NONDISTORTIONARY ELICITATION OF BELIEFS IN DYNAMIC PROBLEMS

MARCIN PĘSKI AND COLIN STEWART
UNIVERSITY OF TORONTO

Abstract. We study elicitation of beliefs that are related to choices in a dynamic decision problem. We consider three cases with respect to the relationship between actions and beliefs: (i) beliefs are independent of actions, (ii) actions affect information but not the payoff-relevant state, and (iii) actions affect the distribution of states. In each case, we provide necessary and sufficient conditions for incentivizing truthful reports of beliefs without distorting behavior in the original decision problem. For two-period decision problems, questions asked in the first period can be incentivized if and only if they are about expectations of payoffs plus some function of the belief, where the class of functions that can be added varies across the three cases. In contrast, incentivizing truthful reports in the second period always distorts incentives in the first period unless the first-period action affects neither the information nor the incentives in the second period.

Keywords: belief elicitation, experimental design, dynamic problems

1. Introduction

In many experimental and empirical studies, researchers are interested in eliciting subjects' beliefs about the consequences of their actions. These beliefs may change over time due to the arrival of new information or because actions affect the quality of information or the distribution of some payoff-relevant state of the world. For example, a researcher may want to compare a student's expectation about their score before and after taking a test. Alternatively, a researcher may want to evaluate what the subject anticipates about the quality of information by asking about their expected improvement in a prediction task.

We analyze a two-period model in which a decision-maker (DM) chooses an action in each period, may receive new information between the two periods, and faces a belief elicitation question that can depend on the actions chosen. We vary the model

Date: February 11, 2026.

This research was supported by the Social Sciences and Humanities Research Council of Canada.

along two dimensions: whether the elicitation question takes place in the first period or the second (i.e., before or after the arrival of new information), and how the first-period action can affect the decision-maker's belief about the state. In each case, we identify the class of elicitation questions that are incentivizable, meaning that the DM has strict incentives to answer the question truthfully and the optimal actions in both periods are the same as in the problem without the elicitation question.

Our framework allows for a wide range of questions the researcher may ask. In the first period, for example, the researcher may ask the DM how likely it is that his choice in the second period will be a particular action. Alternatively, if the main task involves a test in period two and the new information is the content of the test, the researcher may ask the DM about his expected score on the test, or about how likely he thinks it is that his prediction of his score after taking the test will be within some fixed error band around his actual score. In the second period, the researcher may, for example, ask questions about the DM's expected payoff, or about how likely the DM believes it is that his actions are optimal ex post.

To match a variety of decision problems, we consider three different cases with respect to how the first-period action can affect the information the DM receives:

(1) The information is not affected by the first-period action. Applications include the above examples where the main task is a test.
(2) The first-period action may affect the quality of information but not the distribution of states. Applications include information acquisition experiments and bandit problems.
(3) The first-period action can affect both the quality of information and the distribution of states. Applications include dynamic investment and contracting problems in which the first-period action can affect the success probabilities associated with each second-period action, or dynamic games in which the DM's uncertainty is about actions played by the other player, which is, in equilibrium, affected by the DM's action choice.

The problem of nondistortionary belief elicitation was introduced in Pęski and Stewart (2025) (henceforth PS). Here, we extend the definition to dynamic decision problems. In our model, a DM (whom we sometimes refer to as the subject in an experiment) takes actions over two periods, after which he receives payoffs that depend on his choices together with an unknown state of the world θ. The DM enters the problem with a privately known belief about the state of the world and about the information he will receive between the two periods. The information the DM receives may affect his beliefs in the second stage in accordance with Bayesian updating. To elicit some property of the DM's beliefs, a researcher designs an elicitation mechanism. The mechanism can ask the DM to report the expected value of some function of the state, and can also ask the DM to make some additional choices. What the mechanism asks for can depend on the actions chosen by the subject so far. The mechanism rewards the DM according to the actions chosen in the decision problem, the report and action choices in the elicitation mechanism, and the state of the world. The mechanism is *nondistortionary* if the optimal choices in the original problem without elicitation remain optimal in the combined problem with the elicitation mechanism.

The elicitation of dynamic beliefs with strict incentives for truth-telling has been studied previously, most notably in Chambers and Lambert (2021). The previous literature focuses on elicitation problems in isolation, without a related decision problem in which actions could be distorted. To understand the relevance of potential distortions, consider the above example of a DM who is asked to predict his score before taking a test. Rewards for a correct prediction could create an incentive to intentionally underperform on the test to ensure the accuracy of the prediction. Such a distortion may skew the results of the test and make their interpretation more difficult. In other settings, such as field experiments in health and education, distortions could conflict with ethical principles or generate liability issues. Designing payments for belief elicitation that do not distort the incentives in the original problem allows the researcher to honestly tell the subject that he will maximize his expected payment by choosing the action he believes is optimal in the decision problem and then reporting his belief truthfully.

Our results provide necessary and sufficient conditions for a question to be incentivizable. For questions that satisfy our sufficient conditions, we provide a simple construction of payments satisfying both of our incentivizability criteria. This construction is based on the classic Becker-DeGroot-Marschak method.

If actions do not affect information, the incentivizability of questions in the first period can be analyzed using the results from PS. A key difference is that the payoff-relevant state—the second-period belief—is not observed. To deal with this complication, we allow for supporting actions that play the same role as the experimental protocols studied in Chambers and Lambert (2021) (where the elicitation problem faces the same issue). Our results explain how and when such protocols can be used without distorting behavior in the first period.

	Action-independent information	Actions may affect the quality of information	Actions may affect information
$t = 1$	expected payoffs enhanced by arbitrary function of beliefs $u_1^*(a_1, p) + f(p)$	expected payoffs enhanced by affine function of beliefs $u_1^*(a_1, p) + l(p)$	expected payoffs $u_1^*(a_1, p)$
	PS conditions	nontrivial continuation for almost all period 1 actions	Always
$t = 2$	expected payoffs if payoffs are separable; action-independent affine functions otherwise	nothing	nothing
	separable payoffs		

TABLE 1. Results: Top row: sufficient conditions for incentivizability. Bottom row: conditions when the sufficient conditions are also necessary.

The analysis of two other cases differs substantively from the previous literature and is a novel contribution of this paper. When the actions can affect the quality of information, we show that information about expected payoffs enhanced with questions about the (true) payoff states can still be elicited. The latter are elicited by asking the DM about the expected value of an affine function of beliefs. But, unlike in the previous case, one cannot incentivize elicitation of general beliefs about second-period information. For example, a question about the option value to stop or continue the experiment at a fee cannot be incentivized without distorting the incentives governing the first-period actions. Perhaps more surprising is the fact that one cannot incentivize questions like "What is the probability that you will choose a_2 as the second-period action?" unless this probability is either 0 or 1.

The necessary and sufficient conditions for incentivizability coincide when all (or at least sufficiently many) first-period actions have a nontrivial continuation: the second-period decision is not predetermined by the first-period action. In such a case, a question is incentivizable if and only if it is a combination of expected payoffs and an action-independent affine function.

When actions affect the distribution of states, the only incentivizable questions are about expected payoffs.

The analysis of the incentivization of second-period questions is another original contribution of this paper. Elicitation of such questions may distort incentives in both periods. In particular, the anticipated rewards from truthfully reporting second-period information increase the value of information from the point of view of the first-period decision making. If the first-period actions may affect the quality of information, the increase in value of information may tip the balance towards more informative actions. Even in the action-independent case, if the first-period actions have a differential impact on the second-period value of information, then any second-period elicitation will affect the first-period incentives. To demonstrate the latter, we use an example with two first-period actions: safe, which ensures a state-independent payoff, and risky, where the payoff depends on how informed the second-period choice is. We show that in this example, the second-period question about expected payoffs cannot be incentivized without distorting first-period incentives.

We use the interaction between the incentive constraints to show three results. When information is action-independent and the payoffs in the original problem are separable, i.e., they can be represented as a sum of two functions each of which depends only on the action from one period, all incentivizable second-period questions are equivalent to questions that depend only on the second-period action. (The notion of equivalence means that there is another incentivizable question that does not depend on the first-period action, and, using the observed actions, one can use the second question to compute the value of the original question.) Under separable payoffs and action-independent information, the first-period choice does not interact with the second-period value of information, and if the question does not depend nontrivially on the first-period choice, it can be appropriately incentivized.

In the opposite case, for generic payoffs, the first-period choice affects the second-period incentives. We show that, in such a case, the only incentivizable questions are equivalent to questions that are action-independent.

Finally, when the first-period actions may affect the quality of information or distribution of states, there are no incentivizable questions. In such a case, the increase in the value of information due to incentivization will tip the balance towards actions that generate more information.

See Table 1 for a summary of the results.

1.1. **Related literature.** Chambers and Lambert (2021) study belief elicitation in dynamic environments, where a DM, privately, and subjectively, anticipate obtaining new information. They develop a framework to design simple protocols and elicitation mechanisms to induce the DM to reveal his beliefs, including anticipated information flow. They prove that their protocols elicit essentially all information that can be elicited in an incentivized way. The key difference with our paper is that Chambers and Lambert (2021) focuses on situations where information is independent of any decisions, whereas we study situations where information is naturally used by the agent in decision problems, including lab or field experiments, and/or the information may itself be affected by the DM's actions.

Tsakas (2020) shows that it is always possible to design a mechanism for eliciting the prior belief of a DM who can acquire information at a cost (as in rational inattention). In our framework, his problem is one in which the action—i.e., the information acquisition choice—affects the quality of information but not the state, and there is a unique optimal action in the original decision problem (not acquiring any information).

PS analyzes nondistortionary elicitation in static problems. In that paper, we provide sufficient conditions in simple, BDM-type mechanisms. We also provide necessary conditions and show that, in three important classes of decision problems, the necessary and sufficient conditions coincide. In those problems, our results fully characterize all incentivizable questions: those are, up to equivalence and problem-specific details, questions about expected payoffs enhanced with an action-independent question about states.

The key difference with the current work is that here, the agent chooses actions in two periods, and importantly, there is new information arriving between the two periods. We consider different scenarios that vary the period in which the question is asked and the effect of the first-period action on the information. When the information is action-independent and the question is asked in the first period, the analysis of the dynamic problem can be done using a combination of results from PS with protocols from Chambers and Lambert (2021). In all other cases, the analysis goes substantively beyond the past literature.

2. Examples

The following two examples are used throughout the text for illustration purposes.

2.1. **Student.** In a university class, a student's final grade comes from a combination of a score on a test and a current grade of $h = 50\%$ on homework assignments. The weight on the test, which is either $\frac{1}{3}$ or $\frac{2}{3}$, is chosen by the student before the test. The test consists of a single multiple-choice question with at least three answers, of which the student chooses one. Assume that the student's test score is marked as 1 if his answer is correct and 0 otherwise. The student's payoff is proportional to the final grade.

Here, the uncertainty is about which answer is correct. When the student sees the test, he learns how likely he is to know the answer to the question. Thus the uncertainty resolves in two stages. Importantly, the uncertainty is not affected by the first-period action.

A researcher running an experiment wants to analyze the student's beliefs, including how much the student will learn from seeing the test. Here are some questions she may ask:

(1) "What is your expected test score?";
(2) "What final grade reduction would you accept for the ability to revise the grading scheme after the test?";
(3) "What would your expected final grade be if, when you choose the high test weight, the marking scheme was reversed so that you get mark of 1 if your chosen answer is incorrect and 0 otherwise?"

Questions (2) and (3) are asked before the test. Notice that the answers to questions (2) and (3) depend on the grading scheme, i.e., the first-period action. Question (1) can be asked before or after the test.

The researcher would like to incentivize the student to provide truthful answers. The university research ethics board requires the researcher to ensure that the student's behavior in the class is not affected by the experiment.

We show that question (1) asked before or after the test and question (2) are incentivizable without creating distortions. Question (3) is not incentivizable: any attempt to incentivize elicitation of beliefs will distort the incentives in the first period.

