

MIT Open Access Articles

Dynamic Belief Elicitation

The MIT Faculty has made this article openly available. **Please share** how this access benefits you. Your story matters.

Citation: Chambers, Christopher P. and Lambert, Nicolas S. 2021. "Dynamic Belief Elicitation." *Econometrica*, 89 (1).

As Published: <http://dx.doi.org/10.3982/ecta15293>

Publisher: The Econometric Society

Persistent URL: <https://hdl.handle.net/1721.1/140535>

Version: Author's final manuscript: final author's manuscript post peer review, without publisher's formatting or copy editing

Terms of Use: Article is made available in accordance with the publisher's policy and may be subject to US copyright law. Please refer to the publisher's site for terms of use.



Dynamic Belief Elicitation*

Christopher P. Chambers[†] Nicolas S. Lambert[‡]

August 2020

Abstract

At an initial time, an individual forms a belief about a future random outcome. As time passes, the individual may obtain, privately or subjectively, further information, until the outcome is eventually revealed. How can a protocol be devised that induces the individual, as a strict best response, to reveal at the outset his prior assessment of both the final outcome and the information flows he anticipates and, subsequently, what information he privately receives? The protocol can provide the individual with payoffs that depend only on the outcome realization and his reports. We develop a framework to design such protocols, and apply it to construct simple elicitation mechanisms for common dynamic environments. The framework is general: we show that strategyproof protocols exist for any number of periods and large outcome sets. For these more general settings, we build a family of strategyproof protocols based on a hierarchy of choice menus, and show that any strategyproof protocol can be approximated by a protocol of this family.

Keywords: Elicitation device; scoring rule; BDM mechanism; dynamic information; second-order beliefs; high-order beliefs.

*A previous version of the paper circulated under the title “Dynamically Eliciting Unobservable Information.” We are particularly grateful to David Kreps, to the coeditor, and to four anonymous referees for helpful comments, suggestions and discussions. We also thank Glenn W. Harrison, Paul J. Healy, Matthew Jackson, Edi Karni, Norio Takeoka, Juuso Välimäki, seminar participants at many institutions, as well as the participants and organizers of several conferences and workshops. We gratefully acknowledge financial support from the National Science Foundation through grants SES-1426867 and CCF-1101209. Lambert thanks Microsoft Research and the Cowles Foundation at Yale University for their hospitality and financial support.

[†]Department of Economics, Georgetown University; christopher.chambers@georgetown.edu.

[‡]Department of Economics, Massachusetts Institute of Technology; nicolas.lambert@mit.edu.

1 Introduction

Imagine an experimenter (she) who believes her subject (he) conforms to the Bayesian model of uncertainty: the subject has probabilistic beliefs over some set of uncertain outcomes, and uses Bayes' rule to update when new information is available. However, the experimenter recognizes that the subject may condition on information which is either subjectively perceived, or privately observable. How can we design an elicitation device to understand how these beliefs evolve?

Probabilistic beliefs are commonly measured by experimenters. The classical tool for doing so is a *scoring rule*. This device offers a menu of state-contingent payoffs to a subject. The menu is chosen so that the subject's optimal choice uniquely reveals his subjective belief about states of the world. Scoring rules apply in situations in which the state of the world is eventually observed by the experimenter.

Here, instead, the subject is involved in a dynamic experiment in which private information resolves gradually over time. The experimenter wants to understand the subject's perception on how this information is to be revealed. She also wants to know, after the information is revealed, what he learned. To fix ideas, consider a simple experiment to test overconfidence, motivated by [Moore and Healy \(2008\)](#).¹ A subject is to take a pass-fail test. There are three time periods of interest. In period 0, the subject has not yet taken the test, and forms a prior belief about the likelihood he will pass the test. Then, the subject takes the test and in period 1, after having taken the test, he forms an updated, posterior belief about whether he passed. In period 2, the test is graded and the subject is told the outcome.

Of course, prior and posterior beliefs are expected to differ: as the subject tries the test, he gains new information about how difficult he finds the test. So, in the initial period, the subject anticipates that he will update his probability assessment, and forms a belief about his own posteriors. This belief reflects what he anticipates learning about his own performance by taking the test. We refer to it as a second-order belief, to distinguish it from the first-order beliefs that are probability assessments on test outcomes. Suppose the experimenter has interest in such a belief and, in period 0, asks the subject to report a distribution over the posteriors he may have. She then asks the subject, in period 1, to report his believed likelihood that he passed the test. This paper is about understanding how the experimenter can induce the subject to report both beliefs truthfully, as a strict best response, when payoffs to the subject can only depend on the reports and the outcome of the test.

We stress that the subject's distribution over the posteriors gives substantially more information than just the prior likelihood of passing the test, and so allows to answer a host of new questions. For instance, Moore and Healy distinguish between overconfidence as the overestimation of one's performance—in our example, when the posterior is biased upward—and overconfidence as the excessive precision of

¹We are grateful to Paul J. Healy and Matt Jackson for bringing to our attention the connection to this stream of the literature.

one’s signal—in our example, when the subject reports extreme posteriors too often.² However, the subject may well display low levels of “precision overconfidence,” while at the same time holding the strong belief that he will be able to guess his score after taking the test. This interpersonal notion of overconfidence cannot be measured through the elicitation of the prior and posterior beliefs. It is rooted in the subject’s second-order beliefs. Of course, once we know how to elicit these beliefs, we can also ask if participants properly anticipate how much they will learn from the test. We can ask if they have a bias, for example ask if they believe to learn successes more than failures. We can measure subjective overconfidence, by finding the quantile of the posterior in the reported prior distribution. And as Moore and Healy, we can ask related questions across subjects, for instance, if participants anticipate the test to be more informative about their own performance than about the performance of others.

Our point is not to provide an exhaustive list of experiments to be conducted with high-order beliefs. Rather, our goal is to operationalize the acquisition of such beliefs. The current state of the art is to offer a standard scoring rule in both periods, to elicit first the prior and then the posterior probabilities of passing the test. The prior assessment reflects, via the law of iterated expectations, the subject’s period-0 mean posterior, but that is the only statistic one gets on the distribution of posteriors. Scoring rules elicit these probability assessments because they concern whether the subject passes the test, an event directly observable by the experimenter. In contrast, the salient feature of our example is that the experimenter is unable to observe how difficult the test is to a subject. Consequently, scoring rules do not elicit the subject’s initial distribution over his posteriors.

In Section 2, we lay down the foundation for our approach, and explain how the elicitation of second-order beliefs can be done in the simple context of the above experiment. It is based on revealed preference. To illustrate, consider two possible menus of outcome-contingent payoffs, from which the subject is permitted to choose in period 0. One menu gives the subject \$6 for sure, independently of the outcome of the test. The other menu offers a choice, in period 1, between two options: the first is \$10 in the event of failure and \$0 otherwise, and the second gives \$10 in the event of passing and \$0 otherwise. Consider two risk-neutral subjects, who, in period 0, both believe that they will pass the test with probability .5, but hold different second-order beliefs: subject *A* believes that he will not learn anything from taking the test, while subject *B* believes that he will learn perfectly. Subject *A* would take the \$6 for sure, because his expected payout with the other menu is \$5. Subject *B*, on the other hand, would prefer to leave his options open by choosing the other menu. As we show, this phenomenon is general. Any subjects with differing beliefs can be behaviorally distinguished via a choice between some pair of menus. Thus, if the

²In practice, nonbinary outcomes are preferred to prevent the confounding of overestimation and overprecision, which can occur when probability assessments are always above 50%. In our example, probability assessments can be as high as 100% and as low as 0%, which makes it possible to disentangle these two notions of overconfidence.

experimenter could elicit a subject’s choice from sufficiently many pairs of menus, she could in principle back out the second-order beliefs.

In Section 3, we leverage this methodology to design simple elicitation protocols for more complex beliefs in special dynamic environments. Our mechanisms can be intuitively grasped as follows. The elicitor selects a collection of elementary decision problems, carefully chosen so that observing an individual’s choice behavior on every one of these problems permits the identification of the individual’s belief. By appropriate randomization, the elicitor can ask the individual to announce his beliefs and pay him *as if* he had to confront every one of these decision problems. This makes truthful communication a strict best response.

This indirect “revealed-preference” approach is simple and powerful. It is also general: it extends to essentially arbitrary dynamic environments. The challenge is that we must rely only on the observed outcome to elicit as a strict best response potentially complex subjective information (the high-order beliefs). To illustrate, suppose there is a coin toss whose outcome is only revealed at some future date. An individual holds a prior assessment on the outcome and is to observe, privately, new information at m different dates, where each observation is the realization of some k -dimensional signal. We want the individual to tell, as the outset, the full joint probability distribution over the m signals and the outcome, and then reveal the signals he observes as he receives them. As m and k grow large, the object to elicit includes a vast amount of *subjective* information, yet to enforce truthfulness, the only *objective* information is a single outcome that takes only two possible values.

We elaborate and formalize the general theory in Section 4. We introduce a family of protocols, where each protocol is identified with a probability distribution over choice menus. We establish that these protocols induce truth telling as a strict best response, and that any strategyproof mechanism can be approximated by a protocol of the family.

In recent years, a large body of work has been devoted to how people learn over time, how beliefs evolve, and how it affects their decisions. Our theoretical framework is relevant to experiments that meet three conditions: (i) the environment is dynamic, (ii) the subject’s beliefs are of interest, and (iii) the signals presented to him either are not controlled by the experimenter and unobservable to her, or are open to interpretation. In the experiment of Moore and Healy, taking the test generates the unobservable signal. Other common cases of unobservable signals are social cues or cheap talk.

Many experimental designs fit these three conditions. A recent stream of the literature devotes special attention to the question of how people learn in repeated games (as in, for example, Nyarko and Schotter, 2002, Palfrey and Wang, 2009, or Hyndman et al., 2012 and Danz et al., 2012). These studies elicit a player’s beliefs about the actions of the other players using classical probability scoring rules. In those environments, the actions taken by a player provide a signal to another player. The signal is observed by the experimenter, but is open to interpretation. In games

with incomplete information, actions continue to provide signals, and beliefs involve both the actions and the private information of the other players. Our framework can be used to estimate how the players anticipate their belief to change, and how it affects their own play.

More broadly, the elicitation of high-order beliefs helps us refine our understanding of people’s learning process and its interplay with observed decisions, and explain violations of equilibrium predictions. For example, it is widely documented that, in games of imperfect information, our ability to learn from strategies is limited—the textbook example being the winner’s curse in common-value auctions. Tools such as the concept of Cursed Equilibrium (Eyster and Rabin, 2005) have been introduced for the purpose of explaining these facts. Knowing high-order beliefs would help push the analysis further by enabling the experimenter to measure how much information a bidder expects to obtain from observing the bids of their opponents. Even in simpler games, the relation between a player’s beliefs and his actions poses interesting questions. Not only equilibrium play is often not observed, but there is evidence of inconsistencies, a player’s belief revealing deeper strategic thinking than his action, as Costa-Gomes and Weizsäcker (2008) demonstrates for normal-form games using probability scoring rules. Being able to elicit beliefs of second or higher order can help us understand how much of these effects can be linked to the complexity of the dynamics. Some other works examine the evolution of subject beliefs in response to signals and stimuli, as the study of information cascades (Ziegelmeyer et al., 2010) or belief polarization (Fryer et al., 2019). In these instances, the signals are controlled by the experimenter, but left open to interpretation. These works track the change of first-order beliefs over time. With second-order beliefs, it becomes possible to explain how much of what is observed is due to an error in how the subject updates his belief (e.g., the subject overreacts to information), versus how much is due to a misspecified cognitive model (e.g., the subject incorrectly believes that future signals will carry significant information).

While we use belief elicitation in experiments as our leading example, probability scoring rules have been applied to a large range of settings to induce honest or accurate reports of information.³ They are also a main tool by which to evaluate, in theory as in practice, learning models, predictions and forecasters. To the extent that our approach develops the foundation for the dynamic analog of probability scoring rules, we believe that our theory can be applied for the same purposes of elicitation and performance evaluation, but in dynamic environments, in which forecasts arrive over time and what matters is not only the quality of those forecasts, but also how fast uncertainty is anticipated to resolve.⁴

³For example, in contract theory (e.g., Thomson, 1979, Osband, 1989, or more recently Carroll, 2019), prediction markets (Ostrovsky, 2012), problems of strategic distinguishability (Bergemann et al., 2017), the testing of forecasters (Stewart, 2011), or the literature rational inattention (Steiner et al., 2017). Harrison and Phillips (2014) expand on the practical aspects of using scoring rules as incentive devices in risk management.

⁴Knowing the information structure—forming expectations on what information will arrive and

Related Literature

Foremost, our paper relates to the literature on scoring rules and belief or preference elicitation, that goes back to [Brier \(1950\)](#) and [Good \(1952\)](#), who establish the first two proper probability scoring rules. [McCarthy \(1956\)](#) and later [Savage \(1971\)](#) offer a general method to construct these scoring rules, which has been extended and exploited extensively. The literature is vast and spans several fields, it is impossible to do it justice (for a survey, see [Gneiting and Raftery, 2007](#)). Importantly, the literature assumes a static setting. Our work departs from the static benchmark, providing a general rule for making scoring rules that apply to dynamic settings.

The experiment conducted by [Manski and Neri \(2013\)](#) demonstrates the practical feasibility of eliciting probabilistic beliefs of one subject on the probabilistic beliefs on the other. To do so, they use the Brier score to elicit the first-order beliefs of a subject A , and a sum of Brier scores to elicit the beliefs of another subject B regarding subject A 's first-order beliefs. These are also “second-order beliefs,” but there, the subject forecasts the beliefs of someone else, whereas here, the subject forecasts his own future beliefs. The distinction is crucial: subject B has no ability to manipulate the reports of subject A , they are, from the viewpoint of subject B , states that the experimenter can observe, so that standard probability scoring rules apply. In elegant recent works, [Karni \(2018, 2020\)](#) uses a similar structure to elicit the second-order beliefs of the same subject. Karni argues that this structure is useful when the subject's behavior conforms to nonstandard decision models. In this case, however, the mechanism is not incentive compatible, because the subject would manipulate his future reports (see [Appendix A](#)).

