int \pi_t(M_t, q_t) dp_{t-1}(q_t)$, is also jointly continuous in M_t and p_{t-1} .*

Lemma 3. *The map $(q_1, \dots, q_T, x, M_1) \mapsto \Pi^1(q_1, \dots, q_T, x; M_1)$ is jointly measurable in the product σ -algebra.*

We can now state and prove our main result.

Theorem 1. *If ξ has full support, then the T -period randomized protocol just described, initiated by ξ , is strategyproof.*

Before we lay down the proof, it is useful to discuss the intuition behind the design of the elicitation protocol. To recover the information structure of the expert, we essentially ask him to solve simultaneously a large number of “simple” dynamic decision problems. In our mechanisms, these decision problems correspond to the exercise decisions at different times for a decision maker who is offered a particular finite menu at time 1.

Observing the exercise decisions for every one of these problems give us some information about what the expert knows, but not all of it, because there are only a finite number of choices at every time, but infinitely many information structures, or “probability trees”. Therefore, we must ask the expert to solve sufficiently many simple decision problems so as to recover the entire information structure. This is the role of the randomization. Because the expert does not know which menu he faces, he is induced to reveal enough information so that the individual who decides on his behalf gets to solve all, or nearly all possible simple decision that may come up. In the risk-neutral framework, this is equivalent to offering the expert a small fraction of every simple decision problem.

However, we must also ensure the expert has enough incentives to produce a strict best response. If offered a “too small fraction” of a decision problem, the incentives vanish. If offered a “too large fraction” of a decision problem, or if there are “too many” such problems, the payoffs are no longer finite, and incentives vanish again. For this reason, we must randomize over a relatively small collection of simple decision problems, which at the same time must be rich enough to uncover the object of interest, i.e., the induced belief tree.

A difficulty arises as this object can become extraordinarily complex. For concreteness, take the Alice and Bob example with $T + 1$ periods. At time T , we only require a distribution over outcomes, which reduces to the probability of snow—a real number. At time $T - 1$, we require a distribution over distributions over outcomes—a finite-dimensional distribution. Thus, if we care to have only two periods, the objects to extract are well-known and simple. But starting at time $T - 2$, we require a distribution over finite-dimensional distributions—i.e., a distribution on a function space. At time $T - 3$, we require a distribution over distributions over real functions, and so forth. The danger is that, as we add more time periods, this object becomes increasingly richer, so rich that a tractable collection of simple decision problems may not be enough to retrieve it.

The key to solve this problem is to control the amount of information that is effectively being retrieved, by specifying the σ -algebra of events. By choosing the σ -algebra generated by the weak- $*$ topology, we ensure that the amount of data needed to encode these “probability trees” remains manageable; in fact, it stops growing very quickly, due to a mathematical result that establish that any such tree can be described entirely by a vector of $[0, 1]^\infty$ (often interpreted as the “smallest” infinite dimensional object). Of course, one could retrieve even less information, but as we show in Appendix C this particular σ -algebra captures just enough information so as to solve any decision problem as well as if we had direct access to the expert’s information structure and private observations.

This choice turns out to have crucial consequences of the design of our elicitation protocol. It allows to focus on finite menus—a very small class of dynamic decision problems—and to deal with extremely simple randomization procedures. It is also worth noting that if the choice is important to properly speak about the incentives properties of our protocols, it is entirely transparent in the protocol itself.

To finish, we also note that if considering a small class of information structures, for example a finite number of them, can eliminate the need to recourse to the randomization procedure, it can still be difficult to identify the right decision problem that allows to retrieve this information structure—the operation we performed in our Alice and Bob example is not one that be generalized easily as we increase the outcome space or the number of periods. The randomization procedure is, in that sense, useful here as well, since it does not require to know which simple decision problem elicits what—all is needed is to have enough of them that together elicit every information structure.

We now detail the proof of our main theorem.

Proof of Theorem 1. The protocol always acts in the expert’s best interest, so the expert is guaranteed to maximize his time- t value at every time t by responding truthfully. The difficulty is to show that, if the expert misreports at time t , he does not maximize his time- t value.

Assume without loss of generality that $[\underline{L}, \overline{U}] = [0, 1]$. We start by defining two particular menus, $\mathbf{0} \in \mathcal{M}_t$, which denotes the menu containing a single security whose value is always 0, independently of the outcome, and similarly $\mathbf{1} \in \mathcal{M}_t$, which denotes the menu containing a single security whose value is always 1.

The proof of the theorem relies on the fact that menus discriminate between experts, in the sense that experts with different information at a particular time will have different preferences over the securities or sub-menus at that time:

Lemma 4. *Let $t \in \{1, \dots, T\}$, and $q', q'' \in \Delta^{T-t+1}(\mathcal{X})$, $q' \neq q''$. If $t = T$, then there exist two securities, S^1, S^2 with values in $(0, 1)$, such that*

$$\begin{aligned} \int_{\mathcal{X}} S^1(x) dq'(x) &> \int_{\mathcal{X}} S^2(x) dq''(x), \\ \int_{\mathcal{X}} S^2(x) dq'(x) &> \int_{\mathcal{X}} S^1(x) dq''(x). \end{aligned}$$

If $t < T$, then there exist two menus $M^1, M^2 \in \mathcal{M}_{t+1}^{(0,1)}$ such that

$$\begin{aligned} \int_{\Delta^{T-t}} \pi_{t+1}(M^1, p) dq'(p) &> \int_{\Delta^{T-t}} \pi_{t+1}(M^2, p) dq''(p), \\ \int_{\Delta^{T-t}} \pi_{t+1}(M^2, p) dq'(p) &> \int_{\Delta^{T-t}} \pi_{t+1}(M^1, p) dq''(p). \end{aligned}$$

Proof. The first part of the statement is immediate. The second part of the statement owes to our separating lemma, Lemma 5 below. The separation lemma asserts existence of a menu $M \in \mathcal{M}_{t+1}$ such that two experts who hold posteriors q' and q'' respectively at time t , and who are offered the menu M at time $t + 1$, expect to make a different payoff from that menu:

$$\int_{\Delta^{T-t}} \pi_{t+1}(M, p) dq'(p) \neq \int_{\Delta^{T-t}} \pi_{t+1}(M, p) dq''(p).$$

With appropriate shifting and scaling, we can choose M in the set $\mathcal{M}_{t+1}^{(0,1)}$. If the left-hand side is larger than the right-hand side, let $M^1 = M$, $M^2 = c\mathbf{1}$, otherwise $M^1 = c\mathbf{1}$, $M^2 = M$, with

$$c = \frac{1}{2} \left[\int_{\Delta^{T-t}} \pi_{t+1}(M, p) dq'(p) + \int_{\Delta^{T-t}} \pi_{t+1}(M, p) dq''(p) \right].$$

■

By a similar argument as in the example of §2, detailed in the proof of Proposition 4, there exist securities S^1, S^2 with values in $(0, 1)$ such that if the individual selects the menu $M_1 = \{\dots \{\dots \{S^1, S^2\} \dots\} \dots\}$ in the preliminary step, the expert does not maximize his time- T value if he misreports at time T —no matter what he reports before time T .

Next, for a given time $t < T$, suppose the expert reports q_t , while he privately observes $p_t \neq q_t$. Lemma 4 asserts existence of menus M^1 and M^2 of $\mathcal{M}_{t+1}^{(0,1)}$ such that if the individual must decide between M^1 and M^2 at time t , it will select M^1 while the expert would have been strictly better off if offered M^2 . Therefore, if the individual selects $M_1 = \{\dots \{\dots \{M^1, M^2\} \dots\} \dots\}$ in the preliminary step, the expert does not maximize his time- t value if he misreports at time t , conditional on the protocol working with M_1 , and independently of the previous reports of the expert.

Now we show that the suboptimality that occurs when the expert lies at time t continues to hold when the individual draws M_1^* in a vicinity of M_1 . More precisely, suppose that when the individual selects $M_1 = \{\dots \{\dots \{M^1, M^2\} \dots\} \dots\}$, where M^1, M^2 are menus of order $T - t$, the expert's time- t value is suboptimal when misreporting, for $t < T$ (if $t = T$, the same argument applies with securities S^1, S^2 instead). By a continuity argument, if M^A and M^B are two menus of order $T - t$ chosen ϵ -close to M^1 and M^2 respectively, for some ϵ small enough, then the expert's time- t value remains suboptimal when lying if the individual selects $M_1' = \{\dots \{\dots \{M^A, M^B\} \dots\} \dots\}$. Now if M_1^* is such that $d(M_1, M_1^*) < \epsilon$ and $M_1^* \in \mathcal{M}_T^{[0,1]}$, then, by an induction argument, every submenu M_t of M_1^* of order $T - t + 1$ is ϵ -close to $\{M^1, M^2\}$. In particular, every submenu of M_t is ϵ -close to either M^1 or M^2 , and further M_t includes at least two submenus respectively closer to M^1 and M^2 . Thus if the individual selects M_1^* in the preliminary step such that M_1^* is ϵ -close to M_1 , the

time- t value to the expert is not maximized when he misreport at time t , no matter the reports at the preceding periods. Because ξ is chosen to have full support, the probability of selecting such a menu M_1^* is positive. Hence overall the time- t value of the expert is less than it would have been had the expert reported the truth at time t .

The remainder of the proof is devoted to the separating lemma:

Lemma 5. *For every $t \in \{1, \dots, T\}$, and every $q', q'' \in \Delta^{T-t+1}(\mathcal{X})$, if*

$$\int_{\Delta^{T-t}} \pi_{t+1}(M, p) dq'(p) = \int_{\Delta^{T-t}} \pi_{t+1}(M, p) dq''(p) \quad \forall M \in \mathcal{M}_{t+1},$$

then $q' = q''$.

Proof. We start with some preliminaries. We endow every set of menus \mathcal{M}_t with an algebraic structure and the following operations:

- Addition: For any $M^A, M^B \in \mathcal{M}_T$, define the menu $M^A + M^B \in \mathcal{M}_T$ by $\{S^A + S^B; S^A \in M^A, S^B \in M^B\}$; if $t < T$ and $M^A, M^B \in \mathcal{M}_t$, define recursively $M^A + M^B = \{m^A + m^B; m^A \in M^A, m^B \in M^B\}$.
- Scalar multiplication: For any $\alpha \geq 0$, for any $M \in \mathcal{M}_T$, define $\alpha M = \{\alpha S; S \in M\}$; if $t < T$, and $M \in \mathcal{M}_t$, define recursively $\alpha M = \{\alpha m; m \in M\}$.

We note that for every $M^A, M^B \in \mathcal{M}_t$, $\alpha \geq 0$, and $p \in \Delta^{T-t+1}(\mathcal{X})$, the following equalities hold:

$$\begin{aligned} \pi_t(\mathbf{1}, p) &= 1, \\ \pi_t(M^A + M^B, p) &= \pi_t(M^A, p) + \pi_t(M^B, p), \\ \pi_t(\alpha M^A, p) &= \alpha \pi_t(M^A, p), \\ \pi_t(M^A \cup M^B, p) &= \max\{\pi_t(M^A, p), \pi_t(M^B, p)\}. \end{aligned}$$

For notational convenience, we let \mathcal{M}_{T+1} be the set of securities, i.e., $\mathcal{M}_{T+1} = \mathcal{C}(\mathcal{X}, \mathbb{R})$, and $\Delta^0(\mathcal{X})$ be the set of outcomes \mathcal{X} . For a security S and an outcome x , we let $\pi_{T+1}(S, x) = S(x)$.

Let \mathcal{B}_t be the set of continuous and bounded real functions on $\Delta^{T-t}(\mathcal{X})$, and let \mathcal{L}_t be the linear span of the set of functions $\{p \mapsto \pi_{t+1}(M, p); M \in \mathcal{M}_{t+1}\} \subset \mathcal{B}_t$. We endow \mathcal{B}_t with the topology of uniform convergence. Recall that every $\Delta^k(\mathcal{X})$ is equipped with the weak-* topology. If a space \mathcal{S} is compact and metrizable, then $\Delta(\mathcal{S})$ endowed with the weak-* topology is compact and metrizable, by the Banach-Alaoglu theorem and the Riesz-Radon representation theorem (for example, Aliprantis and Border, 2006, Theorem 15.11). It follows that every $\Delta^k(\mathcal{X})$ is a compact metrizable space.

The proof proceeds by induction. For $t = T$, $q', q'' \in \Delta(\mathcal{X})$, if the equality

$$\int_{\mathcal{X}} S(x) dq'(x) = \int_{\mathcal{X}} S(x) dq''(x)$$

holds for every security S , then $q' = q''$. Let us assume the lemma's claim is correct for $t + 1 \in \{1, \dots, T\}$. Let $q', q'' \in \Delta^{T-t+1}(\mathcal{X})$ and suppose that, for every $M \in \mathcal{M}_{t+1}$, we have

$$\int_{\Delta^{T-t}(\mathcal{X})} \pi_{t+1}(M, p) dq'(p) = \int_{\Delta^{T-t}(\mathcal{X})} \pi_{t+1}(M, p) dq''(p).$$

Then for every $f \in \mathcal{L}_t$,

$$\int_{\Delta^{T-t}(\mathcal{X})} f(p) dq'(p) = \int_{\Delta^{T-t}(\mathcal{X})} f(p) dq''(p).$$

We have the following two lemmas, proved in Appendix D.

Lemma 6. *The family \mathcal{L}_t contains the function 1, and if $f, g \in \mathcal{L}_t$, then $\alpha f \in \mathcal{L}_t$ for every $\alpha \in \mathbb{R}$, $f + g \in \mathcal{L}_t$, and $\max\{f, g\} \in \mathcal{L}_t$.*

Lemma 7. *The family \mathcal{L}_t is a separating family: If $q', q'' \in \Delta^{T-t}(\mathcal{X})$ and if, for every $f \in \mathcal{L}_t$, $f(q') = f(q'')$, then $q' = q''$.*

In summary, we have that $\Delta^{T-t}(\mathcal{X})$ is a non-empty compact metrizable (hence Hausdorff) space and that \mathcal{L}_t is a subset of the bounded continuous real functions defined on $\Delta^{T-t}(\mathcal{X})$ such that (i) \mathcal{L}_t is a separating family, and (ii) \mathcal{L}_t contains the function 1, and if $f, g \in \mathcal{L}_t$, then $\alpha f \in \mathcal{L}_t$, $f + g \in \mathcal{L}_t$, and $\max\{f, g\} \in \mathcal{L}_t$ (so \mathcal{L}_t forms a Boolean ring). We can then apply the version of the Stone-Weirstrass theorem for Boolean rings described in Theorem 7.29 of Hewitt and Stromberg (1997) which implies that \mathcal{L}_t is dense in \mathcal{B}_t in the topology of uniform convergence. In particular, for every $f \in \mathcal{B}_t$,

$$\int_{\Delta^{T-t}(\mathcal{X})} f(p) dq'(p) = \int_{\Delta^{T-t}(\mathcal{X})} f(p) dq''(p).$$

That $\Delta^{T-t}(\mathcal{X})$ is metrizable implies $q' = q''$ by Aleksandrov's theorem (Aliprantis and Border, 2006, Theorem 15.1). Thus, if the lemma's claim is true for $t + 1$, it is also true for t . ■

This completes the proof of the first main result. ■

As in the Alice and Bob example, the above randomized protocol requires that the administrator has the ability to commit to choosing the best option on behalf of the expert without ever revealing what the options are. However, the deterministic equivalent protocol does not require such commitment ability.

