

NON-DISTORTIONARY BELIEF ELICITATION

MARCIN PEŃSKI AND COLIN STEWART
UNIVERSITY OF TORONTO

ABSTRACT. A researcher wants to ask a decision-maker about a belief related to a choice the decision-maker made. When can the researcher provide incentives for the decision-maker to report her belief truthfully without distorting her choice? We identify which questions can be incentivized in this way in three canonical classes of problems. For these questions, we construct a simple mechanism the researcher can use.

1. INTRODUCTION

Researchers who conduct experiments in economics frequently ask subjects to make a choice facing some uncertainty and then wish to elicit subjects' beliefs related to their choice. For example, the researcher may want to ask the subject about how confident they are that they made the best possible choice *ex post* (i.e., the probability the subject assigns to the unknown state being one in which their choice is optimal). Or, following completion of a multiple-choice test, the researcher may want to ask the subject how likely they believe it is that their score exceeds a certain cutoff, or, more generally, that it lies within a certain range.

To ensure that the subjects' reported beliefs are reliable, researchers typically provide incentives that make truthful reporting optimal.¹ However, when the belief to be elicited is tied to an action choice, doing so could distort the incentives governing that choice. To take a simple example, suppose the subject must answer a multiple-choice question and then is asked the probability that they gave the correct answer. Suppose moreover that the subject is rewarded at the belief elicitation stage with a payment that is increasing in the probability the subject assigns to the correct event (namely, whether their answer was correct or not). Then a subject who is not confident about the correct answer but is confident that one of the answers is incorrect may be able to increase their overall expected payment by choosing the obviously incorrect answer and then reporting a low probability that it is correct, thereby obtaining a high expected payoff at the belief elicitation stage.

We introduce a model combining a general decision problem with a belief elicitation stage. The model allows us to consider a wide variety of belief elicitation questions: a question is described by a function $X(a, \theta)$ that, following an action choice a in the decision problem, asks the subject

Date: December 5, 2024.

¹See Healy and Leo (2024) for a discussion of incentivized vs. unincentivized belief elicitation.

the expectation of $X(a, \theta)$ according to their subjective belief about the unknown state θ . We say that a question is *incentivizable* if there exists a payment scheme at the belief elicitation stage for which (i) truthfully reporting the expectation of $X(a, \theta)$ always maximizes the subject’s expected payment, and (ii) the incentives in the decision problem are not distorted, meaning that for any belief the subject may have about θ , the optimal action in the decision problem remains optimal in the combined problem that includes the belief elicitation stage. When incentives are provided that satisfy these two criteria, the researcher can honestly tell the subject that she will maximize the payment she can expect to receive by choosing the action she believes is optimal in the decision problem and then reporting her belief truthfully; these instructions should minimize any attempts by the subject to distort their behavior.²

Two related features distinguish our approach from previous work on belief elicitation. First, the researcher asks the subject only to report a single number.³ Second, the quantity of interest to the researcher—as described by the question $X(a, \theta)$ —depends non-trivially on the subject’s choice of action in the decision problem. In the absence of either of these features, any question $X(a, \theta)$ is incentivizable using standard methods. For example, if the researcher could ask the subject to report her entire belief, it would be enough to incentivize truthful reports and randomly reward the subject either for her choice in the decision problem or for her reported belief. From a practical perspective, however, this approach could be burdensome for subjects if there are more than a few states to report on; if the researcher is only interested in a one-dimensional statistic of the belief, asking about it directly could reduce noise in the reports.

We first identify questions that are incentivizable regardless of the decision problem; we refer to these questions as being aligned with the utility $u(a, \theta)$ in the decision problem. Alignment allows for all questions that can be obtained by applying two operations to $u(a, \theta)$: (i) the addition of a function $d(\theta)$ that does not depend on a , and (ii) an affine transformation of the resulting function $u(a, \theta) + d(\theta)$ with coefficients that can depend on a but not on θ . Thus, for example, asking the subject about the payoff she expects to receive in the decision problem is incentivizable, as is her willingness to pay to have her action replaced with an *ex post* optimal one.⁴ However, the question that asks the subject about the probability that her choice is *ex post* optimal does not generally take this form (and indeed is, in many problems, not incentivizable). A researcher interested in eliciting a measure of cognitive uncertainty (Enke and Graeber, 2023) may therefore do better to ask about the subject’s expected payoff relative to the optimum rather than the probability that their action is optimal.

²Danz, Vesterlund, and Wilson (2022) find that instructions along these lines increase truthful reporting in a belief elicitation problem relative to explicit incentives.

³We discuss in Section 8.4 how to extend our methods to multiple questions.

⁴The latter question directly extends a simpler one used by Hu (2023) to elicit whether subjects were uncertain about the optimality of their choices.

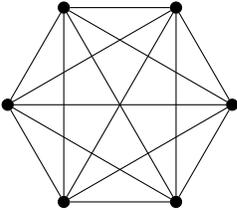
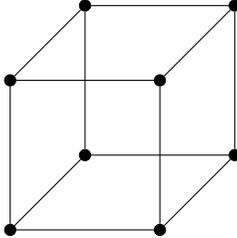
	Tree	Complete graph	Product problem
Adjacency graph			
Examples	Cognitive uncertainty (Enke and Graeber, 2023), one-dimensional supermodular problems	Multiple choice question	Random problem selection, test with ≥ 3 questions
Representation	Piecewise aligned	Aligned	Product aligned

FIGURE 1.1. Summary of main results.

Whether and which other questions are incentivizable depends on the structure of the decision problem. A particularly important role is played by what we call the “adjacency graph.” Two actions are adjacent if there is some belief at which they are both optimal and no other action is. We show that each adjacency places restrictions on how the questions following the adjacent actions are related to one another. Problems with more adjacencies therefore tend to involve stronger restrictions on which questions are incentivizable.

We fully characterize the set of incentivizable questions in three canonical classes of decision problems that differ in the structure of the adjacency graph: complete adjacency, adjacency trees, and product adjacency. (For complete and product adjacency, we also make some mild richness assumptions regarding linear independence of payoffs.) Figure 1.1 summarizes the results.