The standard approach to build scoring rules, explained in [Savage \(1971\)](#), is to take the subgradients of convex functions. This “direct” approach relates to the “payoff equivalence” characterizations in mechanism design ([Krishna and Maenner, 2001](#)). We take a different route. Our approach is inspired by an idea developed in [Allais \(1953\)](#) and also attributed to W. Allen Wallis ([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 over the entire collection, choose two objects at random, and then give the individual the object that is preferred according to his announcement. [Azrieli et al. \(2018\)](#) show that the mechanisms that are incentive compatible under minimal assumptions on the subject's preference reduce to randomized mechanisms of the form given by Allais.

In the static benchmark of the literature, several works relate indirectly to the Allais idea. In their seminal work, [Becker et al. \(1964\)](#) introduce, as an alternative to the Brier score, a method for eliciting an expert's belief via a second-price auction with a random reserve price. [Matheson and Winkler \(1976\)](#) propose a scoring rule to elicit

when it will arrive—as captured by high-order beliefs enables a decision maker to solve any sort of dynamic problem. In contrast, probability assessments on payoff-relevant outcomes are only relevant to static decisions, or degenerate dynamic problems. We demonstrate these facts in Sections S.3 and S.4 of the Supplemental Material ([Chambers and Lambert, 2020](#)).

the distribution function of real-valued random variables. [Schervish \(1989\)](#) proposes a characterization of strictly proper scoring rules for binary outcomes as integrals with positive weights, this type of characterization is also central in scoring rules for distribution properties ([Lambert, 2018](#)). Although these works are independent of each other, they have in common that they can be interpreted as implicit applications of the Allais idea, either by randomization, or by a mechanism that is equivalent to introducing a randomization.⁵ One contribution of our work is thus the formalization of the connection to the Allais idea, and the illustration of its effectiveness beyond the static benchmark.

Another strand of literature compares reports across individuals to obtain honest opinions on subjective matters, by assuming a form of consensus among the individuals (e.g., [Prelec, 2004](#) or [Miller et al., 2005](#)). This consensus allows to discard observable outcomes. In mechanism design, the Crémer-McLean mechanism ([Crémer and McLean, 1988](#)) is a classical example of such a construct. In contrast, in this paper we elicit information individually, but rely on the observability of the final outcome.

Finally, our work connects to the problem of identification in the decision theory literature. To design the elicitation protocols of Sections 2–4, we show how to identify high-order beliefs from behavior across a set of decision problems, and then combine these decision problems into a single elicitation task. The identification that we perform differs from the identification of beliefs in decision theory in that, in the latter, one usually has access to the entire preference relation of the decision maker to identify the parameters of the decision model. By contrast, in our framework, the set of decision problems used for identification must be kept relatively simple. This simplicity is first needed to obtain the simple protocols in the special cases described in the next two sections. But at a more general level, working with a small enough set of decision problems enables us to combine these problems so as to preserve their strict incentive properties, so that a strict best response in any one of the decision problems continues to induce a strict best response in the combined problem—a central aspect of our design. This aggregation can fail when identification is supported by too many decision problems. Thus, knowing that beliefs are identified in the classical sense can help but is not enough to say that beliefs can be elicited; we discuss this point in more details in Section 4.2.

The identification of beliefs (in the classical sense) is known in special cases of the decision theory literature. In the dynamic models of [Takeoka \(2007\)](#), [Dillenberger et al. \(2014\)](#) and [Lu \(2016\)](#), the decision maker observes an interim subjective signal, or acts as if she observed such a signal, as does the subject in our motivating example.

⁵Other elicitation mechanisms, such as the mechanisms of [Roth and Malouf \(1979\)](#), [Grether \(1981\)](#), or [Karni \(2009\)](#), introduce randomization. The purpose of this randomization is different: these mechanisms reward individuals with lottery tickets to create an incentive compatibility in absence of risk neutrality. The key knowledge that the experimenter must have in these contexts is that the subject’s preferences exhibit monotonicity with respect to stochastic dominance of the imputed lotteries.

	Binary outcomes	Number of interim periods	Random time of information arrival	Restricted beliefs
Protocol (I)	yes	1	no	yes
Protocol (II)	yes	1	no	no
Protocol (III)	no	1	no	no
Protocol (IV)	yes	1	yes	no
Protocol (V)	yes	2	no	yes
Protocol (VI)	yes	2	no	no

Table 1: Features of the environments for the protocols of Sections 2 and 3.

In particular, identification for the decision models of [Takeoka \(2007\)](#) and [Dillenberger et al. \(2014\)](#) implies that second-order and third-order beliefs are identified. On the other hand, our results imply that, in any decision model consistent with our behavioral assumptions, higher-order beliefs continue to be identified, hence complementing those works. Note that in most protocols of this paper, all payoffs occur after all uncertainty is resolved. Under risk neutrality and without discounting, redistributing the payoffs over multiple time periods is possible, sometimes allowing for some simplification, but in general, allowing for payoffs in interim periods requires to account for time-related preferences such as intertemporal substitution, thus adding other dimensions to preferences which can complicate the task of elicitation. This is relevant for the models of [Kreps and Porteus \(1978\)](#), and, more recently, [Krishna and Sadowski \(2014\)](#).

Organization

Sections 2 and 3 provide protocols that elicit beliefs in special environments in which the number of time periods, the beliefs of the subject, or the information to be elicited are restricted. These special cases are presented in order of increasing complexity and help build up the general theory. Table 1 outlines the features of these environments. The full, unrestricted framework is presented in Section 4. Section 5 concludes. Appendix A discusses protocols obtained by combining elicitation mechanisms for first-order beliefs. Appendices B and C include the proofs omitted from the main text. The Supplemental Material ([Chambers and Lambert, 2020](#)) presents several extensions of the main results.




Figure 1: Probability tree if the subject anticipates to learn all or nothing.

2 A Simple Example

In this section, we explore a simple example to illustrate the theory of this paper. Throughout, the elicitor is an experimenter who wants to elicit beliefs from her subject regarding a random event. For concreteness, we work with the experiment presented in the Introduction, though the application to other domains is straightforward. The outcome of interest is whether the subject passes or fails the test. The goal is to elicit, in period 1, the subject’s probability assessments on the outcome, and, in period 0, the subject’s initial belief over these probability assessments.

2.1 The Case of Restricted Beliefs

As a first step, consider a restricted information structure: the experimenter hypothesizes that, after taking the test, the subject either fully learns whether he passed or failed, or learns nothing new. Thus, in period 0 (the initial period), the subject is asked to report an element $p \in [0, 1]$, reflecting an initial probability assessment that he will pass the test, together with a probability $\alpha \in [0, 1]$ that he will learn whether he passed after having taken the test (with probability $1 - \alpha$ of learning nothing new). Under the experimenter’s assumption, these two numbers describe fully the subject’s distribution over his posterior beliefs. Then, the subject takes the test and, in period 1 (the interim period), is asked to report how likely he believes to have passed the test, an element $q \in [0, 1]$.⁶ Finally, in period 2 (the final period), the subject is told whether he passed or failed the test. In Figure 1 we draw the probability tree associated to the subject’s belief in period 0. The leaves of the tree correspond to the possible beliefs the subject may form in period 1, regarding whether or not he passed the test, while the branches indicate the ex-ante likelihood attributed to these beliefs in period 0.

The experimenter delivers a payoff as a function of the reported probabilities α , p and q , when x is the outcome that realizes; by convention, $x = 1$ if the subject passes the test, and $x = 0$ if he fails. Following the literature, the experimenter must motivate the subject with strict incentives: the subject must be willing to respond

⁶While we could compel the subject to be consistent with his initial report by requiring $q \in \{0, p, 1\}$, it is not necessary. The belief at the second stage need not be in the support of the belief at the first stage, so that a subject who misreports in period 0 is still induced to report truthfully in period 1.

truthfully, and the truth must be the unique best response. To simplify matters even further, let us suppose the subject is risk neutral.

To elicit the prior and posterior probabilities of passing the test, one can sequentially use probability scoring rules, such as a quadratic scoring rule giving the subject payoffs $1 - (p - x)^2$ and $1 - (q - x)^2$ for respective assessments p and q (Selten, 1998). It is then a strict best response for the subject to act truthfully at each stage. On the other hand, probability scoring rules do not elicit the probability that the subject assigns to learning fully.

Instead, the elicitation of α relies on the following idea. Suppose the experimenter were to use the above-mentioned quadratic scoring rules. As can be easily checked, a subject who believes they are more likely to learn fully also expects to earn more from the elicitation of his posterior. Indeed, the second score yields as expected payoff $1 - p(1 - p)$ to the subject who learns nothing, and 1 to the subject who learns fully. This fact makes it possible to discriminate between the subjects who are more likely to learn fully and those who are not. The higher is α , the higher is the option value associated with delaying a decision. We leverage exactly this idea: individuals with a relatively high α will find it best to delay their decision; in contrast, individuals with a relatively low α will be more willing to commit. The goal is to appropriately choose commitment payoffs and option payoffs in order to distinguish between the different individuals.

For example, the experimenter could fix a baseline payoff $B \in (0, 1)$ and, before administering the test, offer the subject the following choice: in the interim period, he can be paid according to the quadratic score, or he can choose to forgo this elicitation payment and instead get compensated with payoff B . Initially, the subject expects to earn $1 - (1 - \alpha)p(1 - p)$ from the quadratic score, so should decide to use the quadratic score only if α is above some threshold, and otherwise leave with B . Observing the subject's choice enables the experimenter to infer information about the probability of learning fully. By repeating this procedure on the same subject infinitely many times while increasing B smoothly from 0 to 1, the experimenter could, in principle, collect enough data points to infer exactly α . Doing so would be impractical, but the same effect can be achieved through randomization and delegation: *after* the subject communicates his belief, the experimenter draws B at random and chooses between the scoring rule and payoff B on behalf of the subject. The subject strictly prefers to respond honestly as long as a wrong choice is costly with nonzero probability—as when B is uniformly distributed.

This shows how to elicit α . We can elicit all of α , p and q with a small modification: keep the first quadratic scoring rule to get p and postpone the experimenter's actions to the final period, so that in the interim period, the subject is still unsure whether he is being paid as a function of his posterior assessment, inducing a strict best response for q . The overall protocol is summarized below.

Protocol (I) *In the initial period, the subject is asked to estimate the likelihood α that he will learn fully, along with the prior probability p that he will pass the test. In*

the interim period, the subject is asked to assess the posterior probability q that he passed the test. In the final period, the test outcome is revealed, and the experimenter draws $B \in [0, 1]$ uniformly at random. If $1 - (1 - \alpha)p(1 - p) < B$, then the subject is paid $1 + B - (p - x)^2$. Otherwise, the subject is paid $2 - (p - x)^2 - (q - x)^2$.

This construction highlights the key idea of this paper. To elicit an individual's (dynamic) beliefs, we divide the elicitation task into many small parts. We consider a collection of basic decision problems—in this instance, whether to take B or go with the payoffs of a quadratic scoring rule. The collection is designed so that observing the individual's choices from each of the decision problems uncovers the beliefs entirely. But taken separately, each decision problem only reveals a small piece of the individual's beliefs. To elicit these beliefs as a single decision, we combine all the simple problems by suitably randomizing.

As in [Allais \(1953\)](#), randomization is a natural device in this context. However, it is not a necessity. We can design a nonrandom elicitation scheme if, as opposed to drawing one decision problem from a set, we give the subject an infinitesimal fraction of every decision problem from that set. Here, it can be done by computing the average payoffs. For example, averaging the payoffs of Protocol [\(I\)](#) over B yields a nonrandom payoff that represents a scoring rule for the elicitation of α , p and q , which before simplification is written

$$\pi(\alpha, p, q, x) = 1 - (p - x)^2 + \int_0^{1 - (1 - \alpha)p(1 - p)} (1 - (x - q)^2) dB + \int_{1 - (1 - \alpha)p(1 - p)}^1 B dB.$$

In this instance, and in several others that we examine below, we find that randomization makes it possible to have intuitive and relatively simple elicitation schemes. The absence of randomization can be preferred when the payoff, or score, is used for the purpose of evaluating a learning model (as opposed to the elicitation of a subject's beliefs) as in [Feltovich \(2000\)](#). In this context, the complexity of the scoring rule is irrelevant.

2.2 Unrestricted Beliefs

We now depart from the simplifying assumption that the subject fully learns the outcome after taking the test. The subject continues to hold a posterior belief, in the interim period, about whether he passed. In the initial period, the subject forms a belief about this posterior, now captured by a distribution function over $[0, 1]$ that we refer to as second-order belief. The protocols of [Section 2.1](#) do not enable us to elicit such a belief, because the class of decision problems employed in the construction is too coarse. Rather than randomize over quadratic scoring rules, we use a richer set of simpler decision problems.

Protocol (II) *In the initial period, the subject is asked to announce his second-order belief F . The experimenter then draws two numbers A and B independently and uniformly from $[0, 1]$. If*

$$A \geq E^F[\max(B, P)],$$

where the expectation is taken for P distributed according to F , then the protocol stops and the subject gets the payoff A . Otherwise, in the interim period, the subject chooses between getting the fixed payoff B and getting the payoff 1 conditional on him passing the test (and nothing otherwise).

Proposition 1 *In Protocol (II), the subject announces his second-order belief as a strict best response.*

Of course, payoffs can be shifted and scaled as the experimenter sees fit. Throughout let $\varphi(B, F) = E^F[\max(B, P)]$. Straightforward calculations yield

$$\varphi(B, F) = 1 - \int_B^1 F(p) \, dp.$$

The intuition behind Proposition 1 is simple. Observe that $E^F[\max(B, P)]$ is the expected payoff of a subject who is to be given the choice in the interim period, with P the random posterior belief as seen from period 0. Therefore, the experimenter makes the decision that is the best for the subject (given the information the subject provides) and truthful reporting is, at least, a weak best response.

In this protocol, the “simple decision problems” are whether to stop the experiment to get an immediate payoff or continue to the next stage. There are as many decision problems as there are values of A and B . If F is the true second-order belief, but the subject communicates $\tilde{F} \neq F$ instead, then we argue that there are many values of the parameters A and B , thus many simple decision problems—sufficiently many so that, on aggregate, these values generate a positive mass—such that the experimenter who acts on behalf of someone with second-order belief \tilde{F} makes the wrong choice, either stopping the protocol while the subject would have been better off continuing, or conversely. The subject, who is unaware of which decision problem will be selected for him, is at risk of losing some payoff when he deviates from the truth. He can only guarantee himself the maximum payoff with probability 1 when he tells the truth.

