Dynamic Belief Elicitation*

Christopher P. Chambers† Nicolas S. Lambert‡

Monday April 24, 2017

Abstract

At $t = 1$, an expert has (probabilistic) information about a random outcome X. The expert obtains further information as time passes, up to $t = N$ at which X is 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 outcome realization and his reports.) We show that this can be done with the following sort of protocol: At the penultimate time $t = N - 1$, the expert chooses a payoff function from a menu of such functions, where the menu available to him was chosen by him at time $t = N - 2$ 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.

Keywords: Elicitation device; Scoring rule; Dynamic information.

---

*A previous version of the paper circulated under the title "Dynamically Eliciting Unobservable Information" (February 2014). We are particularly grateful to David Kreps for his invaluable help and advice. We also thank Matthew Jackson, Edi Karni, Norio Takeoka, Juuso Välimäki, and seminar participants at ASU, Berkeley, Caltech, Chicago, Georgetown, Glasgow, Harvard/MIT, ITAM, Microsoft Research, NUS Singapore, NYU, Oxford, Rochester, Stanford, the Toulouse Network on Information Technologies, UCLA, UC Davis, UCSD, UC Riverside, Wisconsin, Yale, as well as the participants and organizers of ACM EC, the William Thomson 65th birthday, SAET, the International Symposium on Algorithmic Game Theory, Social Choice and Welfare 2014 conferences, the World Congress of the Econometric Society and ESSET 2015 conferences, for comments and suggestions that greatly improved the paper. We gratefully acknowledge financial support from the National Science Foundation through grants SES-1426867 and CCF-1101209.

†University of California San Diego, La Jolla, CA 92093, USA; cpchambers@ucsd.edu.

‡Stanford Graduate School of Business, Stanford, CA 94305, USA; nlambert@stanford.edu.

1 Introduction

We revisit the classical question of how to elicit expert beliefs regarding a random outcome of interest to some individual. In the standard paradigm, an expert holds probability assessments about the various possible outcomes of a random variable, to be realized publicly at some later time. The literature asks: can we devise a protocol that rewards the expert based on reports of his own assessment, and the eventually realized outcome, which induces him to report truthfully as a strict best response?

In this paper, we depart from the static benchmark examined in the literature to incorporate learning dynamics. There is now a sequence of times $t_1 < \cdots < t_N$. The random outcome of interest—call it random variable X—materializes publicly at $t = t_N$. At $t = t_1$, the expert holds a prior assessment regarding the random outcome, and at every interim period t_k, the expert receives information that may change his probability assessment. We answer the following generalization of the classical question. 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 belief and the *structure of information* he anticipates to receive, and then to provide truthful updates about information he receives, as he receives it? This protocol must provide the expert with a payoff at t_N that depends only on the realized outcome of X, and the announcements of the expert prior to t_N.

Note that we are not just eliciting probability assessments about X at different dates. Crucially, we also ask the expert to tell us about how he thinks his beliefs will evolve over time. To illustrate, consider the simple case of a binary X, such as the outcome of a risky project. There are only three time periods, $t = 0, 1, 2$. The project outcome realizes at $t = 2$. In this case, we would like to ask the expert his assessment of the likelihood of success at both $t = 0, 1$, but principally, we seek to learn his belief, at $t = 0$, regarding the (private) assessment he expects to form at $t = 1$.

To understand why such "dynamic information" is relevant, suppose an investor contemplates buying shares of the project. To simplify, the investor is risk neutral and there is no discounting. If the project succeeds, it pays off $B = 100$. But, if the project fails, it pays off 0. It costs $C = 40$ to participate in the project. Let us look at two scenarios. In a first scenario, the investor must decide at $t = 0$. Clearly, she is best off investing whenever she estimates the chances of success, p, such that expected benefits exceed costs: $p \times B > C$, or $p > .4$. In the sequel, take $p = .5$. In a second scenario, the investor can delay her decision to $t = 1$. However, by doing so, she may lose the opportunity to invest, for example, due to competition. Let us say there is 40% chance of losing the investment opportunity. If the investor expects to remain as informed at $t = 1$ as she is at $t = 0$, she should invest at $t = 0$, and make as expected payoff $100 \cdot .5 - 40 = 10$. Suppose instead the investor expects to conduct further investigation. Specifically, she expects to get good news or bad news, with equal chances, and in the case of good news, she will revise her assessment from .5 to .8, while in the case of bad news, she will revise her assessment to .2. Note that these posteriors are consistent with Bayes law, in that accounting for this interim information, her prior assessment at $t = 0$ remains .5. In this case, if she delays the decision at $t = 0$, and decides to invest at $t = 1$ only when she receives good news, she earns on average $.6(100 \cdot .8 - 40)/2 = 12$. This is more than what she would earn if she had taken the decision at $t = 0$. Yet, the prior likelihood of a successful project at $t = 0$ is the same, both with and without interim information. To maximize her payoff, the investor should account for when and how her information unravels over time—an element not captured in probability assessments over final outcomes, and that the protocols of the literature, which focus on the static benchmark, do not elicit.

This example captures a simple but important fact. Dynamic information, or dynamic beliefs—that is, beliefs that concern future beliefs—enable us to solve any sort of dynamic decision problem. In contrast, probability assessments on the final outcome at one point in time are only relevant to static decision problems. Probability assessments on the final outcome at different dates can help solve dynamic decision problems, but these are degenerate, in that they can be decomposed into a collection of independent, static decision problems at different times.[1]

To date, the literature has focused on the elicitation and evaluation of probability estimates, which help solve static decision problems, but not dynamic problems. *However, many real-world decisions exhibit a dynamic structure.* A standard example is 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 an online advertisement), 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 productions[2]), insurance markets or irreversible investments. This sort of problems is extensively studied in the literature on real options (see, for example, Dixit and Pindyck, 1994). For these problems, probability assessments on

[1] We provide a formal statement and proof in the Supplementary Appendix, Section SA.4.

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

the final outcome at different dates is insufficient to decide optimally.

The key challenge of our setup is that any observation or measurement the expert collects in the interim periods is private. Thus, an elicitor cannot condition on this information. Moreover, the expert alone can tell which type of observations he will receive, and how it will affect his prior beliefs. To appreciate the difficulty, let us take a starker example. Suppose the expert is to receive new information at m different dates, and that each observation is the realization of some k-dimensional signal. We want the expert to tell us, at the outset, the full joint probability distribution over the m signals and the project outcome (success or failure), and then tell us the signals he receives as he receives them. As m and k grow large, the object to elicit becomes fairly complex. However, to enforce truthful expert declarations, as a strict best response, the only piece of tangible information is a single outcome taking only two possible values.

Naturally, in theory, a decision maker could bypass the information elicitation step by delegating the decision to the expert, and motivate the expert by offering a share of the rewards. In pratice, "selling the project" to the expert is often not desirable: the decision maker may want to collect and aggregate information from different experts, she may want to keep private the problem she confronts or she may be the only one to understand the details of the problem, the expert may specialize on some particular type of information and sell it to multiple buyers who each confront different decision problems, and so on. In this paper, as in the literature, we do not make assumptions on the use of the information elicited. We simply suppose that an individual has interest in the information. The expert need not know the details as to how his information is about to be used.

Our first main result, Theorem 1, introduces a class of protocols that answers the aforementioned question positively. Each protocol consists in providing a carefully designed large compound option. At time t_{N-1}, 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_{N-1} was chosen by him at time t_{N-2} from a menu of menus. The menu of menus available to him at time t_{N-3} was chosen from a menu of menus of menus available to him at time t_{N-4}, and so forth. Each protocol of the class is constructed from, and uniquely identified to, a probability distribution over a small and "simple" collection of compound options.

Our second main result, Theorem 2, provides a near-characterization. It shows that any protocol that induces the expert to reveal dynamic information can be approximated by a protocol of the class just described.

We investigate the most difficult situation, whereby the individual who consults the expert observes no information prior to the realization of the uncertain outcome.

It is easy to see that our results apply in an intermediate case in which the individual can observe some 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. Our framework can be used to elicit the probabilities associated with temporal lotteries in the sense of Kreps and Porteus (1978), for example.

The paper is organized as follows. After reviewing the literature, we introduce a simple two-period example to build intuition and illustrate our approach in Section 2. In Section 3, we lay down the general model with a finite number of time periods. We explain in Section 3.2 how to construct the strategyproof protocols. Section 3.3 contains our first main result regarding existence, and Section 3.4 contains our second main result on uniqueness. In Section 4, we extend our construction to a continuous-time framework, allowing to track a flow of private signals and to deal with observations that arrive at privately known random times. The proofs of Section 3 are in the Appendix, and the proofs of Section 4 are relegated to the Supplementary Appendix.

The Supplementary Appendix includes additional results. In Section SA.1, we study various common special cases and supply closed-form formulas of protocols. In Section SA.2, we demonstrate that natural mechanisms of the sort "eliciting an expert's belief over X, and using that expert's report to elicit, at an earlier date, the belief over his future belief" fail to elicit truthful assessments. In Section SA.3, we show that, under regularity conditions, the information the protocols of this paper elicit enables to solve all dynamic decision problems—the protocols of this paper are "universal" in that sense. Finally, in Section SA.4, we show that the mechanisms of the literature can only solve a degenerate class of dynamic problems, which can, essentially, be decomposed into separate static problems.

Related Literature

The literature on eliciting expert beliefs goes back to Brier (1950) and Good (1952), who establish the static benchmark and introduce the Brier score and the logarithmic scoring rule, respectively, as the first "strategyproof" deterministic protocols (i.e., protocols that do not randomize and induce truthtelling with strict incentives). These protocols, known as *strictly proper scoring rules*, have an elegant characterization: they are the subgradient of strictly convex functions, a result first discovered by McCarthy (1956) and Savage (1971) for special cases, and subsequently extensively generalized. The characterization makes it easy to construct scoring rules in the finite-dimensional case: one just needs to take a strictly convex function, and compute the partial derivatives. The literature on scoring rules is vast and spans several fields. Importantly, it focuses on the static setting. For a survey and applications of the subgradient characterization, see Gneiting and Raftery (2007). Contributions to, and applications of the scoring rules literature in economics include, for example, Thomson (1979), Karni (2009), Fang, Stinchcombe and Whinston (2010), Stewart (2011), Ostrovsky (2012), Lambert (2013), Steiner, Stewart and Matějka (2016). Scoring rules are commonly used in the experimental literature to elicit beliefs, see for example Nyarko and Schotter (2002). The literature on strategic distinguishability also utilizes scoring rules, see for example Bergemann, Morris and Takahashi (2017). The idea is to ensure that types are economically well-defined concepts: that two different types can be distinguished in terms of their behavior. In related recent work, Karni (2017) introduces a mechanism to elicit approximate second-order beliefs in a setting with incomplete preferences under Knightian uncertainty.

Our approach departs from the standard "subgradient method": while there exists an analog to the subgradient characterization in our setting,[3] the characterization cannot be used to construct strategyproof protocols in the presence of private learning (or even to prove their existence). Instead, our approach relies on an idea originally developed by Allais (1953) and also attributed to W. Allen Wallis (see Savage, 1954) in a revealed-preference context: to elicit an individual's preference over a collection of objects, one can ask the individual for his preference of the entire collection, choose two objects at random, and then give to the individual the object that is preferred according to his announcement.

This simple, yet powerful idea is central to our approach. To elicit the expert's (dynamic) beliefs, we consider a class of simple dynamic decision problems. The class is designed so that watching the expert operate on each of these decision problems uncovers the expert's beliefs entirely; however, each decision problem taken separately only reveals a small part of the expert's beliefs. To elicit the expert's beliefs as a single decision, we combine all of the simple decision problems by a suitable randomization among them, in the spirit of Allais. As we explain in Sections 2 and 3, the crucial step of our approach is to identify what constitutes a suitable class of simple decision problems and a suitable randomization device.

In the static benchmark of the literature, several works relate indirectly to the Allais idea. In their seminal work, Becker, DeGroot and Marschak (1964) introduce, as an alternative to the Brier score, a method for eliciting an expert's belief via a second-price auction with random reserve price. Karni (2009) offers a related mechanism more specifically targeted at experts who depart from risk neutrality.[4] Matheson

[3]See Proposition 1 in Section 3.

[4]In experimental economics, it is well-known that randomization enables experimenters to elicit

and Winkler (1976) propose a scoring rule to elicit the c.d.f. of real-valued random variables. In a similar fashion, Schervish (1989) shows a characterization of strictly proper scoring rules for binary outcomes as an integral of some specific functions with positive measures. Although these works are independent from one another, they share the common feature that they can be interpreted as applications of the Allais idea, either by explicit randomization as in Becker, DeGroot and Marschak (1964) and Karni (2009), or as a mechanism that is equivalent to introducing randomization, as in Matheson and Winkler (1976) and Schervish (1989). One contribution of our work is thus to formalize the connection to the Allais idea, illustrate its effectiveness beyond the static benchmark, and the key difficulties of its implementation.

2 A Simple Example

To begin, we explore a simple example to illustrate the main question of interest and the method we develop in the general framework of Section 3.

Throughout this section, an individual (she) has interest in some random variable X, indicator of a public random event, taking values in $\{0, 1\}$. For instance, the random variable may designate the outcome of a risky project. The individual seeks to elicit information about X from a risk-neutral expert (he). Most importantly, the information that the elicitor seeks to obtain is dynamic and unravels over time.

The setup is as follows. There are three time periods, indexed $t = 0, 1, 2$. At the final period, X publicly realizes. At the interim period, the expert performs some measurements and information gathering, which result in the private observation of a random signal, denoted by the random variable P. To simplify matters, the signal is a summary statistic: the posterior probability that $X = 1$.[5] In the first period, the expert forms a private prior belief about the outcome X, and also, importantly, about the signal he expects to receive at time $t = 1$. Those beliefs are summarized

probability assessments from subjects without risk-neutral preferences (as explained in Smith (1961) and Savage (1971)). However, it is important to note the difference. In the Allais idea, as well as in the works mentioned, randomization is performed at an ex ante stage, in effect randomizing over several possible decision problems. In protocols that elicit beliefs from non-risk-neutral subjects, randomization is typically applied ex post, by paying subjects with "probability currency" in the language of Savage, that is, by issuing a lottery ticket whose probability of winning corresponds to the normalized payment awarded by one of the standard protocols for risk-neutral subjects. In effect, the lottery ticket "linearizes" the subject's utility.

