Supplementary Appendix to "Dynamic Belief Elicitation"

Christopher P. Chambers Nicolas S. Lambert

This supplement includes additional results and proofs omitted from the main paper. In Section S1, we provide an algorithm that computes the payoffs for a simple instance of strategyproof protocol in the general setting of Section 4. In Section S2, we show how to construct elicitation protocols for information structures involving potentially infinitely many time periods using menus with random deadlines. Sections S3–S5 are relevant to situations in which expert knowledge is solicited or evaluated for the purpose of helping decision makers. In Section S3, we show that, subject to regularity conditions, the knowledge of high-order beliefs elicited by the protocols we study is sufficient to solve any dynamic decision problem. In Section S4, we argue that knowledge of these high-order beliefs is much needed when the decision environment is dynamic: we ask what decision problems can be solved using the classical methods that elicit only first-order beliefs, and show they form a degenerate class. Finally, in Section S5, we illustrate our results in the context of simple principal-agent problems.

S1 Randomization Protocols: An Algorithm

While the protocols introduced in the general setting of Section 4 do not have a compact closed form representation, it is straightforward to compute efficiently the payoffs for special cases of menu randomizations. In this section, we show how this payoff computation could be done in practice with a simple algorithm.

Throughout, the set of possible outcomes is $\{1, \ldots, n\}$. In this proposed implementation, beliefs are represented by a standard tree structure that captures a probability tree with finite support at every level. So, for any belief given in this representation, the "children" are the beliefs in the support, and every branch carries a weight which is the probability for the associated belief in the support. Algorithm S1 takes as input the number of periods, the beliefs communicated at every period, and the final outcome. It computes the individual's payoff, by first drawing a menu at random (Algorithm S2), and then calculating the best payoff that can be achieved from the menu by making sequentially optimal choices given the information communicated (Algorithm S3). Securities follow the uniform distribution, and the number of elements in a menu follows the Poisson distribution with a given parameter λ. Because this randomization is full support, Theorem 1 guarantees that the protocol is strategyproof.

Algorithm S1 Compute the individual's payoff

Require: T (number of periods minus one), p_t (reported belief in period t), x (realized outcome)

```
menu ← RandomMenu(T)
t ← 1
while t < T − 1 do
    choice ← BestAction(p_t, menu)
    submenu ← menu[choice]
    menu ← submenu
    t ← t + 1
end while
choice ← BestAction(p_t, menu)
security ← menu[choice]
return security[x]
```

Algorithm S2 Draw a menu of a given level at random

Require: n (number of possible outcomes), λ (average number of submenus in a menu)

```
function RandomMenu(level)
    K ← random draw from the Poisson distribution with parameter λ
    menu ← empty array of K elements
    for all i ∈ {1, ..., K} do
        if level > 1 then
            menu[i] ← RandomMenu(level − 1)
        else
            menu[i] ← array of n elements drawn randomly and uniformly on [0, 1]
        end if
    end for
    return menu
end function
```

Algorithm S3 Compute the best choice in a given menu for a given belief

Require: n (number of possible outcomes)

```
function OptimalChoice(belief, menu)
    C ← number of possible choices in menu
    bestChoice ← ∅
    bestPayoff ← 0
    for all choice ∈ {1, . . . , C} do
        payoff ← MenuPayoff(belief, menu[choice])
        if payoff ≥ bestPayoff then
            bestPayoff ← payoff
            bestChoice ← choice
        end if
    end for
    return bestChoice
end function

function MenuPayoff(belief, submenu)
    if submenu is a security then
        return ∑_{i=1}^{n} belief.Pr[i] × submenu[i]
    else
        K ← number of probability trees in the support of belief
        payoffs ← empty array of K elements
        for all i ∈ {1, . . . , K} do
            choice ← OptimalChoice(belief.tree[i], submenu)
            payoffs[i] ← MenuPayoff(belief.tree[i], submenu[choice])
        end for
        return ∑_{i=1}^{K} belief.Pr[i] × payoffs[i]
    end if
end function
```

We stress that this is only one possible implementation, that takes as input the beliefs given in their "canonical" form, a probability tree. Depending on the context, one may want to account for other types of inputs, in the form of parameters of probability distributions as in Section 3.3, information structures that specify the nature of the information to be received and its distribution, or other convenient representations of dynamic forecasts. For instance, dynamic weather forecasts are more naturally expressed in terms of probability densities over future measurements and observations than high-order beliefs. This adaptation does not pose any particular difficulty, because to compute the payoffs, the algorithm must simply determine the best menu choices given the input provided.

S2 General Information Structures

In this section, we work with a more general class of information structures than in Section 4. The individual may now privately observe information dynamically and continuously over a unit interval of time. The uncertain outcome of interest materializes at $t = 1$. This setup includes two common instances not captured in Section 4.

1. Discrete information arrival with random times: the individual is to receive a given number of signals over the unit time interval, but as opposed to the simpler setup of Section 4, the times of arrival are random. Both the distribution of arrival times and their realization are private information to the individual. The elicitor now wants to learn the information that concerns both the signals, and the arrival times.

2. Information flow that arrives continuously over time: here the individual tracks a continuous signal over time modeled as a stochastic process, for example, a price process modeled as a Brownian motion with varying drift and scale. The elicitor wants to know the individual's assessment of the signal process distribution—such as drifts and diffusions terms—and, at every instant, the up-to-date signal value.

Of course, in principle, one could obtain approximately the information by considering an environment with a large but fixed number of time periods, and then apply the results of Section 4. Doing so complicates the elicitation unnecessarily. The purpose of this section is to show that we can obtain the information exactly and efficiently when adding random deadlines to the menus we consider in Section 4,

S2.1 Information Structures

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

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

Definition S1 *An information structure is a tuple* $(\Omega, \mathbb{F}, \mathrm{P}, X)$ *in which:*

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

- $\mathbb{F} = \{\mathcal{F}_t\}_{t\in[0,1]}$ *is a right-continuous filtration.*[S1] $\mathcal{F}_t$ *captures all events that are known to be either true or false by the individual at time* t*.*

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

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

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

In the sequel, $\Delta(\Omega)$ is the set of all probability measures on Ω, endowed with the weak-* topology. The filtration determines, in every state of the world, what the individual will know, and when he will know it. The prior determines the degree of uncertainty over the various states and events. Specifically, the full prior gives, at every time t, and in every state ω, the individual's posterior distribution P_t^ω over states of the world (which, of course, induces a posterior distribution over final outcomes, but generally also includes more information), accounting for all information available to the individual at that time and in that state. Note that the full prior includes information at every time, as opposed to prior information at $t = 0$ only. This is required to condition on future events in a consistent manner.[S2]

Definition S2 *Given a state space* Ω *and a filtration* $\mathbb{F}$*, a full prior is a stochastic process*[S3] $\mathrm{P} : (t, \omega) \mapsto \mathrm{P}_t^\omega$ *with values in* $\Delta(\Omega)$*, and such that*

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

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

The first condition is a technical requirement consistent with the fact that the individual learns information that is right-continuous (and every conditional probability of an event has a right-continuous version, for example, see Theorem 1.3.13, of Karatzas

[S1] Right continuity of a filtration models the requirement that the individual does not learn anything new in any upcoming infinitesimal length of time. This is a natural assumption in continuous-time dynamics, and a large body of the literature requires right-continuous filtrations (see, for example, §1.2 of Karatzas and Shreve, 1991).