Proof of Proposition 1. Let F be the subject’s second-order belief, and \tilde{F} be the

subject's announcement. We have

$$\begin{aligned}
E^{\tilde{F}}[\max(B, P)] &= \int_0^1 \max(B, p) d\tilde{F}(p) \\
&= \int_0^B B d\tilde{F}(p) + \int_B^1 p d\tilde{F}(p) \\
&= B\tilde{F}(B) + \left(1 - B\tilde{F}(B)\right) - \int_B^1 \tilde{F}(p) dp \\
&= 1 - \int_B^1 \tilde{F}(p) dp.
\end{aligned}$$

Therefore, the expected payoff of the subject is

$$\begin{aligned}
&\int_0^1 \int_{\varphi(B, \tilde{F})}^1 A dA dB + \int_0^1 \int_0^{\varphi(B, \tilde{F})} E^F[\max(B, P)] dA dB \\
&= \int_0^1 \frac{1}{2} \left(1 - \varphi(B, \tilde{F})^2\right) dB + \int_0^1 \varphi(B, \tilde{F})\varphi(B, F) dB \\
&= \int_0^1 \left(\frac{1}{2} \left(1 - \varphi(B, \tilde{F})^2\right) + \varphi(B, \tilde{F})\varphi(B, F)\right) dB. \quad (1)
\end{aligned}$$

This expression is maximized if and only if, for almost all B , $\varphi(B, \tilde{F}) = \varphi(B, F)$. As φ is continuous in its first argument, the expression is maximized if and only if for all B , $\varphi(B, \tilde{F}) = \varphi(B, F)$. Naturally, if $F \neq \tilde{F}$ then by the right-continuity of cumulative distribution functions, for some B , $\int_B^1 F(p) dp \neq \int_B^1 \tilde{F}(p) dp$ and so $\varphi(B, \tilde{F}) \neq \varphi(B, F)$. Hence, the expected payoff of the subject is maximized if and only if he reports F . ■

Rather than provide a general discussion, we conclude this section with several observations.

Eliciting the prior and posterior beliefs

Protocol (II) elicits second-order beliefs only. The prior is not elicited directly, but is included as part of the second-order belief, because it is equal to the mean posterior. The posterior, a first-order belief, can either be elicited separately using a quadratic scoring rule, or be elicited in the same protocol if instead all the decisions are made by the experimenter on behalf of the subject. In the interim period the subject would then be asked to send a probability assessment, as opposed to making a binary choice. In this case, it is important that the values of A and B are only drawn or revealed *after* the subject has communicated his information, to ensure that optimal announcements remain strict.

Moral hazard

This paper is concerned with the design of protocols that prevent the manipulation of reported beliefs. It is not concerned with the manipulation of the beliefs themselves. When individuals can influence final outcomes, the elicitation procedure, by adding state-contingent payments, can generate a moral hazard problem when incentives are misaligned.

The problem is common in experimental setups. In incentivized experiments, it can be reduced or eliminated by increasing the incentives on the experimental task. In the above environment, we presume that subjects simply want to pass the test. In practice, subjects have some control over the likelihood of passing the test, for example they could decide to flunk the test, which creates a potential for moral hazard. Suppose the experimenter wants that subjects try to pass to the best of their ability. Then, to ensure that the subject is motivated to pass, the experimenter can add a payment in case of success. If large enough, the incentives of the experimenter and the subject are aligned in the presence of the elicitation procedure.

Note that the presence of such agency problems does not impact the incentive compatibility of belief elicitation protocols. Specifically, if the subject can take actions that influence the final outcome (or the information observed prior to the final period) then in the environment of this section (as for the more general dynamic environment of the sections below) the subject will find it optimal to announce the belief associated with the optimal anticipated stream of actions at the time of the announcement.

Complexity of the protocol

The second-order beliefs in this example are the distributions of a random variable taking values in $[0, 1]$. In principle, they can be complex, but in practice the experimenter, who has control of the communication device, need not account for all possible distributions. For example, the subject may be asked to choose a density shape among a suggested sample, move sliders to control the shape of the density function (Moore and Healy, 2008), or be asked to provide the probabilities of finitely many ranges of posteriors (Manski and Neri, 2013). Beyond the experimental context, distributions are often parameterized, for example, a forecaster may be asked for the mean and variance of a truncated Gaussian, or may give a discrete probability tree, i.e., a distribution with finite support.

When the ability to report precisely one's belief is limited by the technology, the subject may be unable to reach the theoretically optimal payoff. However, the loss incurred is small. It is bounded by the squared error between the announced belief \tilde{F} and the true belief F : if, for every p , $|F(p) - \tilde{F}(p)| < \varepsilon$, then the subject's expected payoff is at least the optimal payoff he would obtain by reporting F minus $\varepsilon^2/2$.⁷

⁷Equation (1) in the proof of Proposition 1 yields the expected loss incurred when reporting the approximate belief \tilde{F} instead of the exact belief F , $\frac{1}{2} \int_0^1 (\varphi(B, \tilde{F}) - \varphi(B, F))^2 dB$. If $|F - \tilde{F}| < \varepsilon$,

Note that the protocols of this section are “direct.” The alternatives are the “indirect” elicitation protocols, in which the subject makes choices, and these choices inform the experimenter on the subject’s beliefs. One benefit of direct protocols is that they do not require the subject to confront difficult choices: as long as the subject agrees with the incentive-compatible nature of the protocol, he only needs to supply his information, without making any computation on his own.

Relation with the BDM mechanism

Protocol (II) can be viewed as a dynamic extension of the BDM mechanism (Becker, DeGroot, and Marschak, 1964). In the usual version, the subject bids for an object in a second-price auction with a random reserve price. This bid reveals the subject’s willingness to pay for the object. In the context of probability elicitation, the object is an Arrow-Debreu security.

What we show is that to elicit second-order beliefs, we can use two auctions, one embedded into the another. In the main auction, the subject formulates a bid for the right to participate in the secondary auction. If the bid is greater than or equal to a reserve price A , the subject pays A and obtains this right. Otherwise, the subject pays nothing and gets nothing. Then, if the subject won the main auction, the secondary auction takes place. The subject formulates a bid for the Arrow-Debreu security that pays off x . If the bid is greater than or equal to a reserve price B , the subject pays B and gets the security. Otherwise, the subject pays nothing and gets nothing. For given values of A and B , this auction mechanism is equivalent to Protocol (II): the payoffs are identical up to an addition of the amounts A and B . Hence, collecting bidding data (in the main auction only) for many uniformly distributed pairs A, B in the unit square makes it possible to learn exactly the subject’s second-order belief. While the BDM mechanism can elicit the subject’s probability assessment with a single bid, many bids are needed to learn the second-order beliefs: the subject’s willingness to pay in the first auction depends on B . In applications, the experimenter could present the subject with a series of main auctions in which B is increases gradually from 0 to 1, and for each auction, demand the subject’s bid. Once all the bids are received, the experimenter applies one of these auctions at random, also setting the reserve price A at random.

On the impossibility of eliciting dynamic beliefs by combining standard elicitation mechanisms

Suppose that, instead of the subject reporting, in period 0, the distribution over the posteriors he anticipates to have in period 1, we ask the subject to report the distribution of the posteriors of another subject passing the test.

then $|\varphi(B, \tilde{F}) - \varphi(B, F)| \leq \int_B^1 |F - \tilde{F}| < \varepsilon$, so the expected loss is no greater than $\varepsilon^2/2$.

This situation poses no particular theoretical challenge: we can elicit the posterior using a quadratic scoring rule, and we can elicit the distribution of posteriors using a scoring rule designed for distributions of random variables, for example one of the scoring rules defined by [Matheson and Winkler \(1976\)](#) which we will refer as Matheson-Winkler score,⁸ taking as observed state the elicited posterior. This mechanism is natural and preserves incentives because, from the viewpoint of both subjects, the realized value of the variable the experimenter asks to forecast is exogenous.

One may be tempted to continue to apply this mechanism even when the two subjects are, in fact, the same, as in this paper. As it turns out however, the incentive compatibility property ceases to hold. For example, the subject who is honest in the initial period will want to manipulate his probability estimate in the interim period. Without getting into details, the intuition is that, when the subject reports truthfully in both periods and contemplates a small deviation in the interim period, the effect of that deviation is of second-order in the quadratic scoring rule (since he was maximizing that score by being truthful) but is generally of first-order in the Matheson-Winkler score.

As we show in [Appendix A](#), this result is quite general. The elicitation of high-order beliefs always requires the interaction of the various reported information at different times through the payoffs.

Equivalent scoring rule formulation

As in [Section 2.1](#), the analog of probability scoring rules for the case of second-order beliefs can be easily constructed for the purpose of evaluating learning models.

To keep matters simple, we continue to use the language of elicitation. Suppose the subject first announces a second-order belief F in period 0, a probability assessment p in period 1, while x continues to be the outcome. The goal is to design a payoff $\pi(F, p, x)$ such that it is uniquely optimal to report one's second-order belief F in the initial period, and then it is uniquely optimal to report one's first-order belief p in the interim period (even after misreporting initially). When those two conditions are met, let us say that π is *strategyproof*.

To construct this payoff, we can compute what the subject earns averaged over the draws of A and B in [Protocol \(II\)](#), assuming that both choices in the initial and interim periods are made by the experimenter on the subject's behalf. This average payoff is equal to

$$\int_0^1 \int_{\varphi(B,F)}^1 A \, dA \, dB + \int_0^p \int_0^{\varphi(B,F)} x \, dA \, dB + \int_p^1 \int_0^{\varphi(B,F)} B \, dA \, dB,$$

⁸The result is general but to fix ideas, we can take for example the scoring rule $S(F, p) = -\int_0^p F(x)^2 \, dx - \int_p^1 (1 - F(x))^2 \, dx$, as defined in [Section 2](#) of [Matheson and Winkler \(1976\)](#), where F is the reported distribution and p is the realization of the random variable, here the posterior.

which reduces to

$$\int_0^1 F(B)\varphi(B, F)B \, dB + \int_0^p x\varphi(B, F) \, dB + \int_p^1 B\varphi(B, F) \, dB. \quad (2)$$

The payoff just defined—let us write it $\pi(F, p, x)$ —is an analog of the quadratic scoring rule for second-order beliefs. This payoff function is strategyproof.

The argument is simple. In the initial period, the expected payoff to the subject is identical to the expected payoff in Protocol (II), and hence the subject's unique best response in this period is to report truthfully. Then, no matter the second-order belief F announced, in the interim period the subject who believes to pass the test with probability p announces \tilde{p} so as to maximize the residual expected payoff

$$\int_0^{\tilde{p}} p\varphi(B, F) \, dB + \int_{\tilde{p}}^1 B\varphi(B, F) \, dB. \quad (3)$$

As $\varphi(B, F)$ is strictly positive (except possibly for $B = 0$), (3) is strictly increasing for $\tilde{p} \leq p$ and strictly decreasing for $\tilde{p} \geq p$, and so maximized exactly when $\tilde{p} = p$: it is strictly optimal to report truthfully in period 1.

Unlike the original protocol, payments can be spread out over time, which simplifies the mechanism. First, the subject reports second-order belief F , and is immediately paid the amount $\int_0^1 F(B)\varphi(B, F)B \, dB$. Then, in the interim period, the subject reports probability assessment p , and is immediately paid the amount $\int_p^1 B\varphi(B, F) \, dB$. Finally, after the event outcome realizes, the subject is paid $\int_0^p x\varphi(B, F) \, dB$.

Behind the seemingly complex formulation of this two-stage quadratic scoring rule lies a simpler intuition. The logic is as follows. Let us rewrite (2) slightly differently as

$$\pi(F, p, x) = \frac{1}{2} + \int_0^1 \left(\max(p, B) - \frac{1}{2}\varphi(B, F) \right) \varphi(B, F) \, dB + \int_0^p (x - p)\varphi(B, F) \, dB. \quad (4)$$

Ignoring the irrelevant constant $1/2$, let us interpret the first component of (4). Recall that $\varphi(B, F)$ is the average value of the subject's interim payoff for second-order belief F . If the realization of the random interim payoff, $\max(B, P)$, is publicly known, the term

$$\left(\max(p, B) - \frac{1}{2}\varphi(B, F) \right) \varphi(B, F) \quad (5)$$

is a quadratic scoring rule that elicits the subject's assessment of the mean interim payoff.⁹ The elicitation is indirect, because the subject does not report explicitly an assessment of this mean, instead he reports second-order belief F from which an

⁹Up to a factor, the quadratic scoring rule for estimations of the mean of a random variable takes the form $s(m, y) = -my + y^2/2 + h(y)$, where h is arbitrary, y is the realization of the random variable, and m is the mean estimate.

implicit assessment can be derived. Integrating over the range of possible values for B , as in (4), ensures that this assessment is elicited for every B .

Let us now interpret the second component of (4). The term

$$\int_0^p (x - p)\varphi(B, F) dB \tag{6}$$

is a probability scoring rule that elicits the likelihood of the event. It is known as a ‘‘Schervish’’ score with weight function $B \mapsto \varphi(B, F)$ (Schervish, 1989).

An important feature is that the probability scoring rule (6) is not fixed: it varies with the subject’s announced second-order belief of the initial period. This dependence is required because the realization of $\max(B, P)$ is only privately observed, and thus the subject may be tempted to manipulate his report to increase the payoff that comes from the first component of (4), as explained in our previous remark. The adaptive weight $\varphi(B, F)$ in (6) ensures that the benefits of misreporting the first-order belief in the quadratic scoring rule (5) never exceed the cost collected through scoring rule (6).

Note that the general procedure of constructing an equivalent scoring rule can be undertaken with any of our protocols below, simply by taking expectations with respect to all randomization done on the part of the elicitor and leveraging risk neutrality. The protocols discussed will incentivize the announcement of whichever beliefs our results claim, but without any randomization.

3 Protocols for Restricted Environments

In this section we apply the general principle illustrated in Section 2 to specific instances of dynamic environments. In every instance, there are finitely many time periods. The elicitor (e.g., an experimenter) has interest in the outcome of a random variable or random event that materializes publicly in the final period. An individual (e.g., the subject) holds beliefs on the distribution of outcomes in the initial period. Those beliefs may evolve over time, through one or more interim period(s), due to information that either is subjectively perceived or interpreted, or is privately observed by the individual.