Complete adjacency graphs naturally arise in problems in which decision-maker chooses an action to match an unknown state and receives a payoff based on whether or not they succeed, as in a multiple-choice question where the state corresponds to which answer is correct and the subject receives a payment for a correct answer. In these problems, only questions that are aligned with the utility (in the sense described above) are incentivizable.

Adjacency trees naturally arise in some problems with ordered states and actions that are monotone in beliefs. For example, if states and actions are real numbers and the subject incurs a quadratic loss based on the distance between her action and the state, then the adjacency graph forms a line.⁵

⁵A typical example is when the decision problem is itself a belief elicitation problem—such as one of belief updating—with the states representing objective probabilities known to the researcher.

Because there are so few adjacencies, this case is, in a sense, the most permissive in terms of which questions are incentivizable. In particular, alignment with the utility on the full set of actions is no longer necessary; it suffices for the question to be “piecewise aligned,” meaning, in this case, that for each pair of adjacent actions it is aligned with the utility (but with the parameters governing the alignment possibly differing across pairs).

Product adjacencies arise when the decision problem comprises a number of separate tasks with complete adjacency graphs and the subject’s expected reward is a sum of rewards across these tasks. For example, the subject may complete a multiple-choice test and receive a payment proportional to their score. The adjacency graph has a simple structure: two actions are adjacent if only if they differ on a single task. In this case, a question is incentivizable if and only if it is aligned with some weighted sum of the utilities in the various tasks (but not necessarily aligned with the overall utility in the decision problem). Thus, for example, a question that asks the subject about the likelihood that her score is above some fixed cutoff is not incentivizable, while a question that asks about the expected improvement in her score across two parts of the test is.

For all of the questions that we show are incentivizable, we provide a simple construction of payments satisfying both of our incentivizability criteria. This construction is based on the classic Becker-DeGroot-Marschak method. One can first normalize $X(a, \theta)$ to lie in $[0, 1]$, and then elicit the value of $y \in [0, 1]$ at which the subject is indifferent between winning the prize with probability y and winning it with probability $X(a, \theta)$.⁶

We are not the first to observe that incentivized belief elicitation can distort other decisions. Chambers and Lambert (2021) and Healy and Leo (2024) discuss the possibility that a subject would purposefully fail a test to increase their payment from belief elicitation about their likelihood of passing. Möbius et al. (2022) describes how their belief elicitation mechanism is designed to preserve incentives in their main task. Blanco et al. (2010) find evidence that, in some problems, subjects who are paid for both a choice in a game and a reported belief take advantage of hedging opportunities, distorting either choice (or both). We implicitly assume that subjects are randomly paid either for the main task or the reported belief, eliminating such hedging opportunities.

The problem we study is motivated in part by Enke and Graeber (2023) and related and follow-up papers (e.g., Amelio, 2022; Arts, Ong, and Qiu, 2024; Xiang et al., 2021).⁷ In each of these papers, subjects’ cognitive uncertainty is elicited using unincentivized questions. Hu (2023) is the first paper we are aware of that provides strict incentives for subjects to reveal whether they are uncertain about their decision. His mechanism is essentially a simplified version of the Becker-DeGroot-Marschak

⁶A typical method for eliciting y is a multiple-price list that asks the subject to choose between probabilities y and $X(a, \theta)$ for various values of y , then choose one such choice at random for payment.

⁷This work follows a longer history of using questions about confidence in decisions that is not expressed in probabilistic terms, going back at least to Butler and Loomes (2007).

mechanism we employ; in his, subjects make a binary choice of whether to pay a cost to have some chance that their action can be replaced with the optimal one.

Belief elicitation has been widely studied and used in both theory and experiments (see Schlag, Tremewan, and Van der Weele (2015), Charness, Gneezy, and Rasochoa (2021), Haaland, Roth, and Wohlfart (2023), and Healy and Leo (2024) for surveys). The closest theoretical work to this proposal is that of Lambert, Pennock, and Shoham (2008) and Lambert (2019), which ask which properties of distributions can be elicited. Our model shares the feature that the belief elicitation question does not ask the decision-maker to report their entire belief. We sidestep their question of elicibility by restricting attention to questions that always correspond to elicitable properties, and we add the condition that the elicitation must not distort the decisions in the main decision problem.

In experiments, the binarized scoring rule of Hossain and Okui (2013) has become a popular choice for eliciting beliefs. Danz, Vesterlund, and Wilson (2022) find that subjects report more accurate beliefs when they are told that reporting truthfully will maximize the payment they can expect to receive than when the payments in the binarized scoring rule are described explicitly. In keeping with this finding, we would expect to see less distortion in behaviour in belief elicitation settings like ours if subjects are instructed that they have incentives to choose what they believe to be the optimum, which researchers can do only if the belief elicitation question is incentivizable.

Azrieli, Chambers, and Healy (2018) study incentives in a sequence of tasks and find that paying for a randomly selected problem is the only incentive-compatible mechanism when allowing for a general class of preferences. In their model, the sequence of tasks is exogenously given, whereas in ours the belief elicitation task depends on the subjects' choice in the main task; random selection for payment is therefore not sufficient to ensure incentive compatibility.

2. SOME EXAMPLES

Example 1. A decision maker chooses an action, with a payoff that depends on unknown state of the world. An analyst asks “what is the probability that your action is correct,” i.e. that it maximizes the *ex post* payoffs.

Example 2. Similarly as in Example 1, but the analyst wants to know how much the decision maker would be willing to pay to change their action after the state is realized.

Example 3. A Wall Street banker chooses the investments to match her expectations of future stock returns. An analysts wants to know the probability that the banker's choice is within $x\%$ of the realized returns.

Example 4. A student takes a test with multiple true/false questions. His payoff is equal to the score, i.e., the number of questions that he got right. An analyst wants to know what is the student subjective probability that the score was above $y\%$.