[S2] It is well-known that a prior does not always generate consistent conditional probabilities, and even when it does, these are generality not unique. Embedding such information in a full prior avoids the need to put overly restrictive assumptions on the state space, and avoids the problem induced by the multiplicity of conditional probabilities. Our definition of full prior is similar to the Conditional Probability Systems of Myerson (1991). For general conditions that ensure systematic existence and coherence of regular conditional probabilities, see Berti and Rigo (1996).

[S3] The second condition implies by Lemma S1 of Section S2.5 that the full prior is a well-defined stochastic process.

and Shreve, 1991). The second condition means that the conditional probability at time t is known given the information the individual has access to at time t. We lack a condition requiring consistency of the posteriors over time: such a condition will not be needed for our results.

S2.2 Payoffs, Strategies, and Strategyproofness

Unlike in Section 4, we do not ask the individual to report his dynamic beliefs, because there is now an infinite hierarchy of high-order beliefs at every instant. Instead, we consider the more natural protocols in which the elicitor asks the individual to declare his information structure (at time $t = 0$) and, at every instant $t \geq 0$, the updated posterior over states. As before, the individual is rewarded at time $t = 1$ as a function of the data he communicated and the outcome that obtains.

Every protocol yields a payoff (or an expected payoff, if the protocol randomizes) to the individual that can be written directly as a function of the individual declarations and the realized outcome of X. For convenience, payoffs are now specified by a family of payoff rules. There is one payoff rule for every information structure the individual might communicate. We encode the flow of posteriors by $\mathcal{Q}$, the space of maps $Q : t \mapsto Q_t$, from times to probability measures over states of the world, that are such that for every event E, $Q_t(E)$ is right-continuous in time.[S4] A *payoff rule* $\Pi^{(\Omega,\mathbb{F},\mathrm{P},X)}$ maps $\mathcal{Q} \times \mathcal{X}$ to the interval of possible payoffs $[0, 1]$. It is required to be measurable. For every declared information structure $(\Omega, \mathbb{F}, \mathrm{P}, X)$, $\Pi^{(\Omega,\mathbb{F},\mathrm{P},X)}(\{Q_t\}_{t\in[0,1]}, x)$ is the payoff to an individual who reports information structure $(\Omega, \mathbb{F}, \mathrm{P}, X)$ at time $t = 0$ and, at every instant t, reports the posterior Q_t, while $X = x$ materializes.

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

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

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

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

[S4]The space $\mathcal{Q}$ is equipped with the product σ-algebra.

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

Of course, it is generally not possible to motivate the individual to reveal all of his information about Ω—for instance the individual may have information that does not even concern the random variable X. Insofar as the elicitor is exclusively concerned about the outcome of X, the individual's information is relevant only to the extent that it impacts the uncertainty on X. As in the main framework, this relevant information is captured by probability trees. However, because time is continuous, these trees can now have any level, and be associated with an arbitrary sequence of intermediate times. For example, at $t = 0.5$, the individual who believes, given his own information at that time, that states of the world are distributed according to Q, can infer a distribution over the public outcomes—a probability tree of level 1. The individual may also receive information between $t = 0.5$ and $t = 1$, say at $t = 0.8$. Thus, at $t = 0.8$, his beliefs about the public outcome are to be updated. At $t = 0.5$, the individual anticipates the update, and forms a belief about the distribution he is about to infer at $t = 0.8$—a probability tree of level 2, with 0.8 as time of interim distribution.

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

$$\varphi(Q) = Q(X).$$ [S5]

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

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

where P_{t_1} is the random variable of the process P sampled at time t_1. (By Lemma S2 in Section S2.5, the induced belief trees are well-defined and measurable.)

We continue to assume that the elicitor has interest in the probabilities that can be inferred from the individual's private information. Thus, at every time, the elicitor cares to learn about the individual's belief trees of all levels. A strategyproof protocol

[S5] For a random variable X, we let $Q(X)$ denote the law of X induced by the probability measure Q.

must therefore induce the individual, as a strict best response, to disclose enough information for the elicitor to learn the individual's belief trees of all levels and at all times. The individual is then induced to communicate all *relevant* information as a strict best response; however, as noted earlier, there will be many different information structures that will be equally relevant to the elicitor (this was already the case in Section 3.3).

Definition S3 *A protocol is strategyproof when, for each individual information structure* $(\Omega^*, \mathbb{F}^*, \mathrm{P}^*, X^*)$*:*

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

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

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

S2.3 Temporal Menus

To deal with the richness of the individual's information structure, we use an extended menu instrument that we refer to as *temporal menus*. These are menus with deadlines.

A *temporal menu of securities* is a pair $\sigma_0 = (M_0, \tau_0)$, where M_0 is a collection of securities and $\tau_0 \in [0, 1]$ is a fixed deadline. The owner of a temporal menu σ_0 must decide, at any $t \leq \tau_0$, to get one security among the collection M_0 (if $t > \tau_0$, the temporal menu is expired and delivers zero payoff). A temporal menu of securities is a *temporal menu of order 1*. Analogously to Section 4, we define in a recursive fashion a temporal menu of submenus of order k as a pair $\sigma_k = (M_k, \tau_k)$, τ_k is the menu's deadline and M_k is a collection of submenus of order $k-1$ whose deadlines are greater than τ_k. An elicitor who has a temporal menu σ_k must select one temporal submenu from M_k at any time $t \leq \tau_k$. A temporal menu is finite when it includes a finite number of submenus, which are in turn finite. Σ_k designates the collection of finite temporal menus of order k. For notational convenience, let Σ_0 be the space of securities, i.e., the continuous maps from Ω to $[0, 1]$.

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

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

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

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

S2.4 A Class of Strategyproof Protocols

We now turn to the elicitation procedure. In a preliminary step, the elicitor draws a random temporal menu of a random order according to a simple procedure detailed below.[S6] That menu is not disclosed to the individual. The elicitor then asks the individual to provide his information structure and to send an update on his posterior at every instant before the outcome realization. Based on the individual's announcements, the elicitor makes decisions optimally on behalf of the individual, under the assumption that the individual reports truthfully, and without revealing her menu choices to the individual. Eventually, at $t = 1$, the individual owns a security issued from the last decision made by the elicitor, and is paid off accordingly.