Specifically, if $\Pi(q_1, \dots, q_T, x; M_1) = \Pi^1(q_1, \dots, q_T, x; M_1)$ is the payoff to the expert for a given menu M_1 drawn in the preliminary stage of the protocol, then from the expert's viewpoint the randomized protocol described above is equivalent to the deterministic protocol with payoff rule

$$\Pi(q_1, \dots, q_T, x; \xi) = \int_{\mathcal{M}_T} \Pi(q_1, \dots, q_T, x; M_1) d\xi(M_1), \quad (1)$$

as it inherits the same incentive properties as the original protocol. It is in fact how we constructed the closed-form payoff rule in our Alice and Bob example. The procedure can be applied in other cases as well, once the elements of the model, \mathcal{X} , T , and ξ have been specified. Of course the payoff rule is measurable—and thus yields a valid protocol—via the Fubini-Tonelli's theorem.

The following corollary uses the protocols just described to establish existence of strategyproof protocols for a larger space of outcomes, the spaces \mathcal{X} that are Polish, i.e., separable and completely metrizable spaces. In particular, boundedness is no longer required. These spaces include the common case of \mathbb{R}^n with its usual topology, or the space of continuous real functions in a bounded interval of time (which comes in handy in for the continuous-time version of the model, in which the outcome is gradually revealed over time). These extended protocols take essentially the same form as the ones described above.

Corollary 1. *If \mathcal{X} is simply assumed Polish, then strategyproof protocols exist.*

Proof. By the Urysohn metrization theorem (Aliprantis and Border, 2006, Theorem 3.40), \mathcal{X} can be embedded in the Hilbert cube, denoted \mathcal{X}^* . Using that the Hilbert cube is compact metrizable (and thus also Polish) and Theorem 1, we get that the protocol described by equation (1) is strategyproof when the space of uncertainty is \mathcal{X}^* .

Note that for each k , each of $\Delta^k(\mathcal{X})$ and $\Delta^k(\mathcal{X}^*)$ are Polish, by Aliprantis and Border (2006, Theorem 15.15). Let $\iota_0 : \mathcal{X} \rightarrow \mathcal{X}^*$ be the inclusion map $\iota_0(x) = x$ and note that it is an embedding. Then the map $\iota_1 : \Delta(\mathcal{X}) \rightarrow \Delta(\mathcal{X}^*)$ given by $\iota_1(p) = p \circ \iota_0^{-1}$ is well-defined and is also an embedding, by Aliprantis and Border (2006, Theorem 15.14). Inductively define $\iota_k : \Delta^k(\mathcal{X}) \rightarrow \Delta^k(\mathcal{X}^*)$ by $\iota_k(p) = p \circ \iota_{k-1}^{-1}$ and note that these are well-defined embeddings.

Now, let Π be a deterministic strategyproof rule for the space \mathcal{X}^* as given by equation (1), and let $\Pi_{\mathcal{X}}$ be defined by $\Pi_{\mathcal{X}}(q_1, \dots, q_T, x) = \Pi(\iota_T(q_1), \dots, \iota_1(q_T), \iota_0(x))$, and note that $\Pi_{\mathcal{X}}$ induces the same incentives as Π . ■

Menu-based protocols. It worth observing that every deterministic protocols can be interpreted and implemented directly through the instruments introduced at the beginning of this section. That is, at $t = 1$, the individual offers a (carefully designed) menu of order T to the expert, and observes the expert's selection every time the expert exercises the menu. The expert's choices at every stage reveals (indirectly) information on his prior/posteriors, and eventually delivers the individual with the same information as she would have obtained from the protocol described above.

In this case, however, the menu contains infinitely many elements, unlike the menus used in the above randomized protocol. Such a menu, let us call it $M^*(\Pi)$, can be constructed explicitly from any payoff rule Π , according to a simple recursive procedure. For every p_1, \dots, p_{T-1} with $p_t \in \Delta^{T-t+1}(\mathcal{X})$, we define

$$\begin{aligned} M_t^{p_1, \dots, p_{t-1}} &= \{M_{t+1}^{p_1, \dots, p_{t-1}, q_t}; q_t \in \Delta^{T-t+1}(\mathcal{X})\}, \\ M_T^{p_1, \dots, p_{T-1}} &= \{\Pi(p_1, \dots, p_{T-1}, q_T, \cdot) \in \mathcal{C}(\mathcal{X}, \mathbb{R}); q_T \in \Delta(\mathcal{X})\}, \end{aligned}$$

where $\mathcal{C}(\mathcal{X}, \mathbb{R})$ is the space of continuous maps from outcomes to payoffs, i.e., the space of securities. The menu $M^*(\Pi)$ is then defined as the menu M_1 that was just obtained.

4.2 Uniqueness

In this section we present two characterizations of the strategyproof (direct) protocols. The first is an exact and general characterization which is a simple and useful test to check whether a particular protocol is strategyproof, but can be difficult to use for the design of strategyproof protocols in the presence of hidden information. The second characterization is approximate and shows that under some regularity condition, every protocol can be (approximately) expressed as a randomized protocol of the form described above when initialized with an appropriate menu distribution ξ . Thus, when it comes to the design of strategyproof protocols, there is essentially no loss of generality in limiting oneself to the above class of randomized protocols.

The exact characterization is obtained by using the convexity of the time- t value functions. For a protocol for T periods with payoff rule Π , let $\Pi_t(p_1, \dots, p_t)$ denote the time- t value of the expert who has reported truthfully p_1, \dots, p_t up to time t and plans to report truthfully after time t :

$$\begin{aligned}\Pi_T(p_1, \dots, p_T) &= \int_{\mathcal{X}} \Pi(p_1, \dots, p_T, x) dp_T(x), \\ \Pi_t(p_1, \dots, p_t) &= \int_{\Delta^{T-t}} \Pi_{t+1}(p_1, \dots, p_t, q) dp_t(q).\end{aligned}$$

Proposition 1. *Let Π be the payoff rule of a direct protocol for T periods. Then the protocol is strategyproof if and only if the following conditions are satisfied: For every $t = 1, \dots, T$, and every p_1, \dots, p_{t-1} , the map $G_t(q_t) = \Pi_t(p_1, \dots, p_{t-1}, q_t)$ is strictly convex on $\Delta^{T-t+1}(\mathcal{X})$, and the map $s_t(q_t, q_{t+1}) = \Pi_{t+1}(p_1, \dots, p_{t-1}, q_t, q_{t+1})$ is a subgradient of G at point q_t (if $t = T$, the condition is becomes $s_T(q_T, x) = \Pi(p_1, \dots, p_{T-1}, q_T, x)$ is a subgradient of G at point q_T).*

Proof. When the expert participates in a strategyproof protocol, $G_t(p_t)$ can be thought of as the time- t value written as a function of the expert's true time- t posterior p_t . If we consider only one-step deviations from the truth at time t —in the sense that the expert always tells the truth before and after time t , but possibly not at time t —then the strict convexity of G_t becomes necessary and sufficient to allow a strict best response. Additionally, it must be the case that the time- $(t+1)$ value to the expert (who is truthful from time $t+1$ onwards) is a subgradient of G_t . The arguments are similar to Gneiting and Raftery (2007) and thus omitted. We then observe that for a protocol to be strategyproof, it is necessary and sufficient that it be robust to every one-time deviation. ■

The characterization is useful to check whether a particular protocol is strategyproof, however because the characterization is not constructive, it is of limited practical use when it comes to designing new strategyproof protocols in the presence of intermediate hidden information. Indeed the construction reduces to that of designing strictly convex value functions with strictly convex subgradients in large spaces, a task which is not so simple. However, the class of payment rules issued from the strategyproof randomized protocols of Section 4.1 approximate arbitrarily closely any strategyproof mechanism that is subject to some regularity conditions.

Theorem 2. *Let Π' be a payoff rule issued from a strategyproof protocol for T periods. If the payoff $\Pi'(q_1, \dots, q_T, x)$ is jointly continuous in (q_1, \dots, q_T, x) , then for every $\epsilon > 0$ there exists a*

full-support distribution over finite menus, ξ , such that the inequality

$$|\Pi'(q_1, \dots, q_T, x) - \Pi''(q_1, \dots, q_T, x; \xi)| < \epsilon$$

holds uniformly for all q_1, \dots, q_T, x , in which $\Pi''(\cdot; \xi)$ denotes the payment rule associated with the ξ -randomized protocol of Section 4.1 described by equation (1).

Proof. The proof consists in two steps. First, we approximate the payoff rule Π' by a payoff rule associated with a finite menu. Since finite menus only reveal partial information, in a second step we complement the payoff rule by a small fraction of a strategyproof protocol. The overall payoff rule can be implemented via a randomized mechanism of the form studied above.

The difficulty is to produce a finite menu by sampling the initial payoff rule (which can be regarded as a large, infinite menu) finitely many times in such a way that, whenever a selection needs to be made from that finite menu or one of the sub-menus, the payoffs associated with that choice are always close to those of the original payoff rule.

As before, for any finite menu M of order T , let $\Pi(q_1, \dots, q_T, x; M)$ be the associated payoff rule. Let us take w.l.o.g. $[\underline{L}, \bar{U}] = [0, 1]$. We will need a compatible metric d on the spaces $\Delta^k(\mathcal{X})$, for example the Lévy-Prokhorov metric.

Fix $\epsilon > 0$. Because Π' is continuous on $\Delta^T(\mathcal{X}) \times \dots \times \Delta(\mathcal{X}) \times \mathcal{X}$ which is a compact set, it is uniformly continuous, and there exists $\delta_1 > 0$ such that if, for every i , q'_i is δ_1 -close to q''_i , i.e., $d(q'_i, q''_i) < \delta_1$, then $|\Pi'(q'_1, \dots, q'_T, x) - \Pi'(q''_1, \dots, q''_T, x)| < \epsilon/2$ for every x .

Step 1(a). We show that there exists a finite subset Σ_1 of $\Delta^T(\mathcal{X})$ such that, for every p_1 , if

$$q_1^* \in \arg \max_{q_1 \in \Sigma_1} \Pi'_1(q_1, p_1),$$

then q_1^* is δ_1 -close to p_1 .

Let $\{\Sigma_{1,k}\}$ be a sequence of finite subsets of $\Delta^T(\mathcal{X})$ such that $\Sigma_{1,k} \rightarrow \Delta^T(\mathcal{X})$ in the Hausdorff metric topology induced by the Lévy-Prokhorov metric. (The compactness of $\Delta^T(\mathcal{X})$ guarantees existence of such a sequence.) We observe that $(q_1, p_1) \mapsto \Pi'_1(q_1, p_1)$ is continuous. (This observation is seen directly via induction, using that every Π'_i is uniformly continuous.) The correspondence $(\mathcal{P}, p_1) \mapsto \mathcal{P}$, where \mathcal{P} is a compact subset of $\Delta^T(\mathcal{X})$ and $p_1 \in \Delta^T(\mathcal{X})$ is also continuous (for instance, by Theorem 18.10 of Aliprantis and Border (2006)). Using Berge's maximum theorem, we get that the correspondence

$$(\mathcal{P}, p_1) \mapsto \arg \max_{q_1 \in \mathcal{P}} \Pi'_1(q_1, p_1)$$

is upper hemicontinuous. Now suppose that for every k , there exists (q_1^k, p_1^k) such that

$$q_1^k \in \arg \max_{q_1 \in \Sigma_{1,k}} \Pi'_1(q_1, p_1^k)$$

with $d(q_1^k, p_1^k) \geq \delta_1$. Because $\Delta^T(\mathcal{X})$ is compact, there exists a subsequence $\{\sigma(k)\}$ such that $p_1^{\sigma(k)} \rightarrow p_1^\infty$ for some p_1^∞ . Also, $\Sigma_{1,\sigma(k)} \rightarrow \Delta^T(\mathcal{X})$, where the limit is again in the Hausdorff

metric topology. Noting that $\arg \max_{q_1 \in \Delta^T(\mathcal{X})} \Pi'_1(q_1, p_1) = \{p_1\}$, by the upper hemicontinuity of the argmax correspondence, we get that $q_1^{\sigma(k)} \rightarrow p_1^\infty$, thus contradicting that $d(q_1^{\sigma(k)}, p_1^{\sigma(k)}) \geq \delta_1$ for every k .

Step 1(b). Next we show that there exists k^* such that for every finite menu M of order T that satisfies

$$|\Pi_1(q_1, p_1; M) - \Pi'_1(q_1, p_1)| < 1/k^* \quad \forall q_1, p_1 ,$$

then, for every p_1 , if

$$q_1^* \in \arg \max_{q_1 \in \Sigma_1} \Pi_1(q_1, p_1; M)$$

then q_1^* is δ_1 -close to p_1 .

If the claim is wrong, then for every k there exists p_1^k, q_1^k, M^k such that

$$|\Pi(q_1, p_1; M^k) - \Pi'(q_1, p_1)| < 1/k \quad \forall q_1, p_1 ,$$

while

$$q_1^k \in \arg \max_{q_1 \in \Sigma_1} \Pi(q_1, p_1^k; M^k)$$

and $d(q_1^k, p_1^k) \geq \delta_1$. Using again the compactness of $\Delta^T(\mathcal{X})$, we can create a subsequence of indices, $\{\sigma(k)\}$, such that $p_1^{\sigma(k)} \rightarrow p_1^\infty$ and $q_1^{\sigma(k)} \rightarrow q_1^\infty$ for some $p_1^\infty \in \Delta^T(\mathcal{X})$ and $q_1^\infty \in \Sigma_1$.

Then $d(q_1^\infty, p_1^\infty) \geq \delta_1$ and following Step 1(a), it implies that q_1^∞ is not a maximizer of the map $q_1 \in \Sigma_1 \mapsto \Pi'_1(q_1, p_1^\infty)$. Let $q_1^* \in \Sigma_1$ be such a maximizer, then we have $\Pi'_1(q_1^*, p_1^\infty) > \Pi'_1(q_1^\infty, p_1^\infty)$, and by continuity, for large enough k 's, $\Pi'_1(q_1^*, p_1^k) > \Pi'_1(q_1^\infty, p_1^k)$, with both sides of the inequality bounded away from each other. Thus for large enough k 's, we get $\Pi_1(q_1^*, p_1^k; M^k) > \Pi_1(q_1^\infty, p_1^k; M^k)$, which contradicts the fact that for large enough k 's, q_1^∞ should also maximize $q_1 \in \Sigma_1 \mapsto \Pi_1(q_1, p_1^k; M^k)$, since Σ_1 is finite.

Once again we use uniform continuity and set $\delta > 0$ such that, if for every i , q'_i is δ -close to q''_i , then $|\Pi'(q'_1, \dots, q'_T, x) - \Pi'(q''_1, \dots, q''_T, x)| < 1/k^*$ for every x . Let $\delta_2 = \min\{\delta_1, \delta\}$.

