Supplementary Appendix to "Dynamic Belief Elicitation"

Christopher P. Chambers Nicolas S. Lambert

This Supplementary Appendix includes additional results and the proofs omitted from the main text. Throughout, references are to the main text and to the appendix of the main text, unless they include the prefix "SA."

It proceeds as follows. In Section SA.1, we provide several examples that illustrate the general method of the main text to special environments. In Section SA.2, we consider the case of protocols that may be decomposed into a separate subprotocols, one for each stage, each subprotocol being responsible for the elicitation of one signal at one time. We demonstrate that these protocols are not strategyproof, which implies the necessity of interactions across stages. In Section SA.3, we prove that, subject to regularity conditions, the information the protocols described in the main text elicit are enough to solve any dynamic decision problem. In Section SA.4, we ask what dynamic decision problems can be solved using the classical methods that elicit beliefs on the uncertainty of the random variable X (as opposed to dynamic beliefs) and show they form a degenerate class of problems. Finally, Section SA.5 includes the proofs for Section 4 of the main text.

SA.1 Additional Examples

In this section, we present several additional examples in the spirit of the example of Section 2, that illustrate the general approach we follow in Section 3.

SA.1.1 Random Times of Information Arrival

Here, time is continuous and indexed by $t \in [0, 1]$. The random variable X is an indicator of a random event, it is valued in $\{0, 1\}$ and materializes at $t = 1$. At some random time τ in the interval $(0, 1)$, the expert is to observe a private signal. As in Section 3, the random signal P takes value in $[0, 1]$ and designates the expert's updated posterior that $X = 1$. The principal difference with Section 3 is that τ is no longer a commonly known fixed time. Learning when new information arrives and when it is expected to arrive can be of interest, for instance, to an individual who faces an investment decision with gradually increasing costs of investment over time.

At $t = 0$, the expert's information structure consists of a prior belief about the signal, captured by a c.d.f. F over $[0, 1]$, and a prior belief about the time of arrival of the signal, captured by a c.d.f. G, also over $[0, 1]$. Although it is not essential, to simplify formulas, the signal and time of arrival are assumed independent. At $t = \tau$, the expert observes the realization of the random signal P, the up-to-date posterior assessment that $X = 1$. The elicitor observes neither the signal realization nor the time of observation.

The objective is to induce the expert to announce, as a strict best response, his information structure at $t = 0$, and then the signal realization, as soon as he receives it (except possibly when the signal realization turns out to be the same as the expert's prior belief on X, in which case we consider weak best response as acceptable, because the signal brings no valuable information).

Let $\overline{p}$ be the expert prior that $X = 1$, before new information arrives: $\overline{p} = \mathbf{E}^F[P]$. For any given t, the expert at time 0 anticipates to receive updated information by time t with probability $G(t)$. Hence, his prior about his future belief, at t, about the random variable X, is captured by the c.d.f. $H_t(p) := G(t)F(p) + (1 - G(t))\mathbf{1}\{\overline{p} \le p\}$.

Let us consider an α-option as in Section 2, that expires at some *fixed* time t (meaning that its owner must exercise the right to sell by time t at the latest). Suppose the elicitor offers the option for price β. From Section 2, we know that, if the parameters (α, β) are suitably randomized, when delegating all choices to the elicitor on behalf of the expert, the expert is strictly best off communicating H_t at time 0, and then his posterior assessment p_t at time t.

Now consider instead an α-option that expires at some *random* time T, for instance uniformly distributed. We continue to randomize over α, β as before. The expert is then strictly best off providing the c.d.f. H_t for every time t, except possibly on a set of times of zero mass. As any c.d.f. is right-continuous, it is equivalent to providing both G and F as a strict best response. Of course, the expert is also strictly best off sending his new information as soon as it arrives, so that the elicitor makes decisions with the most refined information—as any delay causes a loss of payoff with positive probability, unless the probability assessment that $X = 1$ remains identical after the arrival of interim information.

Taking expectations over α, β, T permits us to reduce the randomized protocol to a deterministic reward scheme that pays off only when the random variable X realizes. If the expert reports his information structure (F, G) at time 0, then reports the time t when he observes the signal, along with the value of p, the expert's payoff is:

$$\Pi(F,G,t,p,x) \coloneqq \frac{1}{4} - \frac{1}{4}\int_0^1 \int_0^1 \left(G(s)\textstyle\int_0^\alpha F + (1-G(s))\max(0,\alpha-p)\right)^2 \mathrm{d}\alpha\,\mathrm{d}s$$
$$+\frac{1}{2}\int_0^t \int_{\overline{p}}^1 \left(1 + G(s)\textstyle\int_0^\alpha F + (1-G(s)\max(0,\alpha-\overline{p})\right)(\alpha - x)\,\mathrm{d}\alpha\,\mathrm{d}s$$
$$+\frac{1}{2}\int_t^1 \int_p^1 \left(1 + G(s)\textstyle\int_0^\alpha F + (1-G(s)\max(0,\alpha-p)\right)(\alpha - x)\,\mathrm{d}\alpha\,\mathrm{d}s,$$

with $\overline{p} \coloneqq \mathbf{E}^F[P]$. It is easily verified that such payoff meets the aforementioned objective.

SA.1.2 Tracking a Continuous Signal

Time is again continuous and indexed by $t \in [0,1]$. The expert now privately observes, at every instant $t \in [0,1]$, a continuous signal S_t that starts at $S_0 = 0$ and evolves as a Brownian motion, with deterministic but time-dependent continuous drift and scale, μ_t and σ_t. That is, the law of evolution of S satisfies $\mathrm{d}S_t = \mu_t\,\mathrm{d}t + \sigma_t\,\mathrm{d}W_t$, for some standard Brownian motion W. The random variable X is binary and connected to the random signal in the following way: $X = 1$ if $S_1 \geq 0$ and $S = 0$ if $S_1 < 0$. For concreteness, one can think of X as the outcome of an election, and S_t as a estimate of excess votes for one side or the other.

The expert's information structure is summarized by the drifts and scales $\{\mu_t\}_t$ and $\{\sigma_t\}_t$. It is privately known to the expert. The objective is to induce the expert to announce as a unique best response, at the outset, his information structure, and subsequently, at every instant, the value of the private signal he observes.