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

(a) Preliminary stage: the elicitor draws at random a finite number K from ξ_K. She then draws K deadlines $\tau_1, \ldots, \tau_K$ at random from $\xi_{\tau,K}$, and a finite menu M of order K at random from $\xi_{M,K}$. A temporal menu of order K is then formed by taking all the menus and submenus associated with the finite menu M, and respectively associating to each menu and submenu order k the deadline τ_k. The resulting temporal menu $\sigma^*_K = (M^*_K, \tau^*_K)$ is never disclosed to the individual.

(b) The individual's actions: at $t = 0$, the individual communicates an information structure $(\Omega, \mathbb{F}, \mathrm{P}, X)$. Then, at all subsequent times t, the individual communicates a posterior over states, $Q_t \in \Delta(\Omega)$.

[S6] The σ-algebra of events of finite temporal menus is defined analogously to that of the finite menus of the main framework of Section 4. Specifically, the space of finite menus of order 1 is equipped with a metric d where, if $\sigma' = (M', \tau')$ and $\sigma'' = (M'', \tau'')$ are two menus of order 1, $d(\sigma', \sigma'') = d(M', M'') + |\tau' - \tau''|$, with $d(M', M'')$ the Hausdorff distance between the sets of securities M' and M'', respectively. Next, in a recursive manner, the space of finite menus of order k is equipped with a metric d where, if $\sigma' = (M', \tau')$ and $\sigma'' = (M'', \tau'')$ are two menus of order k, $d(\sigma', \sigma'') = d(M', M'') + |\tau' - \tau''|$ with $d(M', M'')$ is the Hausdorff distance between the sets of submenus of order $k - 1$, M' and M'', respectively. We can then take the Borel σ-algebra induced by the metric. As earlier, we note that the σ-algebra of events does not depend of the particular metric chosen.

(c) The elicitor's actions: every time a deadline is reached, i.e., $t = \tau_k$, the elicitor privately chooses a temporal submenu $\sigma^*_{k-1} = (M^*_{k-1}, \tau^*_{k-1})$ uniformly at random from M^*_k, among all the submenus that are optimal assuming the individual has revealed and will reveal truthful information—i.e., such that under the declared information structure, $\pi_{k-1}(\sigma^*_{k-1}, Q_t) = \max_{\sigma_{k-1} \in M^*_k} \pi_{k-1}(\sigma_{k-1}, Q_t)$.

At the time of the last deadline, $t = \tau_K$, the elicitor selects for the individual a security S^* from M^*_1 (instead of a temporal submenu) following a similar procedure, i.e., uniformly at random among all the securities of M^*_1 that are optimal for the individual who has been truthful in the past.

The elicitor keeps all her menu choices private until $t = 1$, when the outcome materializes and the individual is offered the security S^*.

The following theorem is the main result of this section.

Theorem S1 *The elicitation protocol described above is strategyproof.*

It is worth pointing out that, as for the protocols of the main paper, this protocol uses randomization, but is equivalent to a nonrandom protocol that averages payoffs over the finite menus drawn in the preliminary stage.[S7] In the remainder of this section, we prove Theorem S1.

S2.5 Proofs

S2.5.1 Some Auxiliary Lemmas

The existence result requires the use of some technical lemmas.

Lemma S1 *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 $\mathbf{R}$ is measurable, then Ψ is measurable.*

[S7] We can write explicitly the payoff $\Pi^{(\Omega,\mathbb{F},\mathrm{P},X)}(\{Q_t\}_{t\in[0,1]}, x; \sigma^*)$ of the protocol for each particular draw of temporal menu σ^*. For a security S, we let $\Pi^{(\Omega,\mathbb{F},\mathrm{P},X)}(\{Q_t\}, x; S) = S(x)$. For a finite temporal menu $\sigma_k = (M_k, \tau_k)$ of order k, we let, recursively,

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

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

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

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

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 $f : \mathcal{B} \mapsto \mathbf{R}$. 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 subbase 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 subbase. 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 S2 *All induced belief trees are well-defined and measurable.*

Proof. The proof proceeds by induction. Let us fix an information structure $(\Omega, \mathbb{F}, \mathrm{P}, X)$. 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 S1, 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, $\omega \mapsto \mathrm{P}^{\omega}_{t_1}(E)$ is measurable for every event E, which by Lemma S1 implies that $\omega \mapsto \mathrm{P}^{\omega}_{t_1}$ is measurable, and so a well-defined random variable with values in $\Delta(\Omega)$. Thus, $\omega \mapsto \varphi_{t_2,\ldots,t_{k+1}}(\mathrm{P}^{\omega}_{t_1})$ 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 $\{\omega \in \Omega : \varphi_{t_2,\ldots,t_{k+1}}(\mathrm{P}^{\omega}_{t_1}) \in E\}$ is a well-defined event of Ω. Applying again Theorems 15.12 and 15.13 of Aliprantis and Border (2006), we get that $Q \mapsto Q(\varphi_{t_2,\ldots,t_{k+1}}(\mathrm{P}_{t_1}) \in E)$ is measurable, and by Lemma S1, we get that $Q \mapsto Q(\varphi_{t_2,\ldots,t_{k+1}}(\mathrm{P}_{t_1}))$ is measurable. ■

Lemma S3 *Given an information structure* $(\Omega, \mathbb{F}, \mathrm{P}, X)$*, every value map* $\pi_k : \Sigma_k \times \Delta(\Omega) \mapsto \mathbf{R}$ *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, $\omega \mapsto \mathrm{P}^{\omega}_{t}$ 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 of

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 $t \mapsto \mathrm{P}^\omega_t$ for every ω implies the joint measurability of $(t, \omega) \mapsto \mathrm{P}^\omega_t$ and thus the joint measurability of $(\sigma_{k+1}, \omega) \mapsto \mathrm{P}^\omega_{\tau_{k+1}}$, where we have decomposed σ_{k+1} as (M_{k+1}, τ_{k+1}). We have thus established that the map

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

is jointly measurable. It follows that the map

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

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 S4 *Every map* $(\{Q_t\}_t, x, \sigma_k) \mapsto \Pi^{(\Omega,\mathbb{F},\mathrm{P},X)}(\{Q_t\}_t, x; \sigma_k)$ *from* $\mathcal{Q} \times \mathcal{X} \times \Sigma_k$ *to* $\mathbf{R}$ *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 S3. 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 $\mathbf{R}$ 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,

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

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 S3 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 of 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

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

Thus $(\sigma_{k+1}, Q) \mapsto \left| \arg\max_{\sigma_k \in M_{k+1}} \pi(\sigma_k, Q) \right|$ is measurable as a pointwise limit of a sequence of measurable functions whose limit exists. Next, we observe that $(\tau, \{Q_t\}_t) \mapsto Q_\tau$ 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

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

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

S2.5.2 Proof of Theorem S1

It is clear that, because the protocol always works in the best interest of the individual, 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 individual and a strategy, following the notation of Definition S3.