Example 5. The test in the previous part has two parts, say, Micro and Macro. The analyst wants to know what is the student expected score in the Micro part. Or how much better the student thinks they did on micro vs. macro (in expectation). Or the questions on the test have different weights and the analyst wants to ask about the fraction that are correct.

Can analyst's questions be simply incentivized? Yes, in Examples 2 and 5. No, in Example 3. In Example 4, it depends on the number of questions and x , but the answer is no if the number of questions is sufficiently high and $x \in (0, 1)$. In Example 1, it depends on the payoffs, but generically, no.

3. MODEL

A decision-maker (DM) chooses an action and then faces a belief elicitation problem posed by a researcher that may depend on the action she chose.

A *decision problem* consists of a tuple (Θ, A, \mathbf{u}) , where Θ is a finite set of states of the world, A is a finite set of actions, and $\mathbf{u} : A \rightarrow \mathbb{R}^\Theta$ is a utility function specifying, for each action, the vector of payoffs across all states. We write $u(a; \theta)$ for the θ -coordinate of the vector $\mathbf{u}(a)$. For each belief $p \in \Delta(\Theta)$, let $A(p) = \arg \max_{a \in A} \sum_{\theta} p(\theta) u(a; \theta)$ denote the set of optimal actions at p . For simplicity, we assume that (i) there are no redundant actions, i.e., no a and a' such that $\mathbf{u}(a) = \mathbf{u}(a')$, and (ii) there are no dominated actions, i.e., for each a , there exists some p such that $A(p) = \{a\}$.

After choosing an action a , the DM faces a *question* about her belief described by a function $X : A \rightarrow \mathbb{R}^\Theta$, with θ -coordinate $X(a; \theta)$. We interpret the question $X(a)$ as asking the DM to report her subjective expected value $\mathbb{E}_p X(a; \cdot)$ given the action a that she chose in the first stage and her belief p . The dependence of X on a allows for the possibility that the researcher seeks information about the DM's beliefs that are related to the chosen action.

As the following examples illustrate, this formulation allows for considerable flexibility in what the DM is asked to report.

Example 6. The question “is the chosen action a correct *ex post*?” corresponds to

$$X(a; \theta) = \begin{cases} 1 & \text{if } a \in \arg \max_{b \in A} u(b; \theta) \\ 0 & \text{otherwise.} \end{cases}$$

Then $\mathbb{E}_p X(a; \cdot)$ is the subjective probability of the chosen action being correct *ex post*.

Example 7. Given any action $a_0 \in A$, the question “is action a_0 correct *ex post* (regardless of what action was chosen)?” corresponds to

$$X(a; \theta) = \begin{cases} 1 & \text{if } a_0 \in \arg \max_{b \in A} u(b; \theta) \\ 0 & \text{otherwise.} \end{cases}$$

Then $\mathbb{E}_p X(a; \cdot)$ is the subjective probability of action a_0 being correct *ex post*.

Example 8. The question “what is the payoff from the chosen action?” corresponds to

$$X(a) = u(a)$$

for all a . Then $\mathbb{E}_p X(a; \cdot)$ is the subjective expected value of the original decision problem (given the chosen action a).

Example 9. The question “what is the regret from the chosen action?” corresponds to

$$X(a; \theta) = u(a; \theta) - \max_{b \in A} u(b; \theta).$$

Then $\mathbb{E}_p X(a; \cdot)$ is the subjective expected loss from the chosen action a relative to the *ex post* optimal action.

The DM announces a report $r \in \mathbb{R}$ and is given a reward that depends her report, her action a , and the realized state θ . Her overall payoff—including the payoff from the decision problem in the first stage—is given by a bounded function $V : \mathbb{R} \times A \times \Theta \rightarrow [0, 1]$ (which we normalize to the unit interval for convenience), the first argument of which is the DM’s report r . We refer to V as an *elicitation method*.

For simplicity, we do not include the payoff $u(a)$ explicitly in the elicitation method; adding it would not change anything as the payoff from the elicitation problem can be adjusted accordingly so as to give the same overall payoff. In the applications we have in mind, the researcher pays the DM for the decision problem with some fixed probability $\alpha \in (0, 1)$ and for the belief elicitation problem with the remaining probability $1 - \alpha$. The elicitation method V therefore takes the form $\alpha u(a; \theta) + (1 - \alpha)V_0(r, a, \theta)$ for some V_0 and some $\alpha \in (0, 1)$, where V_0 is the payoff from the belief elicitation mechanism. As is standard in the recent literature on belief elicitation, we implicitly view V_0 as the probability of winning a fixed prize to avoid any influence of risk preferences on the reported belief.

Definition 1. A question X is *incentivizable* if there exists an elicitation method V such that, for every $p \in \Delta(\Theta)$,

$$\arg \max_{r, a} \mathbb{E}_p V(r, a, \theta) \subseteq \left\{ (\mathbb{E}_p X(a; \theta), a) : a \in \arg \max_{b \in A} \mathbb{E}_p u(b; \cdot) \right\}.$$

Any V satisfying this condition *incentivizes* X .

Incentivizability combines two requirements of the elicitation method: first, the payoffs at the belief elicitation stage must not distort her action choices in the decision problem in the sense that any action a she optimally chooses in the overall problem with payoffs V is also optimal in the original decision problem with payoffs u ; and second, the DM must have strict incentives to report her true subjective

expectation of $\mathbb{E}_p X(a; \cdot)$ given her action choice a . For any question that cannot be incentivized, the researcher must either give up on truthful reporting of the DM's belief or on actions that accurately reflect the subject's belief in the decision problem. In particular, if the researcher is conducting an experiment with the DM as a subject, no matter how she designs the incentives at the belief elicitation stage, she cannot honestly tell the subjects that they will maximize their earnings by considering the decision problem in isolation and by reporting beliefs truthfully.

Say that decision problems (Θ, A, u) and $(\tilde{\Theta}, \tilde{A}, \tilde{u})$ are *equivalent* if $\tilde{\Theta} = \Theta$, $\tilde{A} = A$, and, for every p , the set $A(p)$ of optimal actions is the same in the two problems. Note that if X is incentivizable in some decision problem, it is also incentivizable in every equivalent problem since the payoff u enters only through the optimal actions $A(p)$.