It turns out the same protocol as that of Section SA.1.1, which consists in selling random α-options with random deadlines, can be used to extract such information. Indeed, this protocol induces the expert to reveal his posterior belief on X at every instant t, and, most importantly, to reveal his time-0 prior over each of these posteriors. In the special case of the Brownian motion considered here, it is equivalent to declaring the signal at every instant, and to communicate at the outset its ex ante distribution given by the drift and scale parameters.

To be more precise, the posterior belief at time t—call it P_t—is, ex ante, random, and linked to the expert's signal received at t by

$$P_t = \Phi\left(\frac{S_t + \int_t^1 \mu_s \, \mathrm{d}s}{\sqrt{\int_t^1 \sigma_s^2 \, \mathrm{d}s}}\right)$$

where Φ is the c.d.f. of the standard normal distribution. (This expression is immediate observing that S_1, conditional on the expert's information at t, is normally distributed with mean $S_t + \int_t^1 \mu_s \, \mathrm{d}s$ and variance $\int_t^1 \sigma_s^2 \, \mathrm{d}s$.) Different signal realizations will thus induce different posterior assessments.

The expert's initial information about his signal at t, S_t, is that it is normally distributed, with mean $\int_0^t \mu_s \, \mathrm{d}s$ and variance $\int_0^t \sigma_s^2 \, \mathrm{d}s$. Hence his prior, at the outset, about P_t, is captured by the c.d.f.

$$F_t(p_t) := \Phi\left(\frac{\Phi^{-1}(p_t)\sqrt{\int_t^1 \sigma_s^2 \, \mathrm{d}s} - \int_0^1 \mu_s \, \mathrm{d}s}{\sqrt{\int_0^t \sigma_s^2 \, \mathrm{d}s}}\right).$$

Different drift and scale parameters will thus induce different prior beliefs about posteriors. The expert strictly induced to communicate the "true" belief over future posteriors each time (except possibly on a set of mass zero) is thus also strictly induced to disclose the "true" information structure as described by the time-dependent drifts and scales.

SA.1.3 Two Times of Information Arrivals

The purpose of this setup is to illustrate a simple extension of the method of Section 2 in the time dimension.

The variable X continues to indicate a public event of interest, and there are now four times periods, denoted $t = 0, 1, 2, 3$. The expert is to receive interim information twice, at times $t = 1$, and then $t = 2$. It is convenient to describe the information structure backwards. At $t = 2$, the expert observes a signal that, as in Section 2 and without loss of generality, can be interpreted as a posterior probability that $X = 1$. Let P designate this random signal. At $t = 1$, the expert receives information from which he infers a posterior on X, but importantly also permits the expert to update his belief on P. Thus, let us say the expert observes a random signal which consists of a distribution over P, i.e., a c.d.f. over $[0, 1]$, call it Y. Note this unconventional random variable Y captures both the belief about P, and also the belief about X, by the law of iterated expectations. At $t = 0$, the expert forms private beliefs about the information he anticipates to receive. We restrict attention to information structures for which the expert believes on only two possible values of Y, denoted F and G, with $F \neq G$. The probability that $Y = F$ is denoted q. Thus, the expert's information structure is captured by the triplet (q, F, G). It is privately known to the expert. It should be understood that the restriction to two possible values for the first signal is *not* without loss of generality. Without that restriction, one must apply the more complex protocols of Section 3.

The objective is to induce, with strict incentives, the expert to reveal his information structure—characterized by the triple (q, F, G)—at the outset, and then to reveal his observations at times $t = 1, 2$. We already know from Section 2 how to accomplish the second part of the objective, so let us focus on the first part, i.e., the elicitation of the triple (q, F, G).

To construct the protocol, one can leverage the Allais idea one step further. Consider the following compound (α, β)-option instrument: owning the option at $t = 0$ gives the right to buy an α-option at time $t = 1$, at price β. Following the calculations of Section 2, the value of the compound option to the expert with information structure (q, F, G) is

$$q \max\left\{\int_0^\alpha F - \beta, 0\right\} + (1 - q) \max\left\{\int_0^\alpha G - \beta, 0\right\}.$$

For any two given information structures (q_1, F_1, G_1) and (q_2, F_2, G_2), there exists a compound option whose value differ across these structures. Therefore, one can decide of a price γ for the compound option, chosen mid-way between the two values, that an expert endowed with (q_1, F_1, G_1) would rather buy, while the expert endowed with (q_2, F_2, G_2) would prefer to decline the offer (or conversely). Then, to return to the more general case with no assumption made on the class of triples (q, F, G), one can randomize over (α, β) as in Section 2, and also choose γ randomly, independently and uniformly over $[-1, 1]$, for example. The elicitor would then draw a compound option at random, ask the expert for his beliefs, and make choices on the option optimally on behalf of the expert and given the expert's announcements, in effect inducing strict truthtelling. The details of the protocol and the proof of incentive properties are similar to Section 2 and therefore omitted.

SA.2 Stage-Separated Protocols

The purpose of this section is to show that protocols that attempt to induce truthtelling at later stages, and use such presumably truthful announcements to induce, in turn, truthtelling at earlier stages, fail to be strategyproof.

We focus on the three-period case (the impossibility result extends directly to any number of periods $N \geq 3$), and we borrow notation from Section 2. The random variable X takes values in $\mathcal{X} = \{1, \ldots, n\}$, $n \geq 2$. The elicitor asks the expert to disclose his prior $F \in \Delta(\Delta(\mathcal{X}))$ at $t = 0$ (isomorphic to an $(n-1)$-dimensional c.d.f.), his posterior $p \in \Delta(\mathcal{X})$ at $t = 1$ (isomorphic to an $(n-1)$-dimensional vector), and finally, when X materializes to value x, she rewards the expert with payoff $\Pi(F, p, x)$.

We ask if we can choose a strategyproof payoff rule Π of the form $\Pi(F, p, x) = \Pi_1(F, p) + \Pi_2(p, x)$; that is, we separate stages, the expert gets a first payoff after announcing the prior and posterior, and a second payoff after the random variable X realizes, independently of the prior report. These stage-separated protocols have a natural interpretation: we use the publicly observed outcome of X to elicit the posterior p through Π_2, and then, using p, we attempt to elicit the prior F via Π_1. In particular, Π_1 and Π_2 could be the payoffs of classic probability elicitation methods, such as the Brier score (Brier, 1950) or the Becker-DeGroot-Marschak mechanism (Becker, DeGroot and Marschak, 1964) for Π_2, and the Matheson-Winkler elicitation method (Matheson and Winkler, 1976) for Π_1.