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

We proceed by induction. For $k = 0$, 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 $(t_0, \dots, t_k) \mapsto \varphi_{t_1,\dots,t_k}(Q_{t_0})$ is separately right continuous. We have that $\varphi_{t_1,\dots,t_k}(Q_{t_0}) = Q_{t_0}(\varphi_{t_2,\dots,t_k}(\mathrm{P}_{t_1}))$. The right-continuity assumption on $t_0 \mapsto Q_{t_0}(E)$ for every event E ensures again the right continuity of $t_0 \mapsto Q_{t_0}(\varphi_{t_2,\ldots,t_k}(\mathrm{P}_{t_1}))$.

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

$$t_i \mapsto \int f(q)\, \mathrm{d}Q_{t_0}(\varphi_{t_2,\ldots,t_k}(\mathrm{P}_{t_1}) = q)$$

is right continuous, for every such f. Note that

$$\int f(q)\, \mathrm{d}Q_{t_0}(\varphi_{t_2,\ldots,t_k}(\mathrm{P}_{t_1}) = q) = \int f(\varphi_{t_2,\ldots,t_k}(\mathrm{P}^{\omega}_{t_1}))\, \mathrm{d}Q_{t_0}(\omega).$$

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

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

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 $\{t_0, \ldots, t_k : 0 \leq t_0 < \cdots < t_k \leq 1\}$ 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:

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

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

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

S3 Informational Sufficiency

The protocols that we study elicit, directly or indirectly, high-order beliefs. When the information these protocols elicit is used for the purpose of making better informed decisions, a natural question to raise is whether this information is refined enough to enable the decision maker to optimize as well as if she were as informed as the expert. We study this question in this section, and answer it positively.

We borrow the information structure from Section S2, and use finitely many fixed times of signal arrivals as in the multiperiod setting of Section 4. There are N time periods, $t = t_1, \dots, t_N$, $t_1 < \cdots < t_N$. In period t_N a random outcome X taking values in a compact metrizable set $\mathcal{X}$ materializes publicly. There is an expert whose information structure is $(\Omega, \mathbb{F}, \mathrm{P}, X)$. There is also a less informed utility-maximizing decision maker, who faces a dynamic decision problem: at every time t_k, the decision maker must choose an action a_k from a collection of possible actions $\mathcal{A}_k$, assumed compact metrizable. In the last period t_N, the decision maker receives utility $u(a_1, \dots, a_{N-1}, x)$, where x designates the final outcome that obtains. The decision maker's utility function, $u : \mathcal{A}_1 \times \cdots \times \mathcal{A}_{N-1} \times \mathcal{X} \mapsto \mathbf{R}$, is bounded and jointly continuous.

We demonstrate that, if the decision maker and the expert were the same and only person, then he or she would not be able to get more utility than the decision maker who only gets to observe the expert's high-order beliefs induced by his information structure and subsequent private observations.

To begin, we show that there always exists a solution to the decision problem: if the expert gets to decide which actions to take, he can always choose information-contingent actions that maximize the decision maker's expected utility. A decision policy for the expert is summarized by a tuple $(\alpha_1, \dots, \alpha_{N-1})$, where every α_k is an $\mathcal{F}_k$-measurable map from Ω to $\mathcal{A}_k$.[S8] 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 S5 *There exists a decision policy* $(\alpha_1^*, \dots, \alpha_{N-1}^*) \in \mathcal{D}(\mathbb{F})$ *such that*

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

Proof. We first note that

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

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 $(\alpha_1, \dots, \alpha_{N-1})$. Note that

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

[S8] Every space is tacitly endowed with its Borel σ-algebra.

is continuous in $(a_k, \dots, a_{N-1})$ by the Dominated Convergence Theorem. Thus, we have that the supremum $\sup_{a_{N-1}} \mathrm{E}[u(\alpha_1, \dots, \alpha_{k-1}, a_k, \dots, a_{N-1}, X) \mid \mathcal{F}_{N-1}]$ is continuous in $(a_k, \dots, a_{N-2})$ by Berge's Maximum Theorem. By the Measurable Maximum Theorem (Theorem 18.19 of Aliprantis and Border, 2006) it is also $\mathcal{F}_{N-1}$-measurable. An inductive argument yields that, for every k, and every decision policy $(\alpha_1, \dots, \alpha_{N-1})$,

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

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

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

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

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

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

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

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

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

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

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

which concludes the proof. ■

Now, let us assume that the expert does not take actions in place of the decision maker. Instead, the expert communicates to the decision maker his high-order beliefs 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 $Z_k(\omega) = \varphi_{t_{k+1},\dots,t_{N-1}}(\mathrm{P}_t^\omega)$, 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 S2. The variable Z_k represents the information that the expert communicates to the decision maker. Let $\mathbb{Z} = \{\mathcal{Z}_k\}_k$ be the filtration generated by the discrete process Z_k. $\mathbb{Z}$ represents the dynamic information the decision maker learns from the expert.

A decision policy for the decision maker 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 decision maker. The same argument as in Lemma S5 shows that there exists an optimal policy $(\beta_1^*, \dots, \beta_{N-1}^*)$, in the sense that

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

Because the decision maker only cares about information that is relevant to the random outcome, it is intuitive that information about the probability trees are enough for the decision maker 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 the set of possible states of the world Ω is finite, that intuition can be easily 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 decision maker 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 S1 *The decision maker's optimal actions that follow from communicating with the expert yield the same expected utility as if she had delegated the entire problem to the expert, i.e.,*

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

Proof. Note that if $W : \Omega \mapsto \mathbb{R}$ is $\sigma(X)$-measurable, then we obviously have that $\mathrm{E}[W \mid \mathcal{Z}_{N-1}] = \mathrm{E}[W \mid \mathcal{F}_{N-1}]$. However we also have that if $W : \Omega \mapsto \mathbb{R}$ is bounded and $\sigma(Z_{k+1})$-measurable, then $\mathrm{E}[W \mid Z_k] = \mathrm{E}[W \mid \mathcal{F}_k]$. 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 of Aliprantis and Border, 2006).

In particular, we have

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

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

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

In particular,

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

We conclude by the dynamic programming principle established in Lemma S5. Note that all suprema are measurable by the Measurable Maximum Theorem, as detailed in the lemma. ■

S4 Information Structures and Probabilities

In the context of a decision maker eliciting an expert's beliefs to help her make better decisions, we ask for which type of decision problems learning high-order beliefs is strictly more valuable than learning, over time, the first-order beliefs that the classical methods elicit. We argue that high-order beliefs are valuable for all but a degenerate class of dynamic decision problems, in which the payoffs are essentially additively separable.