Step 2. Next we iterate the preceding two steps for every $t = 2, \dots, T$. Let $t \geq 2$ and $\delta_t > 0$ be given. Fix p_1, \dots, p_{t-1} such that $p_1 \in \Sigma_1, p_2 \in \Sigma_2^{p_1}, p_3 \in \Sigma_3^{p_1, p_2}$, and so forth, where every set of the form $\Sigma_k^{p_1, \dots, p_{k-1}}$ is a finite subset of $\Delta^{T-k+1}(\mathcal{X})$.

By an argument analogous to Step 1(a), we can define $\Sigma_t^{p_1, \dots, p_{t-1}}$ a finite subset of $\Delta^{T-t+1}(\mathcal{X})$ such that, for every p_t , if

$$q_t^* \in \arg \max_{q_t \in \Sigma_t^{p_1, \dots, p_{t-1}}} \Pi'_t(p_1, \dots, p_{t-1}, q_t, p_t)$$

then q_t^* is δ_t -close to p_t .

By an argument analogous to Step 1(b), there exists k^* such that for every finite menu M of order T that satisfies

$$|\Pi_t(p_1, \dots, p_{t-1}, q_t, p_t; M) - \Pi'_t(p_1, \dots, p_{t-1}, q_t, p_t)| < 1/k^* \quad \forall p_t, q_t ,$$

for every p_t , if

$$q_t^* \in \arg \max_{q_t \in \Sigma_t^{p_1, \dots, p_{t-1}}} \Pi_t(p_1, \dots, p_{t-1}, q_t, p_t; M)$$

then q_t^* is δ_t -close to p_t .

We then let δ to be such that if, for every i , q'_i is δ -close to q''_i , then $|\Pi'(q'_1, \dots, q'_T, x) - \Pi'(q''_1, \dots, q''_T, x)| < 1/k^*$ for every x . Let $\delta_{t+1} = \min\{\delta, \delta_t\}$.

Step 3. We build a finite menu M_1^* of order T by sampling the (continuous) menu associated with Π' as follows: For every p_1, \dots, p_{T-1} where for every t , $p_t \in \Sigma_t^{p_1, \dots, p_{t-1}}$, we define

$$\begin{aligned} M_T^{p_1, \dots, p_{T-1}} &= \{ \Pi(p_1, \dots, p_{t-1}, q_T, \cdot); q_T \in \Sigma_T^{p_1, \dots, p_{T-1}} \} , \\ M_t^{p_1, \dots, p_{t-1}} &= \{ M_{t+1}^{p_1, \dots, p_{t-1}, q_t}; q_t \in \Sigma_t^{p_1, \dots, p_{t-1}} \} . \end{aligned}$$

Then, the menu M_1^* can be defined as $\{M_1^{q_1}; q_1 \in \Sigma_1\}$. Let ξ' be the probability measure that allocates full mass on M_1^* . We note that we have, by Steps 1(a), 1(b), and 2,

$$|\Pi(p_1, \dots, p_T, x; \xi') - \Pi'(p_1, \dots, p_T, x)| < \epsilon/2 \quad \forall p_1, \dots, p_T, x$$

Step 4. Finally, let ξ'' be a probability measure over $\mathcal{M}_T^{[0,1]}$ with full support. Take $\xi = (1 - \epsilon/2)\xi' + \epsilon/2\xi''$. Then, $\Pi(p_1, \dots, p_T, x; \xi)$ defines a strategyproof payoff rule, and

$$|\Pi(p_1, \dots, p_T, x; \xi) - \Pi'(p_1, \dots, p_T, x)| < \epsilon \quad \forall p_1, \dots, p_T, x .$$

■

Our characterization requires some regularity conditions on the protocols, but these are easy to satisfy. The sample distribution presented earlier generates continuous payoff rules, and so do many others. We can get a large class of full-support randomization devices that generate continuous payoff rules by defining a reference distribution on finite menus, and then applying an arbitrary positive density function.

For example, let μ be the distribution over finite menus of order $T - t$ generated by the sample procedure described above. Then any distribution of menus with full support that is absolutely continuous w.r.t. μ yields strategyproof protocols with continuous payoffs. Such distribution is easily generated by applying a positive density function f over menus. Moreover, the argument of the preceding result extends directly to show that these protocols, parameterized by a positive density f , are dense in the set of all continuous and strategyproof protocols.

Proposition 2. *If ξ is a full-support distribution absolutely continuous with respect to μ and if $\Pi(\cdot; \xi)$ is the payoff rule issued from a ξ -randomized protocol described above, then $\Pi(q_1, \dots, q_T, x; \xi)$ is continuous in (q_1, \dots, q_T, x) .*

Proof. For a fixed point (p_1, \dots, p_T) , it can be seen that there exists a unique maximizer at every stage of the protocol, for μ -almost every M_1 , when the reports are $(q_1, \dots, q_T) = (p_1, \dots, p_T)$.

By a simple continuity argument, the same is true when (q_1, \dots, q_T) is in a small neighborhood of (p_1, \dots, p_T) . Hence for a fixed point (p_1, \dots, p_T) , it is the case that $(q_1, \dots, q_T, x) \mapsto \Pi(q_1, \dots, q_T, x; M_0)$ is continuous at every point (p_1, \dots, p_T, x) , for μ -almost every M_1 . Therefore, by the dominated convergence theorem, $(q_1, \dots, q_T, x) \mapsto \Pi(q_1, \dots, q_T, x; \xi)$ is also continuous at point (p_1, \dots, p_T, x) . ■

5 Extension: Continuous time

We extend the elicitation mechanism of Section 4.1 to work in a continuous time framework. The expert now privately observes information dynamically over a unit interval of time before the uncertain outcome of interest to the individual materializes at time $t = 1$. The setup we detail below is fairly general. It includes two common situations of interest.

1. Discrete information arrivals with random times: Here the expert is to receive N different signals over the unit time interval, but as opposed to the main setup of Section 3, the times of arrival can be private and random—both their ex-ante distribution and ex-post realization are private information to the agent. The individual now wants to learn the information that concerns both the signals and the arrival times.
2. Information flow that arrives continuously over time: Here the expert is to receive a continuous signal over time that takes the form of a continuous stochastic processes valued in, say, \mathbb{R}^k . The individual seeks to obtain from the expert his assessment of the signal process distribution (for example, drifts and diffusions terms, if parameterized with Brownian motions) and, at every instant, the up-to-date signal value.

Model. Time is continuous and indexed by $t \in [0, 1]$. During this period of time, the individual has interest in information regarding a future outcome. The outcome continues to take values in a compact metrizable space \mathcal{X} . It materializes publicly at $t = 1$.

At time $t = 0$, the expert learns his information structure. As in the discrete-time version, it includes information on when and how the uncertainty on the outcome gets resolved over time.

Definition 6. An *information structure* is a tuple $(\Omega, \mathbb{F}, \mathbb{P}, X)$ in which:

- Ω is a set of states of the worlds relevant to the agent, assumed separable metrizable.
- $\mathbb{F} = \{\mathcal{F}_t\}_{t \in [0, 1]}$ is a right-continuous filtration.
- \mathbb{P} is a full prior.
- $X : \Omega \mapsto \mathcal{X}$ is a random variable that links states of the world to observable outcomes.

Saying that the information the expert receives over time is given by a right-continuous filtration means that he does not learn anything new in any upcoming infinitesimal length of time. This

is a natural assumption in continuous-time dynamics, and a large body of the literature requires right-continuous filtrations (see, for example, Karatzas and Shreve, 1991, §1.2).

Definition 7. Given a state space Ω and a filtration \mathbb{F} , a *full prior* is a stochastic process¹⁴ $\mathbb{P} : (t, \omega) \mapsto \mathbb{P}_t^\omega$ with values in $\Delta(\Omega)$, and such that

- (1) For all ω and all events E , $t \mapsto \mathbb{P}_t^\omega(E)$ is right-continuous.
- (2) For all t and all events E , $\omega \mapsto \mathbb{P}_t^\omega(E)$ is \mathcal{F}_t -measurable.

The first condition is a technical requirement consistent with the fact that the agent learns information that is right-continuous (and every conditional probability of an event has a right-continuous version, for example, Theorem 1.3.13, Karatzas and Shreve (1991)). The second condition means that the conditional probability at time t is known given the information the expert has access to at time t . We lack a condition that states a consistence of the posteriors over time, this condition will not be needed for our results.

As opposed to the benchmark model, we no longer ask the expert to report probability trees, because in continuous time there is an infinite hierarchy of such trees at every instant—an object which can become very complex. So instead we consider the more natural protocols in which the individual asks the expert to declare his information structure (at time $t = 0$) and, at every instant $t \geq 0$, the updated posterior over states. As before, the expert is rewarded at time $t = 1$ as a function of the data he communicated and the outcome that obtains.

Payoffs are specified by a family of *payoff rules*, one payoff rule for every reported information structure. To encode the flow of posteriors, we let $\mathcal{D}([0, 1], \Delta(\Omega))$ be the space of maps $Q : t \mapsto Q_t$ from time indices to probability measures over states, that are such that for every event E , $Q_t(E)$ is right-continuous in time. It is equipped with the product σ -algebra.

Definition 8. A *payoff rule* corresponding to information structure $(\Omega, \mathbb{F}, \mathbb{P}, X)$ is a measurable map $\Pi^{(\Omega, \mathbb{F}, \mathbb{P}, X)} : \mathcal{D}([0, 1], \Delta(\Omega)) \times \mathcal{X} \mapsto [\underline{L}, \overline{U}]$.

$\Pi^{(\Omega, \mathbb{F}, \mathbb{P}, X)}(\{Q_t\}_{t \in [0, 1]}, x)$ is the payoff to an expert who reports information structure $(\Omega, \mathbb{F}, \mathbb{P}, X)$ at time $t = 0$ and, at every instant t , reports the posterior Q_t on states of the world, while the outcome x materializes.

Definition 9. A *reporting strategy* for an expert whose information structure is $(\Omega^*, \mathbb{F}^*, \mathbb{P}^*, X^*)$ consists in two objects:

- An information structure $(\Omega, \mathbb{F}, \mathbb{P}, X)$ declared at time 0.
- A stochastic process $\mathbb{Q} : [0, 1] \times \Omega^* \mapsto \Delta(\Omega)$, in which $\mathbb{Q}_t^{\omega^*}$ is the posterior declared at time t when the true state is ω^* . We require that, for every event E , the process $(t, \omega^*) \mapsto \mathbb{Q}_t^{\omega^*}(E)$ be measurable in ω^* with respect to \mathcal{F}_t and that it be right continuous in t .

¹⁴The second condition implies by Lemma 12 that the full prior is a well-defined stochastic process.

The expert's *time- t value* at state ω^* is the expected payoff given what he knows at time t . For an expert with a given information structure $(\Omega^*, \mathbb{F}^*, \mathbb{P}^*, X^*)$, who participates in a protocol defined by a family of payment rules Π and plays strategy $\langle (\Omega, \mathbb{F}, \mathbb{P}, X), \mathbb{Q} \rangle$, it is defined as

$$V_t^{\omega^*} = \int \Pi^{(\Omega, \mathbb{F}, \mathbb{P}, X)}(\{f(t', \omega')\}_{t' \in [0, 1]}, X(\omega')) dP_t^{*, \omega^*}(\omega').$$

A given strategy $\langle (\Omega, \mathbb{F}, \mathbb{P}, X), \mathbb{Q} \rangle$ is *optimal* at $t = 0$ and state ω^* when, for every alternative strategy $\langle (\Omega', \mathbb{F}', \mathbb{P}', X'), \mathbb{Q}' \rangle$, the time-0 value in state ω^* for the original strategy is at least as large as the time-0 value for the alternative strategy. A given strategy $\langle (\Omega, \mathbb{F}, \mathbb{P}, X), \mathbb{Q} \rangle$ is *optimal* at time $t > 0$ and state ω^* when for every alternative strategy $\langle (\Omega', \mathbb{F}', \mathbb{P}', X'), \mathbb{Q}' \rangle$ with $(\Omega', \mathbb{F}', \mathbb{P}', X') = (\Omega, \mathbb{F}, \mathbb{P}, X)$ and for all $\tau < t$, $Q_\tau^{\omega^*} = Q'_\tau{}^{\omega^*}$, the time- t value at state ω^* of the original strategy is at least as large as the time- t value for the alternative strategy.

We will continue to assume that the individual who runs the elicitation protocol has interest in the probabilities that can be inferred from the expert's private information. It means that for every time t , the individual cares to learn about posteriors on the outcome, of all (finite) orders. *Induced belief trees* are defined as in the benchmark model:

Definition 10. Given an information structure $(\Omega, \mathbb{F}, \mathbb{P}, X)$, the *induced belief tree of order 0* is the map $\varphi : \Delta(\Omega) \mapsto \Delta(\mathcal{X})$ defined as

$$\varphi(Q) = Q(X).$$

The *induced belief tree of order j with intermediate times $t_1 < \dots < t_j$* , is noted $\varphi_{t_1, \dots, t_j} : \Delta(\Omega) \mapsto \Delta^j(\mathcal{X})$ and defined recursively as

$$\varphi_{t_1, \dots, t_j}(Q) = Q(\varphi_{t_2, \dots, t_j}(P_{t_1}))$$

where P_{t_1} is the random variable of the process \mathbb{P} sampled at time t_1 .

Lemma 1 continues to hold and shows that induced belief tree are well-defined and measurable.

We wish to induce the expert to communicate true *relevant* information as a strict best response. As noted earlier, there will be many different information structures that that will be equally relevant to the individual. The strict best response accounts for the assumption that the individual only cares about the information through the probability trees it induces.

Definition 11. A protocol is *strategyproof* when for all information structures $(\Omega^*, \mathbb{F}^*, \mathbb{P}^*, X^*)$:

- The strategy that consists in declaring the true information structure and send the truly updated posterior at all times and for all states is optimal.
- If a strategy $\langle (\Omega, \mathbb{F}, \mathbb{P}, X), \mathbb{Q} \rangle$ is optimal at a given state ω^* and at all times $t \leq \tau$, then for every t_0, \dots, t_j with $t_0 \leq \tau$,

$$\varphi_{t_1, \dots, t_j}(Q_{t_0}^{\omega^*}) = \varphi_{t_1, \dots, t_j}^*(P_{t_0}^{\omega^*})$$

where $\varphi_{t_1, \dots, t_j}$ is the induced belief tree associated with the information structure $(\Omega, \mathbb{F}, \mathbb{P}, X)$ and $\varphi_{t_1, \dots, t_j}^*$ is the tree associated with $(\Omega^*, \mathbb{F}^*, \mathbb{P}^*, X^*)$.

Temporal menu. We will use of an extended menu instrument that we will refer to as *temporal menus*. These are menus with deadlines.

A *temporal menu of securities* is a pair $\sigma_0 = (M_0, \tau_0)$, where M_0 is a collection of securities and $\tau_0 \in [0, 1]$ is a deadline. A person who owns a temporal menu σ not yet expired must decide, at or before time τ_0 , to get one security among the collection M_0 . When that person makes a decision, we say he or she *exercises the menu*. A temporal menu of securities is called temporal menu of order 1.

Similarly we define a temporal menu of sub-menus of order k as a pair $\sigma_k = (M_k, \tau_k)$, τ_k is the menu's deadline and M_k is a collection of sub-menus of order $k - 1$ whose deadlines are greater than τ_k . A person who owns a temporal menu σ_k must choose to own one temporal sub-menu from M_k at any time $t \leq \tau_k$.