To build intuition, we begin with protocols that satisfy some regularities and that, in absence of a first stage, would induce the expert to report truthfully his posterior. Let us focus on the expert's decision at $t = 2$. Assume the expert has reported his true prior F at time 1, and that his true posterior, given a particular signal realization, is p, while 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 $\Pi_2(\widehat{p}, p)$ designates the expert's expected payoff at $t = 2$ when he reports $\widehat{p}$ while his true posterior is p. Because $\Pi_2(\widehat{p}, p)$ is maximized when $\widehat{p} = p$, 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 $\Pi_2(F, \widehat{p})$ is constant in $\widehat{p}$, 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 at $t = 2$ exceed the losses incurred at the second stage: the protocol is not strategyproof.

The following proposition formalizes the claim in its generality without assuming any regularity condition.

Proposition SA.1 *If the payoff rule* Π *of a protocol can be decomposed as* $\Pi(F, p, x) = \Pi_1(F, p) + \Pi_2(p, x)$*, then the protocol is not strategyproof.*

Proof. Consider a protocol whose payoff rule satisfies the above decomposition. For every declared prior $\widehat{F}$, let $g_{\widehat{F}}(\widehat{p}, x)$ be the total payoff to the expert as a function of the announced posterior $\widehat{p}$ and realization x:

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

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

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

where $p \cdot q$ is the dot product between p and q on the simplex $\Delta(\mathcal{X})$ interpreted as a subset of $\mathbb{R}^n$. Thus for all $F, \widehat{F}$, we get that $\overline{g}_F - \overline{g}_{\widehat{F}}$ is constant: at $t = 0$, the expert is best off reporting any $\widehat{F}$ that maximizes $\overline{g}_{\widehat{F}}(\widehat{p})$, for an arbitrary $\widehat{p}$, independently of his true prior. Therefore, the protocol is not stategyproof. ■

It is worth noting, however, that if one cannot elicit the exact belief using such stage-separated protocols, an implication of the mechanism proposed by Karni (2017) is that one can do so arbitrarily closely, in this two-period case, by increasing the magnitude of the payoffs from the second stage Π_2, with respect to the first stage Π_1.

SA.3 Information Sufficiency

We use the general structure of Section 4 with N periods, $t = t_1, \ldots, t_N$, $t_1 < \cdots < t_N$. At $t = t_N$ the uncertain outcome from the compact metrizable set $\mathcal{X}$ materializes publicly. There is an expert whose information structure is $(\Omega, \mathbb{F}, \mathsf{P}, X)$. There is also a less informed utility-maximizing individual, who faces a dynamic decision problem: at every time t_k, the individual must choose an action a_k from a collection of possible actions $\mathcal{A}_k$, assumed compact metrizable. At the last period t_N, the individual receives utility $u(a_1, \ldots, a_{N-1}, x)$, where x designates the outcome of the random variable. The individual's utility function, $u : \mathcal{A}_1 \times \cdots \times \mathcal{A}_{N-1} \times \mathcal{X} \mapsto \mathbb{R}$, is bounded and jointly continuous.

The purpose of this section is to demonstrate the following: if the individual delegates the decisions to the expert, then she gets as much utility as if the expert were to communicate to her the "belief tree" induced by his information structure and subsequent private observations, in the sense of Section 3.1, and where, as explained in the main text, the probability measures use the recursive structure of Borel σ-algebras on the weak-* topology.

To begin, we show that there always exists a solution to the individual's multi-period decision problem: if the expert gets to choose the actions of this decision problem, he can always choose information-contingent actions that maximize the individual's expected utility.

A decision policy for the expert is summarized by a tuple $(\alpha_1, \ldots, \alpha_{N-1})$, where every α_k is an $\mathcal{F}_k$-measurable map from Ω to $\mathcal{A}_k$.[1] Denote by $\mathcal{D}(\mathbb{F})$ the set of all decision policies available to the expert. The following lemma asserts that an optimal decision policy always exists, and yields an expected utility that can be computed via a dynamic programming principle.

Lemma SA.1 *There exists a decision policy* $(\alpha_1^*, \ldots, \alpha_{N-1}^*) \in \mathcal{D}(\mathbb{F})$ *such that*

$$\begin{aligned}\mathbf{E}[u(\alpha_1^*, \ldots, \alpha_{N-1}^*, X)] &= \sup_{(\alpha_1,\ldots,\alpha_{N-1})\in\mathcal{D}(\mathbb{F})} \mathbf{E}[u(\alpha_1, \ldots, \alpha_{N-1}, X)] \\ &= \mathbf{E}[\sup_{a_1} \mathbf{E}[\ldots \sup_{a_{N-1}} \mathbf{E}[u(a_1, \ldots, a_{N-1}, X) \mid \mathcal{F}_{N-1}] \ldots \mid \mathcal{F}_1]].\end{aligned}$$

Proof. We first note that

$$\begin{aligned}&\sup_{(\alpha_1,\ldots,\alpha_{N-1})\in\mathcal{D}(\mathbb{F})} \mathbf{E}[u(\alpha_1, \ldots, \alpha_{N-1}, X)] \\ &\qquad\leq \mathbf{E}[\sup_{a_1} \mathbf{E}[\ldots \sup_{a_{N-1}} \mathbf{E}[u(a_1, \ldots, a_{N-1}, X) \mid \mathcal{F}_{N-1}] \ldots \mid \mathcal{F}_1]],\end{aligned}$$

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

Take an arbitrary decision policy $(\alpha_1, \ldots, \alpha_{N-1})$. Note that

$$\mathbf{E}[u(\alpha_1, \ldots, \alpha_{k-1}, a_k, \ldots, a_{N-1}, X) \mid \mathcal{F}_{N-1}]$$

is continuous in $(a_k, \ldots, a_{N-1})$ by the Dominated Convergence Theorem. Thus, we have that the supremum $\sup_{a_{N-1}} \mathbf{E}[u(\alpha_1, \ldots, \alpha_{k-1}, a_k, \ldots, a_{N-1}, X) \mid \mathcal{F}_{N-1}]$ is continuous in $(a_k, \ldots, a_{N-2})$ by Berge's Maximum Theorem. By the Measurable

[1] Every space is tacitly endowed with its Borel σ-algebra.