The setup is as follows. There is a decision maker and three time periods $t = 0, 1, 2$. The information structure is borrowed from Section 3.1. In the final period $t = 2$, a utility-relevant random variable X realizes, taking values in $\mathcal{X} = \{1, \ldots, n\}$. In the interim period $t = 1$, the decision maker receives information and forms a posterior belief $p \in \Delta(\mathcal{X})$ on X. In the initial period $t = 0$, the decision maker forms a second-order belief $F \in \Delta(\Delta(\mathcal{X}))$, a distribution over the possible posteriors. In this period, the decision maker's prior over outcomes—let us call it $\mu(F)$—is

$$\mu(F) = \int p \, \mathrm{d}F(p). \tag{S1}$$

The decision maker is an expected-utility maximizer who confronts a dynamic choice. There are two actions to be taken in periods 0 and 1 respectively. If the decision maker chooses action $a_t \in \mathcal{A}_t$ in period t, she gets utility $u(a_0, a_1, x)$, where x is the realization of X, and where $\mathcal{A}_t$ is finite.

We define the value function for the decision maker in period 1 as

$$V(a_0, p) = \max_{a_1} \sum_{x \in \mathcal{X}} u(a_0, a_1, x) p(x). \tag{S2}$$

We also define the expected continuation payoff in period 0 as a function of the action taken in that period as

$$\overline{V}(a_0, F) = \int V(a_0, p) \, \mathrm{d}F(p). \tag{S3}$$

To maximize her overall expected payoff, the decision maker should choose action a_0 so as to maximize $\overline{V}(a_0, F)$, and then, after updating her belief to posterior p, choose action a_1 so as to maximize

$$\sum_{x \in \mathcal{X}} u(a_0, a_1, x) p(x). \tag{S4}$$

To simplify matters, assume that every action is at least weakly optimal for *some* posterior which is full support; that is, for every a_0, there is $p \in \Delta(\mathcal{X})$ for which $p(x) > 0$ for all x, and $V(a_0, p) \geq V(b, p)$ for all $b \in \mathcal{A}_0$. A *dynamic decision problem* in this section refers to any tuple $(\mathcal{A}_0, \mathcal{A}_1, u)$ satisfying this property. 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, and so no decision maker ever need to play this action, which, as a result, can be safely removed from the set of possible actions.

We are interested in the structure of the dynamic decision problems—here represented by $u(a_0, a_1, x)$—for which the only relevant information for the period-0 decision is the prior over outcomes $\mu(F)$. That is, for any second-order beliefs 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-0 action optimal under F_1 is also optimal under F_2. Formally, if

$$\overline{V}(a_0, F_1) = \max_{b_0} \overline{V}(b_0, F_1), \tag{S5}$$

then

$$\overline{V}(a_0, F_2) = \max_{b_0} \overline{V}(b_0, F_2). \tag{S6}$$

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 S2 *Let $(\mathcal{A}_0, \mathcal{A}_1, u)$ be a marginal dependent dynamic decision problem. There exists a function $u_1' : \mathcal{A}_1 \times \mathcal{X} \to \mathbf{R}$ and for each $a \in \mathcal{A}_0$, there exists $f_a \in \mathbf{R}^{\mathcal{X}}$ such that $V(a, p) = f_a \cdot p + \sup_{a_1 \in \mathcal{A}_1} \sum_{x \in \mathcal{X}} u_1'(a_1, x)p(x)$.*

The proof is relegated to the end of this section. To better understand Proposition S2, it is helpful to introduce two definitions. First, we will call a dynamic decision problem $(\mathcal{A}_0, \mathcal{A}_1, u)$ *additively separable* if for each $i = 0, 1$, there is $u_i : \mathcal{A}_i \times \Omega \to \mathbf{R}$ for which we can decompose the decision problem as $u(a_0, a_1, x) = u_0(a_0, x) + u_1(a_1, x)$. 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 $(\mathcal{A}_0, \mathcal{A}_1, u)$ and $(\mathcal{A}_0', \mathcal{A}_1', u')$ are *first-period payoff isomorphic* if the sets

$$\{V(a_0, \cdot) : a_0 \in \mathcal{A}_0\}$$

and

$$\{V'(a_0', \cdot) : a_0' \in \mathcal{A}_0'\}$$

coincide. To understand the definition of first-period payoff isomorphism, observe that, in terms of incentives in period 0, two isomorphic problems are equivalent, in that they present exactly the same set of period-1 value functions.

Corollary S1 *Any marginal independent dynamic decision problem $(\mathcal{A}_0, \mathcal{A}_1, u)$ is first-period payoff isomorphic to an additively separable dynamic decision problem.*

Corollary S1 is straightforward: let $\mathcal{A}_0' = \mathcal{A}_0$, and for any $a_0 \in \mathcal{A}_0$, define $u_0'(a_0, x) = f_a(x)$. Let $\mathcal{A}_1' = \mathcal{A}_1$, and u_1' as in Proposition S2, so that $V^*(p) = \sup_{a_1' \in \mathcal{A}_1'} \int u_1'(a_1', x)\, \mathrm{d}p(x)$.

Corollary S1 is an immediate consequence of Proposition S2, but facilitates its interpretation: probabilities over the final outcome are enough only when the available choice in period 0 does not constrain the available choices in period 1, and when the utility in period 1 does not depend on the utility in period 0.

We emphasize that our notion of dynamic decision problem does not directly incorporate any type of intertemporal budget constraint; the available decisions in period 1 (the set $\mathcal{A}_1$) do not depend on the decision made in period 0. Rather, our framework allows intertemporal budget constraints to be captured by modifying the utility function so that $u(a_0, a_1, \omega) = -\infty$ whenever a_1 is not available after having chosen a_0. Additive separability of course rules this type of construction out. Hence, problems obeying marginal dependence are very few indeed: marginal dependence requires that period-0 decisions do not constrain period-1 decisions at all, and further that there is no complementarity across the two time periods.

S4.1 Proof of Proposition S2

Fix $a, b \in \mathcal{A}_0$, and suppose that p^* 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 $H : \Delta(\mathcal{X}) \to \mathbf{R}$ 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 period 0 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 $[\delta_{p_a}, \delta_{p_b}] \subseteq \Delta(\Delta(\mathcal{X}))$.[S9] 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 subsegments of the form $[F_i, F_{i+1}]$, where some action c_i of $\mathcal{A}_0$ is optimal on (F_i, F_{i+1}). This fact owes to that the value function in period 0 (i.e., the function that gives the decision maker's expected payoff under optimal decisions

[S9] Here, δ_p denotes the Dirac measure on p.

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

$$V^*(p) = V(a^*, p) = \sup_{a_1 \in \mathcal{A}_1} \sum_{x \in \mathcal{X}} u(a^*, a_1, x) p(x)$$

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

S5 Application to Principal-Agent Problems

In this section, we apply our framework in a context of contractual design. Naturally, an in-depth analysis is beyond the scope of this paper. Instead, we work with two simple models, whose purpose is to illustrate how the revealed-preference approach of our paper can be used in principal-agent problems.