A temporal menu is finite when it includes a finite number of sub-menus, which are in turn finite. We denote by Σ_k the collection of finite temporal menus of order k , and $\Sigma_k^{[a, b]}$ the sub-collection in which the securities of all menus take values in $[a, b]$. For notational convenience, let $\Sigma_0 = \mathcal{C}(\Omega, \mathbb{R})$ (the space of securities, i.e., the continuous maps from Ω to \mathbb{R}) and $\Sigma_0^{[a, b]} = \mathcal{C}(\Omega, [a, b])$ (the securities that take values in $[a, b]$).

For an expected-value maximizer with no discounting whose information structure is $(\Omega, \mathbb{F}, \mathbb{P}, X)$, we denote by $\pi_0(S, Q)$ the value of the security S when his prior/posterior over states is Q :

$$\pi_0(S, Q) = \int_{\Omega} S(X(\omega)) dQ(\omega) .$$

We also define $\pi_k(\sigma_k, Q)$ to be the value, to the same person, of the finite temporal menu $\sigma_k = (M_k, \tau_k)$ of order k whose deadline has not yet passed. Recursively, we have:

$$\begin{aligned} \pi_1(\sigma_1, Q) &= \int \left[\sup_{S \in M_1} \pi_0(S, \mathbb{P}_{\tau_1}^{\omega}) \right] dQ(\omega), \\ \pi_k(\sigma_k, Q) &= \int \left[\sup_{\sigma_{k-1} \in M_k} \pi_{k-1}(\sigma_{k-1}, \mathbb{P}_{\tau_k}^{\omega}) \right] dQ(\omega). \end{aligned}$$

The following lemma asserts that those values are well-behaved:

Lemma 8. *Given an information structure $(\Omega, \mathbb{F}, \mathbb{P}, X)$, every value map $\pi_k : \Sigma_k \times \Delta(\Omega) \mapsto \mathbb{R}$ is well-defined and jointly measurable, for all $k \geq 0$.*

Proof. See Appendix D. ■

A class of strategyproof protocols. We now describe the mechanism. In a preliminary step, the individual who administers the protocol draws a random temporal menu of a random order

according to a simple procedure detailed below.¹⁵ That menu is kept secret from the expert. The individual then asks the expert to provide his information structure and to send an update on his posterior on states at every instant before the outcome realization. Based on the expert's announcements, the individual makes exercise decisions optimally on behalf of the expert, under the assumption that the expert reports truthfully, and without ever revealing her action to the expert. Eventually, at $t = 1$, the expert owns a security issued from the last decision made by the individual, and is paid off accordingly. The exact protocol is detailed below.

Let ξ_K be a full support distribution on \mathbb{N}^* , $\xi_{\tau,k}$ be a full support distribution on $\{t_1, \dots, t_k : 0 \leq t_1 < \dots < t_k < 1\}$, and $\xi_{M,k}$ be a full support distribution on the set of finite menus of order k , for example the simple randomization procedure of Section 4.1.

- (a) Preliminary stage: The individual first draws at random a finite number K from ξ_K . She then draws K deadlines τ_1, \dots, τ_K at random from $\xi_{\tau,K}$, and a finite menu M of order K at random from $\xi_{M,K}$. A temporal menu of order K is then formed by taking all the menus and sub-menus associated with the finite menu M , and respectively associating to each menu and sub-menu order k the deadline τ_k . The resulting temporal menu $\sigma_K^* = (M_K^*, \tau_K^*)$ is never disclosed to the expert.
- (b) The expert's actions: At $t = 0$, the expert communicates an information structure $(\Omega, \mathbb{F}, \mathbb{P}, X)$. Then, at all subsequent times t , the expert communicates a posterior over states, $Q_t \in \Delta(\Omega)$.
- (c) The individual's actions: Every time a deadline is reached, i.e., $t = \tau_k$, the individual privately chooses a temporal sub-menu $\sigma_{k-1}^* = (M_{k-1}^*, \tau_{k-1}^*)$ uniformly at random from M_k^* , among all the sub-menus that are optimal assuming the expert has revealed and will reveal truthful information—i.e., such that under the declared information structure, $\pi_{k-1}(\sigma_{k-1}^*, Q_t) = \max_{\sigma_{k-1} \in M_k^*} \pi_{k-1}(\sigma_{k-1}, Q_t)$.

At the time of the last deadline, $t = \tau_K$, the individual selects for the expert a security S^* from M_1^* (instead of a temporal sub-menu) following a similar procedure, i.e., uniformly at random among all the securities of M_1^* that are optimal for the expert who has been truthful in the past.

The individual keeps all her actions secret until $t = 1$, when the outcome materializes and the expert is offered the security S^* .

We can write explicitly the payoff $\Pi^{(\Omega, \mathbb{F}, \mathbb{P}, X)}(\{Q_t\}_{t \in [0,1]}, x; \sigma^*)$ of the protocol for each particular draw of temporal menu σ^* . For a security S , we let $\Pi^{(\Omega, \mathbb{F}, \mathbb{P}, X)}(\{Q_t\}, x; S) = S(x)$. For a finite

¹⁵The σ -algebra of events of finite temporal menus is defined analogously to that of the finite menus of the main framework. Specifically, the space of finite menus of order 1 is equipped with a metric d where, if $\sigma' = (M', \tau')$ and $\sigma'' = (M'', \tau'')$ are two menus of order 1, $d(\sigma', \sigma'') = d(M', M'') + |\tau' - \tau''|$, with $d(M', M'')$ be the Hausdorff distance between the sets of securities M' and M'' , respectively. Next, in a recursive manner, the space of finite menus of order k is equipped with a metric d where, if $\sigma' = (M', \tau')$ and $\sigma'' = (M'', \tau'')$ are two menus of order k , $d(\sigma', \sigma'') = d(M', M'') + |\tau' - \tau''|$ with $d(M', M'')$ is the Hausdorff distance between the sets of sub-menus of order $k-1$, M' and M'' , respectively. We can then take the Borel σ -algebra induced by the metric. As earlier, we note that the σ -algebra of events does not depend of the particular metric chosen.

temporal menu $\sigma_k = (M_k, \tau_k)$ of order k , we let, recursively,

$$\Pi^{(\Omega, \mathbb{F}, \mathbb{P}, X)}(\{Q_t\}, x; \sigma_k) = \frac{1}{|\mathcal{K}|} \sum_{\sigma_{k-1} \in \mathcal{K}} \Pi^{(\Omega, \mathbb{F}, \mathbb{P}, X)}(\{Q_t\}, x; \sigma_{k-1}), \quad \text{with} \quad \mathcal{K} = \arg \max_{\sigma_{k-1} \in M_k} \pi(\sigma_{k-1}, Q_{\tau_k}).$$

The equivalent deterministic protocol is then defined by the family of payoff rules

$$\Pi^{(\Omega, \mathbb{F}, \mathbb{P}, X)}(\{Q_t\}, x) = \int_{\cup_k \Sigma_k^{[L, \bar{U}]}} \Pi^{(\Omega, \mathbb{F}, \mathbb{P}, X)}(\{Q_t\}, x; \sigma^*) d\xi(\sigma^*).$$

where ξ is the probability measure associated with the randomized device of the protocol's preliminary step.

The following lemma ensures that such the randomized version of the protocol is well-defined. Measurability of the resulting payoff rules follow from the Fubini-Tonelli theorem.

Lemma 9. *Every map $(\{Q_t\}, x, \sigma_k) \in \mathcal{D}([0, 1], \Delta(\Omega)) \times \mathcal{X} \times \Sigma_k^{[L, \bar{U}]} \mapsto \Pi^{(\Omega, \mathbb{F}, \mathbb{P}, X)}(\{Q_t\}, x; \sigma_k) \in \mathbb{R}$ is jointly measurable.*

Proof. See Appendix D. ■

We show the main result of this section, namely that the protocol just described is strategyproof.

Proposition 3. *The elicitation protocol for continuous-time dynamic information is strategyproof.*

Proof. It is clear that, because the protocol always works in the best interest of the expert, reporting the truth is optimal. Thus part (1) of strategyproofness is satisfied. We will show that part (2) is also satisfied. Our proof relies on a density argument applied to the main result Theorem 1, using the fact that induced belief trees satisfy certain regularity conditions.

Let us fix an information structure of the expert and a strategy, following the notation of Definition 11.

We will show that, for every $\{Q_t\} \in \mathcal{D}([0, 1], \Delta(\Omega))$, the map $(t_0, \dots, t_j) \mapsto \varphi_{t_1, \dots, t_j}(Q_{t_0})$ is right-continuous in the weak-* topology of $\Delta^j(\mathcal{X})$ separately in each variable.

We proceed by induction. For $j = 0$, and every $\{Q_t\} \in \mathcal{D}([0, 1], \Delta(\Omega))$, the map $t_0 \mapsto \varphi(Q_{t_0}) = Q_{t_0}(X)$ is right-continuous, because for every event E of Ω , by assumption $t_0 \mapsto Q_{t_0}(E)$ is right-continuous. Now fix j and suppose that for every $\{Q_t\} \in \mathcal{D}([0, 1], \Delta(\Omega))$, the map $(t_0, \dots, t_j) \mapsto \varphi_{t_1, \dots, t_j}(Q_{t_0})$ is separately right-continuous. We have that $\varphi_{t_1, \dots, t_j}(Q_{t_0}) = Q_{t_0}(\varphi_{t_2, \dots, t_j}(P_{t_1}))$. The right-continuity assumption on $t_0 \mapsto Q_{t_0}(E)$ for every event E ensures again the right-continuity of $t_0 \mapsto Q_{t_0}(\varphi_{t_2, \dots, t_j}(P_{t_1}))$.

Now, let $f : \Delta^j(\mathcal{X}) \mapsto \mathbb{R}$ be a (bounded) continuous function, with respect to the weak-* topology. Saying that the map $t_i \mapsto Q_{t_0}(\varphi_{t_2, \dots, t_j}(P_{t_1}))$ is right-continuous is saying that the map

$$t_i \mapsto \int_{\Delta^j(\mathcal{X})} f(q) dQ_{t_0}(\varphi_{t_2, \dots, t_j}(P_{t_1}) = q)$$

is right-continuous, for every such f . Note that

$$\int_{\Delta^j(\mathcal{X})} f(q) dQ_{t_0}(\varphi_{t_2, \dots, t_j}(\mathbf{P}_{t_1}) = q) = \int_{\Omega} f(\varphi_{t_2, \dots, t_j}(\mathbf{P}_{t_1}^\omega)) dQ_{t_0}(\omega).$$

By the induction hypothesis, $t_i \mapsto \varphi_{t_2, \dots, t_j}(\mathbf{P}_{t_1}^\omega)$ is separately right-continuous, for every ω . The dominated convergence theorem then yields the right-continuity of

$$t_i \mapsto \int_{\Omega} f(\varphi_{t_2, \dots, t_j}(\mathbf{P}_{t_1}^\omega)) dQ_{t_0}(\omega).$$

We conclude that, for every $\{Q_t\} \in \mathcal{D}([0, 1], \Delta(\Omega))$ and every j , the map $(t_0, \dots, t_j) \mapsto \varphi_{t_1, \dots, t_j}(Q_{t_0})$ is separately right-continuous.

Now, consider a strategy that is optimal at state ω^* and time t_0 . Since the randomization device uses the full-support distribution, it means that, for every j , there exists a set of times \mathcal{T} dense in $\{t_0, \dots, t_j : 0 \leq t_0 < \dots < t_j \leq 1\}$ such that, for every $\tau = (\tau_0, \dots, \tau_j) \in \mathcal{T}$, the expected payoff from the protocol that randomizes over menus of \mathcal{M}_j according to ξ_M and uses τ as exercise times is optimal at t_0 and for state ω^* . According to Theorem 1, this means that for every τ , the posterior of order j , formed at time t_0 , with intermediate times t_1, \dots, t_j , is the same under both the truthful strategy and the alternative strategy:

$$\varphi_{\tau_1, \dots, \tau_j}(\mathbf{Q}_{t_0}^{\omega^*}) = \varphi_{\tau_1, \dots, \tau_j}(\mathbf{P}_{t_0}^{\omega^*}).$$

By density of \mathcal{T} and right-continuity of the inference maps with respect to times, we get that, for every $t_0 < \dots < t_j$,

$$\varphi_{\tau_1, \dots, \tau_j}(\mathbf{Q}_{t_0}^{\omega^*}) = \varphi_{\tau_1, \dots, \tau_j}(\mathbf{P}_{t_0}^{\omega^*}).$$

■

6 Related literature

The literature on scoring rules is vast; Gneiting and Raftery (2007) provide a useful recent summary. Savage (1971) formalizes many of the ideas in the literature, in particular he discusses at length the complementarity issue. The mechanism whereby a choice set is chosen randomly was first suggested by Allais (1953) and W. Allen Wallis (see Savage, 1954). Azrieli et al. (2012a,b) investigate this mechanism in an abstract framework and establish that its incentive properties hold under quite broad specifications of behavior under uncertainty. In particular, there it is explained that two clever approaches to eliciting valuations, Becker et al. (1964) and Karni (2009), can be viewed as applications of Allais' idea. Matheson and Winkler (1976) is an explicit application of this idea to the literature on scoring rules. Related contributions to the scoring rules literature in economics include Thomson (1979), Lambert (2013), Fang et al. (2010), Ostrovsky (2012), and Hossain and Okui (2013). Scoring rules are used in the experimental literature to elicit beliefs; e.g., Nyarko and

Schotter (2002).

Following the decision theory literature, our paper can be viewed in a subjectivist light (Savage, 1972). Under this interpretation, the expert’s beliefs are not objective and are merely a subjective part of his preference. Likewise, his understanding of how information evolves at each time would also be subjective. The recent work by Dillenberger et al. (2012) provides behavioral foundations over menu choice for a decision maker who behaves as if there is an intermediate signal realization before the choice of an act must take place; formally, this object is called an *experiment* by Blackwell (1953). In this interpretation, the signal realization and probability measure are all subjectively perceived by the decision maker *ex-ante*. Our mechanism then provides a way for the individual to elicit subjective beliefs from multiple experts; she may then aggregate these as she sees fit; classical methods are described in Genest and Zidek (1986) and Clemen and Winkler (1999).

We wish to distinguish our work from the literature on dynamic mechanism design, the goal of which is usually to use transfers to enforce a specific outcome or property (such as Pareto efficiency) through weak incentives. In our environment, the goal is to use strict incentives in order to get the expert to perfectly reveal his beliefs, however, there is no outcome (other than transfers) which is relevant to the expert. Hence, we could always obtain weak incentive compatibility by offering no transfers at all. By asking for a protocol that makes truth-telling a strict best response, we are asking for a form of protocol that is “robust”. Of course, the protocols that we devise can be scaled up or down; a protocol that is sufficiently “scaled up” will be useful when there are conflicts of interests between the individual who needs the information and the expert who possesses it or can acquire it. However, in this work, it is the nature of the protocol that we feel is of economic interest, rather than its “scale”.