We examine several cases of special but salient types of information structures. In each case, we show that simple protocols enable the elicitor to obtain, as a strict best response, the individual’s relevant dynamic beliefs, or equivalently, the individual’s private subjective information structure. To keep protocols simple, we focus on the elicitation of period-0 beliefs. As in Section 2, elicitation of the subsequent lower-order beliefs can be done by combining the protocol with another elicitation procedure (e.g., a quadratic scoring rule for first-order beliefs), or by delegating the individual’s choices to the elicitor (we use this approach in Section 4). We assume risk neutrality. Extending the protocols to expected utility maximizers is straightforward.¹⁰

¹⁰One can apply the following change to the protocols: shift/scale the payoffs (say, in dollars) to

3.1 Multiple Outcomes

Here, as in Section 2, there is a single interim period at which the individual is able to collect new information. The outcome is now a discrete random variable X taking values in the finite set $\{1, \dots, n\}$. The individual's interim belief is captured by a vector $P = (P_1, \dots, P_n)$ where P_k denotes the assessed probability of $X = k$. In the initial period, the individual holds a belief about P . This is a second-order belief, represented by a multidimensional distribution function. It is worth noting that, perhaps surprisingly, going from two to more than two outcomes makes the elicitation of second-order beliefs substantially more difficult, unlike the case of first-order beliefs. With first-order beliefs, we can always sum n quadratic scores for binary events to elicit the probabilities for each of the n outcomes, whereas we cannot iterate Protocol (II) to extract second-order beliefs. Instead, we develop the following protocol.

Protocol (III) *In the initial period, the individual is asked to announce his second-order belief F . The elicitor then draws two numbers A and B , and n numbers c_1, \dots, c_n independently and uniformly from $[0, 1]$. Let $C_i = c_i / (c_1 + \dots + c_n)$. If*

$$A \geq E^F[\max(B, C_1P_1 + \dots + C_nP_n)],$$

then the protocol stops and the individual gets the payoff A . Otherwise, in the interim period, the individual is offered a choice between getting the fixed payoff B , or getting a contingent the payoff of 1 if $X = I$ (and nothing otherwise), where $I \in \{1, \dots, n\}$ is drawn randomly in the final period, with $\Pr[I = i] = C_i$.

The term $E^F[\max(B, C_1P_1 + \dots + C_nP_n)]$ denotes the expected value of $\max(B, \sum_i C_iP_i)$ when (P_1, \dots, P_n) is distributed according to F . The dimension of the class of “simple decision problems,” in the terms of Section 2, must be proportional to the size of the outcome set. This is because the domain of the second-order beliefs has dimension $n - 1$.

Proposition 2 *In Protocol (III), the individual announces his second-order belief as a strict best response.*

Proof. Let F be the individual's second-order belief, and \tilde{F} be the individual's announcement. For $C = (C_1, \dots, C_n)$, let $\psi(C)$ be the distribution function of the random variable $\sum_i C_iP_i$ when (P_1, \dots, P_n) is distributed according to the true second-order belief F , and similarly let $\tilde{\psi}(C)$ be the distribution function of the random

take values in the normalized interval $[0, 1]$, and then, instead of paying the subject $\$y$, pay the subject $\$1$ (or any fixed amount) with probability y , and pay nothing otherwise (or any smaller fixed amount). The idea of using “probability currency” to overcome the problem of risk aversion is discussed notably in Savage (1971) who attributes it to Smith (1961). It is common in the experimental literature (for more details, see, for example, Harrison et al., 2014 and references therein).

variable $\sum_i C_i P_i$ when (P_1, \dots, P_n) is distributed according to the announced second-order belief \tilde{F} . As the proof of Proposition 1 demonstrates, the expected payoff of the individual is

$$\int \left(\frac{1}{2} \left(1 - \varphi(B, \tilde{\psi}(C))^2 \right) + \varphi(B, \tilde{\psi}(C)) \varphi(B, \psi(C)) \right) dB dC,$$

which is maximized if and only if, for almost all tuples $(B, C) = (B, C_1, \dots, C_n)$, $\varphi(B, \tilde{\psi}(C)) = \varphi(B, \psi(C))$. By a continuity argument, this condition is equivalent to the condition that for all tuples (B, C) , $\varphi(B, \tilde{\psi}(C)) = \varphi(B, \psi(C))$, which in turn is equivalent to the condition that for all C , $\tilde{\psi}(C) = \psi(C)$, as shown in the proof of Proposition 1. By the Cramér-Wold Theorem, the distribution of a finite-dimensional random vector is uniquely determined by the distributions of its one-dimensional projections, and therefore, the condition that for all C , $\tilde{\psi}(C) = \psi(C)$ is equivalent to the condition that $\tilde{F} = F$. Hence, the individual maximizes his expected payoff if, and only if, he announces the true second-order belief. ■

3.2 Information Arriving at a Random Time

We now turn to an environment with more than one interim period. The individual's belief on the final outcome is still refined only once. However, this time of refinement is random and is neither controlled nor observed by the elicitor.

Section 3.1 deals with multiple outcomes, so for simplicity, we assume the final outcome is binary. Time periods are indexed $t = 0, 1, \dots, T$. At $t = 0$, the individual possesses a probability that the event materializes (a prior belief). At some future time $\tau \in \{1, \dots, T - 1\}$, he refines his initial assessment after observing a private or subjective signal, updating his prior to a posterior belief.

In the initial period, the individual's "dynamic belief" is captured by (i) a belief about *how much* he anticipates to learn, which we describe via a second-order belief F over the range $[0, 1]$ of possible posterior assessments, and (ii) a belief about *when* he anticipates to learn, described by a distribution over the possible dates $\{1, \dots, T - 1\}$. We assume F is nondegenerate, i.e., that F does not put full mass on the initial assessment (if it did, it would mean that the individual never refines his prior). Consider the following protocol.

Protocol (IV) *In the initial period, the individual announces a distribution F over posterior assessments, together with a distribution G over the times at which he anticipates updating his prior belief. The elicitor then draws numbers A and B independently and uniformly from $[0, 1]$. In addition, she draws a time t_c from $\{1, \dots, T - 1\}$, uniformly and independently. Let $\bar{p} = E^F[P]$ be the individual's prior assessment of the chance that the event occurs,¹¹ and construct the following distribution function H*

¹¹This assessment can be deduced from the belief reported, or can be reported separately.

over event probabilities:

$$H(p) = \begin{cases} G(t_c)F(p) + (1 - G(t_c)) & \text{if } p \geq \bar{p}, \\ G(t_c)F(p) & \text{if } p < \bar{p}. \end{cases}$$

If

$$A \geq E^H[\max(B, P)],$$

then the individual gets the payoff A and the protocol stops. Otherwise, in period t_c , the individual is offered the choice between getting the fixed payoff B , or getting the contingent payoff of 1 if the event occurs (and nothing otherwise).

In this mechanism, the value $E^H[\max(B, P)]$ equals the individual's expected payoff when the protocol does not immediately stop. The individual who understands this fact also understands that he cannot gain at manipulating his reports. Of course, to elicit the posterior, the elicitor could just ask for a revised probability assessment in each time period, offering a quadratic scoring rule.

Proposition 3 *In Protocol (IV), it is strictly optimal for the individual to honestly announce both his second-order belief and his belief about when he will receive his private signal.*

Proof. Let F be the individual's second-order belief, and G the believed distribution on the time of information arrival. Let \tilde{F} and \tilde{G} be the individual's announcements of these two beliefs, respectively. For every $t = 1, \dots, T - 1$, let

$$\tilde{H}_t(p) = \begin{cases} \tilde{G}(t)\tilde{F}(p) + (1 - G(t_c)) & \text{if } p \geq E^{\tilde{F}}[P], \\ \tilde{G}(t)\tilde{F}(p) & \text{if } p < E^{\tilde{F}}[P], \end{cases}$$

and let

$$H_t(p) = \begin{cases} G(t)F(p) + (1 - G(t)) & \text{if } p \geq E^F[P], \\ G(t)F(p) & \text{if } p < E^F[P]. \end{cases}$$

If we fixed the time t_c , the protocol would reduce to Protocol (II), and the distribution H_{t_c} would be elicited as a strict best response. Here, however, the time t_c can be any time period with positive probability, and so by Proposition 1, the individual's expected payoff is maximized if and only if, for every t , $\tilde{H}_t = H_t$, condition which, in turn, is equivalent to the condition that $\tilde{F} = F$ and $\tilde{G} = G$. (This equivalence is immediate, noting that if, for all t , $\tilde{H}_t = H_t$, then as $\tilde{G}(T - 1) = G(T - 1) = 1$, so $\tilde{F} = F$, which then implies $\tilde{G} = G$.) Overall, the individual's expected payoff is maximized if and only if both $\tilde{F} = F$ and $\tilde{G} = G$. ■

3.3 Two Interim Periods

We conclude with the case of two interim periods. This setting adds one period to the baseline setup of Section 2.

There are now four time periods: the initial period ($t = 0$), two interim periods ($t = 1, 2$), and the final period ($t = 3$). The outcome of interest concerns an event described by the indicator variable X , revealed at $t = 3$. At $t = 0$, the individual forms a first probabilistic appraisal about the event. In the next two interim periods, the individual receives information that may change his assessment. The information is modeled by signals S_1 and S_2 respectively, taking finitely many values. In the initial period, the individual holds a belief about the joint distribution on the triple (S_1, S_2, X) , which defines the individual's information structure. Signal S_2 contains information on random outcome X only, while signal S_1 may be informative on both X and signal S_2 .

The individual who has observed both signals makes a final probability assessment of the event. Similarly, the individual who has observed S_1 forms a belief on his future probability assessment (a second-order belief), and the individual who has not yet observed any signal holds a belief on the second-order belief he anticipates to have the next period (a third-order belief). Of course, at both times $t = 0, 1$ the individual can also appreciate the event likelihood, and, at $t = 0$, the distributions over the probabilistic beliefs he anticipates having, but these beliefs are redundant. Because there are finitely many signals, a second-order belief F can be described as a collection of pairs (f, p) , where f is the likelihood of obtaining final assessment p . A third-order belief μ can also be described as a collection of pairs (q, F) , where q is the likelihood of having second-order belief F in the next period. The goal is to elicit the individual's information structure, or equivalently, the individual's third-order beliefs.

Probability trees afford a simple graphical representation. Figure 2 gives an example when S_1 and S_2 are binary. The overall tree depicts the individual's belief at $t = 0$. The two subtrees express the possible second-order beliefs the individual may have at $t = 1$, with their probability shown on the branches. For instance, with the belief represented the right subtree, the individual's probability assessment of the event is $.5 \times .1 + .5 \times .9 = .5$, and the assessment will be revised to .9 or .1 in the next period: the individual anticipates being able to predict the outcome 90% of the time. In this example, the first signal is uninformative on the event itself, because the probabilistic appraisal remains 50%, but it indicates how informative the second signal will be.

The following protocol is an immediate extension of Protocol (II) of Section 2.

Protocol (V) *In the initial period, the individual is asked to announce his third-order belief μ . The elicitor then draws the numbers A, B and C uniformly and independently*

I am uninformed about the event, and will remain uninformed next period. I am unsure how much I will eventually learn, but I will know next period.




Figure 2: Example of a probability tree representing a third-order belief.

from $[0, 1]$, and computes

$$\pi = \sum_{(q,F) \in \mu} q \max \left(B, \sum_{(f,p) \in F} f \max(C, p) \right).$$


If $A \geq \pi$, the protocol stops and the individual gets the payoff A . Otherwise, in the first interim period, the individual is offered the choice either to stop and get the payoff B , or to continue. If the individual continues, then in the second interim period, the individual is offered the choice between getting the payoff C or getting the contingent payoff of 1 if the event occurs (and nothing otherwise).

The value of π corresponds to the expected payoff the individual would make if given the choice in the first interim period. So, as for the other protocols, the individual is always at least weakly better off announcing his true belief. However, in spite of the analogy with the baseline Protocol (II), this protocol generally fails to elicit third-order beliefs. Elicitation fails, for example, for beliefs as simple as those depicted in the probability trees of Figure 3.


Proposition 4 Protocol (V) does not elicit third-order beliefs as a strict best response.

The reason behind this lack of incentives is that the class of decision problems the protocol randomizes upon is not rich enough to differentiate between the elements of the comparatively larger set of possible third-order beliefs. To elicit beliefs successfully, we can either restrict the set of possible beliefs, or enrich the class of decision problems. We examine both possibilities.

Let us start with the first by considering the following restriction: the beliefs must be so that, for any tuple of second-order beliefs (F_1, \dots, F_n) , there exists some



(a)



(b)

Figure 3: Two probability trees not elicited by Protocol (V).

probability threshold x such that if $i \neq j$ then $E^{F_i}[\max(x, P)] \neq E^{F_j}[\max(x, P)]$, or equivalently $\int_x^1 F_i \neq \int_x^1 F_j$. When this condition is satisfied, we say that second-order beliefs are *restricted*.¹² For example, second-order beliefs are restricted when they can be ordered by second-order stochastic dominance, meaning that the possible signals of the first interim period are informative to different degrees.

Proposition 5 *If second-order beliefs are restricted, then Protocol (V) elicits the individual's third-order belief as a strict best response.*

We now examine the alternative possibility. We abstain from restricting beliefs but enrich the class of “simple decision problems.”

Protocol (VI) *In the initial period, the individual is asked to report his third-order belief μ . The elicitor then draws two numbers A, B uniformly and independently from $[0, 1]$. In addition, she draws J numbers C_1, \dots, C_J and J other numbers d_1, \dots, d_J uniformly and independently from $[0, 1]$. Let $D_i = d_i / (d_1 + \dots + d_J)$. The elicitor then*

¹²Of course, we already know that, for any two distinct and unrestricted second-order beliefs F and \tilde{F} , there always exists an x with $\int_x^1 F \neq \int_x^1 \tilde{F}$. The restriction imposes that this inequality extend beyond pairs of beliefs, to all tuples.

computes

$$\pi = \sum_{(q,F) \in \mu} q \max \left(B, \sum_{\substack{(f,p) \in F \\ 1 \leq i \leq J}} D_i f \max(C_i, p) \right).$$

If $A \geq \pi$, the protocol stops and the individual gets the payoff A . Otherwise, in the first interim period, the individual is offered the choice either to stop and get the payoff B , or to continue. If the individual continues, then in the second interim period, the elicitor selects $C = C_I$, with I drawn independently at random with $\Pr[I = i] = D_i$. The individual is then offered the choice between getting the payoff C or getting the contingent payoff of 1 if the event occurs (and nothing otherwise).