Maximum Theorem (Theorem 18.19, Aliprantis and Border, 2006) it is also $\mathcal{F}_{N-1}$-measurable. An inductive argument yields that, for every k, and every decision policy $(\alpha_1, \ldots, \alpha_{N-1})$,

$$\mathbf{E}[\sup_{a_{k+1}} \mathbf{E}[\ldots \sup_{a_{N-1}} \mathbf{E}[u(\alpha_1, \ldots, \alpha_{k-1}, a_k, \ldots, a_{N-1}, X) \mid \mathcal{F}_{N-1}] \ldots \mid \mathcal{F}_{k+1}] \mid \mathcal{F}_k],$$

is continuous in a_k and is measurable with respect to $\mathcal{F}_k$ (thus is a Cathéodory function). In particular, taking $k = 1$, we get that the map

$$(a_1, \omega) \mapsto \mathbf{E}[\sup_{a_2} \mathbf{E}[\ldots \sup_{a_{N-1}} \mathbf{E}[u(a_1, \ldots, a_{N-1}, X) \mid \mathcal{F}_{N-1}] \ldots \mid \mathcal{F}_2] \mid \mathcal{F}_1](\omega),$$

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

$$\begin{aligned}\mathbf{E}[\sup_{a_2} \mathbf{E}[\ldots \sup_{a_{N-1}} \mathbf{E}[u(\alpha_1^*, a_2, \ldots, a_{N-1}, X) \mid \mathcal{F}_{N-1}] \ldots \mid \mathcal{F}_2] \mid \mathcal{F}_1] \\ = \sup_{a_1} \mathbf{E}[\sup_{a_2} \mathbf{E}[\ldots \sup_{a_{N-1}} \mathbf{E}[u(a_1, \ldots, a_{N-1}, X) \mid \mathcal{F}_{N-1}] \ldots \mid \mathcal{F}_2] \mid \mathcal{F}_1].\end{aligned}$$

We define recursively the remaining α_k^*'s by noting that, by the above result, having defined $\alpha_1^*, \ldots, \alpha_{k-1}^*$ where each α_i^* is a $\mathcal{F}_i$-measurable map from Ω to $\mathcal{A}_i$, the map

$$\begin{aligned}&(a_k, \omega) \mapsto \\ &\quad \mathbf{E}[\sup_{a_k} \mathbf{E}[\ldots \sup_{a_{N-1}} \mathbf{E}[u(\alpha_1^*, \ldots, \alpha_{k-1}^*, a_k, \ldots, a_{N-1}, X) \mid \mathcal{F}_{N-1}] \ldots \mid \mathcal{F}_{k+1}] \mid \mathcal{F}_k](\omega),\end{aligned}$$

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

$$\begin{aligned}&\mathbf{E}[\sup_{a_{k+1}} \mathbf{E}[\ldots \sup_{a_{N-1}} \mathbf{E}[u(\alpha_1^*, \ldots, \alpha_k^*, a_{k+1}, \ldots, a_{N-1}, X) \mid \mathcal{F}_{N-1}] \ldots \mid \mathcal{F}_{k+1}] \mid \mathcal{F}_k] \\ &= \sup_{a_k} \mathbf{E}[\sup_{a_{k+1}} \mathbf{E}[\ldots \sup_{a_{N-1}} \mathbf{E}[u(\alpha_1^*, \ldots, \alpha_{k-1}^*, a_k, \ldots, a_{N-1}, X) \mid \mathcal{F}_{N-1}] \ldots \mid \mathcal{F}_{k+1}] \mid \mathcal{F}_k].\end{aligned}$$

Finally, using with the law of iterated expectations to collapse the $\mathcal{F}_k$'s,

$$\begin{aligned}
&\mathbf{E}[\sup_{a_1} \mathbf{E}[\dots \sup_{a_{N-1}} \mathbf{E}[u(a_1, \dots, a_{N-1}, X) \mid \mathcal{F}_{N-1}] \dots \mid \mathcal{F}_1]],\\
&\quad = \mathbf{E}[\sup_{a_2} \mathbf{E}[\dots \sup_{a_{N-1}} \mathbf{E}[u(\alpha_1^*, a_2, \dots, a_{N-1}, X) \mid \mathcal{F}_{N-1}] \dots \mid \mathcal{F}_2]],\\
&\quad = \cdots\\
&\quad = \mathbf{E}[\sup_{a_k} \mathbf{E}[\dots \sup_{a_{N-1}} \mathbf{E}[u(\alpha_1^*, \dots, \alpha_{k-1}^*, a_k, \dots, a_{N-1}, X) \mid \mathcal{F}_{N-1}] \dots \mid \mathcal{F}_k]],\\
&\quad = \cdots\\
&\quad = \mathbf{E}[u(\alpha_1^*, \dots, \alpha_{N-1}^*, X)],
\end{aligned}$$

which concludes the proof. ■

Now, let us assume that the expert does not take actions in place of the individual. Instead, the expert communicates to the individual the relevant "probability tree" (i.e., his induced belief tree) at every time. Let Z_k be the random variable taking values in $\Delta^{N-k}(\mathcal{X})$—i.e., the space of all "probability trees", all endowed with the Borel σ-algebra generated by the weak-* topology—and defined by $Z_k(\omega) = \varphi_{t_{k+1},\dots,t_{N-1}}(\mathsf{P}_t^\omega)$, where $\varphi_{t_{k+1},\dots,t_{N-1}}$ is the induced belief tree with intermediate times $t_{k+1}, \dots, t_{N-1}$, as defined in Section 4. Z_k represents the information that the expert communicates to the individual. Let $\mathbb{Z} = \{\mathcal{Z}_k\}_k$ be the filtration generated by the discrete process Z_k. $\mathbb{Z}$ represents the dynamic information the individual learns from the expert.