For example, suppose, as described above, that V takes the form $\alpha u + (1 - \alpha)V_0$ for some V_0 and $\alpha \in (0, 1)$. Consider two decision problems with respective payoffs u and \tilde{u} . If X is incentivizable using elicitation method $V(r, a, \theta) = \alpha u(a; \theta) + (1 - \alpha)V_0(r, a, \theta)$ in the problem with payoffs u , then it is also incentivizable using elicitation method $\tilde{V}(r, a, \theta) = \alpha \tilde{u}(a; \theta) + (1 - \alpha)V(r, a, \theta)$ in the problem with payoffs \tilde{u} . In this context, if we have a belief elicitation mechanism V_0 for which V incentivizes X in some decision problem, it is straightforward to construct mechanisms that incentivize X in any equivalent decision problem.

Our formulation imposes two substantive restrictions on the belief elicitation problem. First, the DM is asked to report only a single number rather than, say, a full probability distribution. In practice, collecting more complicated information about beliefs quickly becomes impractical in experiments. If, however, the full probability distribution could be elicited, our problem would reduce to a standard belief elicitation problem since there would be no need to make the question or incentives dependent on the action chosen in the decision problem. Second, the elicited belief is based on the expectation of some question $X(a; \cdot)$. While this formulation captures many relevant cases, in principle, the researcher may want to elicit other properties of the distribution for which our approach may not apply.

Given a decision problem (Θ, A, u) , say that two actions $a, b \in A$ are *adjacent* if there is a belief $p \in \Delta(\Theta)$ such that $A(p) = \{a, b\}$, that is, at belief p , a and b are both optimal and there is no other optimal action. Adjacency of actions a and b implies that there is a $(|\Theta| - 2)$ -dimensional set of beliefs $p \in \Delta(\Theta)$ such that $a, b \in \arg \max_{c \in A} \mathbb{E}_p u(c; \cdot)$. The *adjacency graph* is the undirected graph with vertices A and edges consisting of the adjacent pairs $\{a, b\}$. Note that, since there are no redundant or dominated actions, the adjacency graph is connected for every decision problem.

4. SUFFICIENT CONDITIONS

We begin by identifying simple sufficient conditions under which a question X is incentivizable. In the following section, we show that these conditions are also necessary in some natural applications.

Definition 2. A question X is *aligned with u on $B \subseteq A$* if there exist $\gamma : B \rightarrow \mathbb{R} \setminus \{0\}$, $\kappa : B \rightarrow \mathbb{R}$, and $d \in \mathbb{R}^\Theta$ such that either

$$X(a) \equiv_B \gamma(a)(u(a) + d) + \kappa(a)\mathbf{1}$$

or

$$X(a) \equiv_B \gamma(a)d + \kappa(a)\mathbf{1},$$

where $\mathbf{1} \in \mathbb{R}^\Theta$ is the vector of all ones and \equiv_B denotes equality for all $a \in B$. We say that X is *non-trivially aligned with u on B* in the former case, and *trivially* so in the latter case. If $B = A$, we say simply that X is *aligned with u* (and similarly with the (non)-trivial qualifier).

Relative to the question $X = u$ that asks the DM about the expected utility from her chosen action, questions aligned with u allow for three changes. First, a vector d may be added to payoffs. Since d is independent of a , this change has no effect on the optimal action in the decision problem. Then, for each a , the question $X(a)$ can be rescaled by a (non-zero) constant $\gamma(a)$ and translated by another constant $\kappa(a)$ uniformly across θ . These changes make the question X essentially equivalent to the question $u + d$. The case of $X(a) \equiv_B \gamma(a)d + \kappa(a)\mathbf{1}$ can be viewed as a limit of these operations as d is scaled up and $\gamma(a)$ scaled down by the same constant, causing the $u(a)$ term to vanish.

Note that for X aligned with u , the parameters γ , κ , and d are not uniquely determined in general.

Proposition 1. *The following hold:*

- (1) *For any $d \in \mathbb{R}^\Theta$, $X(a) = u(a) + d$ and $X(a) = d$ are incentivizable.*
- (2) *Suppose that X is incentivizable. Then, for any $\gamma, \kappa : A \rightarrow \mathbb{R}$, $Y(a) = \gamma(a)X(a) + \kappa(a)\mathbf{1}$ is also incentivizable.*
- (3) *If X is aligned with u , then it is incentivizable.*

The proposition indicates that alignment with u is sufficient for incentivizability. The proof proceeds by construction using a standard Becker-DeGroot-Marschak (BDM) mechanism.

Proof. For the first statement, let $L, M \in \mathbb{R}$ be such that, for all a and θ , $L < u(a; \theta) + d(\theta) < M$ and $L < d(\theta) < M$. If $X(a) = u(a) + d$, let

$$V(r, a, \theta) = \int_L^r X(a; \theta) dx + \int_r^M x dx - \frac{M^2}{2} = (u(a; \theta) + d(\theta))(r - L) - \frac{1}{2}r^2,$$

where $d(\theta)$ is the θ -coordinate of d . For this V , simple calculations show that the optimal choice of r given a is $\mathbb{E}_p X(a; \theta)$. Since the optimal r is greater than L , the optimal choice of action a is the same as in the decision problem with utility u . A similar argument applies if $X(a) = d$, in which case we let $V(r, a, \theta) = u(a; \theta) + d(\theta)(r - L) - \frac{1}{2}r^2$.

For the second statement, suppose V incentivizes X . Let $W(r, a, \theta) = V\left(\frac{1}{\gamma(a)}(r - \kappa(a)), a, \theta\right)$. It is straightforward to verify that W incentivizes Y .

The third statement follows from the first two. \square

For a researcher interested in eliciting a measure of the DM's confidence in her choice of action, an immediate implication of the proposition, captured in the following corollary, is that it is possible to elicit the DM's expected regret without distorting her decisions.