[5]More generally, the expert could observe all sorts of private signals, from which he would be able to infer a private posterior belief. Insofar as the elicitor cares exclusively about X, the expert's posterior is the only payoff-relevant variable. Thus, for simplicity, we interpret the signal directly as the expert's posterior belief.

by a c.d.f. F over the interval $[0, 1]$ of possible signal values. Note that, given F, we can compute the prior belief $\overline{p}$ the expert holds at the outset about the likelihood that $X = 1$ by the law of iterated expectations: $\overline{p} = \mathbf{E}^F[P]$, where the expectation is taken under the c.d.f. F. In the remainder of this section, F is referred to as the *prior*, and P as the *posterior*.

The elicitor wants to motivate the expert to reveal his prior at the outset, and subsequently the posterior he observes at the interim period. Note the unconventional nature of the exercise. Here, we are not only interested in extracting probability estimates of a publicly observable random variable, we also seek to extract the expert's belief regarding information he is about to receive privately—plainly, we seek to extract the expert's information structure, summarized by the c.d.f. F. Following the literature, the individual must motivate the expert with strict incentives: the expert should be willing to respond truthfully at every period, and the truth should be the only best response.[6]

It is worth observing that F and P capture the only payoff-relevant information when solving a dynamic decision problem that spans over these three time periods (while as we remark below and in the Supplementary Appendix, Section SA.4, for most decision problems, observing $\overline{p}$ and then P is not enough). We may assume the elicitor is a decision maker relatively less informed and wants to learn F and P to solve such a problem. However, the decision problem the elicitor may confront is outside of the model. We may suppose that there is some decision problem for which F and P are relevant to her, such as to whether and when to invest in a risky project, but we abstract away from the details of the problem: the elicitation method we seek to construct must apply to *any* such decision problem. Importantly, we impose no particular restriction on F and P.

The elicitor delivers to the expert a payoff $\Pi(\widehat{F}, \widehat{p}, x)$, where $\widehat{F}$ is the expert's announced prior at the outset, $\widehat{p}$ is the expert's announced posterior at the interim period, and x is the realization of the random variable at the final period. We seek to construct Π such that the following two conditions are met.

(i) For every $\widehat{F}$ announced at $t = 0$, it is uniquely optimal to report the true

[6]As is the case for the more classical probability elicitation methods, strict incentives yield strictly convex value functions, which permit to leverage the setup in two main types of environments: in the presence of moral hazard, such as when the expert must acquire costly information, and in the presence of adverse selection, in a broad sense, such as when comparing the level of expertise of several experts of varying quality.

posterior at $t = 1$: for all posteriors $p, \widehat{p}$ with $\widehat{p} \neq p$,[7]

$$\mathbf{E}^p\left[\Pi(\widehat{F}, p, X)\right] > \mathbf{E}^p\left[\Pi(\widehat{F}, \widehat{p}, X)\right],$$

where X is distributed according to p.

(ii) It is uniquely optimal to report the true prior at $t = 0$. Denote by

$$V(\widehat{F}; p) = \sup_{\widehat{p}} \mathbf{E}^p\left[\Pi(\widehat{F}, \widehat{p}, X)\right]$$

the "value function" at $t = 1$, after he has reported $\widehat{F}$ and observed signal realization p. The condition of strict best response is that, for all priors $F, \widehat{F}$ with $\widehat{F} \neq F$,

$$\mathbf{E}^F\left[V(F; P)\right] > \mathbf{E}^F\left[V(\widehat{F}; P)\right], \tag{1}$$

where P is distributed according to F.

Inequality 1 illustrates the crucial difficulty that arises with private observations of signals. If signals were public, the elicitor would be free to choose any value function V to elicit F, for example the value function induced by the Matheson-Winkler elicitation method (Matheson and Winkler, 1976). However, that signals are private implies that V cannot be arbitrarily chosen. Indeed, V must correspond to a value function associated with the entire protocol, a significant restriction. Most importantly, $V(F; p)$ must be strictly convex in p. Classical methods, which do not worry about a second stage, are free to choose $V(F; p)$ arbitrarily, were p observable. This is one of the main technical hurdles in our construction.

To build intuition for our approach, we begin with a simplified environment in which there are only two possibilities for the expert's prior: F_1 or F_2, both fixed and known to the elicitor. For instance, F_1 can correspond to a case in which the expert expects to revise his initial probability assessment substantially, while under F_2 the expert remains relatively less informed.

Our construction uses securities and options. A *security* S is an outcome-contingent payoff. Here we consider the security whose value is the realization of the random variable, i.e., $S = X$. Let us call the option to short-sell the security at $t = 1$, and for a given price α, an α-*option*. Suppose the elicitor offers, at $t = 0$, the chance to purchase an α-option for a given price β. The only time the expert can sell the

[7] Alternatively, we could have a slightly weaker requirement that the incentive constraint (i) holds only after the expert has reported the truth at time 0. It would not change the results substantially.

security is therefore at $t = 1$, and only in the case he purchased the option. If he declines to sell at the interim period, the option expires and the expert earns no payoff.

When the expert's posterior is p, his expected value for the security is $\mathbf{E}^p[S] = p$. So, if the expert holds the option at $t = 1$, he is best off exercising the option when $\alpha > p$. He then earns an expected payoff $\alpha - p$. If $\alpha < p$, the expert is best off not exercising the option, and obtaining a zero payoff. Thus, his value of holding the option, at $t = 1$, is $\max(\alpha - p, 0)$. The value of holding the option, at $t = 0$, is then the average over the possible posteriors the expert may have,

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

Here, the last inequality follows from integration by parts.[8] As $F_1 \neq F_2$, there exists a selling price of the security, α, such that $\int_0^\alpha F_1(p)\, \mathrm{d}p \neq \int_0^\alpha F_2(p)\, \mathrm{d}p$. Say, for example, that $\int_0^\alpha F_1(p)\, \mathrm{d}p > \int_0^\alpha F_2(p)\, \mathrm{d}p$. Then the "type" F_1 values the option more than the "type" F_2. We can fix the option price between these two values, for example

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

Then, if the expert has prior F_1, he will strictly prefer to acquire the option at price β, but if he has prior F_2, he will strictly prefer not to. Thus, the prices α and β can be chosen so that we can distinguish between any two given information structures for the expert with a single choice—as long as we know which ones of the only two possible information structures the expert may have.

Let us return to the general case with no restrictions imposed on the prior. To distinguish between all prior distributions, we introduce randomization, in the spirit of Allais (1953) and Becker, DeGroot, Marschak (1964). In this particular example, we will select 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, the expert would always prefer to get the option no matter his

[8]The integration is over $(-\infty, +\infty)$ instead of $[0, 1]$ to avoid a minor technical difficulty of the Riemann-Stieltjes integral when there are mass points at the boundary.

prior information, which in turn motivates him to give his posterior as a strict best response—this objective is secondary, the primary goal is of course to extract the expert's information structure at the first period.)

This is a natural extension of the Allais idea to the dynamic setting. However, to preserve strict incentives, the expert should not observe which option he is facing, nor its price. Choices are therefore delegated to the elicitor. The expert announces a prior $\widehat{F}$ and, later, a posterior $\widehat{p}$. The elicitor randomly selects the option and its price, and, on the expert's behalf, makes the optimal decision to buy or not to buy the option, then to exercise or not to exercise the right to sell the security. The elicitor must never inform the expert of which decision she has made until all uncertainty about the random variable is resolved. The expert eventually gets the final value that results from the elicitor's choices.

For any given randomly selected option and price, specified by parameters α, β, the expert's final payoff is:[9]

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

As described, the dynamic counterpart of the Allais mechanism relies on the individual's ability to *commit* to making the best choices on the expert's behalf *without* the expert having the ability to verify the elicitor's actions. But from the expert's viewpoint, the randomized protocol reduces to a deterministic payoff that averages payoffs of the randomized protocol over α and β:

$$\Pi(\hat{F}, \hat{p}, x) = \mathbf{E}\left[\Pi(\hat{F}, \hat{p}, x; \alpha, \beta)\right],$$

where the expectation is taken over α, β.[10]

Hence, $\Pi(\hat{F}, \hat{p}, x)$ is a deterministic mechanism that we can use in place of the randomized protocol, and notably requires no commitment. Moreover, we can calculate an explicit formula for this mechanism. For a given α, and reports $\widehat{F}$, $\widehat{p}$,

[9] Here the elicitor breaks ties in favor of not purchasing the option (at the first period) and not selling the security (at the second period).

[10] Without assuming risk-neutrality and without even making assumptions on the utility function, one can continue to induce a strict and truthful best response from the expert by interpreting $\Pi(\hat{F}, \hat{p}, x)$ as the probability of winning a fixed prize, i.e., paying the expert in lottery tickets (see Savage, 1971). Rewards in "probability currency" is a common method to deal with risk aversion that applies broadly.

the expert gets on average over β's,

$$P\left(\beta < \int \max(\alpha - p, 0)\,\mathrm{d}\widehat{F}(p)\right)$$
$$\times\, \mathbf{E}\left[(\alpha - x)\mathbf{1}\left\{\alpha \leq \widehat{p}\right\} - \beta \,\middle|\, \beta < \int \max(\alpha - p, 0)\,\mathrm{d}\widehat{F}(p)\right]$$

which simplifies to

$$\frac{1}{4}\left[1 - \left(\int_0^\alpha \widehat{F}\right)^2\right] + \frac{1}{2}\left(1 + \int_0^\alpha \widehat{F}\right)(\alpha - x)\mathbf{1}\left\{\alpha \leq \widehat{p}\right\}.$$

Thus, averaging over α's and β's altogether, when announcing $\widehat{F}$ in the first period and $\widehat{q}$ in the interim period, the expert gets

$$\Pi(\hat{F}, \hat{p}, x) = \frac{1}{4} - \frac{1}{4}\int_0^1 \left(\textstyle\int_0^\alpha \widehat{F}\right)^2 \mathrm{d}\alpha + \frac{1}{2}\int_{\widehat{p}}^1 \left(1 + \textstyle\int_0^\alpha \widehat{F}\right)(\alpha - x)\,\mathrm{d}\alpha. \tag{2}$$

We refer to Π as the *two-stage quadratic scoring rule*. This scoring rule enables an elicitor to elicit an expert's belief over a binary event, and to elicit the expert's beliefs over his future beliefs. Note that it requires neither randomization nor commitment from the elicitor.

If the above derivation is heuristic, it is straightforward to establish that Π has the right incentive properties at every stage. Suppose the expert's true prior at time 0 is F and his observed signal at time 1 is p. We first claim that that it is a strict best response for the expert to report truthfully at $t = 1$, no matter his previous declarations. His expected payoff from announcing $\hat{p}$ given his information at $t = 1$ is:

$$\frac{1}{4} - \frac{1}{4}\int_0^1 \left(\textstyle\int_0^\alpha \hat{F}\right)^2 \mathrm{d}\alpha + \frac{1}{2}\int_{\hat{p}}^1 \left(1 + \textstyle\int_0^\alpha \hat{F}\right)(\alpha - p)\,\mathrm{d}\alpha.$$

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

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

We can also see that it is a best response for the expert to report truthfully at $t = 0$.

His expected payoff at $t=0$ is

$$\frac{1}{4}-\frac{1}{4}\int_0^1\left(\int_0^\alpha \hat{F}\right)^2 \mathrm{d}\alpha+\frac{1}{2}\int_0^1\int_p^1\left(1+\int_0^\alpha \hat{F}\right)(\alpha-p)\,\mathrm{d}\alpha\,\mathrm{d}F(p),$$

which simplifies to

$$\frac{1}{4}+\frac{1}{2}\int_0^1\left(\int_0^\alpha F\right)\mathrm{d}\alpha+\frac{1}{4}\int_0^1\left(\int_0^\alpha F\right)^2\mathrm{d}\alpha-\frac{1}{4}\int_0^1\left(\int_0^\alpha \hat{F}-\int_0^\alpha F\right)^2\mathrm{d}\alpha.$$

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

$$\frac{1}{4}\int_0^1\left(\int_0^\alpha \hat{F}-\int_0^\alpha F\right)^2>0.$$

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

Aside from the three fixed time periods, we make no particular assumption on the information structure: we elicit a fully general c.d.f. over posteriors. To facilitate communication, it may be of interest to assume a parameterized family of distributions—here, for example, truncated Gaussian or Beta distributions. Naturally, the above payoff function remains valid under such more restrictive class of information structures, and may be expressed more directly. To illustrate, let us consider the simplest case of the uniform distribution: the expert's prior over posterior is assumed to be uniform on $[\mu-\delta,\mu+\delta]$, for some parameters μ and δ. Although it is somewhat unwieldy, the quadratic scoring rule can be written directly as a function of the parameters of the distribution:

$$\begin{aligned}&\Pi(\mu,\delta,p,x)=\\&\quad-\frac{\delta^3}{80}+\frac{(\mu-p)^3(\mu+3p-4x)}{48\delta}+\frac{1}{8}\delta(\mu-p)(\mu+p-2x)+\frac{1}{12}\delta^2(\mu-x)\\&\frac{1}{12}\left(3\mu^3-3\mu^2(x+2)+3\mu\left(p^2-2px+4x\right)+p(3(p+4)x-2p(p+3))-18x+14\right).\end{aligned}$$

In summary, to elicit the expert's information structure, we ask the expert to solve a sample of elementary decision problems. To ensure the expert solves these problems simultaneously in single decision, we randomize over the decision problems, without ever telling the expert the actual problem he is actually dealing with. Instead, decisions are delegated to the elicitor. The expert provides his belief or information to the elicitor, committed to perform the best she can on behalf of the expert. The expert is naturally never worse off being truthful. Moreover, because there are a wide variety of problems in the sample, and as the randomization device spans over the entire sample, the expert is strictly best off telling the truth: if he were to lie, he would get suboptimal outcomes on a positive mass of problems. The approach is thus very different from the classical subgradient methods of probability elicitation.