A decision policy for the individual is summarized by a tuple $(\beta_1, \dots, \beta_{N-1})$, where every β_k is an $\mathcal{Z}_k$-measurable map from Ω to $\mathcal{A}_k$. Let $\mathcal{D}(\mathbb{Z})$ be the set of all decision policies available to the individual. The same argument as in Lemma SA.1 shows that there exists an optimal policy $(\beta_1^*, \dots, \beta_{N-1}^*)$, in the sense that

$$\mathbf{E}[u(\beta_1^*, \dots, \beta_{N-1}^*, X)] = \sup_{(\beta_1,\dots,\beta_{N-1}) \in \mathcal{D}(\mathbb{Z})} \mathbf{E}[u(\beta_1, \dots, \beta_{N-1}, X)].$$

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

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

Proposition SA.2 *The individual's optimal decisions yield the same expected utility as if she had delegated the problem to the expert, i.e.,*

$$\mathbf{E}[u(\alpha_1^*, \dots, \alpha_{N-1}^*, X)] = \mathbf{E}[u(\beta_1^*, \dots, \beta_{N-1}^*, X)].$$

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

In particular, we have

$$\mathbf{E}[u(a_1, \dots, a_{N-1}, X) \mid \mathcal{F}_{N-1}] = \mathbf{E}[u(a_1, \dots, a_{N-1}, X) \mid Z_{N-1}],$$

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

$$\begin{aligned}
&\mathbf{E}[\sup_{a_k} \mathbf{E}[\dots \mathbf{E}[\sup_{a_{N-1}} \mathbf{E}[u(a_1, \dots, a_{N-1}, X) \mid \mathcal{F}_{N-1}] \mid \mathcal{F}_{N-2}] \dots \mid \mathcal{F}_k] \mid \mathcal{F}_{k-1}] \\
&\quad = \mathbf{E}[\sup_{a_k} \mathbf{E}[\dots \mathbf{E}[\sup_{a_{N-1}} \mathbf{E}[u(a_1, \dots, a_{N-1}, X) \mid Z_{N-1}] \mid Z_{N-2}] \dots \mid Z_k] \mid \mathcal{F}_{k-1}] \\
&\quad = \mathbf{E}[\sup_{a_k} \mathbf{E}[\dots \mathbf{E}[\sup_{a_{N-1}} \mathbf{E}[u(a_1, \dots, a_{N-1}, X) \mid Z_{N-1}] \mid Z_{N-2}] \dots \mid Z_k] \mid Z_{k-1}].
\end{aligned}$$

In particular,

$$\begin{aligned}
&\mathbf{E}[\sup_{a_1} \mathbf{E}[\dots \mathbf{E}[\sup_{a_{N-1}} \mathbf{E}[u(a_1, \dots, a_{N-1}, X) \mid \mathcal{F}_{N-1}] \mid \mathcal{F}_{N-2}] \dots \mid \mathcal{F}_1]] \\
&= \mathbf{E}[\sup_{a_1} \mathbf{E}[\dots \mathbf{E}[\sup_{a_{N-1}} \mathbf{E}[u(a_1, \dots, a_{N-1}, X) \mid Z_{N-1}] \mid Z_{N-2}] \dots \mid Z_1]] \\
&= \mathbf{E}[\sup_{a_1} \mathbf{E}[\dots \mathbf{E}[\sup_{a_{N-1}} \mathbf{E}[u(a_1, \dots, a_{N-1}, X) \mid \mathcal{Z}_{N-1}] \mid \mathcal{Z}_{N-2}] \dots \mid \mathcal{Z}_1]].
\end{aligned}$$

We conclude by the dynamic programming principle established in Lemma SA.1. Note that all suprema are measurable by the Measurable Maximum Theorem, as detailed in the lemma. ■

SA.4 Information Structures and Probabilities

In this section, we seek to understand the advantage that learning full information structures confer when making dynamic decisions, in comparison to learning only probability estimates over the public random outcomes that are updated over time. While information structures have to be elicited via a menu-type structure of the sort presented in Section 2 and Section 3, probability estimates can be elicited via standard methods, such as the Brier score (Brier, 1950) and other probability scoring rules (see, for example, Savage, 1971).

There is an individual and three periods $t = 1, 2, 3$. To represent the individual's information structure, we borrow notation from Section 2. At the final period, a utility-relevant random variable X realizes publicly. The random variable takes n possible values in $\mathcal{X} = \{1, \ldots, n\}$. At the interim period, the individual receives information and infers a *posterior* $p \in \Delta(\mathcal{X})$. At the first period, the individual observes his information structure, summarized by the *prior* $F \in \Delta(\Delta(\mathcal{X}))$ over the posterior to be (indirectly) observed at $t = 1$.

The individual is an expected-utility maximizer, who faces a dynamic decision. There are two actions to be taken at respective times $t = 1, 2$. If the individual chooses action a_t in period t, she gets utility $u(a_1, a_2, x)$, where x is the realization of a random variable X. The set of possible actions $\mathcal{A}_t$ at $t = 1, 2$ is finite.

The prior over posteriors also yield a prior over outcomes—let us call it $\mu(F)$—defined as

$$\mu(F) = \int_{p \in \Delta(\mathcal{X})} p \,\mathrm{d}F(p). \tag{SA.1}$$

We define the value function for the individual at period two as

$$V(a_1, p) = \max_{a_2} \sum_{x \in \mathcal{X}} u(a_1, a_2, x) p(x). \tag{SA.2}$$

We also define the expected continuation payoff at period one as a function of the action taken at that stage as

$$\overline{V}(a_1, F) = \int_{p \in \Delta(\mathcal{X})} V(a_1, p) \,\mathrm{d}F(p). \tag{SA.3}$$

To maximize his payoff, the individual should choose a_1 so as to maximize $\overline{V}(a_1, F)$, and then, once p realizes, choose action a_2 so as to maximize

$$\sum_{x \in \mathcal{X}} u(a_1, a_2, x)p(x). \tag{SA.4}$$

To simplify matters, we assume that every action is weakly optimal for *some* posterior which is full-support; that is, for every a_1, there is $p \in \Delta(\mathcal{X})$ for which $p(x) > 0$ for all x, and $V(a_1, p) \geq V(b, p)$ for all $b \in \mathcal{A}_1$. Throughout this section, a *dynamic decision problem* refers to any tuple $(\mathcal{A}_1, \mathcal{A}_2, u)$ satisfying this property. Observe that this assumption is without loss of generality, because any action that violates this requirement is weakly suboptimal for *all* posteriors, including those that are not full-support (by a continuity argument), and thus no individual ever need to play this action, which can be safely removed from the set of possible actions.

We want to find out the general structure of dynamic decision problems—here represented by $u(a_1, a_2, x)$—such that the only relevant information for the stage 1 decision of the decision maker is the prior over outcomes $\mu(F)$. That is, for any prior over posteriors F_1, F_2 such that F_1 and F_2 yield the same prior over outcomes, i.e., $\mu(F_1) = \mu(F_2)$, we have that any period-one action optimal under F_1 is also optimal under F_2: if

$$\overline{V}(a_1, F_1) = \max_{b_1} \overline{V}(b_1, F_1) \tag{SA.5}$$

implies

$$\overline{V}(a_1, F_2) = \max_{b_1} \overline{V}(b_1, F_2). \tag{SA.6}$$

Such a dynamic decision problem will be termed *marginal dependent.*

Instead of getting a condition on u directly, it is more convenient to get a condition on the value function V.