Proposition 6 *If third-order beliefs have support of size at most K and $J = 2K^2$, then Protocol (VI) elicits the individual’s third-order belief as a strict best response. Otherwise, if J is chosen randomly and $\Pr[J \leq j] < 1$ for every j , Protocol (VI) elicits the individual’s third-order belief as a strict best response, without restriction.*

4 Multiperiod Environments

In this section, we consider the case of any number of time periods, and larger outcome spaces (for the most general case, we refer the reader to Section S.2 of the Supplemental Material). Our purpose is threefold.

First, it is to demonstrate that the revealed-preference approach is general: with a well-chosen, large-enough class of “simple decision problems” and a suitable randomization over the class, we can elicit dynamic beliefs of any order in essentially arbitrary dynamic environments.

Second, one can use the protocols of this section for the elicitation and evaluation of dynamic beliefs or forecasts which do not conform to any of the instances investigated in Section 3. The fact that the protocols of this section work with general dynamic environments implies that they continue to elicit beliefs for any more specialized environment that may be of interest: one does not have to utilize the full dynamics to benefit from these protocols. In that sense, the most general protocols can be interpreted as “universal protocols,” which can be used to extract more or less refined information, depending on the application. While this use may seem excessive for simple dynamics, we emphasize that if we can assert the existence a class of decision problems that enables us to distinguish between different dynamic beliefs, in many settings, identifying the class of the relevant decision problems for making this distinction can be challenging, as Section 3.3 hints at.

Finally, we show that the family of protocols introduced can also approximate arbitrarily closely the payoffs of any sufficiently regular protocol. We believe that this near characterization can be convenient for the problem of selecting a protocol so as to maximize a given objective, possibly subject to some constraints. While a full analysis

is beyond the scope of this paper, we illustrate this idea in simple principal-agent problems in Section S.5 of the Supplemental Material.

Time periods are indexed $t = 0, \dots, T$. Period 0 is referred to as the *initial* period, period T the *final* period, and periods 1 to $T - 1$ are the *interim* periods. As in the preceding section, a random outcome X materializes in the final period. The random outcome takes values in a compact metrizable space \mathcal{X} , which covers the common cases $[a, b]^k$ and $\{1, \dots, n\}$. The individual privately or subjectively receives information gradually at each period. The individual holds probabilistic beliefs about this information, and the outcome. The goal is to elicit the high-order beliefs of the individual, which inform us about the individual's beliefs about the outcome, and also the individual's beliefs about the beliefs he anticipates having. For simplicity, the individual continues to be risk neutral and without discounting.¹³

Let $\Delta^1(\mathcal{X})$ be the set of distributions over \mathcal{X} , these are the set of first-order beliefs. Recursively let $\Delta^{k+1}(\mathcal{X}) = \Delta(\Delta^k(\mathcal{X}))$. The set $\Delta^k(\mathcal{X})$ is the set of the probability trees of level k , i.e., the k^{th} order beliefs. Endow each $\Delta^k(\mathcal{X})$ with the weak-* topology and the usual Borel σ -algebra.¹⁴ In the sequel we use the symbols p and q to denote probability trees of any level. In period t , the “dynamic belief” is a probability tree of level $T - t$. To avoid confusion, we use the subscript notation p_t to denote the high-order belief relevant in period t , and the superscript notation $p^{(k)}$ to denote a probability tree of level k .

The *elicitation protocol* describes the rules of interactions between the individual and the elicitor. By a revelation principle argument, every protocol is payoff-equivalent to a direct protocol, whereby the individual reveals directly his belief. That means that in every period t , the elicitor asks the individual to announce his $(T - t)^{\text{th}}$ order belief. The *payoff rule* Π of a direct elicitation protocol is the individual's overall expected payoff $\Pi(p_0, \dots, p_{T-1}, x)$, as a function of the final outcome x and the successive reports of high-order beliefs of the individuals p_t in period t . We require that Π be jointly measurable in its arguments, and we normalize payoffs to take values in $[0, 1]$.

The objective is to produce a protocol that induces the individual, as a strict best response, to communicate his dynamic beliefs truthfully in every period. Define an individual *strategy* as a family of maps $\{f_0, \dots, f_{T-1}\}$, where $f_t(p_0, \dots, p_t)$ gives the belief tree declared in period t as a function of the history of beliefs the individual has experienced up to period t (such definition rules out randomized strategies and dependence on other private information; it is, for our purpose, without loss). The

¹³Alternatively, it is sufficient to assume that, at every period, the individual behaves as if he was a subjective expected utility maximizer, believing that he will receive information in future periods. The individual's beliefs, or information, need not be consistent across time periods. For example, the protocols of this section properly elicit the beliefs of an individual who learns information he did not anticipate learning.

¹⁴The weak-* topology refers to the weakest topology for which, given any continuous function, integration with respect to that function is a continuous linear functional.

time- t expected payoff of the individual, under strategy f , is then

$$U(p_0, \dots, p_t; f) = \int \Pi(f_0(p_0), \dots, f_{T-1}(p_0, \dots, p_{T-1}), x) dp_{T-1}(x) \dots dp_t(p_{t+1}).$$

A strategy f is *optimal* for the history of beliefs p_0, \dots, p_t and a protocol with payoff rule Π if the individual who follows strategy f after having the sequence of beliefs p_0, \dots, p_t maximizes his payoff, no matter the strategy followed up to period t . Formally, for every pair of strategies (g, h) , where $g = \{h_0, \dots, h_{t-1}, f_t, \dots, f_{T-1}\}$, we have

$$U(p_0, \dots, p_t; g) \geq U(p_0, \dots, p_t; h).$$

Definition 1 *A protocol is strategyproof if*

- For all histories, an optimal strategy exists.
- For all histories (p_0, \dots, p_t) , and all optimal strategies f , $f_t(p_0, \dots, p_t) = p_t$.

4.1 A Family of Randomized Protocols

Central to our protocols are three instruments: *securities*, *menus of securities*, and *menus of (sub)menus*. A security is a continuous map $S : \mathcal{X} \rightarrow [0, 1]$ (continuity is irrelevant if the set of outcomes is discrete). It gives a payoff for every possible realization of the random outcome. 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, in (or before) period $T - 1$. A menu of order k gives the obligation to its owner to pick one (and only one) submenu among the collection it contains, in (or before) period $T - k$. Thus, an individual endowed with a menu of order k in period $T - k$ makes k choices at successive times $T - k, \dots, T - 1$, 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 submenus, 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 on his dynamic beliefs, captured by belief trees. Let us denote by $\pi_k(M_k, p^{(k)})$ the expected value of the menu M_k of order k in period $T - k$, to an individual who holds, as k^{th} order belief, a probability tree of level k , $p^{(k)}$. Recursively, we have

$$\begin{aligned} \pi_1(M_1, p^{(1)}) &= \max_{S \in \mathcal{M}_1} \int S(x) dp^{(1)}(x), \text{ and, if } k > 1, \\ \pi_k(M_k, p^{(k)}) &= \max_{m_{k-1} \in \mathcal{M}_k} \int \pi_{k-1}(m_{k-1}, p^{(k-1)}) dp^{(k)}(p^{(k-1)}). \end{aligned}$$

Our protocols randomize over large collections of menus.¹⁵ As a preliminary, the elicitor who administers the protocol draws a finite menu M_T of order T at random, according to a probability distribution ξ . The menu is known only to the elicitor. Then, in every period $t = 0, \dots, T - 2$, the elicitor asks the individual to reveal his full dynamic belief at that time—a belief tree of level $T - t$. She then chooses the submenu of M_{T-t} that is best according to the individual’s announcement: she selects a submenu $M_{T-t-1} \in M_{T-t}$ of highest expected value in period t ,

$$M_{T-t-1} \in \arg \max_{m_{T-t-1} \in M_{T-t}} \int \pi_{T-t-1}(m_{T-t-1}, p^{(T-t-1)}) dp^{(T-t)}(p^{(T-t-1)}).$$

Finally, in the penultimate period $T - 1$, the individual communicates a posterior distribution over X —a first-order belief. 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 *randomized menu protocols*. We stress that, while these protocols are presented in their general form, with an arbitrary randomization device—and so do not have the convenient “closed form” expression of the protocols of the previous sections—their implementation does not pose any particular difficulty: for baseline randomization devices, the protocol’s random payoffs can be computed efficiently via a simple algorithm. We give an example of implementation in Section S.1 of the Supplemental Material.

For a given randomized menu protocol with randomization device ξ , we denote by $\Pi(p_0, \dots, p_{T-1}, x; M)$ the payoff to the individual who announces p_t in period t , when the realization of X is x , and if the elicitor draws menu $M \in \mathcal{M}_T$ in the initial period. Then, the payoff rule of the overall protocol is expressed as

$$\Pi(p_0, \dots, p_{T-1}, x; \xi) = \int \Pi(p_0, \dots, p_{T-1}, x; M) d\xi(M). \quad (7)$$

¹⁵To 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} , 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 of [Aliprantis and Border, 2006](#)).

¹⁶At every stage, if there is more than one submenu or one security that is optimal for the individual, the elicitor selects a submenu uniformly at random among all optimal submenus. Selecting a submenu uniformly at random guarantees the measurability of the payoff rule. Alternatively, the individual could get an equal fraction of all optimal submenus, or he could get any optimal submenu according to a measurable selection. In the proof of Lemma 3 in Appendix C, we show that such a measurable selection is guaranteed to exist.

4.2 Existence and Uniqueness

Our next result asserts that strict incentives are implemented by the protocols of the class just described when the probability measure ξ is full support. 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.

Theorem 1 *If a randomized menu protocol randomizes according to a full-support distribution, then it is strategyproof.*

The randomized menu protocols follow the general approach illustrated in Section 2, in which the class of simple decision problems is the class of finite choice menus.

The key challenge in the elicitation of dynamic beliefs in a multiperiod environment is that the beliefs become naturally richer as the number of periods increases. This creates two difficulties. The first one is to find the simple decision problems that allow the distinction between two given beliefs. To overcome this issue, we operate on the relatively large class of decision problems that are the menu choices. The second difficulty is that we must randomize over the simple decision problems in a way that preserves the incentives of the individual, making sure that enough mass is put on the decision problems that matter for the separation of beliefs. Because we operate on a large class of decision problems, it can be intricate to ensure that the randomization is “proper.”

To illustrate, suppose we have a continuum of decision problems indexed by $d \in [0, 1]$, and some problem d_0 turns out to be crucial to separate between some beliefs. We would then want to put a positive weight on d_0 —so the uniform distribution, for example, would not work. We would want to put a mass specifically on d_0 , but we may not be able pinpoint d_0 precisely. And if every problem $d \in [0, 1]$ turned out to be crucial to separate between beliefs—perhaps due to the richness or complexity of the beliefs—then no randomization scheme would work. For that reason, the problem of identification we face is different from the standard problem of identification in decision theory. We surmount this issue by including only finite menus (which also helps with the implementation) and by controlling the amount of data required to encode the high-order beliefs (formally given by the σ -algebra of possible events): we ensure that a $(k + 1)^{\text{th}}$ order belief is no more complex than k^{th} order belief, if k becomes large. Then, perhaps surprisingly, the class of finite menus is sufficiently large to make it possible to recover the full hierarchy of beliefs.

Notice that, in the above protocols, the elicitor does not disclose her menu choices to the individual, as if she did, the property of strict incentives would be lost. Of course, this is not a limitation: the elicitor can first collect the sequence of all the announcements of the individual, and only after the last announcement is received she draws a menu and operates on it as in the original protocol. This is a less literal but more natural interpretation of the above protocols, our example of implementation

in Section S.1 of the Supplemental Material follows this alternative. In addition, if we are only interested in the elicitation of the individual’s belief in the initial period, then there is no loss in disclosing the menu randomly drawn and the subsequent menu choices once the individual has communicated the initial belief.

As a corollary, Theorem 1 implies that beliefs are identified in a class of decision makers who hold preferences over hierarchies of finite menus. Each high-order belief $p \in \Delta^T(\mathcal{X})$ induces a preference over finite menus of order T captured by a binary relation \succeq_p over \mathcal{M}_T defined by $M \succeq_p M'$ if and only if

$$\int \pi_T(M, q) dp(q) \geq \int \pi_T(M', q) dp(q).$$

Lemma 4 in the proof of Theorem 1 then implies the following fact, extending a result of Takeoka (2007) and Dillenberger et al. (2014).

Corollary 1 *Let $p, p' \in \Delta^T(\mathcal{X})$. If $\succeq_p = \succeq_{p'}$, then $p = p'$.*

Our final result argues that under regularity conditions, any protocol that is strategyproof is approximately payoff-equivalent to some randomized menu protocol. Hence, there is no loss of generality in focusing on randomized menu protocols: the mechanisms we describe provide an essential characterization of the class of strategyproof mechanisms.

Theorem 2 *Consider a strategyproof protocol whose payoff rule $\Pi(p_0, \dots, p_{T-1}, x)$ is jointly continuous. Then, for every $\epsilon > 0$, there exists a strategyproof randomized menu protocol whose payoff rule $\Pi'(p_0, \dots, p_{T-1}, x)$ satisfies*

$$|\Pi(p_0, \dots, p_{T-1}, x) - \Pi'(p_0, \dots, p_{T-1}, x)| < \epsilon$$

for all p_0, \dots, p_{T-1}, x .

5 Conclusion

We have considered a dynamic analogue of the probability scoring rules. To induce truthful announcements, we develop a new constructive approach, based on randomly selecting among a sufficiently large number of simple dynamic decision problems, and operating as if we were asking the individual whose information is being elicited to solve all these problems at the same time. This approach applies quite broadly. It enables us to derive simple protocols for a range of common instances of dynamic environments, and it extends to general dynamic environments. We have set ourselves up for the most difficult version of the problem: the elicitor sees nothing along the way. If she can observe some of the information that the individual observes, so that the individual’s information is only partially private, it only makes it easier for her

to solve the incentive problem. Of course, she could use her information to expand the set of possible mechanisms by conditioning the individual’s payoff on this piece of public information, but she need not need to do so: the protocols of this paper, that condition payments on the final outcome only, still fully retain strict incentive compatibility.

Appendices

A Stage-Separated Protocols

The purpose of this appendix is to show that eliciting a first-order belief via some method that uses the final outcome as observable information, and, separately, eliciting the second-order belief via a method that uses the elicited first-order belief as observable information, does not induce truthful responses.