Throughout, a principal and an agent operate over three time periods $t = 0, 1, 2$. The principal has interest in a random event E whose outcome becomes publicly known in the final period. The associated indicator variable is X. The event occurs with 50% probability. In the interim period, the agent privately observes a signal S taking values in $\{-1, +1\}$. The signal contains information on the event, with conditional probabilities $\Pr[S{=}{+1}|X{=}1] = \Pr[S{=}{-1}|X{=}0] = (1+\gamma)/2 \geq 1/2$. The parameter $\gamma \in [0, 1]$ captures the precision of the agent's signal. After observing $S = s$, the agent's assessment of the event likelihood moves from 1/2 to $P[X{=}1|S{=}s] = (1{+}\gamma s)/2$. Given a precision γ, the agent's second-order belief, in the initial period, is therefore captured by the distribution F_γ that assigns probability 1/2 to posterior assessment $(1+\gamma)/2$, and probability 1/2 to posterior assessment $(1-\gamma)/2$.

We consider two standard environments. The first environment features adverse selection. The precision γ is exogenous and privately known to the agent. It captures the agent's ability. The principal knows the distribution of the ability in the population of agents, and wants to recruit agents with high ability. The second environment features moral hazard. The precision of the signal γ is now endogenous. It is privately chosen by the agent, who incurs a cost increasing with the precision. The principal faces a decision problem whose value depends on the event E and, being uninformed, contracts with the agent to buy information. The principal's objective is twofold. She wants to find the precision that best balances the benefit of taking a better action against the cost of employing the agent, and she also wants to minimize labor costs while ensuring that the agent does the due diligence on acquiring a signal of the desired precision.

S5.1 A Principal-Agent Problem with Adverse Selection

Here, the agent observes a signal with an exogenously determined precision γ, interpreted as the agent's ability. This ability is known to the agent, and unknown to the principal, who holds a prior distribution G on γ. The principal must decide whether to hire the agent. She derives a positive utility for hiring the most capable agents, and a disutility for hiring the least capable agents. Specifically, the principal's utility $u(\gamma)$ is continuously increasing in the agent's ability γ, and $u(0) < 0 < u(1)$.

To make her hiring decision, the principal requests that the agent provides the precision of his signal along with the observed value (or equivalently, asks the agent to tell which outcome is most likely). Once the event outcome publicly realizes, the principal hires the agent with some probability $\pi(\gamma, s, x)$, the *hiring policy*, that may depend on the reported precision γ, the reported signal s, and the realized outcome x. The principal commits to a hiring policy, and the agent acts so as to maximize the likelihood of being hired. By a classical revelation principle argument, it is without loss that we can restrict attention to incentive compatible policies in which the agent is best off reporting true information. The principal's objective is to design a policy that maximizes her expected utility.

Written formally, the principal maximizes expected utility

$$\int_0^1 u(\gamma)h(\gamma)\,\mathrm{d}G(\gamma)$$

over policies π, subject to π being incentive compatible, and where

$$h(\gamma) = \frac{1+\gamma}{4}\left(\pi(\gamma,1,1) + \pi(\gamma,0,0)\right) + \frac{1-\gamma}{4}\left(\pi(\gamma,0,1) + \pi(\gamma,1,0)\right)$$

is the probability that an agent with ability γ is hired.

In this simple environment, asking the agent to report his information is the same as asking the agent to assess the event likelihood. Abusing notation, any incentive-compatible policy $\pi(\gamma, s, x)$ corresponds to an incentive-compatible policy $\pi(p, x)$, with $p = (1 + \gamma s)/2$ a probability assessment of the event. Naturally, saying that π is incentive compatible is simply saying that π is a weakly strategyproof payoff rule, which in this case reduces to a proper scoring rule. For example, the policy

$$\pi(\gamma, s, x) = 1 - \left(\frac{1+s\gamma}{2} - x\right)^2$$

is incentive compatible and corresponds to the quadratic scoring rule $1 - (p - x)^2$.

Following the methodology of the main text, we construct hiring policies by averaging over payoffs for elementary decision problems, that are menus of securities as in Section 4. A security is represented as a vector of $\mathbf{R}^2$, the first component is the payoff when the event is false, and the second component is the payoff when the event is true (as opposed to the securities of Section 4, the range of payoffs is not bounded). Let $\mathcal{M}$ be the set of pairs of securities. This set represents menus of two securities. It is useful to order the securities: if $(x, y) = ((x_0, x_1), (y_0, y_1)) \in \mathcal{M}$ then $x_0 \leq y_0$ and $y_1 \leq x_1$ (this ordering is without loss of generality). Let us consider the class $\mathcal{S}$ of functions s that can be written as

$$s(p, 0) = \int_{\mathcal{M}} (x_0 \mathbb{1}_{(1-p)x_0+px_1 \geq (1-p)y_0+py_1} + y_0 \mathbb{1}_{(1-p)x_0+px_1 < (1-p)y_0+py_1}) \, \mathrm{d}\mu(x, y),$$
$$s(p, 1) = \int_{\mathcal{M}} (x_1 \mathbb{1}_{(1-p)x_0+px_1 \geq (1-p)y_0+py_1} + y_1 \mathbb{1}_{(1-p)x_0+px_1 < (1-p)y_0+py_1}) \, \mathrm{d}\mu(x, y),$$

where μ is a probability distribution over binary menus. These are the payoff rules generated by randomizations over the decision problems that require the agent to choose between two securities randomly chosen, in the spirit of Section 4. Let $\mathcal{C}$ be the set of payoff rules of $\mathcal{S}$ may deliver payoffs in the range $[0, 1]$. This set captures a large class of incentive-compatible hiring policies—and in fact, it captures *all* relevant policies.

Recall that Theorem 2 asserts that, subject to certain regularity conditions, averaging over large enough menus of securities with payoffs in $[0, 1]$ approximates arbitrarily closely any proper scoring rule. The above construction randomizes over menus of two securities only, but in this simple environment, larger menus are not needed. As long as the payoffs of the securities of $\mathcal{S}$ are not restricted to be in the range $[0, 1]$, the value functions of any incentive-compatible policy can be replicated by a member of $\mathcal{C}$. This is implied by the Schervish representation (Schervish, 1989). Thus, restricting ourselves the policies in the set $\mathcal{C}$ is without loss. In richer environments, one may want to consider larger menus.