2.2. **Oncologist.** In caring for a cancer patient, an oncologist makes two decisions: the diagnostic approach and the treatment plan. The diagnostic choice is between a liquid biopsy (a specialized blood analysis) and a tissue biopsy; the latter is more invasive but provides more definitive information. Subsequently, based on the estimated tumor burden, the oncologist chooses between chemotherapy and active surveillance, which differ in their impact on the tumor count and the well-being of the patient. The oncologist's objective is to select the path that maximizes the patient's Quality-Adjusted Life Years (QALYs).

Here, the uncertainty is about the initial tumor burden. The choice of the first-period action affects the quality of information but has no direct impact on QALY.

A researcher analyzes how medical decisions are made. She wants to incentivize elicitation of one of the following beliefs:

(1) "What is the expected QALY?" This question can be asked in either period (i.e., alongside the diagnostic decision or the treatment decision).
(2) "How likely are you to prescribe chemotherapy?" This question can be asked in the first period (alongside the diagnostic decision).

As in the previous example, the approval process for the experiment commits the researcher not to distort medical decisions. This implies that the researcher must use nondistortionary elicitation.

We show that question (1) asked alongside the diagnostic decision is incentivizable. Neither question (1) asked alongside the treatment decision nor question (2) are incentivizable. In particular, incentivizing the prediction of the treatment distorts incentives governing the diagnostic decision.

3. Model

3.1. **Decision problem.** A single-period decision problem is a pair (A, u) of a set of actions A and a payoff function $u : \Theta \times A \longrightarrow \mathbb{R}$, where Θ is a fixed finite state space. The decision-maker maximizes expected payoffs given beliefs $p \in \Delta\Theta$.

This paper focuses on dynamic decision problems, with two periods. An agent participates in a dynamic experiment (A_1, A_2, u). In each period t, the agent chooses an action $a_t \in A_t$. The payoffs $u(a_1, a_2, \theta)$ depend on the two actions as well as an unknown state of the world $\theta \in \Theta$. The agent is a Bayesian expected utility maximizer. If the action sets are clear from the context, we refer to the experiment using its payoff function u. We assume that the sets of actions in the decision problem u are finite. Below, we consider modifications of u that expand the set of actions to, possibly, infinite sets.

The decision-maker receives information about the state of the world. We assume that the information itself is not observable to the researcher. The information in period 2 is defined as a belief, i.e., a probability distribution over the state space $p \in \Delta\Theta$. Define the expected optimal payoffs following the choice of action a_1 and given belief $p \in \Delta\Theta$ as

$$u_1^*(a, p) = \max_{a_2} \mathbb{E}_p u(a_1, a_2, \theta).$$

The beliefs in period 1 are more complicated because the decision-maker anticipates possibly receiving some information between the two periods. It is without loss to describe the agent's beliefs as a probability distribution over posterior beliefs, i.e., as an element of $\Delta\Delta\Theta$. In general, the information may depend on the first-period action. Thus, we define the information structure as a mapping $\mu \in \mathcal{U} \subseteq (\Delta(\Delta\Theta))^{A_1}$, where we interpret $\mu_{a_1} = \mu(a_1)$ as the anticipated information after taking action a_1. The set $\mathcal{U}$ contains all admissible information structures. We consider the following three cases:

- action-independent information:

$$\mathcal{U}_{\text{ind}} = \{\mu \in (\Delta(\Delta\Theta))^{A_1} : \mu_{a_1} = \mu_{b_1} \text{ for all } a_1, b_1 \in A_1\},$$

- actions affect the quality of information:

$$\mathcal{U}_{\text{quality}} = \{\mu \in (\Delta(\Delta\Theta))^{A_1} : \mathbb{E}_{\mu(a_1)}p = \mathbb{E}_{\mu(b_1)}p \text{ for all } a_1, b_1 \in A_1\},$$

- actions affect states:

$$\mathcal{U}_{\text{general}} = (\Delta(\Delta\Theta))^{A_1}.$$

The optimal decisions in experiment u are given by

$$A_2^u(a_1, p) = \arg\max_{a_2} \mathbb{E}_p u(a_1, a_2, \theta),$$
$$A_1^u(\mu) = \arg\max_{a_1} \mathbb{E}_{\mu(a_1)} u_1^*(a_1, p) = \arg\max_{a_1} \mathbb{E}_{\mu(a_1)} \left(\max_{a_2} \mathbb{E}_p u(a_1, a_2, \theta) \right)$$

where the expectation in the first line and the second expectation in the second line are over the state θ and the first expectation in the second line is over p. For later use, it is also convenient to define the inverse operators: for each $a_1 \in A_1, a_2 \in A_2$,

$$P_1^u(a_1) = \{\mu \in \mathcal{U} : a_1 \in A_1^u(\mu)\},$$
$$P_2^u(a_1, a_2) = \{p \in \Delta\Theta : a_2 \in A_2^u(a_1, p)\}.$$

Sets $P_1^u(a_1)$ and $P_2^u(a_1, a_2)$ are closed and convex.

Say that a_2 is an *essential best response* following a_1 if $P^u(a_1, a_2)$ has a nonempty interior. Let

$$\mathcal{P}(a_1) = \{P^u(a_1, a_2) : a_2 \text{ is essential b.r. after } a_1\}$$

be the collection of closed subsets of beliefs $P \subseteq \Delta\Theta$ induced by second-period essential best responses after a_1. Each collection $\mathcal{P}(a_1)$ consists of closed convex sets; their union is equal to the space of all beliefs and the intersection of any two of them has empty interior. (The latter is a consequence of the fact that if there are two actions a_2, b_2 and an open set of beliefs where a_2, b_2 are best responses following a_1, then $P(a_1, a_2) = P(a_1, b_2)$.) Slightly abusing language, we refer to $\mathcal{P}(a_1)$ as the belief partition after action a_1.

In order to avoid breaking the second-period ties, some of our results are restricted to beliefs for which there is always a unique best response no matter what the continuation is. Let $\Delta^u = \{p : |A^u(a_1, p)| = 1\}$ be the set of such beliefs. We assume that Δ^u is a dense subset of all beliefs $\Delta\Theta$ (this is an assumption on the decision problem payoffs and, because the action sets are finite, it is generically satisfied). Let

$$\mathcal{U}^u = \{\mu \in \mathcal{U} : \mu_{a_1}(\{p : |A^u(a_1, p)| > 1\}) = 0 \; \forall a_1\}.$$

3.2. **Expansions.** In order to elicit beliefs, a researcher designs an expanded version of the original decision problem, with a larger space of actions. Formally, an *expansion* of problem (A_1, A_2, u) is any decision problem $(A_1 \times S_1, A_2 \times S_2, v)$. The expansion is *nondistortionary* if

$$A_1^u(\mu) = \{a_1 : \exists s_1 \text{ st. } (a_1, s_1) \in A_1^v(\mu)\} \text{ for each } \mu,$$
$$A_2^u(a_1, p) = \{a_2 : \exists s_1, s_2 \text{ st. } (a_2, s_2) \in A_2^v((a_1, s_1), p)\} \text{ for each } p \text{ and } a_1.$$

In a nondistortionary expansion, the optimal choices of the original experiment remain optimal in the expansion.

We are interested in a special type of expansion called an *elicitation mechanism.* In either period $t = 1$ or 2, the DM is asked to report the value of a random variable. Additionally, the DM may be asked to choose other actions, with the idea that additional choices may help with elicitation. Formally, assume that $S_t = \mathbb{R} \times S'_t$ where S'_t and $S'_{-t} = S_{-t}$ (where $-t$ is the other period) are *finite*. The interpretation is that the t-period action (a_t, r, s_t) in the elicitation mechanism consists of the action in the original problem a_t, the elicitation report r, and the supporting action s_t.

If the elicitation mechanism does not have any supportive actions, i.e., the sets $S'_t = \{*\}$ are trivial, we say that the mechanism is *simple.*

3.3. **Comments.** We compare the model in this paper to the setting from PS. There, we consider static decision problems (A, u). We study elicitation of affine properties of beliefs, which correspond to the expected values $\mathbb{E}_p X(a, \theta)$ of action-dependent random variables $X : A \times \Theta \longrightarrow \mathbb{R}$. We refer to X as a *question profile* and $X(a, \cdot) \in \mathbb{R}^\Theta$ as a *question.* We say that X is incentivizable if there exists a simple (static) nondistortionary expansion $(A \times \mathbb{R}, v)$ such that, if $(a, r) \in A^v(p)$, then $r = \mathbb{E}_p X(a, \theta)$.

There are two main differences with respect to the current paper. First, the decision problem is dynamic. The key difference is that beliefs may evolve between two periods, either because of the arrival of new information or the impact of actions on the distribution of states. Without evolving beliefs, the two-period decision problem would be equivalent to the static one and the question of incentivizability would reduce to the one studied in PS.

We limit ourselves to two-period problems for simplicity. Similarly, we assume that new information arrives only after the first period, and, in the case of actions affecting the distribution of states, only the first-period (but not the second-period) action affects the distribution. Extensions to more than two periods or additional arrivals of information are straightforward and would generate analogous results.

Second, PS's definition of incentivizability does not allow for any other actions in the expansion apart from the report of the expected value of X. In contrast, we allow for (finitely many) such actions in expansions of dynamic problems. In the second period, such actions are crucial to elicit period 1 questions about dynamic beliefs, including beliefs about learning (see Chambers and Lambert (2021), for example). At the same time, we find that such actions in the first period do not expand the set of incentivizable period 1 questions, and neither period 1 nor period 2 supportive actions play any role in incentivizability in the second period.

As in PS, we assume that belief elicitation concerns a single one-dimensional random variable. This assumption captures the idea that the DM has limits on his cognitive load and capability to process and answer questions. It makes the model cleaner. It is also straightforward to relax.

4. Elicitation in the first period

4.1. **Questions.** In this section, we study elicitations of properties of period 1 beliefs μ. The researcher asks the subject to report an expected value of a random variable. To accommodate a range of questions asked in practice, we allow the value of the random variable to depend on the chosen action. Formally, a *period* 1 *question profile* is defined as a mapping $X : A_1 \times \Delta\Theta \longrightarrow \mathbb{R}$. For technical reasons, we assume that $X(a_1, p)$ is continuous at beliefs p for which the second-period best response is unique, i.e., $p \in \Delta^u$.

We say that X is *incentivizable* if there exists a nondistortionary elicitation mechanism $(A_1 \times \mathbb{R} \times S_1, A_2 \times S_2, v)$ such that, if $(a_1, r, s_1) \in A^v(\mu)$, then $r = \mathbb{E}_{\mu(a_1)} X(a_1, p)$ for each $\mu \in \mathcal{U}^u$.

Consider the following examples.

Example 1. The questions from the example in Section 2.1 are formalized as follows:

(1) "What is your expected test score?" corresponds to

$$X^{(1)}(a_1, p) = \max_\theta p(\theta),$$

which is an action-independent question

(2) "What final grade reduction would you accept for the ability to revise the grading scheme after the test?" corresponds to

$$X^{(2)}(a_1, p) = \max_{a_1' = \frac{1}{3}, \frac{2}{3}} \left(a_1' \max_\theta p(\theta) + (1 - a_1')h\right) - \left(a_1 \max_\theta p(\theta) + (1 - a_1)h\right)$$

(3) "What would be your final grade if, when you choose the high test weight, the marking was changed so that you get payoff 1 if your chosen answer is incorrect and 0 otherwise (cross-out one incorrect)?" corresponds to

$$X^{(3)}(a_1, p) = \begin{cases} a_1 \max_\theta p(\theta) + (1 - a_1)h & a_1 = \frac{1}{3} \\ a_1(1 - \min_\theta p(\theta)) + (1 - a_1)h & a_1 = \frac{2}{3}. \end{cases}$$

Example 2. We formalize the example from Section 2.2. Let $a_1 \in A_1 = \{b, a\}$, $a_2 \in A_2 = \{c, s\}$, and $\theta \in [0, 1]$ be the tumor count. Assume that the survival rate and QALY are equal to

$$s(a_1, a_2, \theta) = 1 - \theta(1 - \kappa(a_2)),$$
$$q(a_1, a_2, \theta) = s(a_1, a_2, \theta) - \rho_1(a_1) - \rho_2(a_2),$$

Here, $\kappa(a_2)$ is the effectiveness of treatment a_2, and ρ_1 and ρ_2 are the costs of the diagnostics and the treatment, respectively.