We focus on the three-period case (the impossibility result extends directly to any number of periods), and we borrow notation and terminology from Section 3.1. The random variable X takes values in \mathcal{X} defined as $\{1, \dots, n\}$, $n \geq 2$. The elicitor asks the individual to disclose his second-order belief $F \in \Delta(\Delta(\mathcal{X}))$ in the initial period, his first-order belief $p \in \Delta(\mathcal{X})$ in the interim period, and finally, when the random variable materializes to value x , she rewards the individual with a payoff equal to $\Pi(F, p, x)$ (on average, if the protocol is randomized).

We ask if we can choose a strategyproof payoff rule Π (following the definition of strategyproofness in Section 4) of the form $\Pi(F, p, x) = \Pi_1(F, p) + \Pi_2(p, x)$; that is, we separate stages, the individual gets a first payoff after announcing the second- and first-order beliefs, and a second payoff after the random variable realizes that depends only on the reported first-order belief and the realization. When the payoff rule of a protocol satisfies this condition, we say it is *stage separated*. Stage-separated protocols have a natural interpretation: they use the publicly observed outcome of X to elicit the posterior p through Π_2 , and then, using p , they attempt to elicit the prior F via Π_1 . For example, Π_1 and Π_2 could be the payoffs of classic probability elicitation methods, such as the quadratic score or the BDM mechanism for Π_2 , and one of the Matheson-Winkler elicitation methods (Matheson and Winkler, 1976) for Π_1 .

Proposition 7 *If a protocol is stage separated, then the protocol is not strategyproof.*

To understand the intuition behind this result, suppose that the protocol satisfies some smoothness conditions and that, absent a first stage, the protocol would induce the individual to report truthfully his first-order belief. Let us focus on the individual’s decision in the interim period. Assume that the individual has reported his true second-order belief F , that his true first-order belief is p , but that he reports $p + \Delta p$. The

expected payoff difference due to his deviation is

$$\Pi_1(F, p + \Delta p) - \Pi_1(F, p) + \Pi_2(p + \Delta p, p) - \Pi_2(p, p),$$

where $\Pi_2(\tilde{p}, p)$ designates the individual's expected payoff in the interim period when he reports \tilde{p} while his true belief is p . Because $\Pi_2(\tilde{p}, p)$ is maximized when $\tilde{p} = p$, we expect the second term $\Pi_2(p + \Delta p, p) - \Pi_2(p, p)$ to be of order at most $\|\Delta p\|^2$, under smoothness conditions. However, unless $\Pi_2(F, \tilde{p})$ is constant in \tilde{p} , we also expect the first term $\Pi_1(F, p + \Delta p) - \Pi_1(F, p)$ to be of order $\|\Delta p\|$, for at least some instances of p . Thus, there are situations in which the gains realized from the first stage when deviating from the truth in the interim period exceed the losses incurred in the second stage: the protocol is not strategyproof. The formal proof follows.

Proof of Proposition 7. Consider a stage-separated protocol. For every declared second-order belief \tilde{F} , let $g_{\tilde{F}}(\tilde{p}, x)$ be the total payoff (or average total payoff, if the protocol is randomized) to the individual as a function of the announced first-order belief \tilde{p} and realization x :

$$g_{\tilde{F}}(\tilde{p}, x) = \Pi_1(\tilde{F}, \tilde{p}) + \Pi_2(\tilde{p}, x).$$

Suppose that $g_{\tilde{F}}(p, p) > g_{\tilde{F}}(\tilde{p}, p)$ for every $\tilde{p} \neq p$, where $g_{\tilde{F}}(\tilde{p}, p)$ is the individual's total expected payoff given his realized posterior belief p —this inequality would be required of any strategyproof protocol. Let $\bar{g}_{\tilde{F}}$ be the map on $\Delta(\mathcal{X})$ defined by $\bar{g}_{\tilde{F}}(p) = g_{\tilde{F}}(p, p)$. Note that $\bar{g}_{\tilde{F}}$ is convex, so the preceding inequality can be interpreted saying that the map $x \mapsto \Pi_1(\tilde{F}, \tilde{p}) + \Pi_2(\tilde{p}, x)$ is a subgradient of $\bar{g}_{\tilde{F}}$ at point \tilde{p} . Because the domain of $\bar{g}_{\tilde{F}}$ is the simplex, the map $x \mapsto \Pi_2(\tilde{p}, x)$ is also a subgradient. Thus the convex functions $\bar{g}_{\tilde{F}}$ share the same subgradients. In particular, for every $p', p'' \in \Delta(\mathcal{X})$,

$$\bar{g}_{\tilde{F}}(p'') - \bar{g}_{\tilde{F}}(p') = \int_0^1 (p'' - p') \cdot \Pi_2(\alpha p'' + (1 - \alpha)p', \cdot) d\alpha$$

where $p \cdot q$ is the dot product between p and q on the simplex $\Delta(\mathcal{X})$ interpreted as a subset of \mathbf{R}^n . Thus for all F, \tilde{F} , we get that $\bar{g}_F - \bar{g}_{\tilde{F}}$ is constant: in the initial period, the individual is best off reporting any \tilde{F} that maximizes $\bar{g}_{\tilde{F}}(\tilde{p})$, for an arbitrary \tilde{p} , independently of his true second-order belief. This fact means that the protocol is not strategyproof. ■

It can be seen that the payoff rule associated to the mechanism suggested by [Karni \(2018, 2020\)](#) is equivalent to the sum of a Matheson-Winkler score and a quadratic scoring rule. Therefore, one reading of his work is that, by increasing the magnitude of the payoffs in the second stage comparatively to the payoffs of the first stage, one can get the individual to make reports increasing closer to his true belief for a range of decision models.

B Proofs of Section 3

B.1 Proof of Proposition 4

Let μ and $\tilde{\mu}$ be the two probability trees of Figure 3:

$$\mu = \left\{ \left(\frac{1}{2}, F_1 \right), \left(\frac{1}{2}, F_2 \right) \right\} \quad \text{and} \quad \tilde{\mu} = \left\{ \left(\frac{1}{2}, \tilde{F}_1 \right), \left(\frac{1}{2}, \tilde{F}_2 \right) \right\},$$

with

$$F_1 = \left\{ \left(\frac{1}{4}, 0 \right), \left(\frac{1}{2}, \frac{1}{3} \right), \left(\frac{1}{4}, 1 \right) \right\},$$

$$F_2 = \left\{ \left(\frac{1}{2}, 0 \right), \left(\frac{1}{2}, \frac{2}{3} \right) \right\},$$

and

$$\tilde{F}_1 = \left\{ \left(\frac{1}{2}, 0 \right), \left(\frac{1}{4}, \frac{1}{3} \right), \left(\frac{1}{4}, 1 \right) \right\},$$

$$\tilde{F}_2 = \left\{ \left(\frac{1}{4}, 0 \right), \left(\frac{1}{4}, \frac{1}{3} \right), \left(\frac{1}{2}, \frac{2}{3} \right) \right\}.$$

For any second-order belief F , let $\Pi(F; C)$ be the expected payoff of the individual at $t = 1$ when the choice is to continue. Simple calculations yield

$$\Pi(F_1; C) = \begin{cases} \frac{5}{12} + \frac{1}{4}C & \text{if } 0 \leq C \leq \frac{1}{3} \\ \frac{1}{4} + \frac{3}{4}C & \text{if } \frac{1}{3} \leq C \leq 1 \end{cases}, \quad \Pi(F_2; C) = \begin{cases} \frac{1}{3} + \frac{1}{2}C & \text{if } 0 \leq C \leq \frac{2}{3} \\ C & \text{if } \frac{2}{3} \leq C \leq 1 \end{cases},$$

and

$$\Pi(\tilde{F}_1; C) = \begin{cases} \frac{1}{3} + \frac{1}{2}C & \text{if } 0 \leq C \leq \frac{1}{3} \\ \frac{1}{4} + \frac{3}{4}C & \text{if } \frac{1}{3} \leq C \leq 1 \end{cases}, \quad \Pi(\tilde{F}_2; C) = \begin{cases} \frac{5}{12} + \frac{1}{4}C & \text{if } 0 \leq C \leq \frac{1}{3} \\ \frac{1}{3} + \frac{1}{2}C & \text{if } \frac{1}{3} \leq C \leq \frac{2}{3} \\ C & \text{if } \frac{2}{3} \leq C \leq 1 \end{cases}.$$

Similarly, let $\Pi(\mu; B, C)$ be the expected payoff at $t = 0$ of an individual with third-order belief μ assuming the protocol continues at least to the next period. Let $\Pi(\tilde{\mu}; B, C)$ be the expected payoff for third-order belief $\tilde{\mu}$.

If $0 \leq C \leq 1/3$, then $\Pi(F_1; C) = \Pi(\tilde{F}_2; C)$ and $\Pi(F_2; C) = \Pi(\tilde{F}_1; C)$. So, for all B ,

$$\begin{aligned}\Pi(\mu; B, C) &= \frac{1}{2} \max(B, \Pi(F_1; C)) + \frac{1}{2} \max(B, \Pi(F_2; C)) \\ &= \frac{1}{2} \max(B, \Pi(\tilde{F}_2; C)) + \frac{1}{2} \max(B, \Pi(\tilde{F}_1; C)) \\ &= \Pi(\tilde{\mu}; B, C).\end{aligned}$$

Similarly, if $1/3 \leq C \leq 1$, then $\Pi(F_1; C) = \Pi(\tilde{F}_1; C)$ and $\Pi(F_2; C) = \Pi(\tilde{F}_2; C)$. So, for all B ,

$$\begin{aligned}\Pi(\mu; B, C) &= \frac{1}{2} \max(B, \Pi(F_1; C)) + \frac{1}{2} \max(B, \Pi(F_2; C)) \\ &= \frac{1}{2} \max(B, \Pi(\tilde{F}_1; C)) + \frac{1}{2} \max(B, \Pi(\tilde{F}_2; C)) \\ &= \Pi(\tilde{\mu}; B, C).\end{aligned}$$

Altogether, for all B, C , $\Pi(\mu; B, C) = \Pi(\tilde{\mu}; B, C)$ and the decision of the elicitor, at $t = 1$, is the same whether the announced third-order belief is μ or $\tilde{\mu}$, independently of the draw of A, B, C . Hence, announcing $\tilde{\mu}$ when one's true belief is μ (and conversely) is still a best response.

B.2 Proof of Proposition 5

We have already argued that truthful announcements are a weak best response, at least. We show that the best response is strict. As for Proposition 4, $\Pi(\mu; B, C)$ denotes the expected payoff of an individual, at $t = 0$, whose third-order belief is μ , and who is about to face the choice in the first interim period:

$$\Pi(\mu; B, C) = \sum_{(q, F) \in \mu} q \max \left(B, \sum_{(f, p) \in F} f \max(C, p) \right).$$

Let μ and $\tilde{\mu}$ be distinct third-order beliefs whose supports lie within the restricted class considered. If there exist B, C such that $\Pi(\mu; B, C) \neq \Pi(\tilde{\mu}; B, C)$ then, by continuity of $\Pi(\mu; B, C)$ in B and C , there exists a positive mass of triples (A, B, C) such that

$$\min(\Pi(\mu; B, C), \Pi(\tilde{\mu}; B, C)) < A < \max(\Pi(\mu; B, C), \Pi(\tilde{\mu}; B, C)),$$

so that an individual whose belief is μ and who announces $\tilde{\mu}$ gets a strictly suboptimal payoff with positive probability on the random triple (A, B, C) —and so gets a strictly suboptimal payoff overall.

The proof then reduces to demonstrating the existence of B and C such that $\Pi(\mu; B, C) \neq \Pi(\tilde{\mu}; B, C)$. It is convenient to assume that μ and $\tilde{\mu}$ share the same

probability trees that characterize the second-order beliefs. To do so we write

$$\mu = \{(q_i, F_i)\}_{1 \leq i \leq n} \quad \text{and} \quad \tilde{\mu} = \{(\tilde{q}_i, F_i)\}_{1 \leq i \leq n}$$

with $F_i \neq F_j$ if $i \neq j$ and possibly $q_i = 0$ or $\tilde{q}_i = 0$.

Let C be such that the n payoffs $\Pi(F_1; C), \dots, \Pi(F_n; C)$ can be totally ordered, for example, $\Pi(F_1; C) < \dots < \Pi(F_n; C)$. It is then easily verified that the linear span of the set of vectors of \mathbf{R}^n

$$\{(\max(B, \Pi(F_1; C)), \dots, \max(B, \Pi(F_n; C))) \mid B \in [0, 1]\} \quad (8)$$

is \mathbf{R}^n (for instance, by varying B gradually on its range, one can progressively construct the vectors $(0, \dots, 0, 1)$, $(0, \dots, 0, 1, 1)$, and so forth, to make a basis). As $\mu \neq \tilde{\mu}$, $(q_1, \dots, q_n) \neq (\tilde{q}_1, \dots, \tilde{q}_n)$, and the set of vectors (8) has full rank, there exists B such that

$$\sum_i q_i \max(B, \Pi(F_i; C)) \neq \sum_i \tilde{q}_i \max(B, \Pi(F_i; C)),$$

and hence $\Pi(\mu; B, C) \neq \Pi(\tilde{\mu}; B, C)$.

B.3 Proof of Proposition 6

Let $\Pi(F; C, D)$ be the expected payoff of an individual in the first interim period, with second-order belief F and who chooses to continue, and for the parameters $C = (C_1, \dots, C_J)$ and $D = (D_1, \dots, D_J)$. Let μ and $\tilde{\mu}$ be distinct third-order beliefs and suppose that the support of each of these third-order beliefs has size at most K . Let $\mathcal{F} = \{F_1, \dots, F_M\}$ be the union of the two supports, so that $M \leq 2J$.

The key argument in Proposition 5 relies on the fact that for some draws of the elicitor, *any* two second-order beliefs in \mathcal{F} yield different expected payoffs of the continuing individual in the first interim period. In that proposition, the fact is simply assumed, by having restricted second-order beliefs. In this proposition, we show the fact holds in the modified protocol. The same argument then continues to apply.

Let q_i be the probability, according to μ , of obtaining belief F_i in the first interim period, and let \tilde{q}_i be the analog for $\tilde{\mu}$. As argued in the proof of Proposition 1, for every $i \neq j$, there exists $\alpha_{ij} \in [0, 1]$ such that $E^{F_i}[\max(\alpha_{ij}, P)] \neq E^{F_j}[\max(\alpha_{ij}, P)]$. The remainder of the proof requires the following result.