7 Conclusion

In this work, we have considered a dynamic analogue of the classical scoring rule problem. We recursively construct menus of menus of . . . of menus of outcome-contingent payoffs from which an agent is allowed to choose. We show that such a method can completely elicit the dynamic structure of an agent’s belief or information as it is revealed to her. To establish the existence of such objects, we develop a new constructive approach based on randomly selecting among finite menus. The construction applies to the discrete time and continuous time cases equally well.

The technique we develop is not unique to elicitation of probabilities, but can be employed to elicit a distribution of linear characteristics in more general environments. For example, consider an environment where workers are parameterized by a scalar $\theta \in [0, 1]$ (a cost of effort, say). The utility of a worker of type θ from working x hours and receiving compensation T is $u_\theta(x, T) = T - \theta x$. Constructing a payoff rule or contract $(x(\theta), T(\theta))_\theta$ which allows a firm to completely elicit θ is a standard convex analysis problem, isomorphic to constructing a scoring rule. Now, suppose the firm wants to elicit the *distribution* of worker types in the economy. Suppose the firm negotiates with a union that knows the distribution, $\mu \in \Delta([0, 1])$. The union evaluates contracts by a utilitarian

criterion, so that the utility of incentive compatible contract $(x(\theta), T(\theta))_\theta$ is given by

$$\int_{[0,1]} T(\theta) - \theta x(\theta) d\mu(\theta).$$

Our technique can be employed to show that the firm can, by offering the union to choose from a *menu* of contracts, completely elicit the distribution of worker types.

We wish to emphasize that we have set ourselves up for the most difficult possible version of the problem: the individual sees nothing along the way. If she can observe some of the information that the expert observes, it only makes it easier for her to solve the incentive problem. For example, let us consider a simple case in which there are two possible outcomes, and an intermediate signal realization which is observable both to the individual and to the expert. The individual can simply use a classical scoring rule to elicit the joint beliefs of the expert over the outcome and signal realization. Upon observing the signal realization, the individual can then form her own updated belief; and in particular, the joint belief of the expert can be used to construct a probability over probabilities. Intermediate cases in which the individual can observe some of the information which the expert can observe can be similarly studied; the point is the individual does not *need* to condition her payoff on the information which is observable in the interim: she can use our menu-based mechanisms to condition her payoff *only* on the observed outcome and still fully retain strict incentive compatibility.

References

- C.D. Aliprantis and K.C. Border. *Infinite dimensional analysis: A hitchhiker's guide*. Springer, third edition, 2006.
- M. Allais. Le comportement de l'homme rationnel devant le risque: Critique des postulats et axiomes de l'école américaine. *Econometrica*, 21(4):503–546, 1953.
- F. Allen. Discovering personal probabilities when utility functions are unknown. *Management Science*, 33(4):542–544, 1987.
- F.J. Anscombe and R.J. Aumann. A definition of subjective probability. *The annals of mathematical statistics*, 34(1):199–205, 1963.
- Y. Azrieli, C.P. Chambers, and P.J. Healy. Incentives in experiments: A theoretical analysis. Working paper, 2012a.
- Y. Azrieli, C.P. Chambers, and P.J. Healy. Incentives in experiments with objective lotteries. Working paper, 2012b.
- G.M. Becker, M.H. DeGroot, and J. Marschak. Measuring utility by a single-response sequential method. *Behavioral science*, 9(3):226–232, 1964.

- P. Berti and P. Rigo. On the existence of inferences which are consistent with a given model. *The Annals of Statistics*, 24(3):1235–1249, 1996.
- P. Billingsley. *Probability and measure*. John Wiley & Sons, anniversary edition, 2012.
- D. Blackwell. Equivalent comparisons of experiments. *Annals of Mathematical Statistics*, 24(2): 265–272, 1953.
- R.T. Clemen and R.L. Winkler. Combining probability distributions from experts in risk analysis. *Risk Analysis*, 19(2):187–203, 1999.
- D. Dillenberger, J.S. Lleras, P. Sadowski, and N. Takeoka. A theory of subjective learning. Working paper, 2012.
- A.K. Dixit and R.S. Pindyck. *Investment under uncertainty*. Princeton University Press, 1994.
- F. Fang, M.B. Stinchcombe, and A.B. Whinston. Proper scoring rules with arbitrary value functions. *Journal of Mathematical Economics*, 46(6):1200–1210, 2010.
- C. Genest and J.V. Zidek. Combining probability distributions: A critique and an annotated bibliography. *Statistical Science*, 1(1):114–135, 1986.
- T. Gneiting and A.E. Raftery. Strictly proper scoring rules, prediction, and estimation. *Journal of the American Statistical Association*, 102(477):359–378, 2007.
- E. Hewitt and K.R. Stromberg. *Real and abstract analysis*, volume 25 of *Graduate Texts in Mathematics*. Springer-Verlag, 1997.
- T. Hossain and R. Okui. The binarized scoring rule. *Review of Economic Studies*, 80(3):984–1001, 2013.
- I. Karatzas and S.E. Shreve. *Brownian motion and stochastic calculus*. Springer, second edition, 1991.
- E. Karni. A mechanism for eliciting probabilities. *Econometrica*, 77(2):603–606, 2009.
- A.S. Kechris. *Classical descriptive set theory*. Springer-Verlag, 1995.
- D.M. Kreps and E.L. Porteus. Temporal resolution of uncertainty and dynamic choice theory. *Econometrica*, 46(1):185–200, 1978.
- N.S. Lambert. Elicitation and evaluation of statistical forecasts. working paper, 2013.
- J. Lu. Random choice and private information. Working paper, 2013.
- J.E. Matheson and R.L. Winkler. Scoring rules for continuous probability distributions. *Management Science*, 22(10):1087–1096, 1976.

- Y. Nyarko and A. Schotter. An experimental study of belief learning using elicited beliefs. *Econometrica*, 70(3):971–1005, 2002.
- M. Ostrovsky. Information aggregation in dynamic markets with strategic traders. *Econometrica*, 80(6):2595–2647, 2012.
- L.J. Savage. *The foundations of statistics*. Dover, 1954.
- L.J. Savage. Elicitation of personal probabilities and expectations. *Journal of the American Statistical Association*, 66(336):783–801, 1971.
- L.J. Savage. *The foundations of statistics*. Dover Publications, 1972.
- W. Thomson. Eliciting production possibilities from a well-informed manager. *Journal of Economic Theory*, 20(3):360–380, 1979.
- H. Wold. Ordinal preferences or cardinal utility? *Econometrica*, 20(4):661–664, 1952.

A Appendix: Protocols for one-time hidden information

This section extends the Alice and Bob example to the case of any finite number of outcomes.

Using the notation of our Section 2, let us assume that Alice now observes the outcome of a categorical random variable X that takes values in $\{1, \dots, n\}$. The timing is the same as in Section 2, but now, at $t = 2$, Bob forms a posterior belief about X which is summarized by a vector of probabilities $p \in \mathbb{R}^n$. At $t = 1$, Bob forms a prior over the posteriors he anticipates to observe at $t = 2$, now summarized by a multi-dimensional cumulative distribution function F on \mathbb{R}^n . (At $t = 3$, the outcome of the random variable is publicly observed, as before.)

We will continue to use menus and securities as defined in our baseline Alice and Bob example. Let us suppose payoffs from securities take values in some bounded interval $[\underline{L}, \bar{U}]$. A security can be regarded as a vector of $[\underline{L}, \bar{U}]^n$ —a vector of payoffs, one payoff for each possible outcome. Let ξ_S be a probability distribution over $[\underline{L}, \bar{U}]^n$, with full support, and ξ_C be a probability distribution over $[\underline{L}, \bar{U}]$, also with full support—for instance, the uniform distributions.

Consider the following protocol administered by Alice:

- As a preliminary stage, Alice draws two securities S_1, S_2 independently at random according to ξ_S . Let $M = \{S_1, S_2\}$ the menu of these two securities. Alice also draws a fixed payoff C independently at random according to ξ_C . The menu and payoff are kept secret from Bob.
- At $t = 1$, Alice asks Bob to announce his prior over anticipated posteriors. If Bob reports the cumulative distribution \hat{F} , Alice chooses to offer Bob either the menu M or the constant payoff C which maximizes Bob’s present expected value, assuming truthful reports, ties broken in favor of the constant payoff, for example. Whatever choice Alice makes is kept hidden from Bob.

- At $t = 2$, Alice asks Bob to announce his observed posterior. Bob reports a vector of probabilities \hat{p} . If Alice had selected the constant payoff at the previous stage, nothing happens. Otherwise, Alice gives Bob a security S from M that maximizes Bob's expected value at that time, assuming truthful reports, ties broken in favor of S_1 , for example.
- At $t = 3$, the outcome of the random variable, x , is observed, and Bob receives his due payoffs.

For any given choice of securities S_1, S_2 , and constant payoff C , the payoff to Bob is concisely written as

$$\Pi(\hat{F}, \hat{p}, x; S_1, S_2, C) = \begin{cases} C & \text{if } C \geq \int \max \{ \sum_x S_1(x)p(x), \sum_x S_2(x)p(x) \} d\hat{F}(p), \\ S_1(x) & \text{if } C < \int \max \{ \sum_x S_1(x)p(x), \sum_x S_2(x)p(x) \} d\hat{F}(p), \\ & \text{and } \sum_x S_1(x)\hat{p}(x) \geq \sum_x S_2(x)\hat{p}(x), \\ S_2(x) & \text{otherwise.} \end{cases}$$

As in the baseline Alice and Bob example, the protocol relies on Alice's ability to commit to making the best decisions on Bob's behalf. This problem can be avoided by looking at the expected payoff that the protocol generates over all instances of the randomly chosen menus and constant payoff. This expected payoff is

$$\Pi(\hat{F}, \hat{p}, x) = \int_{[\underline{L}, \bar{U}]} \int_{[\underline{L}, \bar{U}]^n} \int_{[\underline{L}, \bar{U}]^n} \Pi(\hat{F}, \hat{p}, x; S_1, S_2, C) d\xi_S(S_1) d\xi_S(S_2) d\xi_C(C).$$

Proposition 4. *The protocol administered by Alice induces truthful reports a strict best response at each time and given Bob's information at that time, no matter what Bob reports in the previous round:*

1. At $t = 1$, for all pairs $F \neq \hat{F}$,

$$\int_{\mathbb{R}^n} \Pi(F, p) dF(p) > \int_{\mathbb{R}^n} \Pi(\hat{F}, p) dF(p),$$

where $\Pi(F, p)$ is Bob's value function at time 2 after he has reported F and observes p , i.e., $\Pi(F, p) = \sup_{\hat{p}} \sum_x \Pi(F, \hat{p}, x)p(x)$.

2. At $t = 2$, for all pairs $p \neq \hat{p}$, and all F ,

$$\sum_x \Pi(F, p, x)p(x) > \sum_x \Pi(F, \hat{p}, x)p(x).$$

Proof. As Alice always acts in Bob's best interest, it never hurts Bob to tell the truth. What must be shown is that, if Bob misreports at any time, then he does not maximize his expected payoff given the private information he possesses at that time. Without loss of generality, we assume $[\underline{L}, \bar{U}] = [-1, 1]$.

Step 1: We first prove the first incentive property: That if the report \hat{F} differs from the true prior F , then Bob does not maximize his time-1 value. We note that, no matter what he chooses to report, given a choice of \hat{F} , the time-2 expected payoff is never greater than when he reveals his information truthfully, since Alice acts in his best interest. Thus,

$$\Pi(F, p) = \sup_{\hat{p}} \sum_x \Pi(F, \hat{p}, x) p(x) = \sum_x \Pi(F, p, x) p(x).$$

Thus throughout this first step, we can safely assume that Bob reports truthfully at $t = 2$.

For a given menu of securities $M = \{S_1, S_2\}$, let $\pi(M, p)$ be the expected payoff obtained by a person who is offered menu M and assess p as outcome probabilities:

$$\pi(M, p) = \max \left\{ \sum_x S_1(x) p(x), \sum_x S_2(x) p(x) \right\}.$$

The following lemma is proved via the Cramér-Wold theorem (Billingsley, 2012, Theorem 29.4), which states that a probability distribution on a vector of n random variables is uniquely defined by the marginal distributions of its linear combinations. According to the lemma, two different types of Bob who own different beliefs about the hidden information they expect to receive at $t = 2$ will value an option to exercise a security differently.

Lemma 10. *If F, G are two cumulative distribution functions over the probability vectors of outcomes (i.e., the simplex of \mathbb{R}^n), and if $F \neq G$ then there exists an affine functional ϕ on the simplex such that*

$$\int_{\mathbb{R}^n} \max\{\phi(p), 0\} dF(p) \neq \int_{\mathbb{R}^n} \max\{\phi(p), 0\} dG(p). \quad (2)$$

Proof. The Cramér-Wold theorem implies that a probability measure on an Euclidian space is entirely determined by the probabilities it assigns to its half-spaces. Hence if $F \neq G$ there exists an affine functional ψ positive on the simplex and for which the probability of $\{\psi \leq a\}$ is different under F and under G , for some $a > 0$. Let F_ψ be the cumulative distribution of ψ (interpreted as a random variable on the simplex) under F , and let G_ψ be the cumulative distribution of ψ under G . Integrating by parts yields

$$\int_0^b F_\psi(z) dz = \int_0^b (b - z) dF_\psi(z) = \int_{\mathbb{R}^n} \max\{b - \psi(p), 0\} dF(p)$$

for every $b > 0$, with the analog equality for G_ψ and G . Since $F_\psi(a) \neq G_\psi(a)$, there exists c such that

$$\int_0^c F_\psi(z) dz \neq \int_0^c G_\psi(z) dz$$

implying that (2) is satisfied with $\phi = c - \psi$. ■

Consider Φ as in Lemma 10 (with $G = \hat{F}$), for example, $\int \max(\Phi(p), 0) dF(p) > \int \max(\Phi(p), 0) d\hat{F}(p)$. For the opposite inequality a symmetric argument applies. Rescaling Φ if necessary, we can choose

Φ to take values in $(-1, 1)$ on the simplex. We can then define the security $S \in (-1, 1)^n$ to be such that, for every p ,

$$\Phi(p) = \sum_x S(x)p(x). \quad (3)$$

Define the outcome-independent payoff

$$C = \frac{1}{2} \left[\int_{\mathbb{R}^n} \max(\Phi(p), 0) dF(p) + \int_{\mathbb{R}^n} \max(\Phi(p), 0) d\hat{F}(p) \right] \in [0, 1)$$

and consider the menu M of the two securities $S^1 = S$, $S^2 = 0$ (degenerate security that yields zero payoff).

Let us assume Alice has drawn this menu M and this outcome-independent payment C in the preliminary step. If she picks menu M instead of fixed payoff C at time 1 when Bob reports his prior, then the time-1 expected payoff to Bob is $\int_{\mathbb{R}^n} \pi(M, p) dF(p)$ —given Alice’s choice. The way we defined C yields the two inequalities

$$\begin{aligned} \int_{\mathbb{R}^n} \pi(M, p) dF(p) &> C, \\ \int_{\mathbb{R}^n} \pi(M, p) d\hat{F}(p) &< C. \end{aligned}$$

Therefore Alice, acting on behalf of Bob, selects the constant payoff C whereas Bob would have been strictly better off with menu M .