(1) "What is the expected QALY?" corresponds to

$$X^{(1)}(a_1, p) = \max_{a_2} \mathbb{E}_p q(a_1, a_2, \theta) = \mathbb{E}_p q(a_1, A^q(a_1, p), \theta),$$

(2) "What is the probability that you choose chemotherapy?"

$$X^{(2)}(a_1, p) = \mathbb{1}\{c \in A^u(a_1, p)\}.$$

As the above examples demonstrate, our definition of a question profile is quite flexible and it accommodates a large number of practically applicable questions. Nevertheless, it is important to note that the definition is restrictive. In order to describe the limits of our definition, consider first a more fundamental problem of what properties $X(\mu)$ of probability distributions can be elicited in an incentivized way but ignoring the issue of distortions or action-dependence. This problem has been studied in the literature (see Lambert, Pennock, and Shoham (2008) or Lambert (2019)) and necessary and sufficient conditions for incentivizability have been identified. An example of a necessary condition is that the level sets $X^{-1}(r) = \{\mu : X(\mu) = r\}$ must be convex. [1] The convexity restriction eliminates a number of potentially interesting distribution properties, like questions about variance or higher-moments are not incentivizable.

In addition to convexity, our notion adds some additional restrictions. First, we require the property $X(\mu)$ to be affine in μ, which is a stronger property. In the case of action-dependent information, standard functional analytic results show that our definition captures all continuous and affine properties, as such properties can be represented as expected values of continuous random variables. For the other two cases (actions affecting the quality of information and/or distribution of states), an affine property takes the form $X(\mu) = \sum_a X^a(\mu_a)$, where each $X^a(\mu_a)$ has a representation as an expectation over a continuous variable.

At the same time, it is worth noting that we are not aware of any other potentially interesting or relevant questions that satisfy our necessary conditions for incentivizability but are not captured by our definition.

The following simple observation is useful in characterizing incentivizability. We say that two first-period questions $X(a_1, \cdot)$ and $Y(a_1, \cdot)$ are *equivalent* if there are $\alpha \in \mathbb{R} \setminus \{0\}$ and $\beta \in \mathbb{R}$ such that

$$Y(a_1, p) = \alpha X(a_1, p) + \beta.$$

We say that two question profiles X and Y are equivalent if $X(a_1, \cdot)$ and $Y(a_1, \cdot)$ are equivalent for all first-period actions a_1. If X and Y are equivalent, then for each a_1 the researcher can easily compute the answer to either question from the answer to the other one. The (straightforward) proof of the following result can be found in Appendix B.

Lemma 1. *If question profile X is incentivizable, and Y is an equivalent profile, then it is incentivizable as well.*

4.2. **Sufficient conditions.** In this section, we discuss sufficient conditions for incentivizability.

The next result presents the sufficient conditions for incentivizability. Let $L(\Theta) = \mathbb{R} \times \mathbb{R}^{\Theta}$ be the set of affine functions on $\Delta\Theta$ with the interpretation that, for each $l \in L(\Theta), p \in \Delta\Theta$, we have $l(p) = l_0 + \sum_{\theta} l(\theta)p(\theta)$. For any finite set $L_0 \subseteq L(\Theta)$,

[1] To see why, notice that if r is optimal report at μ and μ', the value of information function must be affine in-between the two profiles; see also Lemma 4 below.

function $f_L(p) = \max_{l \in L_0} l(p)$ is convex as a maximum over affine functions. We refer to such functions as *finitely convex* to emphasize that they are constructed from finitely many affine functions. We say that a function is *finitely continuous* if it can be obtained as a difference between two finitely convex functions. Standard results show that any continuous function can be approximated by a sequence of finitely continuous functions.

Theorem 1. *The following question profiles are incentivizable in the first period:*

(1) $X(a_1, p) = \gamma u_1^*(a_1, p) + f(p)$ *for any finitely continuous* $f(p)$ *and* $\gamma \in \mathbb{R}$ *if information is action-independent, i.e.,* $\mathcal{U} = \mathcal{U}^u_{ind}$*,*
(2) $X(a_1, p) = \gamma u_1^*(a_1, p) + l(p)$ *for any affine* l *and* $\gamma \in \mathbb{R}$ *if actions affect the quality of information, i.e.,* $\mathcal{U} = \mathcal{U}^u_{quality}$*,*
(3) $X(a_1, p) = u_1^*(a_1, p)$ *if actions affect the distribution of states, i.e.,* $\mathcal{U} = \mathcal{U}^u_{general}$*.*

The theorem says that questions about enhanced payoffs are incentivizable. The scope of enhancement depends on the restriction on how the beliefs change with first-period actions. In general, the more restrictive are the beliefs, the more permissive are the sufficient conditions. When the information is independent of actions, any question that is a sum of payoffs and a finitely continuous (action-independent) function of beliefs (or is equivalent to one) is incentivizable.

Because γ can be equal to 0, the theorem shows that certain action-independent questions are incentivizable (given Lemma 1, this is the only role of γ in the statement of the theorem).

The proof constructs appropriate elicitation mechanisms. The construction has two elements. The first element is the standard Becker-DeGroot-Marschak (BDM) mechanism. Assume for simplicity that the value of the question is between 0 and 1 (if not, we can replace the question by a normalized equivalent one). After learning the subject's report, r, of his expectation $\mathbb{E}_{\mu(a_1)} X(a_1, p)$, the researcher draws a number x uniformly from $[0, 1]$. If $r > x$, the subject receives a payoff equal to r, and, if $r < x$, equal to $X(a; p)$. (The payoffs can be easily binarized to avoid problems with risk aversion; see Hossain and Okui (2013)), for example). Notice that we need to explain how to pay $X(a_1, p)$ given that the second-period beliefs are not observed. If $X = u_1^*$, then the payoff is generated as in the original decision problem. If the payoffs are additionally enhanced, the details differ in each case. For example, in the case of action-independent information, suppose that X is enhanced with $f(p) = \max_{l \in L_+} l(p) - \max_{l \in L_-} l(p)$ for some finite sets $L_+, L_- \subseteq L(\theta)$ of affine functions. We construct two parallel decision problems, where the DM chooses $l \in L_+$ and $l \in L_-$, respectively and, in each problem, receives a payoff $l_0 + l(\theta)$ from the realized state of the world θ. We use the choices in the two problems to generate payoffs $X(a_1, p)$. The weight of payoffs in the second problem is sufficiently high not to mess up incentives with the negative component of X. Finally, in any case, the value of enhanced X is maximized by choosing the same first-period action as the one that maximizes the payoffs u_1^* from the original decision problem, as needed for incentivizability. The details can be found in Appendix C.

The construction relies on the supporting action in the second period of the expanded decision problem. The number of those actions, hence the added complexity, depends on the enhancement function. Any question about payoffs, possibly enhanced with an affine function (like in the case of actions affecting the quality of information) can be incentivized with simple elicitation mechanisms without any supporting actions.

Together with Lemma 1, the theorem describes a large class of incentivizable questions. Consider the following example.

Example 3. Question profiles $X^{(1)}$ and $X^{(2)}$ from Example 1 are incentivizable. Indeed, $X^{(1)} = \max_\theta p(\theta)$ is equivalent to $u_1^*(a_1, p) = a_1 \max_\theta p(\theta) + (1 - a_1)h$. For each a_1' and θ', define affine functions $l(a_1', \theta') \in L(\Theta)$ so that $l(p|a_1', \theta') = a_1' p(\theta') + (1 - a_1')h$. Then, $X^{(2)}$ is equivalent to $u_1^*(a_1, p) - \max_{a_1', \theta'} l(p|a_1', \theta')$.

4.3. **Necessary conditions.** Next, we describe some necessary conditions. Say that two essential best responses $a_1, b_1 \in A_1$ are *adjacent* if there is a belief $\mu \in \mathcal{U}$ such that $A^u(\mu) = \{a, b\}$: both actions are optimal in the first-period and there is no other optimal action. As in PS, incentivizability imposes strong restrictions on the value of questions for adjacent actions.

Lemma 2. *Suppose that X is incentivizable in the first-period. Then, for any pair of adjacent first-period actions a_1, b_1, there exist non-zero constants $\alpha^{a_1}, \alpha^{b_1} \neq 0$, as well as $\beta^{a_1}, \beta^{b_1}, \gamma^0$, and a function $f(p)$ such that, for all $p \in \Delta^u$,*

$$\alpha^{a_1} X(a_1, p) - \beta^{a_1} - \gamma^0 u_1^*(a_1, p) = \alpha^{b_1} X(b_1, p) - \beta^{b_1} - \gamma^0 u_1^*(b_1, p) = f(p). \quad (1)$$

Additionally,

- *if actions may affect the quality of information, i.e., $\mathcal{U} = \mathcal{U}^u_{quality}$, then we can assume that f is affine,*
- *if actions may affect the distribution of states, i.e., $\mathcal{U} = \mathcal{U}^u_{general}$, then we can assume that $f \equiv 0$.*

The proof relies on the value of information argument (see Lemma 4 below; see also PS) and a careful exploitation of the affine structure of the set of information structures $\mathcal{U}$. The larger is the space of information structures, the tighter are the conditions. The details can be found in Appendix D.

If the original decision problem has two essential best responses, then the necessary conditions of Lemma 2 are very close to the sufficient conditions of Theorem 1.

Otherwise, if there are more than two actions, there is a gap between the necessary and sufficient conditions. We bridge this gap in the rest of this section. From now on, we study each of the three cases separately.

Action-independent information. In the case of action-independent information, the necessary conditions for incentivizability stated in Lemma 9 are equivalent to the necessary conditions for a static decision problem (A_1, u_1^*) analyzed in PS. PS contains three different sets of conditions for a decision problem under which the sufficient and necessary conditions for incentivizability coincide.

It is worth pointing out two possible differences between PS and the problem considered here. First, PS assumes that the state space is finite, whereas here, the relevant state space $\Delta\Theta$ is infinite. Second, there is a gap between finitely continuous functions in the sufficient conditions of Theorem 1 and arbitrary functions in Lemma 9. We leave the exploration of these two differences for future research.

Instead, we illustrate the limits to incentivizability using an example.

Example 4. Question profile $X^{(3)}$ from Example 1 is not incentivizable if $|\Theta| > 2$. Indeed, Lemma 9 implies the existence of constants $\alpha^1, \alpha^2 \neq 0$ and β^1, β^2, γ such that, for each p

$$\alpha^1 \min_\theta p(\theta) + \gamma \max_\theta p(\theta) + \beta^1 = (\alpha^2 + \gamma) \max_\theta p(\theta) + \beta^2.$$

But that's clearly impossible if there are more than two states.

Actions may affect the quality of information. Next, we assume that actions may affect the quality of information, $\mathcal{U} = \mathcal{U}^u_{\text{quality}}$.

We start with a distinction that plays an important role in this case. Say that the first-period action a_1 has a *nontrivial continuation* if the belief partition $\mathcal{P}(a_1)$ contains at least two elements. In other words, an action has a nontrivial continuation if there are at least two essential best responses following the first-period action.

We first consider decision problems where the second-period behavior is entirely determined by the first-period choice.

Corollary 1. *Suppose that* $\mathcal{U} = \mathcal{U}^u_{quality}$. *Suppose that none of the first-period actions has a nontrivial continuation. Then, if* X *is incentivizable, it is affine:* $X(a_1, \cdot) \in L(\Theta)$ *for each* a_1.

Proof. If a_1 has no nontrivial continuation, $u^*_1(a_1, \cdot)$ is affine. Further, Lemma 9 implies that there exist constants α, β, γ and an affine function l such that

$$X(a_1, p) = \alpha u^*_1(a_1, p) + \beta + l(p).$$

Hence, $X(a_1, \cdot)$ is affine as the sum of affine functions. □

Affine questions can be used to ask about the distribution of states. However, they cannot be used to elicit any dynamic beliefs, including beliefs about learning. The Corollary says that if there are no nontrivial second-period continuations, and if the first-period actions may affect the quality of information, it is not possible to elicit dynamic beliefs without creating distortions.