We end this section with four remarks.

First, in spite of its apparent simplicity, the Allais randomization explained in this example is a powerful approach. It permits to elicit various types of information structures where the classical subgradient method fails. In Section SA.1 of the Supplementary Appendix, we apply the method on several other examples: we allow for a random time of signal arrival (in addition to the prior, we want to elicit the expert's belief about when new information is to arrive), we consider the problem of tracking a Brownian signal, and a case of an expert who observes new information at two different interim periods, as opposed to only one.

Second, the approach relies crucially on a careful selection of *simple decision problems* (here, the choice to buy or not buy α-options at a given price), and an appropriate randomization among them. Failure to select or randomize properly results in weak incentives, and thus the inability to distinguish between certain types of experts and observations. Calculating the appropriate classes and randomizations, and proving these have the right incentive properties, can be a difficult task even in simple cases. To illustrate, consider a slightly more general version of the example of this section: let us assume X is a categorical random variable, i.e., a variable that takes value in a finite set $\{1, \ldots, n\}$. The same type of option-based elicitation protocol continues to work, however it is not as simple: the incentive properties now rely on the fact that a distribution over a random vector is uniquely determined by the distributions of the linear combinations of its components, a fairly nontrivial result known as the Cramér-Wold Theorem (Cramér and Wold, 1936). In addition, if X takes values in a larger set, the Cramér-Wold Theorem no longer holds and the option-based elicitation protocol generally fails to properly elicit the expert's belief. As another example, let us keep X binary but instead add one more period: consider an expert who gets two observations at different times. One may attempt a direct extension of the above protocol, where instead of selling an α-option to the expert, we sell an option for the right to sell an α-options. However, such construction does not provide enough incentives. As we show in the next two sections, it is nevertheless possible to achieve the desired incentive properties under these last two cases by using more complex options.

Third, we note that there is a fundamental difference between eliciting dynamic beliefs as whole, and eliciting at different times up-to-date probabilistic assessments about X. If the latter is easily achieved using a probability scoring rule or the Becker-DeGroot-Marschak mechanism, it lacks crucial information about when and how uncertainty gets resolved over time, because it does not elicit the expert's assessment about his future private observations. Such dynamic information is generally valuable: we show in the Supplementary Appendix, Section SA.4, that the class of decision problems that can be resolved using probability estimates at different point in time consists in a degenerate subset of dynamic decision problems. The class essentially contains concatenated static problems.

Fourth and finally, one may wonder if an appropriate combination of classical probability elicitation methods can suffice to uncover dynamic information: it is tempting to use the public realization x to elicit the expert's posterior, and going backward, using the presumably truthful report of the expert's posterior, to elicit, in turn, the expert's prior. As it turns out, we show in the Supplementary Appendix, Section SA.2, that the expert is always induced to deviate in such protocols. More generally, we show that any two-stage procedure which pays the expert a sum $\Pi_1(\hat{F}, \hat{p}) + \Pi_2(\hat{p}, x)$ is unable to provide strict truthtelling incentives.

3 Main Results

3.1 Model

There are now finitely many fixed time periods $t_1 < t_2 < \cdots < t_N$.[11] As in the preceding section, there is an individual (a female), a risk-neutral expert (a male), and a random variable X that materializes publicly at the last period t_N. The random variable takes values in a compact metrizable space $\mathcal{X}$ (e.g., an interval $[a, b]$ or a finite set). The individual wants to learn dynamic information about X. To do so, she solicits the expert's opinion.

At the first period t_1, the expert holds a prior assessment about the distribution of X. Then, as time unravels, the expert may receive additional private information that refines his beliefs about future uncertainty. To be more precise, the expert is to observe a private random signal at every time $t_2, \ldots, t_{N-1}$. The signal takes the form of a "probability tree" or "belief tree". That is, the random signal S_{N-1} observed at

[11] The focus on a finite number of fixed periods permits to focus the discussion on the key features of the elicitation protocols, and abstract away from technical considerations. In Section 4, we extend our main result to more general information structures, which allows to track a continuous signal or to have random times of information arrivals.

the penultimate time t_{N-1} is a probability distribution over the random variable of interest, X; the random signal S_{N-2} observed at t_{N-2} is a probability distribution over S_{N-1}; and so forth, the random signal S_k observed at t_k is a probability distribution over S_{k+1}.

The expert learns his information structure privately at $t = t_1$. This information structure would typically consist of a joint distribution over the sequence of private signals and the public outcomes, but under our assumption that signals are probability trees, it reduces without loss to a distribution over the first private signal to be received at $t = t_2$, S_2. Hence, to simplify notation, we denote this distribution by an extra "signal" S_1.

Of course, such setup is not to be taken literally. In practice, experts may receive all sorts of information, and this information induces a probability tree, that captures the "dynamic belief" of the expert regarding X. Probability trees capture all relevant uncertainty about X and the future signals, and to the extent that the individual cares about X only, restricting attention to probability trees is without loss of generality. For example, if X is the only payoff-relevant variable for a decision maker endowed with an arbitrary dynamic signal structure, then the only payoff-relevant information is captured by the probability trees induced by this signal structure. [12] In that sense, one can think of probability trees as a universal signal space. In our formal results, it is convenient to focus on this summary statistic. [13]

We let $\Delta^1(\mathcal{X})$ be the set of distributions over X, and define the set of probability trees $\Delta^k(\mathcal{X})$ recursively, letting $\Delta^{k+1}(\mathcal{X}) = \Delta(\Delta^k(\mathcal{X}))$. We refer to an element of $\Delta^k(\mathcal{X})$ as a *probability or belief tree of level k*. We endow each $\Delta^k(\mathcal{X})$ with the weak-* topology and the usual Borel σ-algebra.[14] As we argue in the discussion of

[12] Say, for instance, that the individual is a bidder for the right to drill an oil parcel. The expert is an engineer who observes a rough estimate about the amount of oil in the parcel on Day 1, is about to get another estimate on Day 2, both to be used to assess the net-present-value of the parcel. The bidder is unable to interpret the engineer's measurements, but they are irrelevant to her. Instead, the payoff-relevant information is, on Day 2, the engineer's estimate of the distribution of future cash flows, and, on Day 1, the engineer's estimate about the distribution of the cash flow distribution to be assessed subsequently. This "level-2" probability tree captures, in particular, the degree to which the engineer expects to learn valuable information in his second estimate.

[13] As an alternative, one could to ask the expert to communicate, at the outset, the type of measurements he is about to take (i.e., the set of all possible signals), along with the joint probability measure over the future observations and X, and then the realizations that he observes, as he observes them. One difficulty of this approach is that the message of the expert may include information that is not about X, and consequently one cannot ensure that truthful communication is a "strict" best response. Our signal structure takes care of this issue by excluding all information that does not concern the random variable of interest that is publicly observed.

[14] The weak-* topology refers to the weakest topology for which, given any continuous function,

Section 3.5, the σ-algebra must be specified for the formal statements, but can be made transparent in the elicitation protocols. The space of possible signals to be realized at time t_k is $\mathcal{S}_k = \Delta^{N-k}(\mathcal{X})$. In the sequel, the letters p and q denote signal realizations, i.e., belief trees of any level. As the signal received at t_k is a tree of order $N-k$, to avoid confusion, we use the subscript notation p_k to denote a realization of the k-th signal, and the superscript notation $p^{(k)}$ to denote a tree of level k.

The expert communicates information to the elicitor and gets rewarded. The rules of interactions between expert and elicitor are described by the *elicitation protocol*. General protocols can involve randomization, and can extract information indirectly—for example, the expert may be asked to make decisions, and observing the expert's choices may then reveal part of his private belief. By a revelation principle argument, every protocol is payoff-equivalent to a *direct protocol*, whereby the expert reveals directly his signals and information structure. Such protocols are represented by a *payoff rule* $\Pi(p_1, \ldots, p_{N-1}, x)$ that captures the agent's final payoff as a function of the realized outcome x and the successive announcements, $p_k, k = 1, \ldots, N-1$, at every time t_k. Therefore, we may describe a protocol either in full form, or more concisely via its payoff rule, whichever we see as convenient. We require that Π be jointly measurable in its arguments, and we normalize payoffs to take values in $[0, 1]$.

Our objective is to produce a protocol that induces the expert, as a strict best response, to communicate the signals and information structure he observes, as soon as he observes them. Define an expert *strategy* as a family of maps $\{f_1, \ldots, f_{N-1}\}$, where $f_k(p_1, \ldots, p_k) \in \mathcal{S}_k$ gives the belief tree declared at time t_k as a function of the history of signals the expert has observed up to t_k (such definition rules out randomized strategies and dependence on other private information; it is, for our purpose, without loss). The time-t_k expected payoff of the expert, under strategy f, is then

$$\begin{aligned}
&U(p_1, \ldots, p_k; f) \\
&= \int_{\mathcal{S}_{k+1}} \cdots \int_{\mathcal{S}_{N-1}} \int_{\mathcal{X}} \Pi(f_1(p_1), \ldots, f_{N-1}(p_1, \ldots, p_{N-1}), x) \, \mathrm{d}p_{N-1}(x) \ldots \mathrm{d}p_k(p_{k+1}).
\end{aligned}$$

A strategy f is *optimal* for the history of signals $p_1, \ldots, p_k$ and protocol with payoff rule Π if the expert who follows strategy f after observing the sequence of signals $p_1, \ldots, p_k$ maximizes his payoff, no matter the strategy followed up to time t_k. Formally, for every pair of strategies (g, h), where $g = \{h_1, \ldots, h_{k-1}, f_k, \ldots, f_{N-1}\}$, we have

$$U(p_1, \ldots, p_k; g) \geq U(p_1, \ldots, p_k; f).$$

integration with respect to that function is a continuous linear functional.

Definition 1 *A protocol is strategyproof if*

- *For all histories, an optimal strategy exists.*
- *For all histories* $(p_1, \dots, p_k)$*, and all optimal strategies* f*,* $f_k(p_1, \dots, p_k) = p_k$*.*

3.2 Secret Randomization Protocols

Here we introduce a class of protocols that forms the basis of our elicitation method. Central to our protocols are three instruments: *securities, menus of securities, and menus of (sub-)menus.* A security is a continuous map $S : \mathcal{X} \to [0, 1]$. It gives a payoff for every possible realization of the random variable. 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 $k-1$. A menu of securities gives the obligation to its owner to pick one (and only one) security from the menu, at (or before) time t_{N-1}. A menu of order k gives the obligation to its owner to pick one (and only one) sub-menu among the collection it contains, at (or before) time t_{N-k}. Thus, an individual endowed with a menu of order k at time t_{N-k} makes $k-1$ choices at successive times $t_{N-k}, \dots, t_{N-1}$, to eventually end up with a single security. We work mostly with finite menus. 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 denote by $\mathcal{M}_k$ the collection of finite menus of order k.

The value of a menu to an individual depends crucially on the individual's information structure, here captured by belief trees. Let us denote by $\pi_k(M_k, p^{(k)})$ the expected value of menu of order k, M_k, at time t_k, to an individual who has just observed his $(N-k)$-th signal, a probability tree of level k, $p^{(k)}$. Recursively, we have

$$\pi_1(M_1, p^{(1)}) = \max_{S \in M_1} \int_{\mathcal{X}} S(x)\, \mathrm{d}p^{(1)}(x), \text{ and, if } k > 1,$$
$$\pi_k(M_k, p^{(k)}) = \max_{m_{k-1} \in M_k} \int_{\Delta^{k-1}(\mathcal{X})} \pi_{k-1}(m_{k-1}, p^{(k-1)})\, \mathrm{d}p^{(k)}(p^{(k-1)}).$$

Our protocols randomize over large collections of menus.[15] As a preliminary, the elicitor who administers the protocol draws a finite menu M_{N-1} of order $N-1$ at

[15] To ensure the randomization device is well-defined, we endow the set of securities and the set of all menus of a given order with the Borel σ-algebra, where the space of securities is given the usual sup-norm topology, and every space of menus is given the Hausdorff metric topology.

The Hausdorff metric is a standard way to measure distances between sets. If d is a metric on $\mathcal{X}$, random, according to a probability distribution ξ. The menu is known only to the elicitor, and is kept private. Then, at every time t_k, $k = 1, \ldots, N-2$, the elicitor asks the expert to reveal the signal he observes at that time. That signal is a belief tree of level $N-k$. She then chooses sub-menu of M_{N-k} that is best according to the expert's announcement: she selects a sub-menu $M_{N-k-1} \in M_{N-k}$ of highest expected value at the time t_k,

$$M_{N-k-1} \in \operatorname*{arg\,max}_{m_{N-k-1} \in M_{N-k}} \int_{\Delta^{N-k-1}(\mathcal{X})} \pi_{N-k-1}(m_{N-k-1}, p^{(N-k-1)}) \,\mathrm{d}p^{(N-k)}(p^{(N-k-1)}).$$

Finally, at the penultimate time t_{N-1}, the expert communicates a posterior distribution over X. The elicitor then offers a security taken from the last menu selected, M_1, of highest expected value according to the declared posterior.[16] We refer to such protocols as *Secret Randomization Protocols*.

For a given Secret Randomization Protocol with randomization device ξ, let $\Pi(p_1, \ldots, p_{N-1}, x; M)$ be the payoff to the expert who announces p_k at time t_k, when the realization of X is x, and if the elicitor draws menu $M \in \mathcal{M}_{N-1}$ at $t = t_1$. Then, the payoff rule of the overall protocol is expressed as

$$\Pi(p_1, \ldots, p_{N-1}, x; \xi) = \int_{\mathcal{M}_{N-1}} \Pi(p_1, \ldots, p_{N-1}, x; M) \,\mathrm{d}\xi(M). \tag{3}$$

3.3 Existence

We now turn to our first main result.