Bob’s ex-ante expected payoff is suboptimal for this particular choice of securities S_1, S_2 and outcome-independent payoff C . By continuity the ex-ante expected payoff remains suboptimal when Alice draws the securities and constant payoff in a small neighborhood of S_1, S_2, C , respectively (viewed as elements of $\mathbb{R}^n, \mathbb{R}^n$, and \mathbb{R}). Because the distributions ξ_S and ξ_C have full support, the event that the selected securities fall in that neighborhood occurs with positive probability, and announcing \hat{F} yields a lesser time-1 expected payoff overall.

Step 2: Next we show that, if Bob misreports his observed posterior, then his time-2 expected payoff is not maximized.

Suppose Bob reports \hat{p} at time 2 while he truly observes p . Let Φ be an affine functional on the simplex of \mathbb{R}^n taking values in $(0, 1)$ and such that $\Phi(p) > \Phi(\hat{p})$. Define S_1 to be the security associated with Φ in the sense of equation (3). Define S_2 to be the security with outcome-independent value $(\Phi(p) + \Phi(\hat{p}))/2$. If offered the choice between securities S_1 and S_2 at time 2, after observing his private information, Bob strictly prefers S_1 to S_2 . However because he reports \hat{p} , Alice selects security S^2 on his behalf. Hence for this particular menu of securities, Bob’s expected payoff at time 2 is suboptimal. By continuity, Bob’s expected payoff remains suboptimal for a menu M of any two securities in a small neighborhood of S_1 and S_2 (viewed as vectors of \mathbb{R}^n). In addition if the outcome-independent payoff C drawn is negative—an event of positive probability—then Alice always selects menu M instead of C , no matter what Bob declares at time 1. Therefore, Bob’s expected payoff at time 2 is suboptimal when he reports a posterior \hat{p} while observing $p \neq \hat{p}$. ■

B Appendix: Impossibility result for stage-separated protocols

We will use the discrete-time framework established in Section 3. This section examines stage-separated protocols, which form a subclass of (direct) protocols. As in direct protocols, the expert is asked to provide, at every time $t < T$, a distribution on the future posterior (“probability trees”) at time $t + 1$ and, at time $t = T$, his posterior on the outcome. However the expert now receives a payoff at every time $t = 2, \dots, T + 1$. At $t \leq T$ the payoff is as a function of the posterior distribution forecasted at time $t - 1$ and the posterior realization declared at time t . At $t = T + 1$ the final payoff is made as a function of the outcome realization and the declared distribution on the outcome realization at $t = T$.

Why look at stage-separated protocols? The literature on proper scoring rules demonstrates that we can induce as a strict best response the truthful disclosure of a private posterior belief over a public random signal. The difficulty in introducing hidden information is that, now, the protocol must rely on the expert’s honest declaration of his time- t posterior to properly induce truthful reporting of his beliefs at time $t - 1$. But we know how to induce truthful reporting of the posterior on the outcome (at $t = T$). So, if the costs of lying about a posterior at $t = T$ exceeds the benefits of misreporting at $t = T - 1$, we can also induce truthful reports at $t = T - 1$. Then, if the cost of lying at $t = T - 1$ exceeds the benefits of misreporting at $t = T - 2$, we can induce truthful reports at $t = T - 2$, and so forth. Such a protocol—essentially a number of strictly proper scoring rules assembled back-to-back—is what we call a stage-separated protocol.

Such a construction turns out not to be possible. It is enough to show impossibility for the 3-period setting (single intermediate hidden information) and the finite space $\mathcal{X} = \{1, \dots, n\}$. The remainder of this section will focus on this case. A stage-separated protocol then asks the expert to disclose his prior over posteriors to be formed at $t = 2$, then once the expert forms his posterior on the outcome, he is asked to report his posterior. At this first stage the expert gets his first payoff, which may only depend on the declared prior and posterior. Then, in the second stage, the realized outcome is known and the expert is offered a second payoff, which may only depend on the declared posterior and the realized outcome. The total payoff to the expert is the sum of these two rewards.

To build intuition of the impossibility result, we first focus on protocols that satisfy some regularities and that, in absence of a first stage, would induce the expert to report truthfully his posterior. Our formal result applies to arbitrary stage-separated protocols.

Let $\Pi_I(q_1, q_2)$ be the expert’s reward from the first stage, if he declares a prior q_1 over the posteriors he expects to observe at $t = 2$, and declares a posterior realization q_2 at $t = 2$. Let $\Pi_{II}(q_2, x)$ be the expert’s reward from the second stage, given a realized outcome x and the posterior declaration q_2 .

We focus the expert’s decision at time 2. Assume the expert has reported his true prior p_1 at time 1 and that his true posterior, given a particular signal realization, is p_2 . Suppose he reports $p_2 + \Delta p_2$ instead of p_2 , and let us see how much he gains (or loses) from this deviation. The expected

gain, given his information at stage 2, is

$$\Pi_I(p_1, p_2 + \Delta p_2) - \Pi_I(p_1, p_2) + \Pi_{II}(p_2 + \Delta p_2, p_2) - \Pi_{II}(p_2, p_2)$$

where $\Pi_{II}(q_2, p_2)$ is the expert's expected payoff from stage 2 when he reports q_2 . Because $\Pi_{II}(q_2, p_2)$ is maximized when $q_2 = p_2$, we expect the second term $\Pi_{II}(p_2 + \Delta p_2, p_2) - \Pi_{II}(p_2, p_2)$ to be of order at most $\|\Delta p_2\|^2$, under enough regularity conditions. However, unless $\Pi_I(p_1, q_2)$ is constant in q_2 , we also expect the first term $\Pi_I(p_1, p_2 + \Delta p_2) - \Pi_I(p_1, p_2)$ to be of order $\|\Delta p_2\|$, for at least some instances of p_2 . Thus there are situations in which the gains realized from the first stage when deviating from the truth at $t = 2$ exceed the losses incurred at the second stage. The following proposition formalizes the claim in its generality without assuming any regularity condition.

Proposition 5. *There does not exist a stage-separated protocol that is strategyproof.*

Proof. Consider an arbitrary stage-separated protocol. For every declared prior q_1 , let $g_{q_1}(q_2, x)$ be the total payoff to the expert as a function of the announced posterior q_2 and outcome realization x , that is,

$$g_{q_1}(q_2, x) = \Pi_I(q_1, q_2) + \Pi_{II}(q_2, x).$$

Suppose that $g_{q_1}(p_2, p_2) > g_{q_1}(q_2, p_2)$ for every $q_2 \neq p_2$, where $g_{q_1}(q_2, p_2)$ is the expert's total expected payoff given his realized posterior p_2 —this would be required of any strategyproof protocol. Let \bar{g}_{q_1} be the map on $\Delta(\mathcal{X})$ defined by $\bar{g}_{q_1}(q_2) = g_{q_1}(q_2, q_2)$. Note that \bar{g}_{q_1} is convex, so the preceding inequality can be interpreted saying that the map $x \mapsto \Pi_I(q_1, q_2) + \Pi_{II}(q_2, x)$ is a subgradient of \bar{g}_{q_1} at point q_2 . Because the domain of \bar{g}_{q_1} is the simplex, the map $x \mapsto \Pi_{II}(q_2, x)$ is also a subgradient. Thus the convex functions \bar{g}_{q_1} share the same subgradients. In particular, for every $p', p'' \in \Delta(\mathcal{X})$,

$$\bar{g}_{q_1}(p'') - \bar{g}_{q_1}(p') = \int_0^1 (p'' - p') \cdot \Pi_{II}(\alpha p'' + (1 - \alpha)p', \cdot) d\alpha$$

where $p \cdot q$ is the dot product between p and q on the simplex $\Delta(\mathcal{X})$ interpreted as a subset of \mathbb{R}^n . Thus for all p_1, q_1 , we get that $\bar{g}_{p_1} - \bar{g}_{q_1}$ is constant: At time 1, the expert is best off reporting any q_1 that maximizes $\bar{g}_{q_1}(q_2)$, for an arbitrary q_2 , independently of his true prior. This makes it impossible for the protocol to be strategyproof. ■

C Appendix: Information sufficiency in multi-period decision problems

We will use the framework and notation of the baseline model of Section 3. There are $T + 1$ periods. At time $T + 1$ an uncertain outcome from the compact metrizable set \mathcal{X} materializes publicly. There is an expert whose information structure is $(\Omega, \mathbb{F}, \mathbb{P}, X)$. There is also a less informed utility-maximizing individual, who faces a dynamic decision problem: At every period $t \leq T$, the individual must choose an action a_t from a collection of possible actions \mathcal{A}_t , assumed compact

metrizable. At the last period $T + 1$, the individual receives utility $u(a_1, \dots, a_T, x)$, where x is the outcome that materializes. The individual's utility function, $u : \mathcal{A}_1 \times \dots \times \mathcal{A}_T \times \mathcal{X} \mapsto \mathbb{R}$, is bounded and jointly continuous.

We will show that, if the individual delegates the decisions to the expert, she gets as much as if the expert communicates to her the “probability tree” induced by his information structure and subsequent private observations, in the sense of Section 3.

We begin by showing that, if the expert gets to choose the actions of this decision problem, he could always choose information-contingent actions that maximizes the individual's expected utility. In other words, there always exists a solution to the individual's multi-period decision problem. A decision policy for the expert is summarized by a tuple $(\alpha_1, \dots, \alpha_T)$ where every α_t is an \mathcal{F}_t -measurable map from Ω to \mathcal{A}_t .¹⁶ Denote by $\mathcal{D}(\mathbb{F})$ the set of all decision policies available to the expert. The following lemma asserts that an optimal decision policy always exists, and yields an expected utility that can be computed via a dynamic programming principle.

Lemma 11. *There exists a decision policy $(\alpha_1^*, \dots, \alpha_T^*)$ such that*

$$\begin{aligned} \mathbb{E}[u(\alpha_1^*, \dots, \alpha_T^*, X)] &= \sup_{(\alpha_1, \dots, \alpha_T) \in \mathcal{D}(\mathbb{F})} \mathbb{E}[u(\alpha_1, \dots, \alpha_T, X)] \\ &= \mathbb{E} \left[\sup_{a_1} \mathbb{E} \left[\dots \sup_{a_T} \mathbb{E} [u(a_1, \dots, a_T, X) \mid \mathcal{F}_T] \dots \mid \mathcal{F}_1 \right] \right]. \end{aligned}$$

Proof. We first note that

$$\sup_{(\alpha_1, \dots, \alpha_T) \in \mathcal{D}(\mathbb{F})} \mathbb{E}[u(\alpha_1, \dots, \alpha_T, X)] \leq \mathbb{E} \left[\sup_{a_1} \mathbb{E} \left[\dots \sup_{a_T} \mathbb{E} [u(a_1, \dots, a_T, X) \mid \mathcal{F}_T] \dots \mid \mathcal{F}_1 \right] \right]$$

assuming the suprema are all measurable which will be shown below. What remains to be shown is that the right-hand side is attained for at least one decision policy.

Take an arbitrary decision policy $(\alpha_1, \dots, \alpha_T)$. Note that $\mathbb{E}[u(\alpha_1, \dots, \alpha_{t-1}, a_t, \dots, a_T, X) \mid \mathcal{F}_t]$ is continuous in (a_t, \dots, a_T) by the dominated convergence theorem. Thus, we have that the supremum $\sup_{a_T} \mathbb{E}[u(\alpha_1, \dots, \alpha_{t-1}, a_t, \dots, a_T, X) \mid \mathcal{F}_t]$ is continuous in (a_t, \dots, a_{T-1}) by Berge's maximum theorem, and by the measurable maximum theorem (Aliprantis and Border, 2006, Theorem 18.19) it is also \mathcal{F}_t -measurable. An inductive argument yields that, for every t , and every decision policy $(\alpha_1, \dots, \alpha_T)$,

$$\mathbb{E} \left[\sup_{a_{t+1}} \mathbb{E} \left[\dots \sup_{a_T} \mathbb{E} [u(\alpha_1, \dots, \alpha_{t-1}, a_t, \dots, a_T, X) \mid \mathcal{F}_T] \dots \mid \mathcal{F}_{t+1} \right] \mid \mathcal{F}_t \right],$$

is continuous in a_t and is measurable with respect to \mathcal{F}_t (thus is a Cathéodory function).

In particular, taking $t = 1$, we get that the map

$$(a_1, \omega) \mapsto \mathbb{E} \left[\sup_{a_2} \mathbb{E} \left[\dots \sup_{a_T} \mathbb{E} [u(a_1, \dots, a_T, X) \mid \mathcal{F}_T] \dots \mid \mathcal{F}_2 \right] \mid \mathcal{F}_1 \right] (\omega),$$

¹⁶Recall that every space is tacitly endowed with its Borel σ -algebra.

is a Cathéodory function. Because \mathcal{A}_1 is compact metrizable, by the measurable selection theorem, there exists a map $\alpha_1^* : \Omega \rightarrow \mathcal{A}_1$, which is \mathcal{F}_1 -measurable, such that

$$\begin{aligned} \mathbb{E} \left[\sup_{a_2} \mathbb{E} \left[\dots \sup_{a_T} \mathbb{E} [u(\alpha_1^*, a_2, \dots, a_T, X) \mid \mathcal{F}_T] \dots \mid \mathcal{F}_2 \right] \mid \mathcal{F}_1 \right] \\ = \sup_{a_1} \mathbb{E} \left[\sup_{a_2} \mathbb{E} \left[\dots \sup_{a_T} \mathbb{E} [u(a_1, \dots, a_T, X) \mid \mathcal{F}_T] \dots \mid \mathcal{F}_2 \right] \mid \mathcal{F}_1 \right]. \end{aligned}$$

We define recursively the remaining α_t^* 's by noting that by the above result, having defined $\alpha_1^*, \dots, \alpha_{t-1}^*$ where each α_s^* is \mathcal{F}_s -measurable map from Ω to \mathcal{A}_s , the map

$$(a_t, \omega) \mapsto \mathbb{E} \left[\sup_{a_t} \mathbb{E} \left[\dots \sup_{a_T} \mathbb{E} [u(\alpha_1^*, \dots, \alpha_{t-1}^*, a_t, \dots, a_T, X) \mid \mathcal{F}_T] \dots \mid \mathcal{F}_{t+1} \right] \mid \mathcal{F}_t \right] (\omega),$$

is a Cathéodory function, and using the compact metrizable property of \mathcal{A}_t , we get that there exists a map $\alpha_t^* : \Omega \rightarrow \mathcal{A}_t$ which is \mathcal{F}_t -measurable and such that

$$\begin{aligned} \mathbb{E} \left[\sup_{a_{t+1}} \mathbb{E} \left[\dots \sup_{a_T} \mathbb{E} [u(\alpha_1^*, \dots, \alpha_t^*, a_{t+1}, \dots, a_T, X) \mid \mathcal{F}_T] \dots \mid \mathcal{F}_{t+1} \right] \mid \mathcal{F}_t \right] \\ = \sup_{a_t} \mathbb{E} \left[\sup_{a_{t+1}} \mathbb{E} \left[\dots \sup_{a_T} \mathbb{E} [u(\alpha_1^*, \dots, \alpha_{t-1}^*, a_t, a_{t+1}, \dots, a_T, X) \mid \mathcal{F}_T] \dots \mid \mathcal{F}_{t+1} \right] \mid \mathcal{F}_t \right]. \end{aligned}$$