Because actions do not affect the distribution of the states, in this case, the problem of incentivizability is equivalent to the static problem studied in PS.

Next, we look at the opposite case where either all, or at least sufficiently many first-period actions have a nontrivial continuation. We show that, in such a case, the sufficient conditions of Theorem 1 are also essentially necessary in the sense that all incentivizable questions are equivalent to a question about enhanced payoffs.

Let G^* be the graph, where the nodes are first-period actions with nontrivial continuations, and any pair of two actions is connected if they are adjacent.

Theorem 2. *Suppose that* $\mathcal{U} = \mathcal{U}^u_{quality}$. *Suppose that* G^* *is connected and, either it contains all first-period actions, or, for any action* a_1 *without nontrivial continuation,* a_1 *is adjacent to some action with nontrivial continuation.*

Then, question profile X *is incentivizable if and only if it is equivalent to a question profile*

$$Y(a_1, p) = \gamma u^*_1(a_1, p) + l(p).$$

about expected payoffs enhanced with an affine function $l \in L(\Theta)$ *for some* $\gamma \in \mathbb{R}$.

The proof of the theorem relies on the observation that, if an action a_1 has nontrivial continuation, then the expected payoffs $u^*_1(a_1, \cdot)$ are not affine. In such a case, the characterization from Lemma 9 implies that the question $X(a_1, \cdot)$ must be equivalent to a question about enhanced payoffs. We have to make sure that the coefficients of equivalence are properly preserved along the adjacency graph. The details of the proof can be found in Appendix E.

Finally, we use an example to show that, in general, asking for the probability of a nontrivial continuation is not incentivizable. This is a surprising observation given that, in a static decision problem, asking about the probability of an action is trivially incentivizable (such a question has a zero-one answer).

Example 5. Question profile $X^{(2)}$ from Example 2 is not incentivizable if both treatments are essential best responses. Indeed, on the contrary, assume that both chemotherapy and surveillance are uniquely chosen at some beliefs. If $X^{(2)}$ were incentivizable, then Lemma 9 would imply the existence of constants γ, β and affine l such that

$$\mathbb{1}\{c \in A^u(a_1, p)\} = \gamma u_1^*(a_1, p) + l(p) + \beta.$$

However, while the right-hand side is continuous, the left-hand side is not. The contradiction demonstrates that $X^{(2)}$ is not incentivizable.

Actions may affect the distribution of states. We turn to the last case, when actions may affect the distribution of the states. An immediate corollary of Lemma 9 is that the sufficient conditions of Theorem 1 are essentially necessary, in the sense that any incentivizable question is equivalent to a question about payoffs.

Corollary 2. *Suppose that $\mathcal{U} = \mathcal{U}^u_{general}$. Then, question X is incentivizable if and only if it is equivalent to the question about expected payoffs $u_1^*(a_1, p)$.*

Proof. The result follows from the characterization in Lemma 9. □

5. Elicitation in the second period

In this section, we study elicitation of the properties of second-period beliefs $p \in \Delta\Theta$. The main results are that nondistortionary elicitation is very limited. If information is action-independent, nontrivial elicitation can happen only if payoffs are essentially separable. Otherwise, there are no incentivizable nontrivial questions.

5.1. **Questions.**

As in Section 4, we focus on affine properties of beliefs. This means that we ask the DM to report the expected value of an action-dependent random variable. Formally, a second-period question is a mapping $X : A_1 \times A_2 \times \Theta \longrightarrow \mathbb{R}$. We say that question X is *incentivizable in period* 2 if there is a nondistortionary expansion $(A_1 \times S_1, A_2 \times \mathbb{R} \times S_2, v)$ such that, for each a_1, a_2, s_1, s_2, if $(a_2, r, s_2) \in A_2^v(a_1, s_1, p)$, then $r = \mathbb{E}_p X(a_1, a_2, \theta)$.

Consider the following examples:

Example 6. A second-period question about the expected final grade in the example from Section 2.1 is given by

$$X(a_1, a_2, \theta) = a_1 \mathbb{1}\{a_2 = \theta\} + (1 - a_1)h.$$

As for their first-period analogues, we define equivalent questions:

Say that two second-period question profiles X, Y are *equivalent* if there are mappings $\alpha : A_1 \times A_2 \longrightarrow \mathbb{R} \setminus \{0\}$ and $\beta : A_1 \times A_2 \longrightarrow \mathbb{R}$ such that, for each a_1, a_2,

$$Y(a_1, a_2, \theta) = \alpha(a_1, a_2) X(a_1, a_2, \theta) + \beta(a_1, a_2).$$

If X and Y are equivalent, then the researcher can easily compute the answers to the first one from the answers to the second one. The proof of the next result can be found in Appendix F.

Lemma 3. *If a question profile X is incentivizable and Y is an equivalent profile, then Y is also incentivizable.*

5.2. **Sufficient conditions.** We start with some definitions. Say that the decision problem (A_1, A_2, u) has *separable* payoffs if there are functions $u_t : A_t \times \Theta \longrightarrow \mathbb{R}$ such that

$$u(a_1, a_2, \theta) = u_1(a_1, \theta) + u_2(a_2, \theta).$$

We say that the payoffs are *essentially separable* if for any two essential best responses in the first period $a_1, b_1 \in A_1$, the continuation best responses in the second period are identical: for each p,

$$A^u(a_1, p) = A^u(b_1, p).$$

In such a case, we ignore the first-period actions and write $A_2^u(p)$ to denote the set of second-period best responses.

Finally, we say that payoffs are *weakly separable* if $\mathcal{P}^u(a_1)$ does not depend on the first-period action. Separability implies essential separability which implies weak separability. If a decision problem is weakly separable, one can make it essentially separable by relabeling actions.

The next result presents our strongest sufficient conditions. Say that a function $v : A_2 \longrightarrow \mathbb{R}$ is *weakly aligned* with the second-period payoffs if (A_2, v) is a single-period decision problem such that $A^v(p) = A_2^u(p)$. Note that if v is weakly aligned then so is $v + d$ for any (action-independent) function $d : \Theta \longrightarrow \mathbb{R}$.

Theorem 3. *Suppose that the information is action-independent, i.e.,* $\mathcal{U} = \mathcal{U}_{ind}$*. If the payoffs are essentially separable, and* v *is weakly aligned with the second-period payoffs, then* $X(a_1, a_2, \theta) = v(a_2, \theta)$ *(or any equivalent question) is incentivizable.*

The Theorem shows that a question about the expected value of a function that weakly aligned with the second period payoffs is incentivizable. Observe that all such questions depend only on the second-period actions.

The proof is standard and relies on the BDM. The details can be found in Section G.

If payoffs are essentially separable, then an example of a question that satisfies the assumptions of the theorem is a question about expected payoffs *as if* the DM chooses a fixed action a_1^* in the first period: $X(a_1, a_2, \theta) = u(a_1^*, a_2, \theta)$.

Example 7. Consider the example from Section 2.1. Notice that the optimal answer in the test a_2 does not depend on the weight that the final grade puts on the test. Moreover, the optimal answer maximizes $\mathbb{E}_p X^{(1)}(a_2, \theta) = \text{Prob}_p(a_2 = \theta)$. Hence, the decision problem is essentially separable. As a result, the question $Y(a_1, a_2, \theta) = \mathbb{1}\{a_2 = \theta\}$ is incentivizable.

Because the question X from Example 6 is equivalent to Y, it is also incentivizable.

5.3. **Necessary conditions: case $\mathcal{U} = \mathcal{U}_{\text{ind}}$.** Assume that information is action-independent. We present two sets of necessary conditions for incentivizability. The conditions apply to two different classes of problems: with separable payoffs and the opposite extreme case, when payoffs are generic.

Theorem 4. *Suppose the decision problem has essentially separable payoffs and* X *is incentivizable in the second period. Then,* X *is equivalent to a question profile that depends only on the second-period actions.*

The theorem shows that, with essentially separable payoffs, questions that only depend on the second-period action can be incentivizable. That includes questions about weakly aligned functions, but, in general, other functions as well. That leaves a gap between the sufficient conditions from Theorem 3 and the above necessary conditions.

The proof of the theorem only uses the first-period restrictions on incentivizability. Additional information can be gained from looking at the second-period restrictions. In the second period, the necessary conditions for incentivizability have been analyzed in PS. That paper contains three different sets of conditions for the second-period decision problem under which the necessary and sufficient conditions coincide.

For each first-period action a_1, let

$$D(a_1) = \{u(a_1, a_2, \cdot) - u(a_1, b_2, \cdot) : a_2, b_2 \in A_2 \text{ st. } a_2 \neq b_2\} \subseteq \mathbb{R}^\Theta$$

be a finite set of payoff difference vectors after choosing the first-period action a_1. Say that payoffs are *generic* if sets $D(a_1)$ and $D(a_2)$ are disjoint for any pair of essential best responses a_1, b_1. It is easy to see that the set of so-defined generic payoffs is an open and dense subset of $\mathbb{R}^{A_1 \times A_2 \times \Theta}$, which justifies the name. On the other hand, if payoffs are separable, sets $D(a_1)$ are identical for all first-period actions, hence payoffs are very much *not* generic.

Theorem 5. *For generic payoffs u, if X is incentivizable in the second-period, then X is equivalent to an action-independent question.*

For all such decision problems, a question is incentivizable in the second period if and only if it is action-independent. Because such questions are trivially incentivizable in the second-period, for such problems the necessary and sufficient conditions coincide.

5.4. **Proof intuition for Theorems 4 and 5.** We explain the logic of the proof of the theorem. The details can be found in Appendix H.2.

We start with a reminder of two basic facts about the value of information function $u_1^*(a_1, \cdot)$. This function is convex. Moreover, it is strictly convex over the interval of beliefs $\alpha p + (1-\alpha)q$ if and only if there is no single common best response at p and q, i.e., if and only if $A^u(a_1, p) \cap A^u(a_1, q) = \emptyset$ (see Lemma 4 in the Appendix).

Consider the following example to show how the above observation imposes restrictions on incentivizable questions.

Example 8. There are two states $\Theta = \{0, 1\}$. In period 1, the agent chooses between save and risky options $A_1 = \{s, r\}$. In the second period, the agent chooses between actions in $A_2 = \{0, 1\}$. If the risky option is chosen, then the agent gets payoff 1 if the second-period action matches the state, and -2 otherwise. If the safe option is chosen in the first period, the agent receives payoff 0 if action 0 is taken in the second period and payoff -1 otherwise. In particular, there is a single optimal decision after action s, and the expected payoffs are constant and equal to $u_1^*(s, p) = 0$ for each p. The expected payoffs after risky action

$$u_1^*(r, p) = \max\{1 - 3p, 3p - 2\} \tag{2}$$

are convex at $p = \frac{1}{2}$ and affine otherwise.

Suppose that $X(a_1, a_2, \theta) = u(a_1, a_2, \theta)$ is a question about payoffs. We show that this question is not incentivizable in period 2.

For simplicity, we only consider simple nondistortionary elicitation mechanisms $(A_1, A_2 \times \mathbb{R}, v)$. On the contrary, suppose that v incentivizes X. Let $v_1^*(a_1, p) = \max_{a_2, r} \mathbb{E}_p v(a_1, a_2, r, \theta)$ be the value of information of the second-period decision problem.

Because there is a single optimal decision $(a_2, r) = (0, 0)$ after action s, it must be that $v_1^*(s, p)$ is affine in p. At the same time, because the expected payoffs vary with p (see equation (2)), there is a different optimal decision at each p. Hence, $v_1^*(r, p)$ is strictly convex at each belief p. In particular, $v_1^*(r, p) - v_1^*(s, p)$ is strictly convex at each p.