Theorem 1 *If a Secret Randomization Protocol randomizes according to a full-support distribution, then it is strategyproof.*

---

the Hausdorff metric on every $\mathcal{M}_k$ is defined recursively by

$$d(M, M') = \max\left\{\max_{m \in M} \min_{m' \in M'} d(m, m'), \max_{m' \in M'} \min_{m \in M} d(m, m')\right\} \text{ for } M, M' \in \mathcal{M}_k,$$

where m and m' denote securities when $k = 1$. Because 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 (Theorem 3.91, Aliprantis and Border, 2006).

[16] At every stage, if there is more than one sub-menu or one security that is optimal for the expert, the administrator selects a sub-menu uniformly at random among all optimal sub-menus. Selecting a sub-menu uniformly at random guarantees the measurability of the payoff rule. Alternatively, the expert could get an equal fraction of all optimal sub-menus, or he could get any optimal sub-menu according to a measurable selection. In the proof of Lemma 2 in Appendix A.1, we show that such a measurable selection is guaranteed to exist.

Thus, the protocols of the class just described are strategyproof, provided that the probability measure ξ has full support.[17]

The proof is in Appendix A.2, we provide a high-level intuition below. The reader not interested in the intuition may skip the remainder of this subsection.

To gain intuition, it is useful to think of Secret Randomization Protocols as a way to induce the expert to solve simultaneously a large number of "simple decision problems". In our case, these decision problems correspond to the menu choices that must be made at every time that precedes the realization of X.

If we were to give the expert a finite menu, and we observed his actions over time, we would be able to draw an inference regarding what he believes or has learned over time. Of course, it would only be partial inference: there are a finite number of choices at every time, but infinitely many possible information structures, or "probability trees." However, if we were to go back in time and start again but with different menus, we would be able to combine our inferences, and get increasingly better pictures of what the information structure looks like. Our approach is thus to ask the expert to solve many such simple decision problems, enough so as to recover the complete information structure.

To avoid issues of complementarity across tasks, and avoid asking the expert to deal with, not only a sequence, but uncountably many decision problems, we recourse to the secret randomization device. Because the expert does not know which one of the "simple decision problem" he confronts, he has to optimize over all these problems simultaneously.

Randomization would be useless if the expert knew which menu was drawn. So, decisions are delegated to the protocol administrator on the expert's behalf, and it is the expert's best interest to communicate enough information so that the administrator does well. In a risk-neutral framework, it is equivalent to offering the expert an infinitesimal fraction of every simple decision problem.

The class of simple decision problems must therefore be chosen so as to be able to infer the expert's private beliefs, and also, one must ensure the expert is motivated to reveal his beliefs as a *strict* best response. This creates a tradeoff. Adding problems to the class permits a better inference, but potentially reduces incentives: as the class of decision problems gets richer, the expert must get a smaller fraction of each decision problems, for payoffs to be finite (or equivalently, for probabilities to sum to one). If the class grows too large, there is the risk that incentives will vanish on some subset of those decision problems, thereby hindering belief extraction. Therefore, the

[17]Here, a full support distribution over menus of order k means that for every finite menu $M \in \mathcal{M}_k$ and every $\epsilon > 0$, the probability of drawing a menu at most ϵ-close to M is positive, with respect to the Hausdorff distance.

elicitor needs to randomize over a class of simple decision problems that is not too large, yet rich enough to reveal the expert's belief.

A difficulty arises as the belief can be complex. As we add more time periods, probability trees become richer objects, so that the collection of simple decision problems may need to grow very large, thereby compromising the expert's incentives. A key ingredient of the proof is to show that the class of finite menus is a rich enough class of "simple decision problems" to randomize upon. The cornerstone of the proof is a duality between the space of finite menus and the space of value functions associated with finite menus. We show that the space of value functions has the structure of a *boolean ring*, when endowed with the appropriate operations. A boolean ring is powerful structure that enables to construct approximation arguments, here used to claim that any real-valued mapping on belief trees can be approximated arbitrarily closely by the difference between two suitably chosen value functions on finite menus. A direct consequence is that finite menus can distinguish between any two probability trees—and thus that the class of finite menus is a suitable class of "simple decision problems".

To help us in this task, and in particular for the approximation argument to apply, we must control the amount of information that is effectively being carried in the probability trees. This amount of information is controlled by the collection of events, the σ-algebra of sets to which a probability is associated. Our choice of σ-algebra ensures that the amount of data needed to encode these probability trees remains manageable. Specifically, it ensures that a probability tree of level $k+1$ requires no more data to encode than a tree of level k, for $k \geq 2$, a key fact to keep the class of "simple decision problems" small enough.

Of course, one could retrieve even less information by choosing an even coarser σ-algebra, but we then run the risk of leaving out potentially useful information. As we show in the Supplementary Appendix, Section SA.3, our choice of σ-algebra captures just enough information so as to solve standard dynamic decision problems. It is worth noting that if the choice is important to specify the incentives properties of our protocols, it is transparent in the administration of the protocol itself.

To finish, we emphasize that the difficulty we face, in comparison to the simple example of Section 2, can only be attributed to a limited extent to the richest of the information structures we allow in this multi-period, multi-outcome setting. For example, *it does not facilitate the problem to restrict attention to the fully discrete case, in which we require a finite number of possible signal and outcome realizations.* This class of information structures induces probability trees with a finite number of branches at every level, and thus can be described as finite-dimensional objects, however the same class of "simple decision problems" must be used as in the general case, and the proof argument is similar.

3.4 Uniqueness

Here we describe the class of strategyproof protocols by the payoff functions they induce. There are two results. The first one is an exact characterization. It is best used as a test that checks whether a given protocol is strategyproof.[18] Because the proof is not constructive, it generally cannot be used to create strategyproof protocols. The second result is the other main result of this paper, and addresses this shortcoming. It argues that Secret Randomization Protocols are essentially unique: under regularity conditions, any protocol that is strategyproof is approximately payoff-equivalent to some Secret Randomization Protocol. Hence, there is no loss of generality in focusing on Secret Randomization Protocols.

We first need to extend notation. Given a protocol with payoff rule Π, with a slight abuse of notation we denote by $\Pi(p_1, \ldots, p_k)$ the value of the *truthful* expert at time t_k, as a function of the expert announcements up to time t_k. These are defined in a straightforward recursive fashion:

$$\Pi(p_1, \ldots, p_{N-1}) = \int_{\mathcal{X}} \Pi(p_1, \ldots, p_{N-1}, x) \, \mathrm{d}p_T(x),$$
$$\Pi(p_1, \ldots, p_{k-1}) = \int_{\Delta^{N-k}} \Pi(p_1, \ldots, p_{k-1}, p_k) \, \mathrm{d}p_{k-1}(p_k).$$

Having defined these value functions, the test to check if a protocol is strategyproof is a classic subgradient test, which generalizes the common characterization of the scoring rule literature (see, for example, McCarthy (1956) and Savage (1971)).

Proposition 1 *Given a protocol with payoff rule Π, the protocol is strategyproof if and only if the following conditions are satisfied:*

1. *For every $k \leq N-1$, and every $p_1, \ldots, p_{k-1}$, the map $G_k(p_k) \coloneqq \Pi(p_1, \ldots, p_k)$ is strictly convex, and the map $s_k(p_k, p_{k+1}) \coloneqq \Pi(p_1, \ldots, p_{k+1})$ is a subgradient of G at point p_k.*

2. *For every $p_1, \ldots, p_{N-2}$, the map $G_{N-1}(p_{N-1}) \coloneqq \Pi(p_1, \ldots, p_{N-1})$ is strictly convex, and the map $s_{N-1}(p_{N-1}, x) \coloneqq \Pi(p_1, \ldots, p_{N-1}, x)$ is a subgradient of G at point p_{N-1}.*

[18] It extends directly the "subgradient" characterization of the classical scoring rule literature.

The proof of Proposition 1 is in Appendix A.4. Next we present our second main result.

Theorem 2 *Consider a strategyproof protocol whose payoff rule* $\Pi(p_1, \ldots, p_{N-1}, x)$ *is jointly continuous. Then, for every* $\epsilon > 0$*, there exists a strategyproof Secret Randomization Protocol whose payoff rule* $\Pi'(p_1, \ldots, p_{N-1}, x)$ *satisfies*

$$|\Pi(p_1, \ldots, p_{N-1}, x) - \Pi'(p_1, \ldots, p_{N-1}, x)| < \epsilon$$

for all $p_1, \ldots, p_{N-1}, x$.

The proof of Theorem 2 is in Appendix A.5.

3.5 Discussion

On Commitment and Risk Neutrality. Secret Randomization Protocols use a randomization device which helps capture the central idea behind the construction of strategyproof protocols. As such, they require the elicitor's ability to commit to act in the expert's best interest. It is important to note, however, that this is not necessary. The payoff rule associated to each protocol prescribes a payoff for each sequence of expert announcements, and each realization of X. *It requires neither randomization, nor commitment.* The examples of Section 2 and Section SA.1 of the Supplementary Appendix supply closed form solutions of these payoff rules for various special cases. Risk neutrality is not required either, as one can, using the terms of Savage, pay the expert with "probability currency," which here means to give the expert a fixed monetary amount with the probability equal to the payoff value of any strategyproof protocol. This classical method enables to use classical mechanisms originally designed for risk-neutral experts on risk-averse experts, while keeping the strategyproof property.

About the Randomization Device. Although 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, consider $\mathcal{X} = \{1, \ldots, n\}$, let F_S be a full-support probability distribution over $[0, 1]^n$, then identified with the set of securities,[19]and F_N be a full-support probability distribution over the positive integers.

One can draw a menu of order $N - 1$ with full support according to a simple recursive procedure: to draw a menu of order $k \geq 2$ at random, one draws a random

[19] When $\mathcal{X}$ is infinite, one can take a countable dense subset of securities, which exists by Lemma 3.99 of Aliprantis and Border (2006), ordering them as $s_1, s_2, \ldots$, and giving probability 2^{-k} to s_k.

number of menus of order $k-1$ independently at random according to F_N; to draw a menu of order 1 at random, one draws a random number of securities at random (according to F_S) independently. This simple procedure generates a probability measure ξ with full support.

Large Outcome Spaces. Suppose X takes values on the real line. This space of possible outcomes fails the compactness condition, and so falls outside the framework of Theorem 1. But if we consider instead the transformed variable $Y = 1/(e^X + e^{-X})$, then Y falls into the range $[-1, 1]$, which is compact. Moreover Y conveys the same information as X. So, to elicit dynamic information about X, we can transform the expert's reports along with the realization of the random variable, and feed these as input to a protocol that elicits information about Y.

This is a classical example of "one-point compactification", but the method applies more broadly. If the space of possible outcomes is only assumed to be Polish (i.e., a separable and completely metrizable space) then we can "compactify" the space and apply Theorem 1. This allows for some useful generalizations. For example, outcomes can then take values in $\mathbb{R}^n$, and we can consider a continuity of outcomes revealed gradually over time, such as stock prices—since both $\mathbb{R}^n$ and the space of continuous real functions over a bounded time interval are Polish.

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

The key to construct such protocols is to work with Secret Randomization Protocols as above with a transformed variable $Y = f(X)$ instead, for some appropriate map f. A short proof of Corollary 1 is presented in Appendix A.3.

Indirect Protocols and Menu Protocols. Every deterministic (non-random) protocol can be interpreted and implemented directly through the instruments introduced at the beginning of this section. That is, at $t = 1$, the elicitor offers a (carefully designed) menu of order $N-1$ to the expert, and observes the expert's selection every time the expert "exercises the menu," i.e., every time the expert chooses one element of the menu. The expert's choices at every stage reveal, indirectly, information on his belief tree. *Eventually, the elicitor gets the same information as she would have obtained from the protocol described above, without requiring the expert to communicate explicitly a belief tree.*

In this case, however, *the menu must contain infinitely many elements carefully designed.* Such a menu, let us call it $M^\star(\Pi)$, can be constructed explicitly from the payoff rule Π of any strategyproof protocol, according to a simple recursive procedure.

For every $p_1, \ldots, p_{N-1}$ with $p_k \in \Delta^{N-k}(\mathcal{X})$, we define

$$\begin{aligned} M_k^{p_1,\ldots,p_{k-1}} &= \left\{ M_{k+1}^{p_1,\ldots,p_{k-1},q_k}; q_k \in \Delta^{N-k}(\mathcal{X}) \right\}, \\ M_{N-1}^{p_1,\ldots,p_{N-2}} &= \left\{ \Pi(p_1, \ldots, p_{N_2-1}, q_{N-1}, \cdot) \in \mathcal{C}(\mathcal{X}, \mathbb{R}); q_{N-1} \in \Delta(\mathcal{X}) \right\}, \end{aligned}$$

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 $M^\star(\Pi)$ is then defined as the menu M_1.

Costly Information. Strict incentives are useful in the presence of moral hazard. To illustrate, suppose the expert incurs a cost to acquire information. Consider the simplest case of Section 2. Recall there are three periods $t = 0, 1, 2$. There is a binary random variable whose outcome materializes at $t = 2$. The expert observes a private signal at $t = 1$, interpreted as a posterior assessment over X. His information structure at $t = 0$ reduces to a c.d.f. F that gives the prior distribution over the signal.

Let us look at the following variation. There exist n possible priors over signals, denoted $F_1, \ldots, F_n$—for example, different priors might correspond to different signal precisions.[20] Neither the elicitor nor the expert know which F_i is the correct prior, but they share a common belief that the prior F_i is correct with probability $q_i > 0$. So, ex ante, they both believe the expert posterior to be distributed according to $q_1 F_1 + \cdots + q_n F_n$.

Suppose the expert, by exerting effort at cost $C > 0$, is able to learn which F_i is the correct one. Can the elicitor use the protocol described in Section 2 to induce the expert to learn the true prior, and then report truthfully?

The answer is yes, and it owes exclusively to the strict incentives of the protocol. Let $V(F)$ be the expert value, at $t = 0$, of such a protocol, summarized by the payoff rule (2) of Section 2, when the expert's prior belief is F:

$$V(F) = \mathbf{E}^F \left[\mathbf{E}^P \Pi(F, P, X) \right],$$

where X drawn from P in the inside expectation, and P drawn from F in the outside expectation. That the protocol induces a unique best response implies that $V(F)$ is strictly convex. Hence, by Jensen's inequality,

$$\sum_i q_i V(F_i) > V\left(\sum_i q_i F_i \right).$$

[20]The assumption of a finite list of possible priors is for simplicity of exposition. The argument we sketch continues to hold with more general collections of priors.