Then, using with the law of iterated expectations to collapse the \mathcal{F}_t 's,

$$\begin{aligned} \mathbb{E} \left[\sup_{a_1} \mathbb{E} \left[\dots \sup_{a_T} \mathbb{E} [u(a_1, \dots, a_T, X) \mid \mathcal{F}_T] \dots \mid \mathcal{F}_1 \right] \right], \\ = \mathbb{E} \left[\sup_{a_2} \mathbb{E} \left[\dots \sup_{a_T} \mathbb{E} [u(\alpha_1^*, a_2, \dots, a_T, X) \mid \mathcal{F}_T] \dots \mid \mathcal{F}_2 \right] \right], \\ = \dots \\ = \mathbb{E} \left[\sup_{a_t} \mathbb{E} \left[\dots \sup_{a_T} \mathbb{E} [u(\alpha_1^*, \dots, \alpha_{t-1}^*, a_t, \dots, a_T, X) \mid \mathcal{F}_T] \dots \mid \mathcal{F}_t \right] \right], \\ = \dots \\ = \mathbb{E} [u(\alpha_1^*, \dots, \alpha_T^*, X)]. \end{aligned}$$

■

Now let us assume that the expert does not take actions in place of the individual. Instead, the expert communicates to the individual the relevant “probability tree” (i.e., his induced belief tree) at every time. Let Z_t be the random variable taking values in $\Delta^{T-t+1}(\mathcal{X})$ —i.e., the space of all “probability trees”, all endowed with the Borel σ -algebra generated by the weak-* topology—and defined by $Z_t(\omega) = \varphi_{t+1, \dots, T}(\mathbb{P}_t^\omega)$, where $\varphi_{t+1, \dots, T}$ is the induced belief tree with intermediate times $t+1, \dots, T$. Z_t represents the information that the expert communicates to the individual. Let $\mathbb{Z} = \{Z_t\}$ be the filtration generated by the discrete process Z_t . \mathbb{Z} represents the dynamic

information the individual learns from the expert. A decision policy for the individual is summarized by a tuple $(\beta_1, \dots, \beta_T)$ where every β_t is an \mathcal{Z}_t -measurable map from Ω to \mathcal{A}_t . Let $\mathcal{D}(\mathbb{Z})$ be the set of all decision policies available to the individual. The same argument as in Lemma 11 shows that there exists an optimal policy $(\beta_1^*, \dots, \beta_T^*)$, in the sense that

$$\mathbb{E}[u(\beta_1^*, \dots, \beta_T^*, X)] = \sup_{(\beta_1, \dots, \beta_T) \in \mathcal{D}(\mathbb{Z})} \mathbb{E}[u(\beta_1, \dots, \beta_T, X)].$$

Because the individual only cares about information that is relevant to the outcome, it is intuitive that information about the probability trees are enough for the individual to make optimal decisions—decisions that yield an expected utility as large as if she had direct access to the expert’s information. In the case where Ω is finite, that intuition is immediately verified. In the general case however, one must explicitly define a σ -algebra of events on probability trees, choice which is not innocuous, as it determines the amount of information that is effectively communicated. If the σ -algebra is too coarse, it can be that the information communicated is not enough for the individual to optimize his expected utility.

We show that our choice of σ -algebra contains sufficiently many events so that there is no loss of relevant information when the expert only communicates posterior information.

Proposition 6. *The individual’s optimal decisions yield the same expected utility as if she had delegated the problem to the expert:*

$$\mathbb{E}[u(\alpha_1^*, \dots, \alpha_T^*, X)] = \mathbb{E}[u(\beta_1^*, \dots, \beta_T^*, X)].$$

Proof. Note that if $W : \Omega \mapsto \mathbb{R}$ is $\sigma(X)$ -measurable, then we obviously have that $\mathbb{E}[W \mid \mathcal{Z}_T] = \mathbb{E}[W \mid \mathcal{F}_T]$. However we also have that if $W : \Omega \mapsto \mathbb{R}$ is bounded and $\sigma(Z_{t+1})$ -measurable, then $\mathbb{E}[W \mid Z_t] = \mathbb{E}[W \mid \mathcal{F}_t]$. This is a consequence of the fact that every $\Delta^k(\mathcal{X})$ is separable and metrizable, and that if one knows every weak-* event of the set of all probability measures on a separable metric space, then one can compute the expectation of every bounded Borel-measurable function on that space (Aliprantis and Border, 2006, Theorem 15.13).

In particular, we have

$$\mathbb{E}[u(a_1, \dots, a_T, X) \mid \mathcal{F}_T] = \mathbb{E}[u(a_1, \dots, a_T, X) \mid Z_T]$$

and inductively, from the above observation, for every t ,

$$\begin{aligned} & \mathbb{E}[\sup_{a_t \in \mathcal{A}_t} \mathbb{E}[\dots \mathbb{E}[\sup_{a_T} \mathbb{E}[u(a_1, \dots, a_T, X) \mid \mathcal{F}_T] \mid \mathcal{F}_{T-1}] \dots \mid \mathcal{F}_t] \mid \mathcal{F}_{t-1}], \\ &= \mathbb{E}[\sup_{a_t \in \mathcal{A}_t} \mathbb{E}[\dots \mathbb{E}[\sup_{a_T} \mathbb{E}[u(a_1, \dots, a_T, X) \mid Z_T] \mid Z_{T-1}] \dots \mid Z_t] \mid \mathcal{F}_{t-1}], \\ &= \mathbb{E}[\sup_{a_t \in \mathcal{A}_t} \mathbb{E}[\dots \mathbb{E}[\sup_{a_T} \mathbb{E}[u(a_1, \dots, a_T, X) \mid Z_T] \mid Z_{T-1}] \dots \mid Z_t] \mid Z_{t-1}]. \end{aligned}$$

In particular,

$$\begin{aligned}
& \mathbb{E} \left[\sup_{a_1 \in \mathcal{A}_1} \mathbb{E} \left[\sup_{a_2 \in \mathcal{A}_2} \mathbb{E} \left[\dots \mathbb{E} \left[\sup_{a_T \in \mathcal{A}_T} \mathbb{E} [u(a_1, \dots, a_T, X) \mid \mathcal{F}_T] \mid \mathcal{F}_{T-1}] \dots \mid \mathcal{F}_2 \mid \mathcal{F}_1 \right] \right] \right], \\
&= \mathbb{E} \left[\sup_{a_1 \in \mathcal{A}_1} \mathbb{E} \left[\sup_{a_2 \in \mathcal{A}_2} \mathbb{E} \left[\dots \mathbb{E} \left[\sup_{a_T \in \mathcal{A}_T} \mathbb{E} [u(a_1, \dots, a_T, X) \mid \mathcal{Z}_T] \mid \mathcal{Z}_{T-1}] \dots \mid \mathcal{Z}_2 \mid \mathcal{Z}_1 \right] \right] \right], \\
&= \mathbb{E} \left[\sup_{a_1 \in \mathcal{A}_1} \mathbb{E} \left[\sup_{a_2 \in \mathcal{A}_2} \mathbb{E} \left[\dots \mathbb{E} \left[\sup_{a_T \in \mathcal{A}_T} \mathbb{E} [u(a_1, \dots, a_T, X) \mid \mathcal{Z}_T] \mid \mathcal{Z}_{T-1}] \dots \mid \mathcal{Z}_2 \mid \mathcal{Z}_1 \right] \right] \right].
\end{aligned}$$

We conclude by the dynamic programming principle enunciated in Lemma 11. Note that all suprema are measurable by the measurable maximum theorem, as detailed in the lemma. ■

D Appendix: Additional proofs

The following lemma is technical, but useful to show that some variables are well-defined (i.e., measurable) random variables.

Lemma 12. *Let \mathcal{A} be a separable metrizable space and \mathcal{B} be a measurable space, with $\Delta(\mathcal{B})$ the set of probability measures on \mathcal{B} equipped with the weak-* topology. Both \mathcal{A} and $\Delta(\mathcal{B})$ are given their respective Borel σ -algebras. Let $\Psi : a \mapsto \Psi_a$ be a map from \mathcal{A} to $\Delta(\mathcal{B})$. If, for every event E of \mathcal{B} , the map $a \mapsto \Psi_a(E)$ from \mathcal{A} to \mathbb{R} is measurable, then Ψ is measurable.*

Proof. Because \mathcal{B} is metrizable, standard approximation arguments apply to show that if $a \mapsto \psi_a(E)$ is measurable for every event E of \mathcal{B} , then the map $a \mapsto \int f d\psi_a$ is also measurable for every continuous and bounded function $f : \mathcal{B} \mapsto \mathbb{R}$. We remark that the sets of the form $\{\mu \in \Delta(\mathcal{B}) \mid \int f d\mu \in I\}$ for I an open interval, and f a continuous bounded function, form a sub-base of the weak-* topology on $\Delta(\mathcal{B})$. Because $\Delta(\mathcal{B})$ is separable and metrizable (Aliprantis and Border, 2006, Theorem 15.12), every open set in $\Delta(\mathcal{B})$ is a countable unions of finite intersections of elements of the sub-base. Thus the σ -algebra generated by the sets $\{\mu \in \Delta(\mathcal{B}) \mid \int f d\mu \in I\}$ is the Borel σ -algebra on $\Delta(\mathcal{B})$, which makes ψ measurable. ■

D.1 Proof of Lemma 1

Fix a given information structure $(\Omega, \mathbb{F}, \mathbb{P}, X)$. We observe that the inference map of order 1, $Q \mapsto \varphi(Q)$, associates to every state probability measure Q the law of X , and is thus well-defined. We also observe that $\Delta(\Omega)$ is separable measurable (Aliprantis and Border, 2006, Theorem 15.12). Thus, for every event E of \mathcal{X} , $Q \mapsto Q(X \in E)$ is measurable (Aliprantis and Border, 2006, Theorem 15.13). Applying Lemma 12, we get that $Q \mapsto \varphi(Q)$ is measurable.

Now take a given $j \geq 0$ and suppose that the associated inference map $\varphi_{t_1, \dots, t_j}$ is well-defined and measurable, for every t_1, \dots, t_j . Then we show that $\varphi_{t_1, \dots, t_{j+1}}$ is well-defined and measurable, for every t_1, \dots, t_{j+1} .

It is well-defined: By assumption, $\omega \mapsto P_{t_1}^\omega(E)$ is measurable for every event E , which by Lemma 12 implies that $\omega \mapsto P_{t_1}^\omega$ is measurable and thus a well-defined random variable with values

in $\Delta(\Omega)$. Thus, $\omega \mapsto \varphi_{t_2, \dots, t_{j+1}}(\mathbf{P}_{t_1}^\omega)$ is a well-defined random variable, with values in $\Delta^{j+1}(\mathcal{X})$.

It is measurable: By the induction hypothesis, for every event E of $\Delta^{j+1}(\mathcal{X})$, the set $\{\omega \in \Omega : \varphi_{t_2, \dots, t_{j+1}}(\mathbf{P}_{t_1}^\omega) \in E\}$ is a well-defined event of Ω . Applying again Theorems 15.12 and 15.13 of Aliprantis and Border (2006), we get that $Q \mapsto Q(\varphi_{t_2, \dots, t_{j+1}}(\mathbf{P}_{t_1}) \in E)$ is measurable, and by Lemma 12, we get that $Q \mapsto Q(\varphi_{t_2, \dots, t_{j+1}}(\mathbf{P}_{t_1}))$ is measurable.

D.2 Proof of Lemma 2

Let $f_{T+1}(S, p_T) = \int S dp_T$. We note that f_{T+1} is jointly continuous, as securities have a compact domain, and $\Delta(\mathcal{X})$ is given the weak-* topology. For $t \leq T$, let $f_t(M_t, p_{t-1}) = \int \pi_t(M_t, q) dp_{t-1}(q)$. We will show that if f_{t+1} is jointly continuous, then both π_t and f_t are jointly continuous.

Let g_t be the correspondence from $\mathcal{M}_t^{[\underline{L}, \bar{U}]} \times \Delta^{T-t+1}(\mathcal{X})$ to $\mathcal{M}_{t+1}^{[\underline{L}, \bar{U}]}$ defined by $g_t(M_t, p_t) = M_t$ (with $\mathcal{M}_{T+1}^{[\underline{L}, \bar{U}]}$ the set of securities.) Because g_t has non-empty compact values, and is continuous when interpreted as a map from $\mathcal{M}_t^{[\underline{L}, \bar{U}]} \times \Delta^{T-t+1}(\mathcal{X})$ to $\mathcal{M}_t^{[\underline{L}, \bar{U}]}$, the correspondence is continuous (Aliprantis and Border, 2006, Theorem 17.15). Since f_{t+1} is continuous, we can invoke Berge's maximum theorem (Aliprantis and Border, 2006, Theorem 17.31) to get that the map

$$(M_t, p_t) \mapsto \max_{m_{t+1} \in g_t(M_t, p_t)} f_{t+1}(m_{t+1}, p_t)$$

is continuous. This proves the joint continuity π_t . If, in addition, $M_t \mapsto \pi_t(M_t, \cdot)$ is continuous in the sup-norm topology, then f_t is jointly continuous (Aliprantis and Border, 2006, Corollary 15.7).

What remains to be shown is the continuity of the maps $M_t \mapsto \pi_t(M_t, \cdot)$. Letting $\pi_{T+1}(S, x) = S(x)$ we get that $S \mapsto \pi_{T+1}(S, \cdot)$ is continuous in the sup-norm topology. We will show that if $M_{t+1} \mapsto \pi_{t+1}(M_{t+1}, \cdot)$ is continuous for every M_{t+1} of order $T - t + 1$, then so is $M_t \mapsto \pi_t(M_t, \cdot)$.

To do so, let \mathcal{C}_t be the space of continuous real functions on $\Delta^{T-t+1}(\mathcal{X})$ endowed with its sup-norm. Let $\mathcal{K}_t(M_t) \subset \mathcal{C}_t$ be the convex hull of $\{\pi_{t+1}(m_{t+1}, \cdot); m_{t+1} \in M_t\}$, which, being the finite union of points, is closed and bounded in \mathcal{C}_t . Let \mathcal{C}'_t be the norm dual of \mathcal{C}_t , which consists of all norm-continuous linear functionals. Let \mathcal{U}_t be the closed unit ball of \mathcal{C}_t , and $\mathcal{U}'_t \subset \mathcal{C}'_t$ be its polar, so that $f \in \mathcal{U}'_t$ if $|f(x)| \leq 1$ for all $x \in \mathcal{U}_t$. For a given closed, bounded set C of \mathcal{C}_t , let h_C defined by $h_C(f) = \sup_{x \in C} f(x)$ denote its support function. Using the induction hypothesis, we remark that the map $M_t \mapsto \mathcal{K}_t(M_t)$ is continuous, if the set of closed bounded subsets of \mathcal{C}_t is given the Hausdorff metric induced by the sup-norm topology. Let us suppose that a sequence $\{M^k\}$ converges to M^∞ , all being finite menus of order $T - t + 1$. Then $\sup_{u' \in \mathcal{U}'} |h_{\mathcal{K}_t(M^k)}(u') - h_{\mathcal{K}_t(M^\infty)}(u')| \rightarrow 0$ (see Lemma 7.58 of Aliprantis and Border, 2006). By the Riesz-Radon representation (Aliprantis and Border, 2006, Corollary 14.15), every $p_t \in \Delta^{T-t+1}(\mathcal{X})$ can be identified with a member of \mathcal{U}' , so that $\pi_t(M_t, \cdot)$ can be viewed as the support function of $\mathcal{K}_t(M_t)$ restricted to $\Delta^{T-t+1}(\mathcal{X})$. Thus, $\sup_{p_t \in \Delta^{T-t+1}(\mathcal{X})} |\pi_t(M^k, p_t) - \pi_t(M^\infty, p_t)| \rightarrow 0$, which makes $M_t \mapsto \pi_t(M_t, \cdot)$ continuous.