Corollary 1. *For any decision problem, the question about payoffs in Example 8 and that about regret in Example 9 are incentivizable, as is any question that does not depend on the chosen action (such as that in Example 7).*

It is useful to have an explicit formula for an elicitation method that incentivizes a question X aligned with u . When X is non-trivially aligned with u , using the BDM construction from the proof of Proposition 1 gives

$$\begin{aligned}
(4.1) \quad V^{BDM}(r, a, \theta; \gamma, \kappa, \mathbf{d}) &= (u(a; \theta) + d(\theta)) \left(\frac{1}{\gamma(a)}(r - \kappa(a)) - L \right) - \frac{1}{2} \left(\frac{1}{\gamma(a)}(r - \kappa(a)) \right)^2 \\
&= \frac{1}{\gamma(a)^2} \left[(X(a; \theta) - \kappa(a))(r - \kappa(a) - \gamma(a)L) - \frac{1}{2}(r - \kappa(a))^2 \right] \\
&= \frac{1}{\gamma(a)^2} \left(X(a; \theta) - \frac{1}{2}r - \frac{1}{2}\kappa(a) \right) (r - \kappa(a)) - \frac{1}{\gamma(a)}(X(a; \theta) - \kappa(a))L.
\end{aligned}$$

Similarly, if X is trivially aligned with u , adding $u(a; \theta)$ to either of the last two lines gives V^{BDM} .

When X is non-trivially aligned with u , the expected payoff of a DM who chooses action a and then chooses r optimally is equal to

$$\begin{aligned}
(4.2) \quad \max_r \left[\mathbb{E}_p[u(a; \cdot) + d(\cdot)] \left(\frac{1}{\gamma(a)}(r - \kappa(a)) - L \right) - \frac{1}{2} \left(\frac{1}{\gamma(a)}(r - \kappa(a)) \right)^2 \right] \\
= \max_x \left[\mathbb{E}_p[u(a; \cdot) + d(\cdot)](x - L) - \frac{1}{2}x^2 \right] \\
= \frac{1}{2} (\mathbb{E}_p[u(a; \cdot) + d(\cdot)])^2 - L \mathbb{E}_p[u(a; \cdot) + d(\cdot)].
\end{aligned}$$

Similarly, the expected payoff is equal to $\mathbb{E}_p[u(a; \cdot)] + (\mathbb{E}_p[d(\cdot)])^2 / 2 - L \mathbb{E}_p[d(\cdot)]$ if X is trivially aligned with u .

Depending on the structure of the decision problem, there can be incentivizable questions that are not aligned with u . In particular, depending on the adjacency graph, we can extend the sufficient condition from Proposition 1 to certain other questions.

Say that action a is *splitting* if removing it from the adjacency graph makes the graph disconnected. If $B_0, B_1 \subset A$ are disjoint sets such that $B_0 \cup B_1 = A \setminus \{a\}$ and the adjacency graph contains no edges between B_0 and B_1 , we say that a *splits the adjacency graph into B_0 and B_1* .⁸

Splitting actions have the following intuitive separation property.

Lemma 1. *Suppose action a splits the adjacency graph into B_0 and B_1 . Then, for any pair of beliefs p_0 and p_1 such that, for each i , some action $b_i \in B_i$ is optimal at p_i in the decision problem, there is a convex combination $p = \alpha p_1 + (1 - \alpha) p_0$ such that a is optimal at p .*

Proof. Suppose not. Then all actions that are optimal at convex combinations of p_0 and p_1 must be either from B_0 or B_1 . Hence, for some such convex combination p' , the set of optimal actions contains at least one element from each of B_0 and B_1 . Then in any neighbourhood of p' , there must be a belief at which there is exactly one member of each B_i that is optimal, contradicting the fact that no action in B_0 is adjacent to any action in B_1 . \square

A *splitting collection* $\{A_0, \dots, A_k\}$ is a collection of subsets $A_i \subseteq A$ such that $\bigcup A_i = A$ and, for each i and j , either $A_i \cap A_j = \emptyset$ or $A_i \cap A_j = \{a\}$ for some splitting action a .

Example 10. Suppose the decision problem involves guessing the correct state, with the option of opting out and choosing a safe action. The action set is $A = \Theta \cup \{a_s\}$, with payoffs

$$u(a; \theta) = \begin{cases} 1 & \text{if } a = \theta, \\ s & \text{if } a = a_s, \\ 0 & \text{otherwise,} \end{cases}$$

where $s \in (0, 1)$. If $s \geq 1/2$, then the adjacency graph is a star with action a_s in the centre. Letting $\{B_0, B_1\}$ be any partition of Θ , $\{B_0 \cup \{a_s\}, B_1 \cup \{a_s\}\}$ forms a splitting collection. Another splitting collection is given by the set of pairs $\{\theta, a_s\}$ for all $\theta \in \Theta$.

Example 12 below describes an example with ordered actions a_0 through a_n in which the adjacency graph forms a line, with a_i adjacent to a_j if and only if $|i - j| = 1$. Thus $\{\{a_0, a_1\}, \dots, \{a_{n-1}, a_n\}\}$ forms a splitting collection.

We say that X is *piecewise aligned with u* if it is aligned with u on each element of a splitting collection.

Theorem 1. *If X is piecewise aligned with u , then it is incentivizable.*

⁸Note that B_0 and B_1 need not be connected, and hence that the sets into which a splitting action splits the adjacency graph are not uniquely determined in general.

Proof. To keep the notation simple, we present the argument only for the case in which the splitting collection contains two elements, A_0 and A_1 , with splitting action a_0 . Extending the argument to the general case is straightforward.

We construct an elicitation method on each A_i , then show that they agree on a_0 and therefore give rise to well-defined elicitation method on the full action set, A .