The left-hand side is the expected payoff obtained by the expert after exerting the effort to learn the true prior, absent the cost of learning. The right-hand side is the expected payoff to the expert without learning. By appropriately scaling the payoffs of the protocol, the expert is best off paying the cost to acquire the information. (Individual rationality and/or limited liability constraints can be handled by shifting the payoffs by a fixed amount.)

Parameterized Information Structures. The information structures we work with are very general. One may consider useful specializations. Typically, one will want to choose a few parameters to describe the structure of information, as we do in Section SA.1 of the Supplementary Appendix. Naturally, in these cases, the Secret Randomization Protocols continue to be strategyproof, though the use of randomization devices that do not have full support can also yield strategyproofness, and occasionally lead to closed-form payoff rules.

4 Results for General Information Structures

In this section, we work with a more general class of information structures. The expert may now privately observe information dynamically and continuously over a unit interval of time. The uncertain outcome of interest materializes at $t = 1$. In particular, the setup allows to work with two common cases.

1. Discrete information arrival with random times: the expert is to receive a given number of signals over the unit time interval, but as opposed to the simpler setup of Section 3, the times of arrival are random. Both the distribution of arrival times and their realization are private information to the expert. The elicitor 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 tracks a continuous signal over time modeled as a stochastic process, for example, a price process modeled as a Brownian motion with varying drift and scale. The elicitor wants to know the expert's assessment of the signal process distribution—such as drifts and diffusions terms—and, at every instant, the up-to-date signal value.

4.1 Information Structures

Time is continuous and indexed by $t \in [0, 1]$. At the beginning of the time interval, and then during the time interval, the elicitor solicits the expert for information regarding an uncertain outcome that realizes publicly at $t = 1$. The outcome continues to take values in a compact metrizable space $\mathcal{X}$.

At $t = 0$, the expert learns his information structure privately. As in the discrete setup, it includes information on when and how the uncertainty on the outcome resolves over time. However, instead of working with belief and probability trees, we model dynamic information as filtered probability spaces. These are convenient to deal with information flow, and to handle the two special cases described above.

Definition 2 *An information structure is a tuple* $(\Omega, \mathbb{F}, \mathsf{P}, X)$ *in which:*

- Ω *is a (separable, metrizable) set of states of the world. The state of the world captures every aspect of the world that is relevant to the expert.*

- $\mathbb{F} = \{\mathcal{F}_t\}_{t \in [0,1]}$ *is a right-continuous filtration.*[21] $\mathcal{F}_t$ *captures all events that are known to be either true or false by the expert at time* t.

- P *is a full prior, described below.*

- $X : \Omega \mapsto \mathcal{X}$ *is a random variable that links (hidden) states of the world to (publicly observable) outcomes of interest.*

We do not assume that the individual consulting the expert knows the state space Ω, and she does not observe the state of the world. However, this elicitor knows the set of all possible outcomes $\mathcal{X}$, and observes the realization of the random variable X. Together with the expert's declarations, this data point is the only information she can use in the elicitation procedure.

In the sequel, $\Delta(\Omega)$ is the set of all probability measures on Ω, endowed with the weak-* topology. The filtration determines, in every state of the world, what the expert will know, and when he will know it. The prior determines the degree of uncertainty over the various states and events. Specifically, the full prior gives, at every time t, and in every state ω, the expert's posterior distribution P_t^ω over states of the world, accounting for all information available to the expert at that time and in that state. Note that the full prior includes information at every time, as opposed to prior information at $t = 0$ only. This is required to condition on future events in a consistent manner.[22]

[21] Right continuity of a filtration models the requirement that the expert 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, §1.2 of Karatzas and Shreve, 1991).

[22] It is well-known that a prior does not always generate consistent conditional probabilities,

Definition 3 *Given a state space Ω and a filtration $\mathbb{F}$, a full prior is a stochastic process*[23] $\mathsf{P} : (t, \omega) \mapsto \mathsf{P}_t^\omega$ *with values in $\Delta(\Omega)$, and such that*

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

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

The first condition is a technical requirement consistent with the fact that the expert learns information that is right-continuous (and every conditional probability of an event has a right-continuous version, for example, see Theorem 1.3.13, of 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 requiring consistency of the posteriors over time. Such a condition will not be needed for our results.

4.2 Payoffs, Strategies, and Strategyproofness

As opposed to the discrete setup, 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. Instead, we consider the more natural protocols in which the elicitor 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.

As before, every protocol yields a payoff (or an expected payoff, if the protocol randomizes) to the expert that can be written directly as a function of the expert declarations and the realized outcome of X. For convenience, payoffs are now specified by a family of payoff rules. There is one payoff rule for every information structure the expert might communicate. We encode the flow of posteriors by $\mathcal{Q}$, the space of maps $Q : t \mapsto Q_t$, from times to probability measures over states of the world, that are such that for every event E, $Q_t(E)$ is right-continuous in time.[24] A *payoff rule* $\Pi^{(\Omega,\mathbb{F},\mathsf{P},X)}$ maps $\mathcal{Q} \times \mathcal{X}$ to the interval of possible payoffs $[0, 1]$, it is required to be measurable.

---

and even when it does, these are generality not unique. Embedding such information in a full prior avoids the need to put overly restrictive assumptions on the state space, and avoids the problem induced by the multiplicity of conditional probabilities. Our definition of full prior is similar to the Conditional Probability Systems of Myerson (1991). For general conditions that ensure systematic existence and coherence of regular conditional probabilities, see Berti and Rigo (1996).

[23]The second condition implies by Lemma SA.2 of the Supplementary Appendix that the full prior is a well-defined stochastic process.

[24]The space $\mathcal{Q}$ is equipped with the product σ-algebra.

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

The expert announces, at the outset, an information structure, and, at every subsequent time, an up-to-date posterior over states of the world. Therefore, for an expert whose (true) information structure is $(\Omega^*, \mathbb{F}^*, \mathsf{P}^*, X^*)$, a *reporting strategy* consists in two objects:

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

The expert's *time-t value* at state ω^* is the average payoff he anticipates to receive given what he knows at time t. For an expert with a given information structure $(\Omega^*, \mathbb{F}^*, \mathsf{P}^*, X^*)$ who plays strategy $\langle(\Omega, \mathbb{F}, \mathsf{P}, X), \mathsf{Q}\rangle$, it is defined as

$$\mathsf{V}_t^{\omega^*} = \int_\Omega \Pi^{(\Omega,\mathbb{F},\mathsf{P},X)}\big(\{\mathsf{Q}_s^{\omega}\}_{s\in[0,1]}, X(\omega)\big)\, \mathrm{d}\mathsf{P}_t^{*,\omega^*}(\omega).$$

Experts seek to maximize their values at every time. A strategy $\langle(\Omega, \mathbb{F}, \mathsf{P}, X), \mathsf{Q}\rangle$ is *optimal* at $t = 0$ and state ω^* when, for every alternative strategy $\langle(\Omega', \mathbb{F}', \mathsf{P}', X'), \mathsf{Q}'\rangle$, 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 $\langle(\Omega, \mathbb{F}, \mathsf{P}, X), \mathsf{Q}\rangle$ is *optimal* at time $t > 0$ and state ω^* when for every alternative strategy $\langle(\Omega', \mathbb{F}', \mathsf{P}', X'), \mathsf{Q}'\rangle$ with $(\Omega', \mathbb{F}', \mathsf{P}', X') = (\Omega, \mathbb{F}, \mathsf{P}, X)$ and for all $\tau < t$, $\mathsf{Q}_\tau^{\omega^*} = \mathsf{Q}_\tau'^{\omega^*}$, 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.

Of course, 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 random variable X. Insofar as the individual consulting the expert is exclusively concerned about the outcome of X, the expert's information is valuable to the individual only to the extent that it impacts her beliefs on X. As in the baseline framework, the relevant information is captured by belief or probability trees.

However, because time is continuous, belief trees can now have any level, and be associated with an arbitrary sequence of intermediate times. For example, at $t = 0.5$, the expert who believes, given his own information at that time, that states of the world are distributed according to Q, can infer a distribution over the public outcomes—a probability tree of level 1. The expert may also receive information between $t = 0.5$ and $t = 1$, say at $t = 0.8$. Thus, at $t = 0.8$, his beliefs about the public outcome are to be updated. At $t = 0.5$, the expert anticipates the update, and forms a belief about the distribution he is about to infer at $t = 0.8$—a probability tree of level 2, with 0.8 as time of interim distribution.

More generally, fixing any finite sequence of times corresponding to interim updated beliefs, and given a posterior distribution over states of the world, we can define the belief tree induced by the expert information structure and the posterior. Formally, for information structure $(\Omega, \mathbb{F}, \mathsf{P}, X)$, the *induced belief tree of level 1* for posterior Q is defined as

$$\varphi(Q) = Q(X).^{25}$$

The *induced belief tree of level $k+1$ with intermediate times $t_1 < \cdots < t_k$*, is noted $\varphi_{t_1,\ldots,t_k} : \Delta(\Omega) \mapsto \Delta^{k+1}(\mathcal{X})$ and defined recursively as

$$\varphi_{t_1,\ldots,t_j}(Q) = Q(\varphi_{t_2,\ldots,t_j}(\mathsf{P}_{t_1}))$$

where P_{t_1} is the random variable of the process P sampled at time t_1. (By Lemma SA.3 of the Supplementary Appendix, the induced belief trees are well-defined and measurable.)

We continue to assume that the elicitor has interest in the probabilities that can be inferred from the expert's private information. Thus, at every time, the elicitor cares to learn about the expert's belief trees of all levels. A strategyproof protocol must therefore induce the expert, as a strict best response, to disclose enough information for the elicitor to learn the expert's belief trees of all levels and at all times. The expert is then induced to communicate all *relevant* information as a strict best response; however, as noted earlier, there will be many different information structures that will be equally relevant to the elicitor.

Definition 4 *A protocol is strategyproof when, for each expert information structure* $(\Omega^*, \mathbb{F}^*, \mathsf{P}^*, X^*)$*:*

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

$$\varphi_{t_1,\ldots,t_j}(\mathsf{Q}_{t_0}^{\omega^*}) = \varphi^*_{t_1,\ldots,t_j}(\mathsf{P}_{t_0}^{\omega^*})$$

[25] For a random variable X, we let $Q(X)$ denote the law of X induced by the probability measure Q.

where the left-hand side of the equality refers to the induced belief tree associated with the announced information structure, and the right-hand of the equality refers to the true information structure.

4.3 Temporal Menus

To deal with the richness of the expert's information structure, we use an extended menu instrument that we 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 fixed deadline. The owner of a temporal menu σ_0 must decide, at any $t \leq \tau_0$, to get one security among the collection M_0 (if $t > \tau_0$, the temporal menu is expired and delivers zero payoff). A temporal menu of securities is a *temporal menu of order 1.* Analogously to Section 3, we define in a recursive fashion 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. An individual who owns a temporal menu σ_k must select 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. Σ_k designates the collection of finite temporal menus of order k. For notational convenience, let Σ_0 be the space of securities, i.e., the continuous maps from Ω to $[0, 1]$.

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

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

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

$$\pi_1(\sigma_1, Q) = \int \left[\sup_{S \in M_1} \pi_0(S, \mathsf{P}^\omega_{\tau_1}) \right] \mathrm{d}Q(\omega),$$

$$\pi_k(\sigma_k, Q) = \int \left[\sup_{\sigma_{k-1} \in M_k} \pi_{k-1}(\sigma_{k-1}, \mathsf{P}^\omega_{\tau_k}) \right] \mathrm{d}Q(\omega).$$

4.4 A Class of Strategyproof Protocols

We now turn to the elicitation procedure. In a preliminary step, the protocol administrator draws a random temporal menu of a random order according to a simple procedure detailed below.[26] That menu is kept secret from the expert. The elicitor then asks the expert to provide his information structure and to send an update on his posterior at every instant before the outcome realization. Based on the expert's announcements, the elicitor makes decisions optimally on behalf of the expert, under the assumption that the expert reports truthfully, and without ever revealing her choices to the expert. Eventually, at $t = 1$, the expert owns a security issued from the last decision made by the elicitor, and is paid off accordingly.

The formal protocol is detailed below. Let ξ_K be a full support distribution on positive integers, $\xi_{\tau,k}$ be a full support distribution on $\{t_1, \ldots, t_k : 0 \leq t_1 < \cdots < t_k < 1\}$, and $\xi_{M,k}$ be a full support distribution on the set of finite menus of order k.

(a) Preliminary stage: the elicitor draws at random a finite number K from ξ_K. She then draws K deadlines $\tau_1, \ldots, \tau_K$ 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 $(\Omega, \mathbb{F}, \mathsf{P}, X)$. Then, at all subsequent times t, the expert communicates a posterior over states, $Q_t \in \Delta(\Omega)$.

(c) The elicitor's actions: every time a deadline is reached, i.e., $t = \tau_k$, the elicitor 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, $\pi_{k-1}(\sigma^*_{k-1}, Q_t) = \max_{\sigma_{k-1} \in M^*_k} \pi_{k-1}(\sigma_{k-1}, Q_t)$.

[26] The σ-algebra of events of finite temporal menus is defined analogously to that of the finite menus of the main framework of Section 3. 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 $d(M', M'')$ 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.

At the time of the last deadline, $t = \tau_K$, the elicitor 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 elicitor keeps all her actions secret until $t = 1$, when the outcome materializes and the expert is offered the security S^*.

The following theorem is the main result of this section.

Theorem 3 *The elicitation protocol described above is strategyproof.*

The proof is in the Supplementary Appendix. Remark that, as in the preceding sections, the protocol uses randomization and requires commitment from the elicitor, but is also equivalent to a deterministic protocol that averages payoffs over the finite menus drawn in the preliminary stage, and for which *commitment is not required.*[27]

5 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.