As the expansion is nondistortionary, the first-period decision problems with payoffs $u_1^*(a_1, p)$ and $v_1^*(a_1, p)$ must have the same optimal actions. This implies that the expected payoff differences between save and risky actions must be collinear (see Lemma 5 in the Appendix). But, in the original decision problem, the difference payoff vector

$$u_1^*(r, p) - u_1^*(s, p) = u_1^*(r, p) = \max\{1 - 3p, 3p - 2\}$$

is affine in $p \in (0, \frac{1}{2})$. We obtain a contradiction with strict convexity of $v_1^*(r, p) - v_1^*(s, p)$ at each p.

The contradiction shows that the question about payoffs cannot be incentivized in a nondistortionary way.

The key feature of the above example is that there are pairs of beliefs such that the value of the question (i.e., the payoffs, in the example) for the optimal second-period action is the same at the two beliefs after one first-period action, but not after some other adjacent first-period action. If so, a similar argument shows that any attempt to incentivize second-period elicitation will distort first-period incentives. Using a simple linear algebra argument, we further show that if the question profile X is incentivizable, a_1 and b_1 are two adjacent actions, and a_2 and b_2 are respective continuation best responses at the same nonempty set of beliefs (i.e., $P(a_1, a_2) \cap P(b_1, b_2) \neq \emptyset$), then the questions $X(a_1, a_2)$ and $X(b_1, b_2)$ must be equivalent.

We now further illustrate how this argument works. Consider Figure 1, which shows the space of beliefs $\Delta\Theta$ divided into sets of beliefs where actions a_2 or b_2 are continuation best responses after action a_1 (red, higher rectangles) and after action

b_1 (blue, lower rectangles). The belief sets overlap, as in the generic payoff case of Theorem 5.

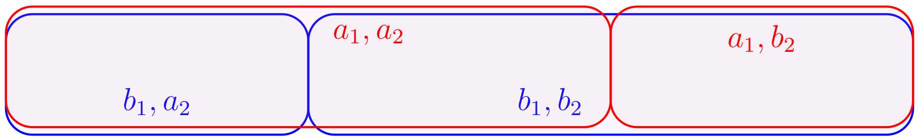

FIGURE 1. Belief sets $P(x_1, x_2)$ for $x_1 \in \{a_1, b_1\}$ and $x_2 \in \{a_2, b_2\}$.

If the question profile X is incentivizable, then the overlap of the belief sets in the center of the Figure implies that questions $X(a_1, a_2)$ and $X(b_1, b_2)$ must be equivalent. Further, the same argument implies that pairs of questions $X(a_1, b_2)$ and $X(b_1, b_2)$ as well as $X(a_1, a_2)$ and $X(b_1, a_2)$ are equivalent as well. Because "equivalence" is an equivalence relation, it must be that all the questions are equivalent to, say, $X(b_1, b_2)$. In other words, in this case, X is equivalent to an action-independent question.

More generally, consider the belief partitions $\mathcal{P}(a_1)$ and define $\mathcal{P}$ as the *join* of the belief partitions $\mathcal{P}(a_1)$ across all first-period essential best responses: $\mathcal{P}$ is the finest collection of closed convex sets with nonempty interior such that, for each essential a_1, each $P \in \mathcal{P}(a_1)$, P is contained in one of the elements of $\mathcal{P}$. For each a_1, each essential a_2, let $\pi(a_1, a_2) \in P$ be the uniquely defined element of collection P that contains $P(a_1, a_2)$. For example, in the case illustrated on Figure 1, the join is trivial: it consists of a single element encompassing all beliefs.

We show in the Appendix that any incentivizable question profile X is equivalent to a question that only depends on the elements of the join of the partition. In the generic case, we show that the join is trivial. As a conclusion, we obtain the thesis of Theorem 5.

On the other hand, if payoffs are essentially separable, the join is equal to the collection $\mathcal{P}(a_1)$ for any essential best response. As a result, we obtain the thesis of Theorem 4: any incentivizable question is equivalent to a question that depends only on the elements of the join, hence, only on the second-period actions.

5.5. **Necessary conditions: other cases.** Here, we show that if actions either affect the quality of information, $\mathcal{U} = \mathcal{U}_{\text{quality}}$, or the distribution of states, $\mathcal{U} = \mathcal{U}_{\text{general}}$, then no nontrivial question is incentivizable. More precisely, we will show that there is no nondistortionary expansion in which nontrivial second-period questions are elicited. We say that a question is nontrivial if it is not equivalent to a constant. More precisely, say that $X : A_1 \times A_2 \times \Theta \longrightarrow \mathbb{R}$ is a nontrivial question if there are essential best responses a_1 and a_2 (after a_1) such that the question is not constant in the state of the world: there are states $\theta \neq \theta'$ such that $X(a_1, a_2, \theta) \neq X(a_1, a_2, \theta')$.

Theorem 6. *Suppose $\mathcal{U}_{quality} \subseteq \mathcal{U}$. There is no non nontrivial question profile that is incentivizable in the second period.*

The theorem says that we cannot incentivize any question, whether it is action-dependent or not. The proof is in Appendix I.

To get the intuition for the result, notice first that, due to the convexity of the value of information function, any incentivization must increase the value of information in the first period relative to the original problem. If actions affect the quality of information, this increase distorts the incentives toward more informative actions.

Appendix A. Static decision problems

In the following discussion, we rely on the following two facts about static decision problems: Let A, u where $u : A \times \Theta$ be an arbitrary (static) decision problem. Let $A^u(p) = \arg\max_a \mathbb{E}_p u(a, \theta)$ be the optimal set and let $V(p) = \max_a \mathbb{E}u(a, \theta)$ be the value of information function. Then, the value of information is affine between beliefs for which there is a common optimal action and strictly convex otherwise.

Lemma 4. *For any two beliefs p, q,*

$$A^u(p) \cap A^u(q) = \emptyset \Longrightarrow \forall_{\alpha \in (0,1)} V(\alpha p + (1-\alpha)q) < \alpha V(p) + (1-\alpha)V(q),$$
$$A^u(p) \cap A^u(q) \neq \emptyset \Longrightarrow \forall_{\alpha \in (0,1)} V(\alpha p + (1-\alpha)q) = \alpha V(p) + (1-\alpha)V(q).$$

Say that actions a_1, b_1 are *adjacent* if there is an interior belief $\mu \in \Delta(\Delta\Theta)$ such that the two actions are uniquely optimal: $A_1^u(\mu) = \{a_1, b_1\}$.

For the second fact, suppose that (A, v) is a decision problem with the same optimal action sets $A^u(.) = A^v(.)$ (i.e., a trivial nondistortionary expansion of u. Then, at least for adjacent actions (i.e., actions such that the two actions are uniquely optimal for some belief), the payoff difference vectors must be collinear:

Lemma 5. *Suppose that decision problems u and v have the same optimal action sets $A^u(.) = A^v(.)$. If a, b are adjacent in decision problem u, then, there exists $c > 0$ st.*

$$u(a, \theta) - u(b, \theta) = c\,(v(a, \theta) - v(b, \theta)) \text{ for each } \theta.$$

The two facts are straightforward and well-known (see also PS).

Appendix B. Proof of Lemma 1

Suppose Y is equivalent to X and $\alpha(a_1), \beta(a_1)$ are the parameters from the definition of equivalence. Suppose that $(A_1 \times \mathbb{R} \times S_1, A_2 \times S_2, v_X)$ is the elicitation mechanism incentivizing X. Let

$$v_Y(a_1, r, s_1, a_2, s_2) = v_X(a_1, \frac{r - \beta(a_1)}{\alpha(a_1)}, s_1, a_2, s_2).$$

Then, $(A_1 \times \mathbb{R} \times S_1, A_2 \times S_2, v_Y)$ incentivizes Y. Indeed, the construction ensures that

- $(a_2, s_2) \in A^{v_Y}(a_1, r, s_1, p)$ is v_Y-optimal in the second period after actions a_1, r, s_1 at beliefs p if and only if $(a_2, s_2) \in A^{v_X}(a_1, \frac{r-\beta(a_1)}{\alpha(a_1)}, s_1, p)$ is v_X-optimal after $(a_1, \frac{r-\beta(a_1)}{\alpha(a_1)}, s_1, p)$.
- The expected continuation payoffs after the first-period in v_Y problem after actions (a_1, r, s_1) are identical to the payoffs in v_X problem after $(a_1, \frac{r-\beta(a_1)}{\alpha(a_1)}, s_1)$, the incentives in the first-period are unchanged. Together with the fact that X is incentivized in the first-period, the latter implies that $r = \alpha(a_1)\mathbb{E}_{\mu(a_1)}X(a_1, \cdot) + \beta(a_1)$ is v_Y-optimal.

Appendix C. Proof of Theorem 1

We start with the case of action-independent information. Suppose that $X(a_1, p) = \gamma u_1^*(a_1, p) + f(p)$, where f is a finitely continuous function and $\gamma \in \mathbb{R}$. Because of equivalence, we can assume hat $\gamma \geq 0$. By the definition of finitely continuous functions, there exist finite sets $L_+, L_- \subseteq L(\Delta\Theta)$ such that

$$f(p) = \max_{l \in L_+} l(p) - \max_{l \in L_-} l(p).$$

Assume that $x_{\min} < X(a_1, p) < x_{\max}$ for each a_1, p.

Construct an elicitation mechanism $(A_1 \times \mathbb{R}, A_2 \times L_+ \times L_-, v)$ with payoffs:

$$\begin{aligned} & v(a_1, r, a_2, l_+, l_-, \theta) \\ =& u(a_1, a_2, \theta) + \frac{r - x_{\min}}{x_{\max} - x_{\min}} \left(u(a_1, a_2, \theta) + l_+(\theta) - l_-(\theta)\right) + \frac{1}{2}\frac{x_{\max}^2 - r^2}{x_{\max} - x_{\min}} + 2l_-(\theta). \end{aligned} \tag{3}$$

Given beliefs p, the first-period actions a_1, r, and the optimal choice of second-period actions leads to expected payoffs equal to

$$u_1^*(a_1, p) + \frac{r - x_{\min}}{x_{\max} - x_{\min}} \mathbb{E}_{\mu(a_1)} X(a_1, p) + \frac{1}{2}\frac{x_{\max}^2 - r^2}{x_{\max} - x_{\min}} + 2\mathbb{E}_p \max_{l \in L_-} l(p).$$

(Notice that the incentives to maximize the $l(p)$ for $l \in L_-$ are due to the fact that the coefficient on such ls is strictly positive.) The optimal choice is $r_{\text{opt}} = \mathbb{E}_{\mu(a_1)} X(a_1, p)$. This leads to expected payoffs given action a_1:

$$u_1^*(a_1, p) + \frac{1}{x_{\max} - x_{\min}} (\mathbb{E}_{\mu(a_1)} X(a_1, p) - x_{\min})^2 + \frac{1}{2}(x_{\max} + x_{\min}) + 2\mathbb{E}_p \max_{l \in L_-} l(p). \tag{4}$$

The expected payoff is increasing in

$$\mathbb{E}_\mu X(a_1, p) = \gamma \mathbb{E}_\mu u_1^*(a_1, p) + \mathbb{E}_\mu f(p).$$

Because the last component of (4) does not depend on a_1, the expected payoffs are maximized by $a_1 \in A^u(\mu)$. This concludes the proof in the case of action-independent information.

Suppose that actions affect the quality of information and assume that $X(a_1, p) = u_1^*(a_1, p) + l(p)$ for some affine l. Notice that $\mathbb{E}_{\mu(a_1)} l(p)$ does not depend on the action for any a_1. In such a case, the same elicitation mechanism, but with l_+ replaced by l in formula (3), and no supporting actions is the appropriate nondistortionary expansion. Finally, the same mechanism, but with $l \equiv 0$ works for the general case.

Appendix D. Proof of Lemma 2

Suppose that question profile X is incentivizable in the first period and actions a, b are adjacent.

D.1. Basic observation.

We start with a straightforward consequence of Lemma 4:

Lemma 6. *For any $\mu \in \mathcal{U}^u$ such that $A^u(\mu) = \{a, b\}$, any $\psi \in \mathcal{U}$, if $\psi_a \cdot u(a, \cdot) = \psi_b \cdot u(b, \cdot)$ and $\mu_a \cdot X(a, \cdot) = \psi_a \cdot X(a, \cdot)$, then $\mu_b \cdot X(b, \cdot) = \psi_b \cdot X(b, \cdot)$.*

Proof. Suppose that $(A_1 \times \mathbb{R} \times S_1, A_2 \times S_2, v)$ is an elicitation mechanism that incentivizes X. For each α, define $\mu_\alpha = \alpha\mu + (1 - \alpha)\psi$.