The principal's prior over signals and abilities yields a prior over the agent's posterior p. Abusing notation, we continue to use G for the distribution function associated with that prior, and we continue to denote by $u(p)$ the utility the principal gets when hiring an agent with posterior p. The principal maximizes

$$\int_0^1 u(p)\pi(p) \, \mathrm{d}G(p) \tag{S7}$$

over policies $\pi \in \mathcal{C}$, where $\pi(p)$ denotes the expected probability of being hired for an agent whose posterior probability assessment is p. It is the same as the principal maximizing

$$\int_{\mathcal{M}} \int_0^1 u(p) \max\{(1-p)x_0 + px_1, (1-p)y_0 + py_1\} \, \mathrm{d}G(p) \, \mathrm{d}\mu(x, y) \tag{S8}$$

over all probability distributions μ over binary menus, and subject to the constraints

$$\int_{\mathcal{M}} x_0 \, \mathrm{d}\mu(x,y) \geq 0 \qquad \int_{\mathcal{M}} y_0 \, \mathrm{d}\mu(x,y) \leq 1$$
$$\int_{\mathcal{M}} y_1 \, \mathrm{d}\mu(x,y) \geq 0 \qquad \int_{\mathcal{M}} x_1 \, \mathrm{d}\mu(x,y) \leq 1$$

to capture the fact that π must take values in $[0,1]$.

Representing policies π as probability distributions μ is a key ingredient in the derivation of an optimal policy, which illustrates the benefit of the revealed-preference approach. The principal's original problem is transformed into a basic linear optimization problem, easily solved by standard methods.[S10]

The solution depends on the signs of $\mathrm{E}[u(\gamma)|S{=}-1, X{=}1]$ and $\mathrm{E}[u(\gamma)|S{=}1, X{=}1]$. The first expectation is the expected utility of a principal hiring an agent who makes the wrong prediction, and the second expectation is the expected utility of a

[S10] Consider, for example, the case $\mathrm{E}[u(\gamma)|S{=}-1, X{=}1] < 0 < \mathrm{E}[u(\gamma)|S{=}1, X{=}1]$. These inequalities are the same as

$$\int_0^{1/2} pu(p) \, \mathrm{d}G(p) < 0 \qquad \text{and} \qquad \int_{1/2}^1 pu(p) \, \mathrm{d}G(p) > 0. \tag{S9}$$

Relaxing the linear optimization problem to internalize the constraints, we maximize

$$\int_{\mathcal{M}} \int_0^1 u(p) \max\{(1-p)x_0 + px_1, (1-p)y_0 + py_1\} \, \mathrm{d}G(p) \, \mathrm{d}\mu(x,y)$$
$$- \int_{\mathcal{M}} (x_0 + y_1) \int_0^{1/2} pu(p) \, \mathrm{d}G(p) \, \mathrm{d}\mu(x,y) + \int_{\mathcal{M}} (x_1 + y_0) \int_{1/2}^1 pu(p) \, \mathrm{d}G(p) \, \mathrm{d}\mu(x,y), \tag{S10}$$

over probability distributions over binary menus, without further restrictions. The coefficients $\pm \int_0^{1/2} pu(p) \, \mathrm{d}G(p)$ and $\pm \int_{1/2}^1 pu(p) \, \mathrm{d}G(p)$ are Lagrange multipliers. By linearity, the problem reduces to maximizing over $(x,y) \in \mathcal{M}$ the value

$$\int_0^1 u(p) \max\{(1-p)x_0 + px_1, (1-p)y_0 + py_1\} \, \mathrm{d}G(p)$$
$$- (x_0 + y_1) \int_0^{1/2} pu(p) \, \mathrm{d}G(p) + (x_1 + y_0) \int_{1/2}^1 pu(p) \, \mathrm{d}G(p), \tag{S11}$$

which yields maximizer $x_0 = y_1 = 0$ and $x_1 = y_1 = 1$, for which (S11) evaluates to zero. This solution is easily verified: if $0 \leq x_0 < y_0 \leq 1$ and $0 \leq y_1 < x_1 \leq 1$, letting $q = (y_0 - x_0)/((x_1 - x_0) + (y_0 - y_1))$, (S10) is equal to

$$\int_0^q u(p) \, ((1-p)y_0 + py_1) \, \mathrm{d}G(p) + \int_q^1 u(p) \, ((1-p)x_0 + px_1) \, \mathrm{d}G(p)$$
$$- (x_0 + y_1) \int_0^{1/2} pu(p) \, \mathrm{d}G(p) + (x_1 + y_0) \int_{1/2}^1 pu(p) \, \mathrm{d}G(p). \tag{S12}$$

principal hiring an agent who makes the correct prediction. As $\mathrm{E}[u(\gamma)|S{=}{-1}, X{=}1] \leq \mathrm{E}[u(\gamma)|S{=}1, X{=}1]$, there are three possible cases. In spite of the rich structure of the incentive-compatible policies, the optimal policy takes a very simple form.

If $\mathrm{E}[u(\gamma)|S{=}{-1}, X{=}1] < 0$ and $\mathrm{E}[u(\gamma)|S{=}1, X{=}1] > 0$, then the solution to the constrained linear programming problem with objective (S8) is a probability measure μ^* that puts full mass on the binary menu $((0,1),(1,0))$. In words, the optimal contract only requires the agent to tell which event outcome is most likely. The agent is hired only if the predicted outcome occurs, independently of the agent's actual ability.

If $\mathrm{E}[u(\gamma)|S{=}{-1}, X{=}1] > 0$ and $\mathrm{E}[u(\gamma)|S{=}1, X{=}1] > 0$ then the solution is probability measure μ^* that puts full mass on the binary menu $((1,1),(1,1))$: it is optimal to always hire the agent, disregarding any information the agent might report—even though the principal incurs a disutility for hiring agents with low ability. Conversely, if instead $\mathrm{E}[u(\gamma)|S{=}{-1}, X{=}1] < 0$ and $\mathrm{E}[u(\gamma)|S{=}1, X{=}1] < 0$, it is optimal to never hire any agent—even though the utility of hiring agents with high ability is positive.

S5.2 A Principal-Agent Problem with Moral Hazard

The principal now faces a decision problem in the interim period. The utility that the problem generates depends on the event outcome. The exact problem the principal confronts is irrelevant: analogously to Section S5.1, we simply assume that, if the principal were to observe a signal of precision γ, she would obtain utility $u(\gamma)$, with u continuously increasing. The principal gets informed via the agent, whose expertise enables the acquisition of a signal of arbitrary precision γ at cost $c(\gamma)$, with $c(\cdot)$ a nonnegative, strictly convex and strictly increasing function. In the initial period, the agent chooses the precision of the signal that realizes in the interim period. Thus, this decision determines the agent's second-order belief.