Proposition SA.3 *Let $(\mathcal{A}_1, \mathcal{A}_2, u)$ be a marginal dependent dynamic decision problem. Then there exists a function $u_2' : \mathcal{A}_2 \times \mathcal{X} \to \mathbb{R}$ and for each $a \in \mathcal{A}_1$, there exists $f_a \in \mathbb{R}^{\mathcal{X}}$ such that $V(a, p) = f_a \cdot p + \sup_{a_2 \in \mathcal{A}_2} \sum_{x \in \mathcal{X}} u_2'(a_2, x)p(x)$.*

To interpret this proposition, let us introduce two definitions.

First, we will call a dynamic decision problem $(\mathcal{A}_1, \mathcal{A}_2, u)$ *additively separable* if for each $i = 1, 2$, there is $u_i : \mathcal{A}_i \times \Omega \to \mathbb{R}$ for which we can decompose the decision problem as $u(a_1, a_2, x) = u_1(a_1, x) + u_2(a_2, x)$. Thus, in additively separable decision problems, there are no interactions between choices made at different times: essentially, these are combine two static and independent problems.

Second, let us say that two dynamic decision problems $(\mathcal{A}_1, \mathcal{A}_2, u)$ and $(\mathcal{A}_1', \mathcal{A}_2', u')$

are *first-period payoff isomorphic* if the sets

$$\{V(a_1,\cdot) : a_1 \in \mathcal{A}_1\}$$

and

$$\{V'(a'_1,\cdot) : a'_1 \in \mathcal{A}'_1\}$$

coincide. To understand the definition of first-period payoff isomorphism, observe that, in terms of first period incentives, two isomorphic problems coincide at period one: they present exactly the same set of second-period value functions.

Corollary SA.1 *Any marginal independent dynamic decision problem* $(\mathcal{A}_1, \mathcal{A}_2, u)$ *is first-period payoff isomorphic to an additively separable dynamic decision problem.*

The proof of Corollary SA.1 is straightforward: Let $\mathcal{A}'_1 = \mathcal{A}_1$, and for any $a_1 \in \mathcal{A}_1$, define $u'_1(a_1, \omega) = f_a(\omega)$. Let $\mathcal{A}'_2 = \mathcal{A}_2$, and u'_2 as in Proposition SA.3, so that $V^*(p) = \sup_{a'_2 \in \mathcal{A}'_2} \int_\Omega u'_2(a'_2, \omega) dp(\omega)$.

Corollary SA.1 is an immediate consequence of Proposition SA.3, but facilitates its interpretation: probabilities over the final outcome are enough only when the available choice at period one does not constrain the available choices at period two, and when the utility at period two does not depend on the utility at period one.

We emphasize that our notion of dynamic decision problem does not directly incorporate any type of intertemporal budget constraint; the available decisions at period two (the set $\mathcal{A}_2$) do not depend on the decision made in period one. Rather, our framework allows intertemporal budget constraints to be captured by modifying the utility function so that $u(a_1, a_2, \omega) = -\infty$ whenever a_2 is not available after having chosen a_1. Additive separability of course rules this type of construction out. Hence, problems obeying marginal dependence are very few indeed: marginal dependence requires that first period decisions do not constrain second period decisions at all, and further that there is no complementarity across the two time periods.

Below we lay down the proof of Proposition SA.3.

Proof of Proposition SA.3.. Fix $a, b \in \mathcal{A}_1$, and suppose that p^* is in the relative interior of $\Delta(\mathcal{X})$ (that is, it is full-support). Suppose that a and b are both optimal for a F^* putting probability one on p^*. In particular, by marginal dependence, we know that $V(a, p^*) - V(b, p^*) = 0$.

Define $H : \Delta(\mathcal{X}) \to \mathbb{R}$ by $H(q) = V(a, q) - V(b, q)$. We claim that H is affine.

First, we establish that for all $q_1, q_2 \in \Delta(\mathcal{X})$ and all $\lambda_1, \lambda_2 \geq 0$ for which $\lambda_1 + \lambda_2 = 1$, we have $H(\lambda_1 q_1 + \lambda_2 q_2) = \lambda_1 H(q_1) + \lambda_2 H(q_2)$.

To this end, let $q^* = \lambda_1 q_1 + \lambda_2 q_2$, and suppose first $q^* \neq p^*$. Observe that since p^* is in the relative interior of $\Delta(\mathcal{X})$, for $\alpha > 0$ small, $r^* = p^* + \alpha(p^* - q^*) \in \Delta(\mathcal{X})$.

Now, let $\beta \in (0,1)$ for which $\beta q^* + (1-\beta) r^* = p^*$. For example, $\beta = \alpha/(\alpha+1)$. Take G^1 which puts weight $\beta\lambda_1$ on q_1, $\beta\lambda_2$ on q_2, and $(1-\beta)$ on r^*, and G^2, which puts weight β on q^* and $(1-\beta)$ on r^*. Observe that $\mu(G^1) = \mu(G^2) = p^*$. Hence, a and b are optimal for each of G^1 and G^2. We conclude that $\beta H(q^*) + (1-\beta)H(r^*) = \beta\lambda_1 H(q_1) + \beta\lambda_2 H(q_2) + (1-\beta)H(r^*)$, which implies $H(q^*) = \lambda_1 H(q_1) + \lambda_2 H(q_2)$.

On the other hand, if $q^* = p^*$, it follows that, by taking a lottery F^* putting probability one on p^*, and lottery G^* putting probability λ_1 on q_1 and λ_2 on q_2, that $H(p^*) = \lambda_1 H(q_1) + \lambda_2 H(q_2)$.

Therefore, for any two actions a, b that are optimal in the first stage under some prior p^* in the relative interior of $\Delta(\mathcal{X})$, $V(a,p) - V(b,p)$ is an affine function of p; hence, can be represented as $V(a,p) - V(b,p) = w \cdot p$ for some w (owing to the fact that the domain p consists of elements of $\Delta(\mathcal{X})$.