Because $A^u(\mu) = \{a, b\}$, there is $0 < \alpha_0 < \alpha_1 < 1$ and $s_1 \in S_1$ such that, for each $\alpha \in [\alpha_0, \alpha_1]$, we have

$$A^u(\mu_\alpha) = \{a, b\} \text{ and } (a, \mu \cdot X(a, \cdot), s_1) \in A^v(\mu_\alpha).$$

Indeed, the first claim follows from the upper hemi-continuity of the best response correspondence. The second claim is a consequence of the first, the assumption that $\mu_a \cdot X(a, \cdot) = \psi_a \cdot X(a, \cdot)$, and the finiteness of S_1.

Together with the definition of elicitation mechanism, Lemma 4 implies that there must be some r and $t_1 \in S_1$ such that $(b, r, t_1) \in A^v(\mu_\alpha)$. Because of the definition of incentivizability, it must be that $r = \mu_\alpha(b) \cdot X(b, \cdot)$ for each $\alpha \in [\alpha_0, \alpha_1]$, which implies $\mu_b \cdot X(b, \cdot) = \psi_b \cdot X(b, \cdot)$. □

In what follows, we say that function $f : \Delta^u \longrightarrow \mathbb{R}$ is an affine transformation of $g : \Delta^u \longrightarrow \mathbb{R}$, if there are constants $\rho, \delta \in \mathbb{R}$ such that $f(p) = \rho g(p) + \delta$.

Define function

$$D(p) = u_1^*(a, p) - u_1^*(b, p).$$

D.2. **Case $\mathcal{U}_{\mathbf{ind}}$.** We start with preliminary result:

Lemma 7. *Fix finite subset $P \subseteq \Delta\Theta$. Then, there exist P-dependent constant $\alpha, \gamma, \beta \in \mathbb{R}$ such that, for each $p \in P$,*

$$\alpha X(a, p) + X(b, p) = \gamma D(p) + \beta.$$

Proof. For each finite subset $P \subseteq \Delta\Theta$, define $\mathcal{U}_P = \{\mu \in \mathcal{U} : \mu(P) = 1 \text{ for each } a \in A_1\}$. Let $\mu \in \mathcal{U}_P$ be an interior information structure such that $A^u(\mu) = \{a, b\}$. Let $\mathcal{W}$ be the (convex) set of information structures ψ such that

$$\psi \cdot u(a, \cdot) = \psi \cdot u(b, \cdot) \text{ and } \mu \cdot X(a, \cdot) = \psi \cdot X(a, \cdot).$$

Lemma 6 implies that

$$\mu \cdot X(b, \cdot) = \psi \cdot X(b, \cdot) \text{ for any } \psi \in \mathcal{W}.$$

Let $F = \{D, \mathbb{1}\}$ and consider a (dual) linear subspace

$$W = \left\{w \in \left(\mathbb{R}^P : \forall_{f \in F \cup \{X(a, \cdot)\}} w \cdot f = 0\right\}.$$

Then, W is the affine hull of the set $\{\psi - \mu : \psi \in \mathcal{W}\}$. The above observation implies that, for each $w \in W$, it must be that $w \cdot X(b, \cdot) = 0$. A standard linear algebra argument implies that $X(b, \cdot) \in \text{span}(F \cup \{X(a, \cdot)\})$. The thesis of the Lemma follows.

□

Note that the constants in equations (5) are typically not uniquely defined and may depend on the choice of set P.

The rest of the proof proceeds in smaller steps. Define function

$$D(p) = u_1^*(a, p) - u_1^*(b, p).$$

Lemma 8. *$X(a,\cdot)$ is an affine transformation of D iff $X(b,\cdot)$ is an affine transformation of D.*

Proof. Suppose $x(a,\cdot)$, but not $X(b,\cdot)$ is an affine transformation of D. Then, there is $P\Delta^u$ such that $X(b,\cdot)$ cannot be represented by affine transformation of D on P. But this leads to a contradiction with Lemma 7. The other direction follows by switching the roles of a and b. □

The Lemma shows that the following cases are exhaustive:

(1) $X(a,\cdot)$ and $X(b,\cdot)$ are not affine transformations of D: For sufficiently large P, neither $X(a,\cdot)$ nor $X(b,\cdot)$ can be represented as a affine transformation of D on P. The hypothesis implies that, if α is the constant from Lemma 7, then $\alpha \neq 0$.

We will show that α is unique. Otherwise, there would be α',γ',β' such that

$$\alpha' X(a,p) + X(b,p) = \gamma' D(p) + \beta'.$$

But then, a subtraction of the first equation from the second leads to a contradiction with $X(a,\cdot)$ not having a representation as affine transformation of D.

Additionally, because $D(p)$ is not constant, γ and β are also uniquely determined for all sufficiently large P. Hence, for all $p \in \Delta^u$

$$\alpha X(a,p) - \beta^a - \gamma u(a,p) = -X(b,p) - \gamma u(b,p),$$

which implies the thesis of Lemma 2.

(2) For each $x = a, b$, $X(x,\cdot)$ is affine transformations of D with coefficients ρ^x, δ^x. The Lemma 2 follows from

$$-X(a,\cdot) + \delta^a + (\rho^b + \rho^a)u(a,\cdot) = \rho^b u(a,\cdot) + \rho^a u(b,\cdot) = X(b,\cdot) - \delta^b + (\rho^b + \rho^a)u(b,\cdot).$$

D.3. **Case $\mathcal{U}_{\mathbf{quality}}$.** We start with preliminary result:

Lemma 9. *Fix finite subset $P \subseteq \Delta\Theta$. Then, there exist P-dependent constant $\alpha, \gamma, \beta^a \in \mathbb{R}$ and affine function $l \in L(\Theta)$ such that, for each $p \in P$,*

$$\alpha X(a,p) = \beta^a + \gamma u(a,p) + l(p), \tag{5}$$
$$X(b,p) = -\gamma u(b,p) - l(p).$$

Proof. For each finite subset $P \subseteq \Delta\Theta$, define $\mathcal{U}_P = \{\mu \in \mathcal{U} : \mu_a(P) = 1 \text{ for each } a \in A_1\}$. Let $\mu \in \mathcal{U}_P$ be an interior information structure such that $A^u(\mu) = \{a, b\}$. Let

$\mathcal{W}$ be the (convex) set of information structures ψ such that

$$\psi_a \cdot u(a,\cdot) = \psi_b \cdot u(b,\cdot) \text{ and } \mu_a \cdot X(a,\cdot) = \psi_a \cdot X(a,\cdot).$$

Lemma 6 implies that

$$\mu_b \cdot X(b,\cdot) = \psi_b \cdot X(b,\cdot) \text{ for any } \psi \in \mathcal{W}.$$

For all actions $x, y \in A$, let

$$\begin{aligned} g(x,p) &= \mathbb{1}\{x=a\}u(a,p) - \mathbb{1}\{x=b\}u(b,p), \\ g^{y,\theta}(x,p) &= p(\theta)(\mathbb{1}\{x=y\} - \mathbb{1}\{x=a\}) \text{ for all } \theta \in \Theta \text{ and } y \in A_1 \setminus \{a\}, \\ h^y(x,p) &= \mathbb{1}\{x=y\}, \\ k^y(x,p) &= \mathbb{1}\{x=y\}X(y,p). \end{aligned}$$

Define $F = \{g\} \cup \{g^{y,\theta} : y \neq a, \theta \in \Theta\} \cup \{h^y : y \in A\}$ and consider a (dual) linear subspace

$$W = \left\{ w \in \left(\mathbb{R}^P\right)^A : \forall_{f \in F \cup \{k^a\}} \sum_x w_x \cdot f_x = 0 \right\}.$$

Then, W is the affine hull of the set $\{\psi - \mu : \psi \in \mathcal{W}\}$.

The above observation implies that, for each $w \in W$, it must be that $\sum_x w_x \cdot k^b_x = 0$.

A standard linear algebra argument implies that $k^b \in \operatorname{span}(F \cup \{k^a\})$: there exist constants γ_b,γ_b^a, β_b^y, for each $y \in A_1$, and $\gamma_b^{y,\theta}$ for θ and $y \neq a$ such that

$$k^b = \sum_y \beta_b^y h^y + \gamma_b g + \sum_{\theta \text{ and } y \neq a} \gamma_b^{y,\theta} g^{y,\theta} + \gamma_b^a k^a.$$

We rewrite it as

$$\begin{aligned} 0 &= \beta_b^a + \gamma_b u(a,p) - \sum_\theta \gamma_b^{\theta,b} p(\theta) - \sum_{\theta \text{ and } y \neq a,b} \gamma_b^{\theta,y} p(\theta) + \gamma_b^a X(a,p), \\ X(b,p) &= \beta_b^b - \gamma_b u(b,p) + \sum_\theta \gamma_b^{\theta,b} p(\theta), \\ 0 &= \beta_b^y + \sum_\theta \gamma_b^{\theta,y} p(\theta) \text{ for each } y \neq a,b. \end{aligned}$$

Let $l(p) = \beta_b^y - \sum_\theta \gamma_b^{\theta,b} p(\theta)$ and rename other constants appropriately. The lemma follows.

□

Note that the constants in equations (5) are typically not uniquely defined and may depend on the choice of set P.

The rest of the proof proceeds in smaller steps.

Lemma 10. *$X(a,\cdot)$ is an affine transformation of D iff $X(b,\cdot)$ is an affine transformation of D.*

Proof. Suppose $x(a,\cdot)$, but not $X(b,\cdot)$ is an affine transformation of D. Then, there is $P\Delta^u$ such that $X(b,\cdot)$ cannot be represented by affine transformation of D on P. At the same time, the summation of equations (5) leads to

$$\alpha X(a,p)+X(b,p)=\gamma D(p)+\beta^a.$$

The contradiction shows that $X(b,\cdot)$ is an affine transformation of D.

The other direction follows by switching the roles of a and b. □

Lemma 11. *Suppose that, for each $x=a,b$, $X(x,\cdot)=\rho^x D(p)+\delta^x$ for some coefficients ρ^x,δ^x. Then, either $\rho^a u(b,\cdot)$ and $\rho^b u(a,\cdot)$ are both affine, or both of them are not affine.*

Proof. Suppose that $\rho^b u(a,\cdot)$ is not affine. Then $u(a,\cdot)$ must be not affine and there is $P\subseteq\Delta^u$ sufficiently large so that $u(a,\cdot)$ is not affine on P. Also, we assume that P is sufficiently large that D is not constant on P (D cannot be constant, otherwise a,b are not adjacent).

Summing equations (5) and using the affine representations for the questions yields

$$(\alpha\rho^a+\rho^b-\gamma)D=\beta^a-\alpha\delta^a-\delta^b$$

Because D is not a constant, we have $\gamma=\alpha\rho^a+\rho^b$.

Using the second equation (5) and substituting for γ, we further obtain

$$\rho^b u(a,\cdot)+\alpha\rho^a u(a,\cdot)=-\delta^b-l.$$

The right-hand side is affine. Hence, $\rho^a u(b,\cdot)$ cannot be affine, otherwise we would get the contradiction with the left-hand side being a sum of affine and not affine functions.

The other direction follows from flipping the roles of the two actions. □

The Lemmas shows that the following cases are exhaustive:

(1) $X(a,\cdot)$ and $X(b,\cdot)$ are not affine transformations of D: For sufficiently large P, neither $X(a,\cdot)$ nor $X(b,\cdot)$ can be represented as a affine transformation of D on P. The summation of equations (5) leads to

$$\alpha X(a,p)+X(b,p)=\gamma D(p)+\beta^a.$$

The hypothesis implies that $\alpha \neq 0$.