Lemma 1 *Let $\mathcal{C} \subset \mathbf{R}^M$ be finite. If, for all $i \neq j$, there exists $X \in \mathcal{C}$ such that $X_i \neq X_j$, then there exists a convex combination Y of the vectors of \mathcal{C} such that for all $i \neq j$, $Y_i \neq Y_j$.*

Proof. Let us start with an arbitrary $Y \in \mathcal{C}$ and apply the following iterative procedure. For any pair $i \neq j$ such that $Y_i = Y_j$, we transform Y into $\alpha X + (1 - \alpha)Y$, where X is a vector of \mathcal{C} such that $X_i \neq X_j$ and $\alpha \in (0, 1)$. The transformed Y satisfies $Y_i \neq Y_j$, and if α is chosen small enough, then all pairs of different elements

under the original vector Y remain pairs of different elements under the transformed vector Y . We iterate this process while there is any remaining pair $i \neq j$ with $Y_i = Y_j$. As there are only finitely many pairs, the procedure terminates and generates a vector whose elements are pairwise different. ■

Returning to the proof of Proposition 6, let X^{ij} be the vector of \mathbf{R}^M

$$X^{ij} = (E^{F_1}[\max(\alpha_{ij}, P)], \dots, E^{F_J}[\max(\alpha_{ij}, P)]),$$

and let \mathcal{C} be the collection of the vectors X^{ij} for every pair (i, j) with $i < j$. There are at most $2K(2K - 1)/2 < 2K^2 = J$ elements in \mathcal{C} , and by Lemma 1, there exists a vector Y written as convex combination of elements of \mathcal{C} such that for every $i \neq j$, $Y_i \neq Y_j$. Therefore, for some vectors $C = (C_1, \dots, C_J)$ and $D = (D_1, \dots, D_J)$, with $0 \leq D_\ell \leq 1$, and such that for every ℓ , $C_\ell = \alpha_{ij}$ for some i, j , element Y_k of vector Y is equal to

$$\sum_{\ell=1}^J D_\ell E^{F_k}[\max(C_\ell, P)] = \Pi(F_k; C, D).$$

Hence, for all $F, \tilde{F} \in \mathcal{F}$, $F \neq \tilde{F}$, there exists C and D for which $\Pi(F; C, D) \neq \Pi(\tilde{F}; C, D)$: any two second-order beliefs in \mathcal{F} yield different expected payoffs of the continuing individual in the first interim period.

C Proofs of Section 4

C.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 2 *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 $(M_k, p^{(k+1)}) \mapsto \int \pi_k(M_k, q) dp^{(k+1)}(q)$, 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 $f_0(S, p^{(1)}) = \int S dp^{(1)}$ and for $k \geq 1$ let $f_k(M_k, p^{(k+1)}) = \int \pi_k(M_k, q) dp^{(k+1)}(q)$.

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 $h_{k+1}(M_{k+1}, p^{(k+1)}) = M_k$. 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 $\lim_{i \rightarrow \infty} \sup_{u' \in \mathcal{U}'} |h_{\mathcal{K}_{k+1}(M^{(i)})}(u') - h_{\mathcal{K}_{k+1}(M^\infty)}(u')| = 0$ 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,

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

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 4. 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 3 *The map $(p^{(T)}, \dots, p^{(1)}, x, M_T) \mapsto \Pi^{(T)}(p^{(T)}, \dots, p^{(1)}, x; M_T)$, where $p^{(k)} \in \Delta^k(\mathcal{X})$, $M_T \in \mathcal{M}_T$, and $x \in \mathcal{X}$, is jointly measurable in the product σ -algebra.*

Proof. As in Lemma 2, for every k we define the correspondence $h_k(M_k, p^{(k)}) = M_k$, and the function $f_k(M_k, p^{(k+1)}) = \int \pi_k(M_k, q) dp^{(k+1)}(q)$.

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.¹⁷ We can then apply the Measurable Selection Theorem (Theorem 18.19 of Aliprantis and Border, 2006), and we get that the argmax correspondence

$$\arg \max_{m \in h_{k+1}(M_{k+1}, p^{(k+1)})} \int \pi_k(m, q) dp^{(k+1)}(q)$$

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 $\{\Phi_{k+1}^{(i)}; i = 1, 2, \dots\}$ such that

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

We observe that

$$\left| \arg \max_{m \in M_{k+1}} \int \pi_k(m, q) dp^{(k+1)}(q) \right| = \lim_{j \rightarrow \infty} \sum_{i=1}^j \frac{1}{\sum_{\ell=1}^j \mathbb{1}_{\Phi_{k+1}^{(i)}(M_{k+1}, p^{(k+1)}) = \Phi_{k+1}^{(\ell)}(M_{k+1}, p^{(k+1)})}}$$

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

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

becomes measurable. This concludes the proof. ■

¹⁷First, 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.

C.2 Proof of Theorem 1

The proof consists of two parts. The first part deals with the separation of different individuals *at a given time* when there are *only two possible types*. In the multiperiod case, we use for “simple decision problems” the class of finite menus. Decisions then consist in choosing an element from the menu at the initial period, then an element from the chosen submenu at the next period, and so forth until the penultimate period 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 individuals whose belief trees are of two possible sorts. In the second part of the proof, we apply the Allais randomization idea to discriminate between any two individuals whose belief trees are no longer restricted.

C.2.1 Part 1: Discriminating Between Two Belief Trees

Let $p^{(k)}$ and $q^{(k)}$ be two different probability trees of level k , that represent the dynamic beliefs of two individuals in period $T - k$. We refer to the individual with (dynamic) belief $p^{(k)}$ as *type* $p^{(k)}$, and the individual with belief $q^{(k)}$ 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} in period $T - k$, 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 period $T - 1$, 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 map $f : \mathcal{X} \rightarrow [0, 1]$ 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:

$$\int f dp^{(1)} \neq \int f dq^{(1)}.$$

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, f is bounded, so that we can choose f to take values in $[0, 1]$. For example, suppose $\int f dp^{(1)}$ is greater than $\int f dq^{(1)}$. Then we can set $S^p = f$ and S^q to be the average of $\int f dp^{(1)}$ and $\int f dq^{(1)}$. A symmetric argument holds if $\int f dp^{(1)}$ is less than $\int f dq^{(1)}$. For this argument to work, the key element is to have essentially complete flexibility in the design of the security—which is also the individual’s value function at the next and final period T .

Now consider the problem of separating individuals 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 $M_k \in \mathcal{M}_k$ in period $T - k$ to any individual who holds belief tree $\mu^{(k)}$ at that time (that is, $\pi_{M_k}(\mu^{(k)}) = \pi_k(M_k, \mu^{(k)})$).

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

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

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 an 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 4 *For every $k \in \{1, \dots, T\}$, $p^{(k)}, q^{(k)} \in \Delta^k(\mathcal{X})$ 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*

$$\int \pi_{M_{k-1}} dp^{(k)} \neq \int \pi_{M_{k-1}} dq^{(k)}. \quad (9)$$

Proof. The proof proceeds by induction. As shown above, Equation (9) 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 $M_k^{pq} := \{M_{k-1}^p, M_{k-1}^q\}$ in period $T - k$, 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,

$$\int \pi_{M_{k-1}} dp^{(k)} > \int \pi_{M_{k-1}} dq^{(k)},$$

we set $M_{k-1}^p = M_{k-1}$ and $M_{k-1}^q = \frac{1}{2} (\int \pi_{M_{k-1}} dp^{(k)} + \int \pi_{M_{k-1}} dq^{(k)})$. 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}_k 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}_ℓ 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 $\mathbf{1} \in \mathcal{M}_k$ be the (degenerate) menu that generate the constant payoff 1, and $\mathbf{0} \in \mathcal{M}_k$ be the (degenerate) menu that generate the constant payoff 0. The following equalities hold for each $\mu \in \Delta^k(\mathcal{X})$ and each $M, M' \in \mathcal{M}_k$, :

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

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

$$\max\{\pi_M - \pi_{M'}, \pi_N - \pi_{N'}\} = \max\{\pi_M + \pi_{N'}, \pi_N + \pi_{M'}\} - (\pi_{M'} + \pi_{N'}) \quad (10)$$

$$= \pi_{(\pi_M + \pi_{N'}) \cup (\pi_N + \pi_{M'})} - (\pi_{M'} + \pi_{N'}). \quad (11)$$

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 $\alpha \in \mathbf{R}$.¹⁸
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.¹⁹

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

$$\int \pi_{M_k} dp^{(k+1)} = \int \pi_{M_k} dq^{(k+1)}$$

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

$$\int f dp^{(k+1)} = \int f dq^{(k+1)}$$

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

$$\int \pi_{M_k} dp^{(k+1)} \neq \int \pi_{M_k} dq^{(k+1)},$$

which concludes the proof by induction. ■

¹⁸By the boolean ring property $\alpha\mathcal{L}_k \subseteq \mathcal{L}_k$ if $\alpha \geq 0$, and by definition of \mathcal{L}_k , $-\mathcal{L}_k \subseteq \mathcal{L}_k$.

¹⁹At a general level, we need that the linear span of the value functions associated to menu protocols is dense in \mathcal{B}_k . It is known that the linear span of the convex functions on $\Delta^k(\mathcal{X})$ is dense in \mathcal{B}_k in the topology of uniform convergence (see, for example, [Meyer, 1966](#), p. 221). However, for our purpose, this result is not sufficient because the value functions form a strict subset of convex functions; specifically, they correspond to the convex functions on $\Delta^k(\mathcal{X})$ whose recursive subgradients are also convex. Therefore Step 2 goes a step further and shows that the linear span of the convex functions on $\Delta^k(\mathcal{X})$ whose recursive subgradients are also convex is dense in \mathcal{B}_k .

C.2.2 Part 2: Randomization

In this second part, we show that a full-support randomization over finite menus allows to distinguish between any two individuals 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- T menus \mathcal{M}_T . Fix any two sequences of belief trees $\mathbf{p} = \{p^{(T)}, \dots, p^{(1)}\}$ and $\mathbf{q} = \{q^{(T)}, \dots, q^{(1)}\}$ 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_T drawn at random according to ξ , the individual who is given menu M_T 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 individual who observes the sequence of belief trees \mathbf{q} . We refer to the individual who observes \mathbf{p} as type \mathbf{p} , and the individual who observes \mathbf{q} as type \mathbf{q} .

Fix an arbitrary level k such that $p^{(k)} \neq q^{(k)}$, and let $M_k^* = \{M_{k-1}^{p,*}, M_{k-1}^{q,*}\}$ 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 $M_{k-1}^{p,*}$ and $M_{k-1}^{q,*}$ denote securities. Define the (degenerate) menu of level N , M_N^* , which includes only M_k^* , i.e., either $M_N^* = M_k^*$ if $k = N$, otherwise $M_N^* = \{\dots \{M_k^*\} \dots\}$. For such a menu, there is no decision to be made until period $T - k$ 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_T selected anywhere in small enough neighborhood of M_T^* , type $\mathbf{p} := (p^{(T)}, \dots, p^{(1)})$ is strictly better off choosing a different submenu/security than type $\mathbf{q} := (q^{(T)}, \dots, q^{(1)})$, for every optimal selection of type \mathbf{q} .

By Step 1 and Lemma 4, there exists $\epsilon > 0$ such that for any level- k menus M_k, M'_k with $d(M_k, M_k^{p,*}) < \epsilon$ and $d(M'_k, M_k^{q,*}) < \epsilon$, type $p^{(k)}$ would be strictly better off choosing M_k over M'_k at $t = T - k$, while type $q^{(k)}$ would be strictly better off choosing M'_k over M_k .

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

C.3 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 randomized menu 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 submenus, the payoffs associated with that choice remain close to the payoffs of the original payoff rule.

For any finite menu M of order T , let $\Pi^*(p_0, \dots, p_{T-1}, x; M)$ be the induced payoff rule, and let $\Pi^*(p_0, \dots, p_{T-1}, x; \xi)$ be the payoff rule induced by the randomized menu protocol that randomizes according to ξ . Let us slightly abuse notation and denote by

$$\Pi_t(q_0, \dots, q_t; p_t)$$

the maximum expected value, in period t , of the individual who faces payoff rule Π and who reports q_0, \dots, q_t from period 0 to period k , but holds belief p_t in period t . Similarly,

$$\Pi_t^*(q_0, \dots, q_t; p_t; M)$$

is the maximum expected value of the individual endowed with the finite menu M in the initial period 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 $\Delta^T(\mathcal{X}) \times \dots \times \Delta(\mathcal{X}) \times \mathcal{X}$, which is a compact set, it is uniformly continuous. Thus, there exists $\delta_0 > 0$ such that if, for each i , p_i is δ_0 -close to p'_i , i.e., $d(p_i, p'_i) < \delta_0$, then $|\Pi(p_0, \dots, p_{T-1}, x) - \Pi(p'_0, \dots, p'_{T-1}, x)| < \epsilon/2$ for each $x \in \mathcal{X}$.

Step 1(a). We show that there exists a finite subset Σ_0 of $\Delta^N(\mathcal{X})$ such that, for each p_0 , if

$$q_0^* \in \arg \max_{q_0 \in \Sigma_0} \Pi_0(q_0; p_0),$$

then q_0^* is δ_0 -close to p_0 .

Let $\{\Sigma_{0,k}\}_k$ be a sequence of finite subsets of $\Delta^T(\mathcal{X})$ such that $\Sigma_{0,k}$ converges to $\Delta^T(\mathcal{X})$ in the Hausdorff metric topology induced by the Lévy-Prokhorov metric. The compactness of $\Delta^T(\mathcal{X})$ guarantees existence of such a sequence. We observe that $(q_0, p_0) \mapsto \Pi_0(q_0; p_0)$ is continuous—as can be seen immediately via induction, using that every Π_t is uniformly continuous. The correspondence $(\mathcal{P}, p_0) \rightarrow \mathcal{P}$, where \mathcal{P} is a compact subset of $\Delta^T(\mathcal{X})$ and $p_0 \in \Delta^T(\mathcal{X})$ is also continuous (see Theorem 18.10 of [Aliprantis and Border, 2006](#)). Using Berge's Maximum Theorem, we get that the correspondence

$$(\mathcal{P}, p_0) \rightarrow \arg \max_{q_0 \in \mathcal{P}} \Pi_0(q_0; p_0)$$

is upper hemicontinuous. Now suppose that for every k , there exists (q_0^k, p_0^k) such that

$$q_0^k \in \arg \max_{q_0 \in \Sigma_{0,k}} \Pi_0(q_0; p_0^k)$$

with $d(q_0^k, p_0^k) \geq \delta_0$. Because $\Delta^T(\mathcal{X})$ is compact, there exists a subsequence of indexes, $\{\sigma(k)\}_k$, such that $p_0^{\sigma(k)}$ converges to p_0^∞ for some p_0^∞ . Also, $\Sigma_{0,\sigma(k)}$ converges to $\Delta^T(\mathcal{X})$, where the limit is with respect to the Hausdorff metric. Noting that $\arg \max_{q_0 \in \Delta^T(\mathcal{X})} \Pi_0(q_0; p_0) = \{p_0\}$, by the upper hemicontinuity of the argmax correspondence, we get that $q_0^{\sigma(k)}$ converges to p_0^∞ , thus contradicting that $d(q_0^{\sigma(k)}, p_0^{\sigma(k)}) \geq \delta_0$ for every k .