First suppose \mathbf{X} is non-trivially aligned with \mathbf{u} on each A_i . For each $i = 0, 1$, let $(\gamma_i, \kappa_i, \mathbf{d}_i)$ be such that $\mathbf{X}(a) \equiv_{A_i} \gamma_i(a)(\mathbf{u}(a) + \mathbf{d}) + \kappa_i(a)\mathbf{1}$. Define the elicitation methods

$$(4.3) \quad V_i(r, a, \theta; w_i, \omega_i) = \left(\frac{\gamma_i(a_0)}{\gamma_i(a)} \right)^2 \left(X(a; \theta) - \frac{1}{2}r - \frac{1}{2}\kappa_i(a) \right) (r - \kappa_i(a)) \\ + w_i(\theta) - \omega_i - \frac{\gamma_i(a_0)^2}{\gamma_i(a)} (X(a; \theta) - \kappa_i(a))L,$$

where $w_i \in \mathbb{R}^\ominus$ and $\omega_i \in \mathbb{R}$. Note that this expression differs from the expression for V^{BDM} in (4.1) only by multiplication by a positive constant and addition of a function that depends only on the state. As neither of these changes affects the optimal choices of a and r , it follows from the argument in the proof of Proposition 1 that V_i incentivizes \mathbf{X} on A_i (i.e., in the decision problem with actions restricted to A_i).

Notice that

$$V_1(r, a_0, \theta; w_1, \omega_1) - V_0(r, a_0, \theta; w_0, \omega_0) \\ = \left(X(a_0; \theta) - \frac{1}{2}r - \frac{1}{2}\kappa_1(a_0) \right) (r - \kappa_1(a_0)) - \left(X(a_0; \theta) - \frac{1}{2}r - \frac{1}{2}\kappa_0(a_0) \right) (r - \kappa_0(a_0)) \\ + (w_1(\theta) - w_0(\theta)) - (\omega_1 - \omega_0) - \gamma_1(a_0)(X(a_0; \theta) - \kappa_1(a_0))L + \gamma_0(a_0)(X(a_0; \theta) - \kappa_0(a_0))L \\ = X(a_0; \theta)(\kappa_0(a_0) - \kappa_1(a_0) + \gamma_0(a_0) - \gamma_1(a_0)) + \frac{1}{2} \left((\kappa_1(a_0))^2 - (\kappa_0(a_0))^2 \right) \\ + (w_1(\theta) - w_0(\theta)) - (\omega_1 - \omega_0) + (\gamma_1(a_0)\kappa_1(a_0) - \gamma_0(a_0)\kappa_0(a_0))L.$$

Given any w_0 and ω_0 , let

$$w_1(\theta) = w_0(\theta) - X(a_0, \theta)(\kappa_0(a_0) - \kappa_1(a_0) + \gamma_0(a_0) - \gamma_1(a_0))$$

for each θ , and

$$\omega_1 = \frac{1}{2}(\kappa_1(a_0)^2 - \kappa_0(a_0)^2) + \omega_0 + (\gamma_1(a_0)\kappa_1(a_0) - \gamma_0(a_0)\kappa_0(a_0))L.$$

Then $V_1(r, a_0, \theta; w_1, \omega_1) = V_0(r, a_0, \theta; w_0, \omega_0)$ for all r and θ .⁹

Let

$$V(r, a, \theta) = V_i(r, a, \theta) \text{ for each } a \in A_i.$$

⁹In the general case with splitting collection $\{A_1, \dots, A_k\}$, one can recursively define each w_{i+1} and ω_{i+1} given w_i and ω_i in an analogous fashion.

Because V_0 and V_1 agree for action a_0 , V is well defined.

In case \mathbf{X} is trivially aligned with \mathbf{u} on some A_i , we add $\gamma_i(a_0)^2 u(a; \theta)$ to the expression for $V_i(r, a, \theta; w_i, \omega_i)$ in (4.3) and adjust the definition of $w_1(\theta)$ accordingly to include any such additional $u(a_0; \theta)$ terms.

To verify that V incentivizes \mathbf{X} , it suffices to show that at any belief p at which no action in A_i is optimal in the decision problem, $\arg \max_a \max_r \mathbb{E}_p[V(r, a, \cdot)] \subseteq A_j$, where $j \neq i$. Without loss of generality, take $i = 1$ and $j = 0$, and let p_0 denote such a belief.

Note first that, by (4.2), the expected value from choosing an action $a \in A_i$ followed by an optimal choice of r is equal to

$$(4.4) \quad \max_r \mathbb{E}_p[V_i(r, a, \cdot)] = \gamma_i(a_0)^2 \frac{1}{2} (\mathbb{E}_p[u(a, \cdot)] + \mathbb{E}_p[d(\cdot)])^2 \\ + \mathbb{E}_p[w_i(\cdot)] - \omega_i - L\gamma_i(a_0)^2 \mathbb{E}_p[u(a; \cdot) + d(\cdot)].$$

It follows that the expected value from an action $b \in A_i$ is at least as large as that from an action $a \in A_i$ if and only if

$$\frac{1}{2} (\mathbb{E}_p[u(b; \cdot)] + \mathbb{E}_p[d(\cdot)])^2 - L \mathbb{E}_p[u(b; \cdot)] \geq \frac{1}{2} (\mathbb{E}_p[u(a; \cdot)] + \mathbb{E}_p[d(\cdot)])^2 - L \mathbb{E}_p[u(a; \cdot)],$$

which holds if and only if $\mathbb{E}_p[u(b; \cdot)] \geq \mathbb{E}_p[u(a; \cdot)]$ since the function $(x + y)^2/2 - Lx$ is increasing in x for $x + y > L$ and L satisfies $\mathbb{E}_p[u(a'; \cdot)] + \mathbb{E}_p[d(\cdot)] > L$ for all actions a' .