We will show that α is unique. Otherwise, there would be α', γ', β' such that

$$\alpha' X(a,p) + X(b,p) = \gamma' D(p) + \beta'.$$

But then, a subtraction of the first equation from the second leads to a contradiction with $X(a,\cdot)$ not having a representation as affine transformation of D.

Additionally, because $D(p)$ is not constant, γ and β are also uniquely determined for all sufficiently large P. Hence, for all $p \in \Delta^u$

$$\alpha X(a,p) - \beta^a - \gamma u(a,p) = -X(b,p) - \gamma u(b,p) = l(p),$$

which implies the thesis of Lemma 2.

(2) For each $x = a, b$, $X(x,\cdot)$ is affine transformations of D with coefficients ρ^x, δ^x. There are two subcases:

(a) Neither $\rho^a u(b,\cdot)$ and $\rho^b u(a,\cdot)$ is affine: The proof of Lemma 11 shows that

$$f(p) = \rho^b u(a,p) + \alpha \rho^a u(b,p)$$

must affine, which implies that $\alpha \neq 0$ and it is uniquely determined. The thesis of Lemma 2 follows from the following equations:

$$-\alpha X(a,\cdot) + \alpha \delta^a + (\rho^b + \alpha \rho^a) u(a,\cdot) = f(p) = X(b,\cdot) - \delta^b + (\rho^b + \alpha \rho^a) u(b,\cdot).$$

(b) $\rho^a u(b,\cdot)$ and $\rho^b u(a,\cdot)$ are both affine: Then,

$$-X(a,\cdot) + \delta^a + (\rho^b + \rho^a) u(a,\cdot) = \rho^b u(a,\cdot) + \rho^a u(b,\cdot) = X(b,\cdot) - \delta^b + (\rho^b + \rho^a) u(b,\cdot).$$

which shows that the thesis of Lemma 2 holds.

D.4. **Case $\mathcal{U}_{\mathbf{general}}$.** We have a preliminary result:

Lemma 12. *Fix finite subset $P \subseteq \Delta\Theta$. Then, there exist P-dependent constant $\alpha, \gamma, \beta^a, \beta^b \in \mathbb{R}$ such that, for each $p \in P$,*

$$\begin{aligned} \alpha X(a,p) &= \beta^a + \gamma u(a,p), \\ X(b,p) &= \beta^b - \gamma u(b,p). \end{aligned} \tag{6}$$

Proof. For each finite subset $P \subseteq \Delta\Theta$, define $\mathcal{U}_P = \{\mu \in \mathcal{U} : \mu_a(P) = 1 \text{ for each } a \in A_1\}$. Let $\mu \in \mathcal{U}_P$ be an interior information structure such that $A^u(\mu) = \{a, b\}$. Let

$\mathcal{W}$ be the (convex) set of information structures ψ such that

$$\psi_a \cdot u(a,\cdot) = \psi_b \cdot u(b,\cdot) \text{ and } \mu_a \cdot X(a,\cdot) = \psi_a \cdot X(a,\cdot).$$

Lemma 6 implies that

$$\mu_b \cdot X(b,\cdot) = \psi_b \cdot X(b,\cdot) \text{ for any } \psi \in \mathcal{W}.$$

For all actions $x, y \in A$, let

$$\begin{aligned} g(x,p) &= \mathbb{1}\{x=a\}u(a,p) - \mathbb{1}\{x=b\}u(b,p), \\ h^y(x,p) &= \mathbb{1}\{x=y\}, \\ k^y(x,p) &= \mathbb{1}\{x=y\}X(y,p). \end{aligned}$$

Define $F = \{g\} \cup \cup\{h^y : y \in A\}$ and consider a (dual) linear subspace

$$W = \left\{ w \in \left(\mathbb{R}^P\right)^A : \forall_{f \in F \cup \{k^a\}} \sum_x w_x \cdot f_x = 0 \right\}.$$

Then, W is the affine hull of the set $\{\psi - \mu : \psi \in \mathcal{W}\}$.

The above observation implies that, for each $w \in W$, it must be that $\sum_x w_x \cdot k_x^b = 0$.

A standard linear algebra argument implies that $k^b \in \text{span}(F \cup \{k^a\})$: there exist constants $\gamma, \alpha,$, and β^y, for each $y \in A_1$, such that

$$k^b = \sum_y \beta^y h^y + \gamma g - \alpha k^a.$$

This implies

$$\begin{aligned} 0 &= \beta^a + \gamma u(a,p) - \alpha X(a,p), \\ X(b,p) &= \beta^b - \gamma u(b,p). \end{aligned}$$

The result follows. □

The second equation of (7) implies that, if $u(x,\cdot)$ is a constant function, so is $X(b,\cdot)$. Because the roles of a, b can be reversed, the same observation applies to $u(a,\cdot)$ and $X(a,\cdot)$.

Because actions a, b are adjacent, at least one of $u(a,\cdot)$ or $u(b,\cdot)$ is not a constant function. To fix attention, assume w.l.o.g. that $u(b,\cdot)$ is not a constant. The second equation of (7) implies that, for sufficiently large P, γ, and hence β^b are uniquely defined and do not depend on P.

If $X(a,\cdot)$ and $X(b,\cdot)$ are constant functions, the thesis of Lemma 2 is trivially satisfied. For the rest of the discussion, we assume that at least one of the functions is not constant. The equationes (6) imply that $\gamma \neq 0$.

If $u(a,\cdot)$ is not a constant function, then, it must be that $\alpha \neq 0$ is uniquely determined and fixed for all sufficiently large P. The thesis of Lemma 2 immediately follows from equations (6).

If $u(a,\cdot)$ is a constant function, then we can replace β^a from the first equation (6) with $\beta^{a\prime} = X(a,\cdot) - \gamma u(a,\cdot)$. Then, $X(a,\cdot) = \beta^{a\prime} + \gamma u(a,\cdot)$. The thesis of Lemma 2 is trivially satisfied in this case.

Appendix E. Proof of Theorem 2

Suppose that a_1 has nontrivial continuation. Then, $u_1^*(a_1,p)$ is not affine in p. We consider two cases:

First, suppose that $X(a_1,p)$ is affine in p. In such a case, the characterization in Lemma 9 implies that the constant γ_0 in equation (1) must be equal to $\gamma^0 = 0$ and that $X(b_1,p)$ is affine in p for each adjacent b_1. An analogous argument extends the claim to the entire graph G^*, and through equations (1), to the entire set of first-period actions, including actions without nontrivial continuation. For each a_1, define $X^0(a_1) \in \mathbb{R}^\Theta$ as the unique vector such that $X(a_1,p) = p \cdot X^0(a_1)$.

For each edge $a_1 - b_1$ in graph G^*, equations (1) further imply that $X^0(a_1)$ and $X^0(b_1)$ belong to $\text{span}\{\mathbb{1}, f_{a_1b_1}\}$ for some $f_{a_1b_1} \in \mathbb{R}^\Theta$. If $b_1 - c_1$ is another edge in G^*, then it must be that $f_{b_1c_1} \in \text{span}\{\mathbb{1}, f_{a_1b_1}\}$. A repetition of the same argument from the other side, shows that there exists f such that, for each edge e, $f_e \in \text{span}\{\mathbb{1}, f\}$. It follows from equations (1) that it must be that f_e is collinear with f for each edge e of graph G^*. In particular, each question $X(a_1)$ is equivalent to affine question $p \cdot f$. One more application of the equations extends this observation to actions without nontrivial continuation.

Second, suppose that $X(a_1,p)$ is not affine in p. Then, we can assume that $\gamma^0 = 1$ in equation (1). Because of the characterization from Lemma 9, if b_1 is adjacent and it has nontrivial continuation, it must be that $X(b_1,p)$ is also not affine in p. Hence, this argument spreads throughout the entire graph G^*.

For each action a_1 with nontrivial continuation, because $X(a_1)$ is non-affine and $\gamma_0 = 1$, there are unique α, β such that

$$\alpha X(a_1,p) - \beta - u_1^*(a_1,p)$$

is affine in p. By the characterization from Lemma 9, the value of the above formula does not depend on a_1. Let $f \in \mathbb{R}^\Theta$ be the unique vector such that $\alpha X(a_1, p) - \beta - u_1^*(a_1, p) = p \cdot f$. It follows that each question $X(a_1)$ is equivalent to $u_1^*(a_1, p) + p \cdot f$ for actions with nontrivial continuation. One more application of equations (1) extends this fact to all actions.

Appendix F. Proof of Lemma 3

Suppose Y is equivalent to X and $\alpha(a_1, a_2), \beta(a_1, a_2)$ are the parameters from the definition of equivalence. Suppose that $(A_1 \times S_1, A_2 \times \mathbb{R} \times S_2, v_X)$ is the elicitation mechanism incentivizing X. Let

$$v_Y(a_1, s_1, a_2, r, s_2) = v_X(a_1, s_1, a_2, \frac{r - \beta(a_1, a_2)}{\alpha(a_1, a_2)}, s_2).$$

Then, $(A_1 \times S_1, A_2 \times \mathbb{R} \times S_2, v_Y)$ incentivizes Y. Indeed, the construction ensures that

- $(a_2, r, s_2) \in A^{v_Y}(a_1, s_1, p)$ is v_Y-optimal in the second period after actions a_1, s_1 at beliefs p if and only if $(a_2, \frac{r-\beta(a_1,a_2)}{\alpha(a_1,a_2)}, s_2) \in A^{v_X}(a_1, s_1, p)$ is v_X-optimal. Together with the fact that X is incentivized in the second period, the latter implies that

$$r = \alpha(a_1, a_2)\mathbb{E}_p X(a_1, a_2, \cdot) + \beta(a_1, a_2) = \mathbb{E}_p Y(a_1, a_2, \cdot).$$

- Because the expected continuation payoffs after the first-period in v_Y problem is identical to the payoffs in v_X problem, the incentives in the first-period are unchanged.

Appendix G. Proof of Theorem 3

As in the proof of Theorem 1, the construction of the payoffs in the elicitation mechanism relies on BDM. Suppose that $X(a_1, a_2, \theta) = v(a_2, \theta)$, where $A_2^v(p) = A_2^u(p)$ for each p. Assume that $x_{\min} < X(a_1, a_2, \theta) < x_{\max}$ for each a_1, a_2, θ.

Construct an elicitation mechanism $(A_1, A_2 \times \mathbb{R}, w)$ with payoffs:

$$w(a_1, r, a_2, \theta) = u(a_1, a_2, \theta) + \frac{r - x_{\min}}{x_{\max} - x_{\min}} v(a_2, \theta) + \frac{1}{2} \frac{x_{\max}^2 - r^2}{x_{\max} - x_{\min}}. \tag{7}$$

Given the first-period action a_1, second-period beliefs p, and second-period action a_2, the optimal choice of r is $r = \mathbb{E}_p v(a_2, \cdot) = \mathbb{E}_p X(a_2, \cdot)$. The expected payoffs in the second period are equal to

$$\mathbb{E}_p u(a_1, a_2, \cdot) + \frac{1}{2(x_{\max} - x_{\min})} (\mathbb{E}_p v(a_2, \cdot) - x_{\min})^2 + \frac{1}{2}(x_{\max} + x_{\min}).$$

Because the payoffs are separable and because of the choice of function v, the above is maximized by $a_2 \in A^u(a_1, p) = A^w(p)$. Notice that the second and the third component do not depend on action a_1. Hence, given any first-period beliefs μ, the optimal first-period choice must maximize the expectation of the first-component, and hence $a_1 \in A^u(\mu)$.

It follows that $(A_1, A_2 \times \mathbb{R}, w)$ incentivizes X.