D.3 Proof of Lemma 3

As in Lemma 2, define the correspondence $g_t(M_t, p_t) = M_t$, and the function $f_t(M_t, p_{t-1}) = \int \pi_t(M_t, q) dp_{t-1}(q)$, with $\pi_{T+1}(S, x) = S(x)$. For every t , g_t is measurable (Aliprantis and Border, 2006, Theorem 18.10), f_t is a Carathéodory function, and every space $\mathcal{M}_t^{[a,b]}$ is separable.¹⁷ We can then apply the measurable selection theorem (Theorem 18.19 Aliprantis and Border, 2006), and we get that the argmax correspondence

$$\arg \max_{m_{t+1} \in g_t(M_t, p_t)} \int_{\Delta^{T-t}(\mathcal{X})} \pi_{t+1}(m_{t+1}, q) dp_t(q)$$

is measurable and admits a measurable selector. Moreover, by the Castaing representation theorem (Aliprantis and Border, 2006, Corollary 14.18), we can enumerate the elements of the argmax in a measurable way, in the sense that there exists a sequence of measurable selectors $\{\Phi_t^k\}$ such that $\arg \max_{m_{t+1} \in g_t(M_t, p_t)} f_{t+1}(m_{t+1}, p_t) = \{\Phi_t^k(M_t, p_t); k = 1, 2, \dots\}$. We observe that

$$\left| \arg \max_{m_{t+1} \in M_t} \int_{\Delta^{T-1}(\mathcal{X})} \pi_{t+1}(m_{t+1}, p) dq_t(p) \right| = \lim_{K \rightarrow \infty} \sum_{k=1}^K \frac{1}{\sum_{i=1}^K \mathbb{1}_{\Phi_t^i(M_t, p_t) = \Phi_t^k(M_t, p_t)}},$$

is measurable as a pointwise limit of real-valued measurable functions. Also, if Π_{t+1} is measurable, then Π^t , which can be written

$$\Pi^t(p_t, \dots, p_T, x; M_t) = \frac{1}{\left| \arg \max_{m_{t+1} \in M_t} \int \pi_{t+1}(m_{t+1}, \cdot) dq_t \right|} \lim_{K \rightarrow \infty} \sum_{k=1}^K \frac{\Pi^{t+1}(p_{t+1}, \dots, p_T, x; \Phi_t^k(M_t, p_t))}{\sum_{i=1}^K \mathbb{1}_{\Phi_t^i(M_t, p_t) = \Phi_t^k(M_t, p_t)}}.$$

is also jointly measurable. Of course, $\Pi_{T+1}(x; S) = S(x)$ is measurable. This concludes the proof.

D.4 Proof of Lemma 6

Suppose $f, g \in \mathcal{L}_t$. By linearity of \mathcal{L}_t we have directly that $\alpha f \in \mathcal{L}_t$ and $f + g \in \mathcal{L}_t$. The function 1 belongs to \mathcal{L}_t since $\pi_{t+1}(\mathbf{1}, \cdot) = 1$. To see that $\max\{f, g\} \in \mathcal{L}_t$, we decompose f and g as $f(\cdot) = \sum_{i \in \mathcal{I}^+} \alpha_i \pi_{t+1}(M_i^A, \cdot) - \sum_{i \in \mathcal{I}^-} \alpha_i \pi_{t+1}(M_i^A, \cdot)$ and $g(\cdot) = \sum_{j \in \mathcal{J}^+} \beta_j \pi_{t+1}(M_j^B, \cdot) -$

¹⁷First, note that the set of securities is a separable metric space, by Aliprantis and Border (2006, Lemma 3.99). Then the result follows as the set of finite sets of a separable metric space is itself separable when endowed with the Hausdorff topology. In particular, the set of finite sets of a countable dense subset is countable and dense in the Hausdorff topology.

$\sum_{j \in \mathcal{J}^-} \beta_j \pi_{t+1}(M_j^B, \cdot)$, with all $\alpha_i, \beta_j \geq 0$. Then

$$\begin{aligned}
& \max\{f, g\} \\
&= \max \left\{ \sum_{i \in \mathcal{I}^+} \alpha_i \pi_{t+1}(M_i^A, \cdot) + \sum_{j \in \mathcal{J}^-} \beta_j \pi_{t+1}(M_j^B, \cdot), \sum_{j \in \mathcal{J}^+} \beta_j \pi_{t+1}(M_j^B, \cdot) + \sum_{i \in \mathcal{I}^-} \alpha_i \pi_{t+1}(M_i^A, \cdot) \right\} \\
&\quad - \left(\sum_{i \in \mathcal{I}^-} \alpha_i \pi_{t+1}(M_i^A, \cdot) + \sum_{j \in \mathcal{J}^-} \beta_j \pi_{t+1}(M_j^B, \cdot) \right), \\
&= \pi_{t+1} \left(\left(\sum_{i \in \mathcal{I}^+} \alpha_i M_i^A + \sum_{j \in \mathcal{J}^-} \beta_j M_j^B \right) \cup \left(\sum_{j \in \mathcal{J}^+} \beta_j M_j^B + \sum_{i \in \mathcal{I}^-} \alpha_i M_i^A \right), \cdot \right) \\
&\quad - \pi_{t+1} \left(\sum_{i \in \mathcal{I}^-} \alpha_i M_i^A + \sum_{j \in \mathcal{J}^-} \beta_j M_j^B, \cdot \right),
\end{aligned}$$

hence $\max\{f, g\} \in \mathcal{L}_t$.

D.5 Proof of Lemma 7

Let $q', q'' \in \Delta^{T-t}(\mathcal{X})$. Suppose that for every $f \in \mathcal{L}_t$, the equality $f(q') = f(q'')$ holds. Then for every $M \in \Delta^{T-t-1}(\mathcal{X})$, and with

$$c = \frac{1}{2} \left[\int_{\Delta^{T-t-1}(\mathcal{X})} \pi_{t+2}(M, p) dq'(p) + \int_{\Delta^{T-t-1}(\mathcal{X})} \pi_{t+2}(M, p) dq''(p) \right]$$

we get that $\pi_{t+1}(\{M, c\mathbf{1}\}, q') = \pi_{t+1}(\{M, c\mathbf{1}\}, q'')$ and thus

$$c = \int_{\Delta^{T-t-1}(\mathcal{X})} \pi_{t+2}(M, p) dq'(p) = \int_{\Delta^{T-t-1}(\mathcal{X})} \pi_{t+2}(M, p) dq''(p)$$

which, by the induction hypothesis, implies $q' = q''$.

D.6 Proof of Lemma 8

The map π_0 is measurable. Indeed, the map $(S, \omega) \mapsto S(X(\omega))$ is bounded and measurable, and Ω is separable metrizable, so by Theorem 17.25 of Kechris (1995), the map $\pi_0 : (S, Q) \mapsto \int S(X(\omega)) dQ(\omega)$ is jointly measurable.

Now suppose that π_k is jointly measurable. It implies that π_{k+1} is well-defined, because for every $\tau, \omega \mapsto P_\tau^\omega$ is a well-defined random variable with values in $\Delta(\Omega)$.

We observe that the map $\sigma_{k+1} \mapsto M_{k+1}$ from Σ_{k+1} to 2^{Σ^k} is measurable for the Borel σ -algebra of the Hausdorff metric topology on 2^{Σ^k} . Thus by Theorem 18.10 of Aliprantis and Border (2006), the correspondence $\sigma_{k+1} \mapsto M_{k+1}$ from Σ_{k+1} to Σ_k is measurable. We then use the Castaing representation theorem (Aliprantis and Border, 2006, Corollary 18.14) to generate a sequence $\{\Phi_i : i = 1, 2, \dots\}$ of measurable maps $\Phi_i : \Sigma_{k+1} \rightarrow \Sigma_k$ such that $M_{k+1} = \{\Phi_i(\sigma_{k+1}) : i = 1, 2, \dots\}$.

Thus, we get

$$\max_{\sigma_k \in M_{k+1}} \pi_k(\sigma_k, Q) = \sup_{i=1,2,\dots} \pi_k(\Phi_i(\sigma_{k+1}), Q).$$

Thus the map $(\sigma_{k+1}, Q) \mapsto \max_{\sigma_k \in M_{k+1}} \pi(\sigma_{k+1}, Q)$ is jointly measurable as the pointwise supremum of countably many jointly measurable maps. Besides, the right-continuity of $t \mapsto P_t^\omega$ for every ω implies the joint measurability of $(\tau, \omega) \mapsto P_\tau^\omega$ and thus the joint measurability of $(\sigma_{k+1}, \omega) \mapsto P_{\tau_{k+1}}^\omega$. We have thus established that the map

$$(\sigma_{k+1}, \omega) \mapsto \max_{\sigma_k \in M_{k+1}} \pi(\sigma_{k+1}, P_{\tau_{k+1}}^\omega)$$

is jointly measurable. It follows that the map

$$(\sigma_{k+1}, Q) \mapsto \int \left[\max_{\sigma_k \in M_{k+1}} \pi_k(\sigma_k, P_{\tau_{k+1}}^\omega) \right] dQ(\omega).$$

is jointly measurable, again applying Theorem 17.25 of Kechris (1995). Hence π_{k+1} is jointly measurable.

D.7 Proof of Lemma 9

The map $(\{Q_t\}, x, S) \mapsto \Pi(\{Q_t\}, x; S) = S(x)$ is jointly measurable.

Let us suppose that $(\{Q_t\}, x; \sigma_k) \mapsto \Pi(\{Q_t\}, x; \sigma_k)$ is measurable. Let $\{\Phi_i : i = 1, 2, \dots\}$ be a sequence of measurable maps from $\Sigma_{k+1}^{[\underline{L}, \overline{U}]}$ to $\Sigma_k^{[\underline{L}, \overline{U}]}$ such that $M_{k+1} = \{\Phi_i(\sigma_{k+1}) : i = 1, 2, \dots\}$, whose existence was proved in Lemma 8. Because securities of $\Sigma_k^{[\underline{L}, \overline{U}]}$ take values in $[\underline{L}, \overline{U}]$ we have that for any $\sigma', \sigma'' \in \Sigma_k^{[\underline{L}, \overline{U}]}$, $d(\sigma', \sigma'') < \overline{D}$ for some constant \overline{D} large enough. Let ν be the argmax correspondence $(\sigma_{k+1}, Q) \mapsto \arg \max_{\sigma_k \in M_{k+1}} \pi(\sigma_k, Q)$.

We show that ν is weakly measurable. Let δ be the associated distance function, it is a map from $\Sigma_k^{[\underline{L}, \overline{U}]} \times (\Sigma_{k+1}^{[\underline{L}, \overline{U}]} \times \Delta(\Omega))$ to \mathbb{R} defined by $\delta(\sigma_k, (\sigma_{k+1}, Q)) = d(\sigma_k, \nu(\sigma_{k+1}, Q))$. We remark that for every finite set \mathcal{S} of $\Sigma_k^{[\underline{L}, \overline{U}]}$, $\sigma_k \mapsto d(\sigma_k, \mathcal{S}) = \min_{\sigma' \in \mathcal{S}} d(\sigma_k, \sigma')$ is continuous. Also,

$$\delta(\sigma_k, (\sigma_{k+1}, Q)) = \min_{i=1,2,\dots} \left(d(\sigma_k, \Phi_i(\sigma_{k+1})) \mathbf{1}_{g(\sigma_{k+1}, Q) = \pi(\Phi_i(\sigma_{k+1}), Q)} + \overline{D} \mathbf{1}_{g(\sigma_{k+1}, Q) \neq \pi(\Phi_i(\sigma_{k+1}), Q)} \right)$$

where $g(\sigma_{k+1}, Q) = \max_{\sigma_k \in M_{k+1}} \pi_k(\sigma_k, Q)$. It was shown in the proof of Lemma 8 that g is jointly measurable. Therefore, $(\sigma_{k+1}, Q) \mapsto \delta(\sigma_k, (\sigma_{k+1}, Q))$ is measurable as the pointwise infimum of countably many measurable functions.

We have thus proved that the distance function δ associated to the argmax correspondence ν is Carathéodory function, which establishes its weak measurability (Aliprantis and Border, 2006, Theorem 18.5).

The weak measurability of ν implies in turn that we can enumerate its elements by a sequence of measurable selectors $\{\tilde{\Phi}_i : i = 1, 2, \dots\}$ where $\tilde{\Phi}_i : \Sigma_{k+1} \times \Delta(\Omega) \mapsto \Sigma_k$, in the sense that

$\nu(\sigma_{k+1}, Q) = \{\tilde{\Phi}_i(\sigma_{k+1}, Q) : i = 1, 2, \dots\}$. We can then write

$$\left| \arg \max_{\sigma_k \in M_{k+1}} \pi(\sigma_k, Q) \right| = \lim_{\ell \rightarrow \infty} \frac{1}{\sum_{j=1}^{\ell} \mathbb{1}_{\tilde{\Phi}_i(\sigma_{k+1}, Q) = \tilde{\Phi}_j(\sigma_{k+1}, Q)}}$$

Thus $(\sigma_{k+1}, Q) \mapsto \left| \arg \max_{\sigma_k \in M_{k+1}} \pi(\sigma_k, Q) \right|$ is measurable as a pointwise limit of a sequence of measurable functions whose limit exists. Next, we observe that $(\tau, \{Q_t\}) \mapsto Q_\tau$ is right-continuous in τ and measurable in $\{Q_t\}$, and therefore is jointly measurable, which in turn implies joint measurability of $(\sigma_{k+1}, \{Q_t\}) \mapsto Q_{\tau_{k+1}}$. Finally, observing that

$$\Pi(\{Q_t\}, x; \sigma_{k+1}) = \frac{1}{\left| \arg \max_{\sigma_k \in M_{k+1}} \pi(\sigma_k, Q_{\tau_{k+1}}) \right|} \lim_{\ell \rightarrow \infty} \frac{\sum_{i=1}^{\ell} \Pi(\{Q_t\}, x; \tilde{\Phi}_i(\sigma_{k+1}, Q_{\tau_{k+1}}))}{\sum_{j=1}^{\ell} \mathbb{1}_{\tilde{\Phi}_i(\sigma_k, Q_{\tau_{k+1}}) = \tilde{\Phi}_k(\sigma_j, Q_{\tau_{k+1}})}}$$

we get measurability of $(\{Q_t\}, x, \sigma_{k+1}) \mapsto \Pi(\{Q_t\}, x; \sigma_{k+1})$ again as a point-wise limit of a sequence of measurable functions. This concludes the proof.