Suppose for contradiction that there is some $b \in A_1 \setminus \{a_0\}$ such that $b \in \arg \max_a \max_r \mathbb{E}_{p_0} V(r, a, \cdot)$. It follows that

$$\max_r \mathbb{E}_{p_0} V_1(r, b, \cdot) = \max_r \mathbb{E}_{p_0} V(r, b, \cdot) \geq \max_r \mathbb{E}_{p_0} V(r, a_0, \cdot) = \max_r \mathbb{E}_{p_0} V_1(r, a_0, \cdot).$$

By the observation in the preceding paragraph,

$$\mathbb{E}_{p_0} u(b, \cdot) \geq \mathbb{E}_p u(a_0, \cdot).$$

Let p_1 be a belief at which b is strictly optimal, i.e., $\{b\} = \arg \max_a \mathbb{E}_{p_1}[u(a, \cdot)]$. By Lemma 1, action a_0 must be optimal at some convex combination $p = \alpha p_1 + (1 - \alpha) p_0$. At the same time, the above inequalities imply that

$$\begin{aligned} \mathbb{E}_p[u(a_0, \cdot)] &= \alpha \mathbb{E}_{p_1}[u(a_0, \cdot)] + (1 - \alpha) \mathbb{E}_{p_0}[u(a_0, \cdot)] \\ &< \alpha \mathbb{E}_{p_1}[u(b, \cdot)] + (1 - \alpha) \mathbb{E}_{p_0}[u(b, \cdot)] \\ &= \mathbb{E}_p[u(b, \cdot)], \end{aligned}$$

contradicting the optimality of a_0 at p . □

Example 11. Suppose there are 3 actions $A = \{a, b, c\}$ with adjacency graph

$$a - b - c.$$

The following question is incentivizable but not aligned with u : $X(a) = u(a)$, $X(b) = u(b)$, and $X(c) = \frac{1}{2}u(b) + u(c)$.

5. NECESSARY CONDITION

We now consider necessary conditions for questions to be incentivizable. In certain important classes of problems, these conditions will imply that, depending on the class, only questions aligned or piecewise aligned with u can be incentivized.

From this point on, except where otherwise stated, we assume that $|\Theta| > 3$.

5.1. Key observation. We begin with a key observation that underlies all of our necessary conditions. Consider two adjacent actions, a and b . Suppose an elicitation method V incentivizes some X . Notice that, for any fixed r , the expected value $\mathbb{E}_p[V(r, a, \cdot)]$ is affine in p . Moreover, the set of beliefs at which a given r is optimal following action a consists of the intersection of the simplex with the hyperplane defined by $r = \mathbb{E}_p[X(a, \cdot)]$. If this set intersects with the set of beliefs at which both a and b are optimal, then the value obtained from action b followed by the optimal report must itself be affine within this intersection. This in turn implies that the optimal report $r' = \mathbb{E}_p[X(b, \cdot)]$ following action b must be constant along this intersection since the value for any fixed report is affine. This observation places a strong restriction on how $X(a)$ and $X(b)$ are related, which is formalized in the following lemma.

For any vector $v \in \mathbb{R}^\Theta$, let

$$\bar{v} = v - \frac{1}{|\Theta|} \sum_{\theta' \in \Theta} v(\theta') \mathbf{1}.$$

Thus \bar{v} is the projection of v onto the hyperplane of vectors whose coordinates sum to 0. Recall that two vectors $u, v \in \mathbb{R}^\Theta$ are *collinear* if there exists $\alpha \neq 0$ such that $v = \alpha u$.

Let $\Delta_a^b = \bar{u}(b) - \bar{u}(a)$ be the payoff difference vector.

Lemma 2 (Adjacency Lemma). *Suppose that X is incentivizable. If actions a and b are adjacent, then there exist $\rho \in \mathbb{R}$ and $\sigma \in \mathbb{R} \setminus \{0\}$ such that*

$$\bar{X}(b) = \rho \Delta_a^b + \sigma \bar{X}(a).$$

If $\bar{X}(a)$ or $\bar{X}(b)$ is collinear with Δ_a^b , then we can take $\rho \neq 0$.

The Adjacency Lemma indicates that if X is incentivizable and a and b are adjacent actions, then X must be aligned with u on $\{a, b\}$. To see this, note first that X is non-trivially aligned with u on a set of actions B if and only if $\bar{X}(a) \equiv_B \gamma(a) \bar{u}(a) + \mathbf{d}$ for some $\gamma(a) \in \mathbb{R}$ and $\mathbf{d} \in \mathbb{R}^\Theta$, and is trivially aligned with u on B if and only if $\bar{X}(a) \equiv_A \gamma(a) \mathbf{d}$ for some $\gamma(a) \in \mathbb{R}$ and $\mathbf{d} \in \mathbb{R}^\Theta$. If $\rho = 0$, then X is trivially aligned with u on $\{a, b\}$. If $\rho \neq 0$, then the alignment is nontrivial, with

$$\mathbf{d} = \frac{1}{\rho} \bar{X}(b) - \bar{u}(b) = \frac{\sigma}{\rho} \bar{X}(a) - \bar{u}(a),$$

$\gamma(a) = \rho/\sigma$, and $\gamma(b) = \rho$.

Proof. Let V be an elicitation method that incentivizes \mathbf{X} . Consider two distinct beliefs p_0 and p_1 such that (i) actions a and b are both optimal in the decision problem, i.e., $a, b \in \arg \max_{a'} \mathbb{E}_{p_k}[u(a', \cdot)]$ for $k = 0, 1$, and (ii) the question $\mathbf{X}(a)$ attains the same value at p_0 and p_1 , i.e., $\mathbb{E}_{p_0}[\mathbf{X}(a, \cdot)] = \mathbb{E}_{p_1}[\mathbf{X}(a, \cdot)]$; such a pair of beliefs exists because $|\Theta| > 3$. Letting $r = \mathbb{E}_{p_0}[\mathbf{X}(a, \cdot)]$, it follows that (a, r) is optimal given V at all $p_\alpha = \alpha p_1 + (1 - \alpha)p_0$ for $\alpha \in [0, 1]$. The optimal expected payoff $\mathbb{E}_{p_\alpha}[V(r, a, \cdot)]$ is therefore linear in α . Since b is also an optimal action at each p_α , the optimal expected payoff $\max_s \mathbb{E}_{p_\alpha}[V(s, b, \cdot)]$ must be linear in α as well. Therefore, there exists some r' such that, for each α , $r' \in \arg \max_s \mathbb{E}_{p_\alpha}[V(s, b, \cdot)]$. In particular, $\mathbb{E}_{p_k}[\mathbf{X}(b, \cdot)] = r'$ for each $k = 0, 1$.