Appendix H. Proof of Section 5.3

H.1. **Preliminary results.**

Lemma 13. *Suppose that X is incentivizable in the second period. Then, for any two pairs (a_1, a_2) and (b_1, b_2) such that $a_1 \neq b_1$ are adjacent and the intersection of belief sets $P(a_1, a_2) \cap P(b_1, b_2)$ has nonempty interior, there are coefficients γ, β such that, for each θ,*

$$X(a_1, a_2, \theta) = \gamma X(b_1, b_2, \theta) + \beta.$$

Proof. The proof follows the idea in the example. Take any actions $a_1, b_1 \in A_1$, and $a_2, b_2 \in A_2$, such that a_1 and a_2 are adjacent and there exists an open set of beliefs $P \subseteq P(a_1, a_2) \cap P(b_1, b_2)$ such that $a_2 \in A^u(a_1, p), b_2 \in A^u(b_1, p)$ for each $p \in P$. Then, the choice of belief set implies that the payoffs $u_1^*(a_1, p)$ and $u_1^*(b_1, p)$ as well as the payoff difference $u_1^*(a_1, p) - u_1^*(b_1, p)$ is affine in $p \in P$.

Suppose that $\mathbb{E}_p X(a_1, a_2, \theta) = \mathbb{E}_q X(a_1, a_2, \theta)$ for some $p, q \in P$, but $\mathbb{E}_p X(b_1, b_2, \theta) \neq \mathbb{E}_q X(b_1, b_2, \theta)$.

Take any nondistortionary expansion $(A_1 \times S_1, A_2 \times \mathbb{R} \times S_2, v)$ and suppose that v incentivizes X. Let $v_1^*(a_1, p) = \max_{a_2, r} \mathbb{E}_p v(a_1, (a_2, r), \theta)$ be the second-period value of information of the expansion. Find supporting actions $s_1, t_1 \in S_1$ such that (a_1, s_1) and (b_1, t_1) are adjacent, and second-period supporting actions $s_2, t_2 \in S_2$ and open subset of beliefs $Q \subseteq P$ such that $p \in Q \subseteq P^v(a_1, s_1, a_2, s_2) \cap P^v(b_1, t_1, b_2, t_2)$. If $q \notin Q$, then replace it by a belief along the interval between p and q that belongs to Q.

Lemma 4 implies that the value function $v_1^*(a_1, s_1, \alpha p + (1 - \alpha)q)$ is affine over α and $v_1^*(b_1, t_1, \alpha p + (1 - \alpha)q)$ is strictly convex over α. It follows that the difference $v_1^*(b_1, t_1, \alpha p + (1-\alpha)q) - v_1^*(a_1, s_1, \alpha p + (1-\alpha)q)$ is strictly convex in α. This leads to a contradiction with Lemma 5. The contradiction shows that it must be $\mathbb{E}_p X(b_1, b_2, \theta) = \mathbb{E}_q X(b_1, b_2, \theta)$.

Because the above argument applies for any p, q in an open set of beliefs Q, we conclude that, for any vector $v \in \mathbb{R}^\Theta$ such that $v \cdot \mathbb{1} = 0$, we have $X(a_1, a_2, \theta) \cdot v = 0$

implies $X(b_1, b_2, \cdot) \cdot v = 0$. A linear algebra argument implies that $X(b_1, b_2, \cdot) \in \text{span}\{X(a_1, a_2, \cdot), \mathbb{1}\}$. The claim follows. $\square$

Lemma 14. *Suppose X is incentivizable in the second period. Then, there is a mapping $y : \mathcal{P} \longrightarrow \mathbb{R}^\Theta$ such that for each pair of essential best responses a_1, a_2, there are γ, β such that*

$$X(a_1, a_2, \theta) = \gamma Y(\pi(a_1, a_2), \theta) + \beta.$$

The lemma says that each incentivizable question can be factorized through the appropriate element of the join collection $\mathcal{P}$.

Proof. Fix action a_1^*. For each element of the join $P \in \mathcal{P}$, choose action $a_2(P)$ such that $P = \pi(a_1^*, a_2(P))$. Define $Y(P) = X(a_1^*, a_2(P), \cdot)$.

We construct a graph G^u as follows. The edges are pairs of essential best responses (a_1, a_2). Any two pairs (a_1, a_2) and (b_1, b_2) are connected if $a_1 \neq b_1$ are adjacent and the intersection of belief sets $P(a_1, a_2) \cap P(b_1, b_2)$ has nonempty interior. Then, the connected components of the graph G^u correspond to the elements of the partition $\mathcal{P}$.

Let $\mathcal{A}$ be the set of edges of G^u for which there is γ, β such that the thesis of the lemma holds. Lemma 13 implies that if one pair belongs to $\mathcal{A}$, then all connected pairs belong to $\mathcal{A}$ as well. As a result, the entire connected components either belong or not belong to $\mathcal{A}$. At the same time, the definition of Y implies that each connected component P has at least one pair $(a_1^*, a_2(P)) \in \mathcal{A}$. The claim follows. $\square$

H.2. **Proof of Theorem 4.** If payoffs are essentially separable, then $P(a_1, a_2) = P(b_1, a_2)$ for each $a_1, b_1 \in A_1$ and $a_2 \in A_2$ and those sets are elements of the join.

Fix an action a_1^* and let $Y(a_2, \theta) = X(a_1^*, a_2, \theta)$. Then, question Y does not depend on the first-period action. Moreover, the proof of Lemma 14 implies that X must be equivalent to Y.

H.3. **Proof of Theorem 5.** We first show that, for generic payoffs, the join $\mathcal{P}$ is trivial. We start with remarks about the elements of the join. Each of such elements is a union of finitely many best response sets. The latter are convex sets bounded by finitely many linear indifference spaces:

$$I_{a_1}(a_2, b_2) = (\{p : p \cdot (u(a_1, a_2, \cdot) - u(a_1, a_2, \cdot)) = 0\},$$

which is the set of beliefs which make the DM indifferent between actions $a_2, b_2 \in A_2$ after some first-period action $a_1 \in A_1$. Each of the indifference spaces is $|\Theta - 2|$-dimensional. As a result, each of the element of the join is a finite polytope enclosed by $|\Theta - 2|$-dimensional faces, which are subsets of the indifference spaces.

On the contrary, suppose that the join is non-trivial. Let $\pi, \pi' \in \mathcal{P}$ are two adjacent elements of the join: i.e., elements that share a $|\Theta - 2|$-dimensional face F. Because of the construction, F must belong to a indifference subspace for *each* first-period action a_1:

$$F \subseteq \bigcap_{a_1} \bigcup_{a_2 \neq b_2} I_{a_1}(a_2, b_2). \tag{8}$$

At the same time, if payoffs are generic, two first-period actions a_1, b_1 never lead to identical indifference subspaces: for any two actions a_2, a_2' and b_2, b_2', vectors

$$u(a_1, a_2, \cdot) - u(a_1, a_2', \cdot) \neq u(b_1, b_2, \cdot) - u(b_1, b_2', \cdot)$$

implies that the intersection of $I_{a_1}(a_2, a_2')$ and $I_{b_1}(b_2, b_2')$ has fewer than $|\Theta - 2|$ dimensions. It follows that the set on the right-hand side of (8) has fewer than $|\Theta - 2|$ dimensions. We obtain a contradiction with the fact that F is $|\Theta - 2|$-dimensional. The contradiction shows that we cannot have two different elements in the join, hence, it must be trivial.

The claim follows from Lemma 14.

Appendix I. Proof of Theorem 6

On the contrary, suppose that there exists a nondistortionary expansion $(A_1 \times S_1, A_1 \times \mathbb{R} \times S_1, v)$ that incentivizes a non-constant question profile X in the second period.

Let a_1 be an essential best response and a_2 be an essential best response after a_1 such that $X(a_1, a_2, \cdot)$ is not constant. Take a second-period interior belief p such that $a_2 \in A^u(a_1; p)$. Let ψ be a first-period interior belief such that $A^u(\psi) = \{a_1\}$ and that assigns a strictly positive probability to belief p. Take any essential best response b_1 that is adjacent to a_1 and let ψ' be a belief at which b_1 is the unique best response. By taking a convex combination between ψ and any ψ', we can find μ such that $A^u(\mu) = \{a_1, b_1\}$, it assigns strictly positive probability $\epsilon := \mu_{a_1}(p) > 0$ to p, and such that in a neighborhood of μ, only actions a_1 and b_1 are best responses.

Find q, q' such that $p = \frac{1}{2}q + \frac{1}{2}q'$ and $\mathbb{E}_q X(a_1, a_2, \cdot) \neq \mathbb{E}_p X(a_1, a_2, \cdot)$ (such beliefs exist because $X(a_1, a_2, \cdot)$ is not constant). We can find q, q' close enough to p so that

a_2 remains a best response at both beliefs. It follows that $u_1^*(a_1, \cdot)$ is affine over the interval q and q'. At the same time, because v incentivizes question profile X, it must be that $v_1^*(a_1, s_1, \cdot)$ is strictly convex over the interval between q and q', and, in particular,

$$v_1^*(a_1, x_1, p) < \frac{1}{2} v_1^*(a_1, x_1, q) + \frac{1}{2} v_1^*(a_1, x_1, q'). \tag{9}$$

Define belief μ' so that $\mu'(P) = \mu(P)$ for each $P \subseteq \Delta\Theta \setminus \{p, q, q'\}$ and

$$\begin{aligned} \mu'_{a_1}(p) &= 0, \\ \mu'_{a_1}(q) &= \mu_{a_1}(q) + \frac{1}{2}\epsilon, \\ \mu'_{a_1}(q') &= \mu_{a_1}(q') + \frac{1}{2}\epsilon. \end{aligned}$$

Notice that $\mu' \in \mathcal{U}_{\text{quality}} \subseteq \mathcal{U}$.

We have

$$\begin{aligned} \mathbb{E}_{\mu'(a_1)} u_1^*(a_1, \cdot) &= \mathbb{E}_{\mu(a_1)} u_1^*(a_1, \cdot) + \left(\frac{1}{2}(u_1^*(a_1, q) + u_1^*(a_1, q')) - u_1^*(a_1, p) \right) \epsilon \\ &= \mathbb{E}_{\mu(a_1)} u_1^*(a_1, \cdot) \\ &= \mathbb{E}_{\mu(b_1)} u_1^*(b_1, \cdot), \end{aligned}$$

where the second equality comes from the fact that $u_1^*(a_1; \alpha p + (1-\alpha) q)$ is affine over α and the last equality is from the choice of μ. Hence, $A^u(\mu') = \{a_1, b_1\}$. At the same time, the expected payoff at μ' is equal to

$$\begin{aligned} \max_{a_1', x_1'} \mathbb{E}_{\mu'(a_1)} v_1^*(a_1', x_1', \cdot) &\geq \mathbb{E}_{\mu'(a_1)} v_1^*(a_1, x_1, \cdot) \\ &= \mathbb{E}_{\mu(a_1)} v_1^*(a_1, x_1, \cdot) + \left(\frac{1}{2}(v_1^*(a_1, x_1, q) + v_1^*(a_1, x_1, q')) - v_1^*(a_1, x_1, p) \right) \epsilon \\ &> \mathbb{E}_{\mu(a_1)} v_1^*(a_1, x_1, \cdot) \\ &= \mathbb{E}_{\mu(b_1)} v_1^*(b_1, y_1, \cdot) \\ &= \max_{b_1', y_1'} \mathbb{E}_{\mu'(b_1)} v_1^*(b_1', y_1', \cdot), \end{aligned}$$

where the inequality comes from (9) and the last equality follows from the choice of beliefs μ and the fact that v is nondistortionary. Thus, b_1 is not v-optimal at μ' and, because v is nondistortionary, $b_1 \notin A^u(\mu')$. A contradiction demonstrates the result.

References

Chambers, Christopher P and Nicolas Lambert (2021). "Dynamic belief elicitation". *Econometrica* 89.1, pp. 375–414.

Hossain, Tanjim and Ryo Okui (2013). "The binarized scoring rule". *Review of Economic Studies* 80.3, pp. 984–1001.

Lambert, Nicolas (2019). "Elicitation and evaluation of statistical forecasts". Working paper.

Lambert, Nicolas, David M Pennock, and Yoav Shoham (2008). "Eliciting properties of probability distributions". *Proceedings of the 9th ACM Conference on Electronic Commerce*, pp. 129–138.

Pęski, Marcin and Colin Stewart (2025). "Nondistortionary belief elicitation".

Tsakas, Elias (2020). "Robust scoring rules". *Theoretical Economics* 15.3, pp. 955–987.