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

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

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

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

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

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

Aliprantis, C.D. and K.C. Border (2006): *Infinite Dimensional Analysis: A Hitchhiker's Guide.* Springer, third edition.

Allais, M. (1953): "Le Comportement de l'Homme Rationnel devant le Risque: Critique des Postulats et Axiomes de l'École Américaine," *Econometrica*, 21, 503–

546.

Anscombe, F.J. and R.J. Aumann (1963): "A Definition of Subjective Probability," *Annals of Mathematical Statistics*, 34, 199–205.

Azrieli, Y., C.P. Chambers and P.J. Healy (2014): "Incentives in Experiments with Objective Lotteries," Working paper.

—— (2016): "Incentives in Experiments: A Theoretical Analysis," *Journal of Political Economy* (forthcoming).

Becker, G.M., M.H. DeGroot and J. Marschak (1964): "Measuring Utility by a Single-Response Sequential Method," *Behavioral Science*, 9, 226–232.

Bergemann, D., S. Morris and S. Takahashi (2017): "Interdependent Preferences and Strategic Distinguishability," *Journal of Economic Theory*, 168, 329–371.

Berti, P. and P. Rigo (1996): "On the Existence of Inferences which are Consistent with a Given Model," *The Annals of Statistics*, 24, 1235–1249.

Blackwell, D. (1953): "Equivalent Comparisons of Experiments," *Annals of Mathematical Statistics*, 24, 265–272.

Brier, G.W. (1950): "Verification of Forecasts Expressed in Terms of Probability," *Monthly Weather Review*, 78, 1–3.

Cramér, H. and H. Wold (1936): "Some Theorems on Distribution Functions," *Journal of the London Mathematical Society*, 1, 290–294.

Dixit, A.K. and R.S. Pindyck (1994): *Investment under Uncertainty*. Princeton University Press.

Fang, F., M.B. Stinchcombe and A.B. Whinston (2010): "Proper Scoring Rules with Arbitrary Value Functions," *Journal of Mathematical Economics*, 46, 1200–1210.

Gneiting, T. and A.E. Raftery (2007): "Strictly Proper Scoring Rules, Prediction, and Estimation," *Journal of the American Statistical Association*, 102, 359–378.

Good, I.J. (1952): "Rational Decisions," *Journal of the Royal Statistical Society, Series B (Methodological)*, 14, 107–114.

Hewitt, E. and K.R. Stromberg (1997): *Real and Abstract Analysis*. Graduate Texts in Mathematics, Springer-Verlag.

Karatzas, I. and S.E. Shreve (1991): *Brownian Motion and Stochastic Calculus*. Springer, second edition.

Karni, E. (2009): "A Mechanism for Eliciting Probabilities," *Econometrica*, 77, 603–606.

——(2017): "A Mechanism for Eliciting Second-Order Beliefs and the Inclination to Choose," Working paper.

Kechris, A.S. (1995): *Classical Descriptive Set Theory.* Springer-Verlag.

Kreps, D.M. and E.L. Porteus (1978): "Temporal Resolution of Uncertainty and Dynamic Choice Theory," *Econometrica*, 46, 185–200.

Lambert, N.S. (2013): "Elicitation and Evaluation of Statistical Forecasts," Working paper.

Lu, J. (2016): "Random Choice and Private Information," *Econometrica*, 84, 1983–2027.

McCarthy, J. (1956): "Measures of the Value of Information" *Proceedings of the National Academy of Sciences*, 42, 654–655.

Matheson, J.E. and R.L. Winkler (1976): "Scoring Rules for Continuous Probability Distributions," *Management Science*, 22, 1087–1096.

Myerson, R.B. (1997):*Game Theory : Analysis of Conflict.* Cambridge, Mass: Harvard University Press.

Nyarko, Y. and A. Schotter (2002): "An Experimental Study of Belief Learning using Elicited Beliefs," *Econometrica*, 70, 971–1005.

Ostrovsky, M. (2012): "Information Aggregation in Dynamic Markets with Strategic Traders," *Econometrica*, 80, 2595–2647.

Pavan, A., I. Segal and J. Toikka (2014): "Dynamic Mechanism Design: A Myersonian Approach," *Econometrica*, 82, 601–653.

Savage, L.J. (1954): *The Foundations of Statistics.* Dover Publications.

—— (1971): "Elicitation of Personal Probabilities and Expectations," *Journal of the American Statistical Association*, 66, 783–801.

Schervish, M.J. (1989): "A General Method for Comparing Probability Assessors," *Annals of Statistics*, 17, 1856–1879.

Smith, C.A.B. (1961): "Consistency in Statistical Inference and Decision," *Journal of the Royal Statistical Society, Series B (Methodological)*, 23, 1–37.

Steiner, J., Stewart, C. and F. Matějka (2016): "Rational Inattention Dynamics: Inertia and Delay in Decision-Making," *Econometrica* (forthcoming).

Stewart, C. (2011): "Nonmanipulable Bayesian Testing," *Journal of Economic Theory*, 146, 2029–2041.

Thomson, W. (1979): "Eliciting Production Possibilities from a Well-Informed Manager," *Journal of Economic Theory*, 20, 360–380.

A Proofs of Section 3

A.1 Some Auxiliary Lemmas

We introduce some technical lemmas to show that the payoff rules and value functions associated with menus satisfy some regularity conditions, such as continuity and measurability. These are needed to enable the computation of expectations, and to allow the use of approximation arguments in the proof of Theorem 1.

In the sequel, to simplify notation, let $\pi_0(S, x)$ be the payoff associated to a security S when the outcome of X is x, let $\Delta^0(\mathcal{X})$ designate $\mathcal{X}$, and let $\mathcal{M}_0$ be the set of securities taking values in the normalized interval $[0, 1]$, instances of such securities will be denoted by S or M_0.

Lemma 1 *For every $k \geq 0$, the value map $(M_k, p^{(k)}) \mapsto \pi_k(M_k, p^{(k)})$ for menu $M_k \in \mathcal{M}_k$ and belief tree $p^{(k)} \in \Delta^k(\mathcal{X})$, is jointly continuous. In addition, the step-ahead value map $(M_k, p^{(k+1)}) \mapsto \int \pi_k(M_k, q)\, \mathrm{d}p^{(k+1)}(q)$, is also jointly continuous in $M_k \in \mathcal{M}_k$ and $p^{(k+1)} \in \Delta^{k+1}(\mathcal{X})$.*

Proof. The proof proceeds by induction.

Let $f_0(S, p^{(1)}) = \int S\, \mathrm{d}p^{(1)}$ and for $k \geq 1$ let $f_k(M_k, p^{(k+1)}) = \int \pi_k(M_k, q)\, \mathrm{d}p^{(k+1)}(q)$.

Note that π_0 is jointly continuous and f_0 is also jointly continuous, because securities have a compact domain and $\Delta^1(\mathcal{X})$ is endowed with the weak-* topology. Also, $S \mapsto \pi_0(S, \cdot)$ is continuous in the sup-norm topology.

We show that if f_k is jointly continuous, and if $M_k \mapsto \pi_k(M_k, \cdot)$ is continuous in the sup-norm topology, then both π_{k+1} and f_{k+1} are jointly continuous, and in addition $M_{k+1} \mapsto \pi_{k+1}(M_{k+1}, \cdot)$ is continuous in the sup-norm topology.

Let h_{k+1} be the correspondence from $\mathcal{M}_{k+1} \times \Delta^{k+1}(\mathcal{X})$ to $\mathcal{M}_k$ that is defined by $h_{k+1}(M_{k+1}, p^{(k+1)}) = M_{k+1}$. Because h_{k+1} has nonempty compact values and is continuous when interpreted as a map from $\mathcal{M}_{k+1} \times \Delta^{k+1}(\mathcal{X})$ to $\mathcal{M}_{k+1}$, the correspondence is continuous (Theorem 17.15 of Aliprantis and Border, 2006). Since

f_k is continuous, we can then invoke Berge's Maximum Theorem (see, for example, Theorem 17.31 of Aliprantis and Border, 2006) to get that the map

$$(M_{k+1}, p^{(k+1)}) \mapsto \max_{m \in h_{k+1}(M_{k+1}, p^{(k+1)})} f_k(m, p^{(k+1)})$$

is continuous. This proves the joint continuity of π_{k+1}. If, in addition, $M_{k+1} \mapsto \pi_{k+1}(M_{k+1}, \cdot)$ is continuous in the sup-norm topology, then f_{k+1} is jointly continuous (Corollary 15.7 of Aliprantis and Border, 2006).

What remains to be shown is the continuity of the maps $M_{k+1} \mapsto \pi_{k+1}(M_{k+1}, \cdot)$.

Let $\mathcal{C}_{k+1}$ be the space of continuous real functions on $\Delta^{k+1}(\mathcal{X})$ endowed with its sup-norm. Let $\mathcal{K}_{k+1}(M_{k+1}) \subset \mathcal{C}_{k+1}$ be the convex hull of $\{\pi_m; m \in M_{k+1}\}$, which, being the finite union of points, is closed and bounded in $\mathcal{C}_{k+1}$. Let $\mathcal{C}'_{k+1}$ be the norm dual of $\mathcal{C}_{k+1}$, which consists of all norm-continuous linear functionals. Let $\mathcal{U}_{k+1}$ be the closed unit ball of $\mathcal{C}_{k+1}$, and $\mathcal{U}'_{k+1} \subset \mathcal{C}'_{k+1}$ be its polar, so that $v \in \mathcal{U}'_{k+1}$ if $|v(x)| \leq 1$ for all $x \in \mathcal{U}_{k+1}$. For a given closed, bounded set C of $\mathcal{C}_{k+1}$, let h_C defined by $h_C(v) = \sup_{x \in C} v(x)$ denote its support function. Using the induction hypothesis, we remark that the map $M_{k+1} \mapsto \mathcal{K}_{k+1}(M_{k+1})$ is continuous, if the set of closed bounded subsets of $\mathcal{C}_{k+1}$ is given the Hausdorff metric induced by the sup-norm topology. Let us suppose that a sequence $\{M^{(i)} \in \mathcal{M}_{k+1}; i = 1, 2, \dots\}$ converges to some $M^\infty \in \mathcal{M}_{k+1}$. Then $\lim_{i\to\infty} \sup_{u' \in \mathcal{U}'} |h_{\mathcal{K}_{k+1}(M^{(i)})}(u') - h_{\mathcal{K}_{k+1}(M^\infty)}(u')| = 0$ by Lemma 7.58 of Aliprantis and Border (2006). By the Riesz-Radon representation (Corollary 14.15 of Aliprantis and Border, 2006), every $p^{(k+1)} \in \Delta^{k+1}(\mathcal{X})$ can be identified with a member of $\mathcal{U}'$, so that $\pi_{k+1}(M_{k+1}, \cdot)$ can be viewed as the support function of $\mathcal{K}_{k+1}(M_{k+1})$ restricted to $\Delta^{k+1}(\mathcal{X})$. Thus,

$$\lim_{i\to\infty} \sup_{p^{(k+1)} \in \Delta^{k+1}(\mathcal{X})} |\pi_{k+1}(M^{(i)}, p^{(k+1)}) - \pi_{k+1}(M^\infty, p^{(k+1)})| = 0,$$

which makes $M_{k+1} \mapsto \pi_{k+1}(M_{k+1}, \cdot)$ continuous. ∎

In the lemma below, we slightly generalize the notation introduced in Section 3. For any $k \geq 1$, and M a menu of order k, let $\Pi^k(p^{(k)}, \dots, p^{(1)}, x; M)$ denote the value of such a menu when $X = x$, for a risk-neutral individual with no discounting and who observes probability trees $p^{(k)} \in \Delta^k(\mathcal{X}), \dots, p^{(1)} \in \Delta^1(\mathcal{X})$ at the successive times of exercise of M and its submenus.

Lemma 2 *The map $(p^{(N-1)}, \dots, p^{(1)}, x, M_{N-1}) \mapsto \Pi^{(N-1)}(p^{(N-1)}, \dots, p^{(1)}, x; M_{N-1})$, where $p^{(k)} \in \Delta^k(\mathcal{X})$, $M_{N-1} \in \mathcal{M}_{N-1}$, and $x \in \mathcal{X}$, is jointly measurable in the product σ-algebra.*

Proof. As in Lemma 1, for every k we define the correspondence $h_k(M_k, p^{(k)}) = M_k$, and the function $f_k(M_k, p^{(k+1)}) = \int \pi_k(M_k, q)\, \mathrm{d}p^{(k+1)}(q)$.

For every k, we note that h_k is measurable (Theorem 18.10 of Aliprantis and Border, 2006), that h_k is a Carathéodory function, and that the space $\mathcal{M}_k$ is separable.[28] We can then apply the Measurable Selection Theorem (Theorem 18.19 of Aliprantis and Border, 2006), and we get that the argmax correspondence

$$\underset{m \in h_{k+1}(M_{k+1}, p^{(k+1)})}{\arg\max} \int_{\Delta^k(\mathcal{X})} \pi_k(m, q)\, \mathrm{d}p^{(k+1)}(q)$$

is measurable and admits a measurable selector. Moreover, by the Castaing Representation Theorem (Corollary 14.18 of Aliprantis and Border, 2006), we can enumerate the elements of the argmax in a measurable way, in the sense that there exists a sequence of measurable selectors $\{\Phi^{(i)}_{k+1}; i = 1, 2, \dots\}$ such that

$$\underset{m \in h_{k+1}(M_{k+1}, p^{(k+1)})}{\arg\max} f_k(m_k, p^{(k+1)}) = \{\Phi^{(i)}_{k+1}(M_{k+1}, p^{(k+1)}); i = 1, 2, \dots\}.$$

We observe that

$$\left| \underset{m \in M_{k+1}}{\arg\max} \int_{\Delta^k(\mathcal{X})} \pi_k(m, q)\, \mathrm{d}p^{(k+1)}(q) \right| = \lim_{j \to \infty} \sum_{i=1}^{j} \frac{1}{\sum_{\ell=1}^{j} \mathbf{1}_{\Phi^{(i)}_{k+1}(M_{k+1}, p^{(k+1)}) = \Phi^{(\ell)}_{k+1}(M_{k+1}, p^{(k+1)})}}$$

is measurable as a pointwise limit of real-valued measurable functions.