Now, consider any pair of actions a and b. Let p_a and p_b be posteriors in the relative interior of $\Delta(\mathcal{X})$ where a and b are weakly optimal, respectively. Consider the line segment $[p_a, p_b] \subseteq \Delta(\mathcal{X})$ connecting p_a and p_b. Observe that $[p_a, p_b]$ lies in the relative interior of $\Delta(\mathcal{X})$.

Finally, we can also consider the segment $[\delta_{p_a}, \delta_{p_b}] \subseteq \Delta(\Delta(\mathcal{X}))$.[2] By assumption, a is optimal for δ_{p_a} and b is optimal for δ_{p_b}. Observe that this interval can be divided into a finite number of sub-segments of the form $[F_i, F_{i+1}]$, where some action c_i of $\mathcal{A}_1$ is optimal on (F_i, F_{i+1}). This fact owes to that the value function at the first stage (i.e., the function that gives the decision maker's expected payoff under optimal decisions and as a function of F) is the upper envelope of finite number of linear functions. Hence, c_i is also optimal on $(\mu(F_i), \mu(F_{i+1}))$. By our first claim, it follows that for each i, $V(c_i, p) - V(c_{i+1}, p)$ is affine in p. Summing over the indexes in i, we get that $V(a,p) - V(b,p)$ is affine in p.

Finally pick some arbitrary $a^* \in \mathcal{A}^1$, and define

$$V^*(p) = V(a^*, p) = \sup_{a_2 \in \mathcal{A}_2} \sum_{x \in \mathcal{X}} u(a^*, a_2, x) p(x)$$

and $f_{a^*} = 0$. Otherwise, f_a is such that $V(a,p) - V(a^*, p) = f_a \cdot p$. ■

[2]Here, δ_p denotes the Dirac measure on p.

SA.5 Proofs of Section 4

SA.5.1 Some Auxiliary Lemmas

As for the baseline case, the general version of our existence result requires the use of some technical lemmas.

Lemma SA.2 *Let $\mathcal{A}$ be a separable metrizable space and $\mathcal{B}$ be a measurable space, with $\Delta(\mathcal{B})$ the set of probability measures on $\mathcal{B}$ equipped with the weak-* topology. Both $\mathcal{A}$ and $\Delta(\mathcal{B})$ are given their respective Borel σ-algebras. Let $\Psi : a \mapsto \Psi_a$ be a map from $\mathcal{A}$ to $\Delta(\mathcal{B})$. If, for every event E of $\mathcal{B}$, the map $a \mapsto \Psi_a(E)$ from $\mathcal{A}$ to $\mathbb{R}$ is measurable, then Ψ is measurable.*

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

Lemma SA.3 *All induced belief trees are well-defined and measurable.*

Proof. The proof proceeds by induction. Let us fix an information structure $(\Omega, \mathbb{F}, \mathsf{P}, X)$. First, $Q \mapsto \varphi(Q)$ associates, to every probability measure Q over states, the law of X. It is thus well defined. We also observe that $\Delta(\Omega)$ is separable measurable (Theorem 15.12 of Aliprantis and Border, 2006). Therefore, for every event E of $\mathcal{X}$, $Q \mapsto Q(X \in E)$ is measurable (Theorem 15.13 of Aliprantis and Border, 2006). Applying Lemma SA.2, we get that $Q \mapsto \varphi(Q)$ is measurable.

Now take a given $k \geq 0$ and suppose that $\varphi_{t_1,\ldots,t_k}$ is well-defined and measurable, for every $t_1, \ldots, t_k$. We show that $\varphi_{t_1,\ldots,t_{k+1}}$ is well defined and measurable, for every $t_1, \ldots, t_{k+1}$.

It is well defined: by assumption, $\omega \mapsto \mathsf{P}^{\omega}_{t_1}(E)$ is measurable for every event E, which by Lemma SA.2 implies that $\omega \mapsto \mathsf{P}^{\omega}_{t_1}$ is measurable, and so a well-defined random variable with values in $\Delta(\Omega)$. Thus, $\omega \mapsto \varphi_{t_2,\ldots,t_{k+1}}(\mathsf{P}^{\omega}_{t_1})$ is a well-defined random variable, with values in $\Delta^{k+1}(\mathcal{X})$.

It is measurable: by the induction hypothesis, for every event E of $\Delta^{k+1}(\mathcal{X})$, the set $\{\omega \in \Omega : \varphi_{t_2,\ldots,t_{k+1}}(\mathsf{P}^{\omega}_{t_1}) \in E\}$ is a well-defined event of Ω. Applying again Theorems

15.12 and 15.13 Aliprantis and Border (2006), we get that $Q \mapsto Q(\varphi_{t_2,\dots,t_{k+1}}(\mathsf{P}_{t_1}) \in E)$ is measurable, and by Lemma SA.2, we get that $Q \mapsto Q(\varphi_{t_2,\dots,t_{k+1}}(\mathsf{P}_{t_1}))$ is measurable. ■

Lemma SA.4 *Given an information structure* $(\Omega, \mathbb{F}, \mathsf{P}, X)$*, every value map* $\pi_k : \Sigma_k \times \Delta(\Omega) \mapsto \mathbb{R}$ *is well-defined and jointly measurable, for all* $k \geq 0$*.*

Proof. The proof proceeds by induction. The map $(S, \omega) \mapsto S(X(\omega))$ is bounded and measurable, and Ω is separable metrizable, so by Theorem 17.25 of Kechris (1995), the map $\pi_0 : (S, Q) \mapsto \int S(X(\omega))\, \mathrm{d}Q(\omega)$ is jointly measurable.