The principal hires the agent and requests a probability assessment in the interim period. She pays the agent $\pi(p, x)$ when the agent communicates assessment p and $X = x$. A *contract* (γ, π) specifies the desired precision γ of the signal to be observed, and the payment scheme π. It is *incentive compatible* when the agent is motivated to acquire a signal of precision exactly equal to γ, and to report truthfully her probability

---

If, for example, $q < 1/2$ then (S12) is equal to

$$(x_0 - y_0) \int_q^{1/2} (1-p)u(p)\,\mathrm{d}G(p) + (x_1 - y_1) \int_q^{1/2} pu(p)\,\mathrm{d}G(p),$$

which is nonpositive by monotonicity of $u(p)$ and (S9). An analogous conclusion holds for $q > 1/2$. Thus, a probability measure μ^* that puts full mass on the binary menu $((0,1),(1,0))$ maximizes (S10). The solution of the relaxed problem solves the original constrained problem, as for every μ associated with a policy in $\mathcal{C}$, the value of (S10) is at least as large as the value of (S7), and the two are equal for $\mu = \mu^*$.

assessment. The contract is *individually rational* when, in the initial period, the agent makes a nonnegative expected net utility if he accepts the terms of the contract (the reservation utility is normalized to zero). Finally, all payments made must be nonnegative.[S11] The principal commits to a contract in the initial period.

The principal must offer a contract (γ, π) that satisfies incentive compatibility and individually rationality, and that guarantees nonnegative payments. Let $\mathcal{F}$ denote the class of these feasible contracts. For a payment scheme π, $\pi(p)$ denotes the expected payment to the agent, in the interim period, when the agent supplies assessment p truthfully. The objective for the principal is then to maximize her utility from the decision problem she confronts minus labor costs,

$$u(\gamma) - \int_0^1 \pi(p)\,\mathrm{d}F_\gamma(p), \tag{S13}$$

over feasible contracts $(\gamma, \pi) \in \mathcal{F}$. If, for every fixed $\gamma \in (0, 1)$, an optimal payment scheme π_γ is known, then the principal simply chooses γ so as to maximize (S13) for $\pi = \pi_\gamma$.

We begin with a derivation of an optimal payment scheme that induces the agent to acquire a signal with precision $\gamma \in (0, 1)$ at minimal cost for the principal. As in Section S5.1, a key observation is that the relevant class of incentive compatible payment schemes π is a subset of class $\mathcal{S}$ (as defined in Section S5.1), whose elements are all the mixtures of elementary payment schemes. Its boundaries are determined by the constraint of feasibility.

Having made this observation, the objective of the principal reduces to the choice of probability measure μ over the set of binary menus $\mathcal{M}$, which yield payment scheme π, so as to minimize the expected payments to the agent,

$$\begin{aligned}\int_0^1 \pi(p)\,\mathrm{d}F_\gamma(p) &= \int_{\mathcal{M}} \int_0^1 \max\left\{(1-p)x_0 + px_1, (1-p)y_0 + py_1\right\} \mathrm{d}F_\gamma(p)\,\mathrm{d}\mu(x, y)\\ &= \int_{\mathcal{M}} \frac{1}{4}\left(x_0 + x_1 + y_0 + y_1 + \gamma(x_1 + y_0 - x_0 - y_1)\right) \mathrm{d}\mu(x, y),\end{aligned}$$

and subject to the condition that the payment scheme be feasible, which can be captured by the following linear constraints:

$$\int_0^1 \pi(p)\,\mathrm{d}F_\gamma(p) \geq c(\gamma), \tag{S14}$$

$$\forall \delta, \quad \int_0^1 \pi(p)\,\mathrm{d}F_\gamma(p) \geq \int_0^1 \pi(p)\,\mathrm{d}F_\delta(p), \tag{S15}$$

$$\forall p, x, \quad \pi(p, x) \geq 0. \tag{S16}$$

[S11] Because we allow for $c(0) > 0$ (a fixed cost for working for the principal), nonnegative payments do not imply individual rationality.

Inequality (S14) is the individual rationality constraint. Inequality (S15) makes the contract incentive compatible regarding the precision (note that truthful reporting is already guaranteed since $\pi \in \mathcal{S}$). Inequality (S16) ensures payments are nonnegative. The set of constraints can then be further simplified as

$$\frac{1}{4}\int_{\mathcal{M}} (x_0 + x_1 + y_0 + y_1 + \gamma(x_1 + y_0 - x_0 - y_1))\,\mathrm{d}\mu(x,y) \geq c(\gamma), \tag{S17}$$

$$\frac{1}{4}\int_{\mathcal{M}} (x_1 + y_0 - x_0 - y_1)\,\mathrm{d}\mu(x,y) = c'(\gamma), \tag{S18}$$

$$\int_{\mathcal{M}} x_0\,\mathrm{d}\mu(x,y) \geq 0, \tag{S19}$$

$$\int_{\mathcal{M}} y_1\,\mathrm{d}\mu(x,y) \geq 0, \tag{S20}$$

where (S18) is a first-order condition equivalent to (S15), and we observe that (S16) reduces to the two cases (S19) and (S20). As in Section S5.1, the problem reduces to a simple linear programming problem that is easily solved. There are two cases to consider, as function of how $(1+\gamma)c'(\gamma)$ compares to $c(\gamma)$.

First, suppose that $(1+\gamma)c'(\gamma) > c(\gamma)$. This condition means that the normalized cost $c(\delta)/(1+\delta)$ is locally strictly increasing at $\delta = \gamma$. A solution is then given by the probability measure μ^* that puts full mass on the menu of the two securities $(0, 2c'(\gamma))$ and $(2c'(\gamma), 0)$. Therefore, an optimal payment scheme is

$$\pi(p,x) = \begin{cases} 2(1-x)c'(\gamma) & \text{if } p < 1/2, \\ 2xc'(\gamma) & \text{otherwise.} \end{cases}$$

On average the agent earns $(1+\gamma)c'(\gamma)$, which exceeds his cost $c(\gamma)$. Thus, in this case, the agent derives a positive rent from working for the principal.

Second, suppose that $(1+\gamma)c'(\gamma) \leq c(\gamma)$, so that the normalized cost $c(\delta)/(1+\delta)$ is locally weakly decreasing at $\delta = \gamma$. A solution is given by the probability measure μ^* that puts full mass on the menu of the two securities $(c(\gamma) - (1+\gamma)c'(\gamma), c(\gamma) + (1+\gamma)c'(\gamma))$ and $(c(\gamma) + (1+\gamma)c'(\gamma), c(\gamma) - (1+\gamma)c'(\gamma))$. It corresponds to the optimal scheme

$$\pi(p,x) = c(\gamma) - (1+\gamma)c'(\gamma) + \begin{cases} 2(1+\gamma)(1-x)c'(\gamma) & \text{if } p < 1/2, \\ 2(1+\gamma)xc'(\gamma) & \text{otherwise.} \end{cases}$$

In this case, the agent earns on average $c(\gamma)$—just enough to cover his cost—and the agent gets his reservation payoff.

Finally, plugging the payment scheme just obtained into the principal's objective (S13), we obtain that the principal chooses precision γ that maximizes $u(\gamma) - \max\{(1+\gamma)c'(\gamma), c(\gamma)\}$.

References

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

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

Schervish, M. J. (1989). A general method for comparing probability assessors. *The Annals of Statistics 17*(4), 1856–1879.