The remainder of the proof continues with a brief induction argument. Note that Π^0 defined by $\Pi^0(x; S) = S(x)$ is measurable. Suppose that Π^{k+1} is measurable. Then Π^k, which can be written

$$\Pi^{k+1}(p^{(k+1)}, \dots, p^{(1)}, x; M_{k+1}) =$$
$$\frac{1}{\left| \arg\max_{m \in M_{k+1}} \int \pi_k(m, \cdot)\, \mathrm{d}p^{(k+1)} \right|} \lim_{j \to \infty} \sum_{i=1}^{j} \frac{\Pi^k(p^{(k)}, \dots, p^{(1)}, x; \Phi^{(i)}_{k+1}(M_{k+1}, p^{(k+1)}))}{\sum_{\ell=1}^{j} \mathbf{1}_{\Phi^{(i)}_{k+1}(M_{k+1}, p^{(k+1)}) = \Phi^{(\ell)}_{k+1}(M_{k+1}, p^{(k+1)})}}.$$

becomes measurable. This concludes the proof. ■

[28] First, note that the set of securities is a separable metric space, by Lemma 3.99 of Aliprantis and Border (2006). 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.

A.2 Proof of Theorem 1

The proof consists of two parts. The first part deals with the separation of different experts *at a given time* when there are *only two possible types*, as in our example of Section 2. In the multi-period case, we use as class of "simple decision problems" the class of finite menus. Decisions then consist in choosing an element from the menu at the initial time, then an element from the chosen submenu at the next time, and so forth until the penultimate time when the decision reduces to choosing among a set of securities from the submenu chosen last. In the first part of the proof, we show that this class of decision problems is rich enough to discriminate between any two experts whose belief trees are of two possible sorts. In the second part of the proof, we apply Allais' randomization idea to discriminate between any two experts whose belief trees are no longer restricted.

Part 1: Discriminating Between Two Belief Trees

Let $p^{(k)}$ and $q^{(k)}$ be two different belief trees of level k, that represent two different expert beliefs at time t_{N-k}. We refer to the expert with belief tree $p^{(k)}$ as *type* $p^{(k)}$, and the expert with belief tree $q^{(k)}$ as *type* $q^{(k)}$.

In this first part, we show that there exists a menu M_k^{pq} of level k with two different submenus M_{k-1}^{p} and M_{k-1}^{q} such that if offered M_k^{pq} at time t_{N-k}, type $p^{(k)}$ is strictly better off choosing submenu M_{k-1}^{p} while type $q^{(k)}$ is strictly better off choosing submenu M_{k-1}^{q}.

To understand the proof, it is helpful to start from the penultimate time t_{N-1}, in which case the belief trees have level $k = 1$ and simply represent outcome distributions. The problem aforementioned reduces to choosing two securities S^p and S^q such that type $p^{(1)}$ strictly prefers S^p and type $q^{(1)}$ strictly prefers S^q. It is easy to achieve when observing that, because $p^{(1)} \neq q^{(1)}$, at least one continuous mapping $f : \mathcal{X} \to [0, 1]$ exists that separates $p^{(1)}$ from $q^{(1)}$, in the sense that the expected payoff from f, when interpreted as a security, is different for the two types:

$$\int f \,\mathrm{d}p^{(1)} \neq \int f \,\mathrm{d}q^{(1)}.$$

It is immediate for the case of finite outcome spaces, and more generally holds for metrizable spaces by Aleksandrov's Theorem (Theorem 15.1 of Aliprantis and Border, 2006). Because $\mathcal{X}$ is compact, we can choose f to be bounded. For example, suppose $\int f \,\mathrm{d}p^{(1)}$ is greater than $\int f \,\mathrm{d}q^{(1)}$. Then we can set $S^p = f$ and S^q to be the average of $\int f \,\mathrm{d}p^{(1)}$ and $\int f \,\mathrm{d}q^{(1)}$. A symmetric argument holds if $\int f \,\mathrm{d}p^{(1)}$ is less than $\int f \,\mathrm{d}q^{(1)}$.

For this argument to work, the key element is to have essentially complete flexibility in the design of the security—which is also the expert's value function at the next and final time t_N.

Now consider the problem of separating experts with with different belief trees of some higher level, and so at some earlier time. To do so, for any $k \geq 1$ and any belief tree $\mu^{(k)}$ of level k, with a slight abuse of notation, let $\pi_{M_k}(\mu^{(k)})$ to be the value of menu $M_k \in \mathcal{M}_k$ at time t_{N-k} to any expert who holds belief tree $\mu^{(k)}$ at that time (that is, $\pi_{M_k}(\mu^{(k)}) = \pi_k(M_k, \mu^{(k)})$ as defined in Section 3).

Thus, for $k > 1$, we seek to design submenus M^p_{k-1}, M^q_{k-1} such that type $p^{(k)}$ strictly prefers $M^p_{k-1}$and type $q^{(k)}$ strictly prefers M^q_{k-1}. Note that the expected payoff for any type μ who chooses submenu M^{pq}_{k-1} at time t_{N-k} is the expectation of the value function at the next time period,

$$\int \pi_{M_{k-1}} \, \mathrm{d}\mu.$$

If we can choose the value functions arbitrarily then the argument of the case $k = 1$ continues to apply. However with $k > 1$ the value functions can no longer be chosen arbitrarily, for $k = 2$ they are the space of strictly convex functions over probability distributions, and as k increases they become a increasingly smaller subset of strictly convex functions whose domain is the growing space of belief trees of level $k - 1$.

Nevertheless, and perhaps surprisingly, the space of value functions is rich enough so that the difference between two value functions can approximate arbitrarily closely any continuous function on $\Delta^{k-1}(\mathcal{X})$. We can then apply a similar argument as for the case $k = 1$ to prove type separation for $k > 1$. The proof relies on a duality between the space of menus and the space of value functions, whereby the set of value functions is shown to have the structure of a Boolean ring, which in turn enables the application of a version of the Stone-Weierstrass Theorem for these algebraic structures. We state and prove the result in the following lemma.

Lemma 3 *For every $k \in \{1, \ldots, N-1\}$, $p^{(k)}, q^{(k)} \in \Delta^k(\mathcal{X})$ with $p^{(k)} \neq q^{(k)}$, there exists $M_{k-1} \in \mathcal{M}_{k-1}$ (M_{k-1} is a security if $k = 1$) such that*

$$\int \pi_{M_{k-1}} \, \mathrm{d}p^{(k)} \neq \int \pi_{M_{k-1}} \, \mathrm{d}q^{(k)}. \tag{4}$$

Proof. The proof proceeds by induction. As shown above, Equation (4) is satisfied for $k = 1$ and some security $M_0 = S$. Now let us assume that the statement of the lemma is valid for k, and show it is then valid for $k + 1$.

Step 1. We begin with two direct implications. First, there exist M_{k-1}^p and M_{k-1}^q, both elements of $\mathcal{M}_{k-1}$, such that when type $p^{(k)}$ is offered $M_k^{pq} := \{M_{k-1}^p, M_{k-1}^q\}$ at time period t_{N-k}, he is strictly better off choosing M_{k-1}^p while type $q^{(k)}$ is strictly better off with M_{k-1}^q. The construction is analogous to the case $k = 1$. If, for example,

$$\int \pi_{M_{k-1}} \, \mathrm{d}p^{(k)} > \int \pi_{M_{k-1}} \, \mathrm{d}q^{(k)},$$

we set $M_{k-1}^p = M_{k-1}$ and $M_{k-1}^q = \frac{1}{2}\left(\int \pi_{M_{k-1}} \, \mathrm{d}p^{(k)} + \int \pi_{M_{k-1}} \, \mathrm{d}q^{(k)}\right)$. Second, if M_{k-1}^p and M_{k-1}^q are chosen as such, we note that the value of M_k^{pq} is different for the two types: $\pi_{M_k^{pq}}(p^{(k)}) \neq \pi_{M_k^{pq}}(q^{(k)})$.

Step 2. Let $\mathcal{B}_k$ be the set of continuous and bounded real functions on $\Delta^k(\mathcal{X})$. We endow $\mathcal{B}_t$ with the topology of uniform convergence. Also 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, Theorem 15.11 of Aliprantis and Border (2006)). It follows that every $\Delta^k(\mathcal{X})$ is a compact metrizable space.

Let $\mathcal{L}_k = \{\pi_{M_k} - \pi_{M_k'}, M_k, M_k' \in \mathcal{M}_k\}$. Note that $\mathcal{L}_k$ is a subset of $\mathcal{B}_k$. We show below that $\mathcal{L}_k$ is a boolean ring for the operations "plus" and "max", in the sense that (a) $0 \in \mathcal{L}_k$, and (b) if $f, g \in \mathcal{L}_k$ then $f + g \in \mathcal{L}_k$ and $\max\{f, g\} \in \mathcal{L}_k$.

To do so, it is useful to endow recursively every set of menus $\mathcal{M}_\ell$ with the following operations:

- Minkowski Addition: for any $M, M' \in \mathcal{M}_1$, we define the menu $M + M' \in \mathcal{M}_1$ by $\{S + S'; S \in M, S' \in M'\}$; if $\ell > 1$ and $M, M' \in \mathcal{M}_\ell$, we define recursively $M + M' = \{m + m'; m \in M, m' \in M'\}$.

- Scalar multiplication: for any $\alpha \geq 0$, and for any $M \in \mathcal{M}_1$, we define $\alpha M = \{\alpha S; S \in M\}$; if $\ell > 1$, and $M \in \mathcal{M}_\ell$, we define recursively $\alpha M = \{\alpha m; m \in M\}$.

Let $\mathbf{1} \in \mathcal{M}_\mathbf{k}$ be the (degenerate) menu that generate the constant payoff 1, and $\mathbf{0} \in \mathcal{M}_\mathbf{k}$ be the (degenerate) menu that generate the constant payoff 0. The following equalities hold for each $\mu \in \Delta^k(\mathcal{X})$ and each $M, M' \in \mathcal{M}_k$, :

$$\begin{aligned}
\pi_{\mathbf{0}}(\mu) &= 0, \\
\pi_{\mathbf{1}}(\mu) &= 1, \\
\pi_{M+M'}(\mu) &= \pi_M(\mu) + \pi_{M'}(\mu), \\
\pi_{\alpha M}(\mu) &= \alpha \pi_M(\mu) \ \forall \alpha \geq 0, \\
\pi_{M \cup M'}(\mu) &= \max\{\pi_M(\mu), \pi_{M'}(\mu)\}.
\end{aligned}$$

Thus, $0 \in \mathcal{L}_k$. In addition, for each $\alpha \geq 0$,

$$\alpha(\pi_M - \pi_{M'}) = \pi_{\alpha M} - \pi_{\alpha M'}.$$

Finally, observe that, for M, M', N, N' menus of level k,

$$(\pi_M - \pi_{M'}) + (\pi_N - \pi_{N'}) = \pi_{M+N} - \pi_{M'+N'}$$

and

$$\begin{aligned}
\max\{\pi_M - \pi_{M'}, \pi_N - \pi_{N'}\} &= \max\{\pi_M + \pi_{N'}, \pi_N + \pi_{M'}\} - (\pi_{M'} + \pi_{N'}) && (5) \\
&= \pi_{(\pi_M + \pi_{N'}) \cup (\pi_N + \pi_{M'})} - (\pi_{M'} + \pi_{N'}). && (6)
\end{aligned}$$

In summary, the following conditions are satisfied:

1. $\mathcal{L}_k$ is a boolean ring.
2. $\mathcal{L}_k$ includes the constant function 1, since $1 = \pi_{\mathbf{1}} - \pi_{\mathbf{0}}$.
3. $\mathcal{L}_k$ is stable by scaling: $\alpha \mathcal{L}_k \subseteq \mathcal{L}_k$ for any $\alpha \in \mathbb{R}$.[29]
4. $\Delta^k(\mathcal{X})$ is a compact Hausdorff space.
5. $\mathcal{L}_k$ separates points in the sense that if $f(p) = f(q)$ for every $f \in \mathcal{L}_k$ then $p = q$. It is a direct consequence of the second implication in Step 1 of the proof.

Therefore, we can apply the version of the Stone-Weirstrass Theorem for Boolean rings described in Theorem 7.29 of Hewitt and Stromberg (1997), which implies that $\mathcal{L}_k$ is dense in $\mathcal{B}_k$ in the topology of uniform convergence.

We end the proof by contradiction. If, for every $M_k \in \mathcal{M}_k$, it is the case that

$$\int \pi_{M_k} \, \mathrm{d}p^{(k+1)} = \int \pi_{M_k} \, \mathrm{d}q^{(k+1)}$$

[29] By the boolean ring property $\alpha \mathcal{L}_k \subseteq \mathcal{L}_k$ if $\alpha \geq 0$, and by definition of $\mathcal{L}_k$, $-\mathcal{L}_k \subseteq \mathcal{L}_k$.

then for every $f \in \mathcal{L}_k$,

$$\int f \, \mathrm{d}p^{(k+1)} = \int f \, \mathrm{d}q^{(k+1)}$$

and by application of the Stone-Weirstrass Theorem, the equality remains true for every $f \in \mathcal{B}_k$. That $\Delta^k(\mathcal{X})$ is metrizable implies $p^{(k)} = q^{(k)}$ by Aleksandrov's Theorem.Thus, there exists a menu M_k of level k such that

$$\int \pi_{M_k} \, \mathrm{d}p^{(k+1)} \neq \int \pi_{M_k} \, \mathrm{d}q^{(k+1)},$$

which concludes the proof by induction. ■

Part 2: Randomization

In this second part, we show that full-support randomization over finite menus allows to distinguish between any two experts whose belief trees differ at some point in time, without restriction on the belief trees.