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

$$|\Pi_0^*(q_0; p_0; M) - \Pi_0(q_0; p_0)| < 1/k^* \quad \forall q_0, p_0,$$

then, for each p_0 , if

$$q_0^* \in \arg \max_{q_0 \in \Sigma_0} \Pi_0^*(q_0; p_0; M)$$

then q_0^* is δ_0 -close to p_0 .

By contradiction, if the claim does not hold, then for every k there exists p_0^k, q_0^k, M^k such that

$$|\Pi_0^*(q_0; p_0; M^k) - \Pi_0(q_0; p_0)| < 1/k \quad \forall q_0, p_0,$$

while

$$q_0^k \in \arg \max_{q_0 \in \Sigma_0} \Pi_0^*(q_0; p_0^k; M^k)$$

and $d(q_0^k, p_0^k) \geq \delta_0$. Using the compactness of $\Delta^T(\mathcal{X})$, we generate a subsequence of indexes, $\{\sigma(k)\}$, such that $p_0^{\sigma(k)}$ converges to p_0^∞ and $q_0^{\sigma(k)}$ converges to q_0^∞ for some $p_0^\infty \in \Delta^T(\mathcal{X})$ and some $q_0^\infty \in \Sigma_0$.

Then, $d(q_0^\infty, p_0^\infty) \geq \delta_0$ and following Step 1(a), it implies that q_0^∞ is not a maximizer of the map $q_0 \in \Sigma_0 \mapsto \Pi_0(q_0; p_0^\infty)$. Let $q_0^* \in \Sigma_0$ be such a maximizer, then we have $\Pi_0(q_0^*; p_0^\infty) > \Pi_0(q_0^\infty; p_0^\infty)$, and by continuity, for large enough k 's, $\Pi_0(q_0^*; p_0^k) > \Pi_0(q_0^\infty; p_0^k)$, with both sides of the inequality bounded away from each other. Thus any k large enough, $\Pi_0^*(q_0^*; p_0^k; M^k) > \Pi_0^*(q_0^\infty; p_0^k; M^k)$. This inequality contradicts the fact that for k large enough, q_0^∞ should also maximize $q_0 \in \Sigma_0 \mapsto \Pi_0^*(q_0; p_0^k; M^k)$, since Σ_0 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 $|\Pi(p_0, \dots, p_{T-1}, x) - \Pi(p'_0, \dots, p'_{T-1}, x)| < 1/k^*$ for each $x \in \mathcal{X}$. Let $\delta_1 = \min\{\delta_0, \delta\}$.

Step 2. We now iterate Step 1 for every $t = 1, \dots, T - 1$. Let $t \geq 1$ and $\delta_t > 0$ be given. Fix p_0, \dots, p_{t-1} such that $p_0 \in \Sigma_0, p_1 \in \Sigma_1^{p_0}, p_2 \in \Sigma_2^{p_0, p_1}$, and so forth, where every set of the form $\Sigma_t^{p_0, \dots, p_{t-1}}$ is a finite subset of $\Delta^{T-t}(\mathcal{X})$.

Analogously to Step 1(a), we define $\Sigma_t^{p_0, \dots, p_{t-1}}$ as a finite subset of $\Delta^{T-t}(\mathcal{X})$ such that, for every p_t , if

$$q_t^* \in \arg \max_{q_t \in \Sigma_t^{p_0, \dots, p_{t-1}}} \Pi_t(p_0, \dots, p_{t-1}, q_t; p_t),$$

then q_t^* is δ_t -close to p_t .

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

$$|\Pi_t^*(p_0, \dots, p_{t-1}, q_t; p_t; M) - \Pi_t(p_0, \dots, p_{t-1}, q_t; p_t)| < 1/k^* \quad \forall p_t, q_t,$$

for every p_t , if

$$q_t^* \in \arg \max_{q_t \in \Sigma_t^{p_0, \dots, p_{t-1}}} \Pi_t^*(p_0, \dots, p_{t-1}, q_t; p_t; M)$$

then q_t^* is δ_t -close to p_t .

Finally, we let δ to be such that if, for every i , q'_i is δ -close to q''_i , then $|\Pi(q_0, \dots, q_{T-1}, x) - \Pi(q'_0, \dots, q'_{T-1}, x)| < 1/k^*$ for every x . Let $\delta_{t+1} = \min\{\delta, \delta_t\}$.

Step 3. We build a finite menu M_0^* of order T by sampling the infinite menu associated with Π as follows: for every p_0, \dots, p_{T-2} where for every t , $p_t \in \Sigma_t^{p_0, \dots, p_{t-1}}$, we define

$$\begin{aligned} M_T^{p_0, \dots, p_{T-2}} &= \{ \Pi(p_0, \dots, p_{t-1}, q_{T-1}, \cdot); q_{T-1} \in \Sigma_{T-1}^{p_0, \dots, p_{T-2}} \}, \\ M_t^{p_0, \dots, p_{t-1}} &= \{ M_{t+1}^{p_0, \dots, p_{t-1}, q_t}; q_t \in \Sigma_t^{p_0, \dots, p_{t-1}} \}. \end{aligned}$$

We let $M_0^* = \{M_0^{q_0}; q_0 \in \Sigma_0\}$. Let ξ be the degenerate probability measure that allocates full mass on M_0^* . We note that we have, by Steps 1(a), 1(b), and Step 2,

$$|\Pi^*(p_0, \dots, p_{T-1}, x; \xi) - \Pi(p_0, \dots, p_{T-1}, x)| < \epsilon/2 \quad \forall p_0, \dots, p_{T-1}, x$$

Step 4. This step concludes the proof. Let ξ' be a probability measure over \mathcal{M}_T with full support. Take $\xi'' = (1 - \epsilon/2)\xi + (\epsilon/2)\xi'$. Then, $\Pi^*(p_0, \dots, p_{T-1}, x; \xi'')$ defines a strategyproof payoff rule, and

$$|\Pi^*(p_0, \dots, p_{T-1}, x; \xi) - \Pi(p_0, \dots, p_{T-1}, x)| < \epsilon \quad \forall p_0, \dots, p_{T-1}, x.$$

References

- Aliprantis, C. D. and K. C. Border (2006). *Infinite Dimensional Analysis: A Hitchhiker's Guide* (third ed.). Springer.
- Allais, M. (1953). Le comportement de l'homme rationnel devant le risque: Critique des postulats et axiomes de l'école américaine. *Econometrica* 21(4), 503–546.
- Azrieli, Y., C. P. Chambers, and P. J. Healy (2018). Incentives in experiments: A theoretical analysis. *Journal of Political Economy* 126(4), 1472–1503.
- Becker, G. M., M. H. DeGroot, and J. Marschak (1964). Measuring utility by a single-response sequential method. *Behavioral Science* 9(3), 226–232.
- Bergemann, D., S. Morris, and S. Takahashi (2017). Interdependent preferences and strategic distinguishability. *Journal of Economic Theory* 168, 329–371.
- Brier, G. W. (1950). Verification of forecasts expressed in terms of probability. *Monthly Weather Review* 78(1), 1–3.
- Carroll, G. (2019). Robust incentives for information acquisition. *Journal of Economic Theory* 181, 382–420.
- Chambers, C. P. and N. S. Lambert (2020). Supplement to 'Dynamic belief elicitation'.
- Costa-Gomes, M. A. and G. Weizsäcker (2008). Stated beliefs and play in normal-form games. *The Review of Economic Studies* 75(3), 729–762.
- Crémer, J. and R. P. McLean (1988). Full extraction of the surplus in Bayesian and dominant strategy auctions. *Econometrica* 56(6), 1247–1257.
- Danz, D. N., D. Fehr, and D. Kübler (2012). Information and beliefs in a repeated normal-form game. *Experimental Economics* 15(4), 622–640.
- Dillenberger, D., J. S. Lleras, P. Sadowski, and N. Takeoka (2014). A theory of subjective learning. *Journal of Economic Theory* 153, 287–312.
- Eyster, E. and M. Rabin (2005). Cursed equilibrium. *Econometrica* 73(5), 1623–1672.
- Feltovich, N. (2000). Reinforcement-based vs. belief-based learning models in experimental asymmetric-information games. *Econometrica* 68(3), 605–641.
- Fryer, R. G., P. Harms, and M. O. Jackson (2019). Updating beliefs when evidence is open to interpretation: Implications for bias and polarization. *Journal of the European Economic Association* 17(5), 1470–1501.
- Gneiting, T. and A. E. Raftery (2007). Strictly proper scoring rules, prediction, and estimation. *Journal of the American Statistical Association* 102(477), 359–378.

- Good, I. J. (1952). Rational decisions. *Journal of the Royal Statistical Society, Series B* 14(1), 107–114.
- Grether, D. (1981). Financial incentive effects and individual decision-making. Manuscript No 401, Division of the Humanities and Social Sciences, California Institute of Technology.
- Harrison, G. W., J. Martínez-Correa, and J. T. Swarthout (2014). Eliciting subjective probabilities with binary lotteries. *Journal of Economic Behavior & Organization* 101, 128–140.
- Harrison, G. W. and R. D. Phillips (2014). Subjective beliefs and statistical forecasts of financial risks: The chief risk officer project. In T. Andersen and S. Mayes (Eds.), *Contemporary Challenges in Risk Management: : Dealing with Risk, Uncertainty and the Unknown*, pp. 163–202. Springer.
- Hewitt, E. and K. R. Stromberg (1997). *Real and Abstract Analysis*. Springer-Verlag.
- Hyndman, K., E. Y. Ozbay, A. Schotter, and W. Z. Ehrblatt (2012). Convergence: An experimental study of teaching and learning in repeated games. *Journal of the European Economic Association* 10(3), 573–604.
- Karni, E. (2009). A mechanism for eliciting probabilities. *Econometrica* 77(2), 603–606.
- Karni, E. (2018). A mechanism for eliciting second-order beliefs and the inclination to choose. *American Economic Journal: Microeconomics* 10(2), 275–85.
- Karni, E. (2020). A mechanism for the elicitation of second-order belief and subjective information structure. *Economic Theory* 69(1), 217–232.
- Kreps, D. M. and E. L. Porteus (1978). Temporal resolution of uncertainty and dynamic choice theory. *Econometrica* 46(1), 185–200.
- Krishna, R. V. and P. Sadowski (2014). Dynamic preference for flexibility. *Econometrica* 82(2), 655–703.
- Krishna, V. and E. Maenner (2001). Convex potentials with an application to mechanism design. *Econometrica* 69(4), 1113–1119.
- Lambert, N. S. (2018). Elicitation and evaluation of statistical forecasts. Unpublished manuscript, Stanford University.
- Lu, J. (2016). Random choice and private information. *Econometrica* 84(6), 1983–2027.
- Manski, C. F. and C. Neri (2013). First- and second-order subjective expectations in strategic decision-making: Experimental evidence. *Games and Economic Behavior* 81, 232–254.
- Matheson, J. E. and R. L. Winkler (1976). Scoring rules for continuous probability distributions. *Management Science* 22(10), 1087–1096.

- McCarthy, J. (1956). Measures of the value of information. *Proceedings of the National Academy of Sciences* 42(9), 654–655.
- Meyer, P. (1966). *Probability and Potentials*. Blaisdell Publishing Company.
- Miller, N., P. Resnick, and R. Zeckhauser (2005). Eliciting informative feedback: The peer-prediction method. *Management Science* 51(9), 1359–1373.
- Moore, D. A. and P. J. Healy (2008). The trouble with overconfidence. *Psychological Review* 115(2), 502.
- Nyarko, Y. and A. Schotter (2002). An experimental study of belief learning using elicited beliefs. *Econometrica* 70(3), 971–1005.
- Osband, K. (1989). Optimal forecasting incentives. *Journal of Political Economy* 97(5), 1091–1112.
- Ostrovsky, M. (2012). Information aggregation in dynamic markets with strategic traders. *Econometrica* 80(6), 2595–2647.
- Palfrey, T. R. and S. W. Wang (2009). On eliciting beliefs in strategic games. *Journal of Economic Behavior & Organization* 71(2), 98–109.
- Prelec, D. (2004). A Bayesian truth serum for subjective data. *Science* 306(5695), 462–466.
- Roth, A. E. and M. W. Malouf (1979). Game-theoretic models and the role of information in bargaining. *Psychological Review* 86(6), 574–594.
- Savage, L. J. (1954). *The Foundations of Statistics*. Dover Publications.
- Savage, L. J. (1971). Elicitation of personal probabilities and expectations. *Journal of the American Statistical Association* 66(336), 783–801.
- Schervish, M. J. (1989). A general method for comparing probability assessors. *The Annals of Statistics* 17(4), 1856–1879.
- Selten, R. (1998). Axiomatic characterization of the quadratic scoring rule. *Experimental Economics* 1(1), 43–61.
- Smith, C. A. B. (1961). Consistency in statistical inference and decision. *Journal of the Royal Statistical Society, Series B (Methodological)* 23(1), 1–37.
- Steiner, J., C. Stewart, and F. Matějka (2017). Rational inattention dynamics: Inertia and delay in decision making. *Econometrica* 85(2), 521–553.
- Stewart, C. (2011). Nonmanipulable bayesian testing. *Journal of Economic Theory* 146(5), 2029–2041.
- Takeoka, N. (2007). Subjective probability over a subjective decision tree. *Journal of Economic Theory* 136(1), 536–571.

Thomson, W. (1979). Eliciting production possibilities from a well-informed manager. *Journal of Economic Theory* 20(3), 360–380.

Ziegelmeyer, A., F. Koessler, J. Bracht, and E. Winter (2010). Fragility of information cascades: An experimental study using elicited beliefs. *Experimental Economics* 13(2), 121–145.