Next, suppose that π_k is jointly measurable. It implies that π_{k+1} is well defined, because for every t, $\omega \mapsto \mathsf{P}_t^\omega$ is a well-defined random variable with values in $\Delta(\Omega)$. We observe that the map $\sigma_{k+1} \mapsto M_{k+1}$ from Σ_{k+1} to 2^{Σ_k} is measurable for the Borel σ-algebra of the Hausdorff metric topology on 2^{Σ_k}. Thus by Theorem 18.10 of Aliprantis and Border (2006), the correspondence $\sigma_{k+1} \twoheadrightarrow M_{k+1}$ from Σ_{k+1} to Σ_k is measurable. We then use the Castaing representation theorem (Corollary 18.14, Aliprantis and Border, 2006) to generate a sequence $\{\Phi_i : i = 1, 2, \dots\}$ of measurable maps $\Phi_i : \Sigma_{k+1} \to \Sigma_k$ such that $M_{k+1} = \{\Phi_i(\sigma_{k+1}) : i = 1, 2, \dots\}$. Thus, we get

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

Hence, the map $(\sigma_{k+1}, Q) \mapsto \max_{\sigma_k \in M_{k+1}} \pi(\sigma_{k+1}, Q)$ is jointly measurable as the pointwise supremum of countably many jointly measurable maps. Besides, the right-continuity of $t \mapsto \mathsf{P}_t^\omega$ for every ω implies the joint measurability of $(t, \omega) \mapsto \mathsf{P}_t^\omega$ and thus the joint measurability of $(\sigma_{k+1}, \omega) \mapsto \mathsf{P}_{\tau_{k+1}}^\omega$, where we have decomposed σ_{k+1} as (M_{k+1}, τ_{k+1}). We have thus established that the map

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

is jointly measurable. It follows that the map

$$(\sigma_{k+1}, Q) \mapsto \int \left[\max_{\sigma_k \in M_{k+1}} \pi_k(\sigma_k, \mathsf{P}_{\tau_{k+1}}^\omega) \right] \mathrm{d}Q(\omega).$$

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

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

Lemma SA.5 *Every map* $(\{Q_t\}_t, x, \sigma_k) \mapsto \Pi^{(\Omega,\mathbb{F},\mathsf{P},X)}(\{Q_t\}_t, x; \sigma_k)$ *from* $\mathcal{Q} \times \mathcal{X} \times \Sigma_k$ *to* $\mathbb{R}$ *is jointly measurable.*

Proof. The map $(\{Q_t\}_t, x, S) \mapsto \Pi(\{Q_t\}_t, x; S) = S(x)$ is jointly measurable. Let us suppose that $(\{Q_t\}_t, x; \sigma_k) \mapsto \Pi(\{Q_t\}_t, x; \sigma_k)$ is measurable. Let $\{\Phi_i : i = 1, 2, \dots\}$ be a sequence of measurable maps from Σ_{k+1} to Σ_k such that $M_{k+1} = \{\Phi_i(\sigma_{k+1}) : i = 1, 2, \dots\}$, whose existence was proved in Lemma SA.4. Because securities of Σ_k take values in a bounded interval we have that for any $\sigma', \sigma'' \in \Sigma_k$, $d(\sigma', \sigma'') < \overline{D}$ for some constant $\overline{D}$ large enough. Let ν be the argmax correspondence $(\sigma_{k+1}, Q) \twoheadrightarrow \arg\max_{\sigma_k \in M_{k+1}} \pi(\sigma_k, Q)$.

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

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

where $g(\sigma_{k+1}, Q) = \max_{\sigma_k \in M_{k+1}} \pi_k(\sigma_k, Q)$. It was shown in the proof of Lemma SA.4 that g is jointly measurable. Therefore, $(\sigma_{k+1}, Q) \mapsto \delta(\sigma_k, (\sigma_{k+1}, Q))$ is measurable as the pointwise infimum of countably many measurable functions.

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

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

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

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

Finally, observing that

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

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

SA.5.2 Proof of Theorem 3

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

Let us fix an information structure of the expert and a strategy, following the notation of Definition 4.

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

We proceed by induction. For $k = 0$, and every $\{Q_t\}_t \in \mathcal{Q}$, the map $t_0 \mapsto \varphi(Q_{t_0}) = Q_{t_0}(X)$ is right-continuous, because for every event E of Ω, by assumption, $t_0 \mapsto Q_{t_0}(E)$ is right-continuous. Now fix k, and suppose that for every $\{Q_t\}_t \in \mathcal{Q}$, the map $(t_0, \ldots, t_k) \mapsto \varphi_{t_1,\ldots,t_k}(Q_{t_0})$ is separately right continuous. We have that $\varphi_{t_1,\ldots,t_k}(Q_{t_0}) = Q_{t_0}(\varphi_{t_2,\ldots,t_k}(\mathsf{P}_{t_1}))$. The right-continuity assumption on $t_0 \mapsto Q_{t_0}(E)$ for every event E ensures again the right continuity of $t_0 \mapsto Q_{t_0}(\varphi_{t_2,\ldots,t_k}(\mathsf{P}_{t_1}))$.

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

$$t_i \mapsto \int_{\Delta^k(\mathcal{X})} f(q) \, \mathrm{d}Q_{t_0}(\varphi_{t_2,\ldots,t_k}(\mathsf{P}_{t_1}) = q)$$

is right continuous, for every such f. Note that

$$\int_{\Delta^k(\mathcal{X})} f(q) \, \mathrm{d}Q_{t_0}(\varphi_{t_2,\ldots,t_k}(\mathsf{P}_{t_1}) = q) = \int_{\Omega} f(\varphi_{t_2,\ldots,t_k}(\mathsf{P}^{\omega}_{t_1})) \, \mathrm{d}Q_{t_0}(\omega).$$

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

$$t_i \mapsto \int_\Omega f(\varphi_{t_2,\ldots,t_k}(\mathsf{P}^{\omega}_{t_1}))\,\mathrm{d}Q_{t_0}(\omega).$$

We conclude that, for every $\{Q_t\}_t \in \mathcal{Q}$ and every k, the map $(t_0,\ldots,t_k) \mapsto \varphi_{t_1,\ldots,t_k}(Q_{t_0})$ is separately right continuous.

Now, consider a strategy that is optimal at state ω^* and time t_0. Since the randomization device uses the full-support distribution, it means that, for every k, there exists a set of times $\mathcal{T}$ dense in $\{t_0,\ldots,t_k : 0 \le t_0 < \cdots < t_k \le 1\}$ such that, for every $\tau = (\tau_0,\ldots,\tau_k) \in \mathcal{T}$, the expected payoff from the protocol that randomizes over menus of $\mathcal{M}_k$ according to $\xi_{M,k}$ and uses τ as exercise times is optimal at t_0 and for state ω^*. According to Theorem 1, this means that for every τ, the probability tree of level $k+1$, formed at time t_0, with intermediate times $t_1,\ldots,t_k$, is the same under both the truthful strategy and the alternative strategy:

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

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

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

References

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

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

Brier, G.W. (1950): "Verification of Forecasts expressed in terms of probability," *Monthly Weather Review*, 78, 1–3.

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

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

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

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