Formally, let us fix a full-support distribution ξ over the set of level-$(N-1)$ menus $\mathcal{M}_{N-1}$. Fix any two sequences of belief trees $\mathbf{p} = \{p^{(N-1)}, \dots, p^{(1)}\}$ and $\mathbf{q} = \{q^{(N-1)}, \dots, q^{(1)}\}$ with $\mathbf{p} \neq \mathbf{q}$ (recall the superscript (k) denotes a tree of level k). Proving Theorem 1 reduces to proving the following statement: with positive probability relative to the menu M_{N-1} drawn at random according ξ, the expert who is given menu M_{N-1} at the outset and observes the unraveling sequence of belief trees $\mathbf{p}$ over time is strictly better off making at least one decision different from all optimal decisions of the expert who observes the sequence of belief trees $\mathbf{q}$. We refer to the expert of observes $\mathbf{p}$ as type $\mathbf{p}$, and the expert of observes $\mathbf{q}$ as type $\mathbf{q}$.

Fix an arbitrary level k such that $p^{(k)} \neq q^{(k)}$, and let $M_k^\star = \{M_{k-1}^{p,\star}, M_{k-1}^{q,\star}\}$ be a menu of level k that separates between the two belief trees $p^{(k)}$ and $q^{(k)}$, and whose existence is shown in Part 1 of this proof. We abuse notation in that if $k = 1$, then $M_{k-1}^{p,\star}$ and $M_{k-1}^{q,\star}$ denote securities. Define the (degenerate) menu of level N, $M_N^\star$, which includes only $M_k^\star$, i.e., either $M_N^\star = M_k^\star$ if $k = N$, otherwise $M_N^\star = \{\dots\{M_k^\star\}\dots\}$. For such a menu, there is no decision to be made until time t_{N-k} when the decision maker must choose between either M_k^p or M_k^q.

Because of the full support assumption, to prove the above statement, it is sufficient to show that for any menu M_{N-1} selected anywhere in small enough neighborhood of $M_{N-1}^\star$, type $\mathbf{p} := (p^{(N-1)}, \dots, p^{(1)})$ is strictly better off choosing a different submenu/security than type $\mathbf{q} := (q^{(N-1)}, \dots, q^{(1)})$, for every optimal selection of type $\mathbf{q}$.

By Step 1 and Lemma 3, there exists $\epsilon > 0$ such that for any level-k menus

M_k, M'_k with $d(M_k, M_k^{p,\star}) < \epsilon$ and $d(M'_k, M_k^{q,\star}) < \epsilon$, type $p^{(k)}$ would be strictly better off choosing M_k over M'_k at t_{N-k}, while type $q^{(k)}$ would be strictly better off choosing M'_k over M_k.

Consider any menu M_{N-1} of level $N-1$ such that $d(M_{N-1}, M_{N-1}^\star) < \epsilon$. In this case, by a direct induction argument, every one of the submenus, subsubmenus, etc. of M_{N-1} of level $k-1$ (or securities if $k=1$) is either ϵ-close to $M_{k-1}^{p,\star}$ or, it is ϵ-close to $M_{k-1}^{q,\star}$; moreover, the use of the Hausdorff distance also implies that in every submenu of level k of M_{N-1}, there is at least one submenu closest to $M_{k-1}^{p,\star}$ and another submenu closest to $M_{k-1}^{q,\star}$. Thus the decisions that are optimal for type $\mathbf{p}$ at time t_{N-k} are strictly suboptimal for type $\mathbf{q}$ and inversely.

A.3 Proof of Corollary 1

By the Urysohn Metrization Theorem (Theorem 3.40 of Aliprantis and Border, 2006), $\mathcal{X}$ can be embedded in the Hilbert cube, denoted $\mathcal{X}^*$. Since the Hilbert cube is compact metrizable—and thus also Polish—and by Theorem 1, we get that any Secret Randomization Protocol using $\mathcal{X}^*$ as outcome space and with full support randomization is strategyproof.

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

Now, let Π be a payoff rule of a strategyproof protocol for the space $\mathcal{X}^*$, whose existence is guaranteed by Theorem 1. Let $\Pi_{\mathcal{X}}$ be defined by $\Pi_{\mathcal{X}}(q_1, \ldots, q_T, x) = \Pi(\iota_T(q_1), \ldots, \iota_1(q_T), \iota_0(x))$. Note that $\Pi_{\mathcal{X}}$ induces the same incentives as Π.

A.4 Proof of Proposition 1

Fix $k \leq N-1$. When the expert participates in a strategyproof protocol, $G_k(p_k)$ can be thought of as the time-t_k value, written as a function of the expert's true time-t_k posterior p_k. If we consider only one-step deviations from the truth at time t_k—in the sense that the expert always tells the truth before and after time t_k, but possibly not at time t_k—then the strict convexity of G_k becomes necessary and sufficient for a strict best response. Additionally, it must be the case that the time-t_{k+1} value to the expert (who is truthful from time t_{k+1} onwards) is a subgradient of G_k.

The arguments are standard 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.

A.5 Proof of Theorem 2

The proof is decomposed in two steps. First, we approximate the payoff rule Π by a payoff rule associated with a finite menu. Since finite menus only uncover beliefs partially, 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 Secret Randomization Protocol.

The main difficulty lies in the construction of the finite menu. This menu is obtained by sampling the original payoff rule, finitely many times in such a way that, whenever a selection needs to be made from that finite menu or one its sub-menus, the payoffs associated with that choice remain close to the payoffs of the original payoff rule.

For any finite menu M of order $N-1$, let $\Pi^\star(p_1, \ldots, p_{N-1}, x; M)$ be the induced payoff rule, and let $\Pi^\star(p_1, \ldots, p_{N-1}, x; \xi)$ be the payoff rule induced by the Secret Randomization Protocol that randomizes according to ξ. Let us slightly abuse notation and denote by

$$\Pi_k(q_1, \ldots, q_k; p_k)$$

the maximum expected value, at time t_k, of the expert who faces payoff rule Π and who reports $q_1, \ldots, q_k$ from time t_1 to time t_k, but receives signal p_k at t_k. Similarly,

$$\Pi^\star_k(q_1, \ldots, q_k; p_k; M)$$

is the maximum expected value of the expert endowed with the finite menu M at time t_1 instead. Let $d(\cdot, \cdot)$ denote a compatible metric on each space $\Delta^k(\mathcal{X})$—for example, the Lévy-Prokhorov metric.

Fix $\epsilon > 0$. Because Π is continuous on $\Delta^{N-1}(\mathcal{X}) \times \cdots \times \Delta(\mathcal{X}) \times \mathcal{X}$, which is a compact set, it is uniformly continuous. Thus, there exists $\delta_1 > 0$ such that if, for each i, p_i is δ_1-close to p'_i, i.e., $d(p_i, p'_i) < \delta_1$, then $|\Pi(p_1, \ldots, p_{N-1}, x) - \Pi(p'_1, \ldots, p'_{N-1}, x)| < \epsilon/2$ for each $x \in \mathcal{X}$.

Step 1(a). We show that there exists a finite subset Σ_1 of $\Delta^{N-1}(\mathcal{X})$ such that, for each 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 $\{\Sigma_{1,k}\}_k$ be a sequence of finite subsets of $\Delta^{N-1}(\mathcal{X})$ such that $\Sigma_{1,k}$ converges to $\Delta^{N-1}(\mathcal{X})$ in the Hausdorff metric topology induced by the Lévy-Prokhorov metric. The compactness of $\Delta^{N-1}(\mathcal{X})$ guarantees existence of such a sequence. We observe that $(q_1, p_1) \mapsto \Pi_1(q_1; p_1)$ is continuous—as can be seen immediately via induction, using that every Π_k is uniformly continuous. The correspondence $(\mathcal{P}, p_1) \twoheadrightarrow \mathcal{P}$, where $\mathcal{P}$ is a compact subset of $\Delta^{N-1}(\mathcal{X})$ and $p_1 \in \Delta^{N-1}(\mathcal{X})$ is also continuous (see Theorem 18.10 of Aliprantis and Border, 2006). Using Berge's Maximum Theorem, we get that the correspondence

$$(\mathcal{P}, p_1) \twoheadrightarrow \arg\max_{q_1 \in \mathcal{P}} \Pi_1(q_1; p_1)$$

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 $\Delta^{N-1}(\mathcal{X})$ is compact, there exists a subsequence of indexes, $\{\sigma(k)\}_k$, such that $p_1^{\sigma(k)}$ converges to p_1^∞ for some p_1^∞. Also, $\Sigma_{1,\sigma(k)}$ converges to $\Delta^{N-1}(\mathcal{X})$, where the limit is with respect to the Hausdorff metric. Noting that $\arg\max_{q_1 \in \Delta^{N-1}(\mathcal{X})} \Pi_1(q_1; p_1) = \{p_1\}$, by the upper hemicontinuity of the argmax correspondence, we get that $q_1^{\sigma(k)}$ converges to p_1^∞, 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 $N-1$ that satisfies

$$|\Pi_1^\star(q_1; p_1; M) - \Pi_1(q_1; p_1)| < 1/k^* \qquad \forall q_1, p_1,$$

then, for each p_1, if

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

then q_1^* is δ_1-close to p_1.

By contradiction, if the claim does not hold, then for every k there exists p_1^k, q_1^k, M^k such that

$$\left|\Pi_1^\star(q_1; p_1; M^k) - \Pi_1(q_1; p_1)\right| < 1/k \qquad \forall q_1, p_1,$$

while

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

and $d(q_1^k, p_1^k) \geq \delta_1$. Using the compactness of $\Delta^{N-1}(\mathcal{X})$, we generate a subsequence of indexes, $\{\sigma(k)\}$, such that $p_1^{\sigma(k)}$ converges to p_1^∞ and $q_1^{\sigma(k)}$ converges to q_1^∞ for some $p_1^\infty \in \Delta^{N-1}(\mathcal{X})$ and some $q_1^\infty \in \Sigma_1$.

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 any k large enough, $\Pi_1^\star(q_1^*; p_1^k; M^k) > \Pi_1^\star(q_1^\infty; p_1^k; M^k)$. This inequality contradicts the fact that for k large enough, q_1^∞ should also maximize $q_1 \in \Sigma_1 \mapsto \Pi_1^\star(q_1; p_1^k; M^k)$, since Σ_1 is finite.

Next, by uniform continuity we set $\delta > 0$ such that, if for each i, p_i is δ-close to p'_i, then $|\Pi(p_1, \dots, p_{N-1}, x) - \Pi(p'_1, \dots, p'_{N-1}, x)| < 1/k^*$ for each $x \in \mathcal{X}$. Let $\delta_2 = \min\{\delta_1, \delta\}$.

Step 2. We now iterate Step 1 for every $k = 2, \dots, N-1$. Let $k \geq 2$ and $\delta_k > 0$ be given. Fix $p_1, \dots, p_{k-2}$ 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 $\Delta^{N-k}(\mathcal{X})$.

Analogously to Step 1(a), we define $\Sigma_k^{p_1,\dots,p_{k-1}}$ as a finite subset of $\Delta^{N-k}(\mathcal{X})$ such that, for every p_k, if

$$q_k^* \in \underset{q_k \in \Sigma_k^{p_1,\dots,p_{k-1}}}{\arg\max} \Pi_k(p_1, \dots, p_{k-1}, q_k; p_k),$$

then q_k^* is δ_k-close to p_k.

Then, by a direct generalization of Step 1(b), there exists k^* such that for every finite menu M of order $N-1$ that satisfies

$$|\Pi_k^\star(p_1, \dots, p_{k-1}, q_k; p_k; M) - \Pi_k(p_1, \dots, p_{k-1}, q_k; p_k)| < 1/k^* \qquad \forall p_k, q_k,$$

for every p_k, if

$$q_k^* \in \underset{q_k \in \Sigma_k^{p_1,\dots,p_{k-1}}}{\arg\max} \Pi_k^\star(p_1, \dots, p_{k-1}, q_k; p_k; M)$$

then q_k^* is δ_k-close to p_k.

Finally, we let δ to be such that if, for every i, q'_i is δ-close to q''_i, then $|\Pi(q_1, \dots, q_{N-1}, x) - \Pi(q'_1, \dots, q'_{N-1}, x)| < 1/k^*$ for every x. Let $\delta_{k+1} = \min\{\delta, \delta_k\}$.

Step 3. We build a finite menu M_1^* of order $N-1$ by sampling the infinite menu associated with Π as follows: for every $p_1, \dots, p_{N-2}$ where for every k, $p_k \in \Sigma_k^{p_1,\dots,p_{k-1}}$, we define

$$
\begin{aligned}
M_{N-1}^{p_1,\ldots,p_{N-2}} &= \left\{\Pi(p_1,\ldots,p_{k-1},q_{N-1},\cdot);q_{N-1}\in\Sigma_{N-1}^{p_1,\ldots,p_{N-2}}\right\},\\
M_k^{p_1,\ldots,p_{k-1}} &= \left\{M_{k+1}^{p_1,\ldots,p_{k-1},q_k};q_k\in\Sigma_k^{p_1,\ldots,p_{k-1}}\right\}.
\end{aligned}
$$

We let $M_1^* = \{M_1^{q_1}; q_1 \in \Sigma_1\}$. Let ξ be the degenerate probability measure that allocates full mass on M_1^*. We note that we have, by Steps 1(a), 1(b), and Step 2,

$$
|\Pi^\star(p_1,\ldots,p_{N-1},x;\xi)-\Pi(p_1,\ldots,p_{N-1},x)|<\epsilon/2 \qquad \forall p_1,\ldots,p_{N-1},x
$$

Step 4. This step concludes the proof. Let ξ' be a probability measure over $\mathcal{M}_{N-1}$ with full support. Take $\xi''=(1-\epsilon/2)\xi+(\epsilon/2)\xi'$. Then, $\Pi^\star(p_1,\ldots,p_{N-1},x;\xi'')$ defines a strategyproof payoff rule, and

$$
|\Pi^\star(p_1,\ldots,p_{N-1},x;\xi)-\Pi(p_1,\ldots,p_{N-1},x)|<\epsilon \qquad \forall p_1,\ldots,p_{N-1},x.
$$