Fixing p_0 , notice that the preceding argument applies to all $p_1 \in \mathbb{R}^\Theta$ satisfying the following orthogonality conditions:

- (1) $p_1 - p_0 \perp \mathbf{1}$, ensuring that p_1 is a well-defined belief;
- (2) $p_1 - p_0 \perp u(a) - u(b)$, ensuring that the DM is indifferent between a and b at p_1 ; and
- (3) $p_1 - p_0 \perp \mathbf{X}(a)$, ensuring that $\mathbf{X}(a)$ attains the same value at p_0 as at p_1 .

For any such p_1 , the preceding argument implies that $p_1 - p_0 \perp \mathbf{X}(b)$.

By a standard linear algebra argument, it follows that

$$\mathbf{X}(b) \in \text{span}(\mathbf{1}, u(a) - u(b), \mathbf{X}(a)).$$

Noting that, for any vector \mathbf{v} , $\bar{\mathbf{v}}$ differs from \mathbf{v} by a scalar multiple of $\mathbf{1}$, and that $\mathbf{v} \perp \mathbf{1}$, we obtain

$$\bar{\mathbf{X}}(b) = \rho(\bar{u}(a) - \bar{u}(b)) + \sigma\bar{\mathbf{X}}(a)$$

for some $\rho, \sigma \in \mathbb{R}$.

If $\bar{\mathbf{X}}(b)$ is not collinear with $\bar{u}(b) - \bar{u}(a)$, then we must have $\sigma \neq 0$, as needed. Otherwise, switching the roles of a and b in the preceding argument yields that $\bar{\mathbf{X}}(a)$ is also collinear with $\bar{u}(b) - \bar{u}(a)$, and therefore one can take ρ and σ to be nonzero. \square

The Adjacency Lemma provides a key tool for identifying whether a question is incentivizable. Since it applies to every pair of adjacent actions, it tends to impose stronger restrictions when the adjacency graph is more complete.

The Adjacency Lemma applies only when there are at least four states. If there are only two states, any question $X(a, \theta)$ that is not constant in θ is equivalent to simply asking the DM to report her belief, which can be easily incentivized by adding a standard scoring rule to the utility $u(a)$. With three states, the set of beliefs at which two adjacent actions are optimal forms a line segment. Further restricting to a particular optimal report for one of the actions generically reduces the set to a single point. Matching the values along the line segment does not, therefore, impose restrictions on the

question. Cycles of adjacencies can, however, imply substantive restrictions; we expand on this point in Section 8.

Recall that two adjacent actions share a $(|\Theta| - 1)$ -dimensional set of beliefs at which they are both optimal. If two non-adjacent actions, a and b , “touch” in the sense that there is a nonempty set of beliefs of dimension strictly between 0 and $(|\Theta| - 1)$ at which both are optimal, then following the logic of the proof of the Adjacency Lemma would give rise to a weaker restriction on the relationship between $X(a)$ and $X(b)$.

5.2. Adjacency trees. Suppose that, as in Examples 10 and 12, the adjacency graph is a tree. Let $\mathcal{A} = \{A_0, \dots, A_k\}$ consist of all adjacent pairs of actions; that is, $\tilde{A} \in \mathcal{A}$ if and only if $\tilde{A} = \{a, b\}$ for some adjacent actions a and b . Because the adjacency graph is a tree, \mathcal{A} is a splitting collection.

Theorem 2. *Suppose the adjacency graph is a tree. Then the following are equivalent:*

- (1) X is incentivizable;
- (2) X is piecewise aligned with \mathbf{u} ;
- (3) the conclusion of the Adjacency Lemma holds for every pair of adjacent actions.

Proof. That (2) implies (1) follows immediately from Theorem 1. That (1) implies (3) follows immediately from the Adjacency Lemma. That (3) implies (2) follows from the discussion in the paragraph immediately following the statement of the Adjacency Lemma. \square

The equivalence of (1) and (2) in Theorem 2 is a special case of Theorem 5 below.

Example 12 (Second-order beliefs). The following example is a slightly simplified version of a belief-updating experiment from Enke and Graeber (2023); the same conclusions apply to their original experiment.

The decision problem involves forecasting a binary event. The action set and state space $A = \Theta = \{0, 1/n, \dots, 1\}$ consist of (discretized) probabilities that the event occurs, where $n \geq 3$. One can think of θ as the “true” probability given the available information, which is known to the researcher but about which the DM may be uncertain (for example because she has doubts about how to update her beliefs in light of the information she observes). This uncertainty is captured by her belief $p \in \Delta(\Theta)$.

The DM is rewarded more for forecasts that are closer to the state according to the payoff function $u(a; \theta) = -(a - \theta)^2$. The DM optimally chooses an action closest to $\mathbb{E}_p[\theta]$. The adjacency graph therefore forms a line: a_i and a_j are adjacent if and only if $|a_i - a_j| = 1/n$. The set $\{\{0, 1/n\}, \{1/n, 2/n\}, \dots, \{(n-1)/n, 1\}\}$ forms a splitting collection.

The researcher wishes to elicit the DM’s confidence in her report by asking how likely she believes it is that her action is within x of the true value of θ for some fixed $x \in [0, 1/2]$. This question is

described by

$$X(a, \theta) = \begin{cases} 1 & \text{if } |a - \theta| \leq x \\ 0 & \text{otherwise.} \end{cases}$$

To check whether X is incentivizable, we use Theorem 2.

Consider adjacent actions a and $b = a + 1/n$. For each θ ,

$$u(b, \theta) - u(a, \theta) = - \left(a + \frac{1}{n} - \theta \right)^2 + (a - \theta)^2 = -\frac{1}{n} \left(2a - 2\theta + \frac{1}{n} \right).$$

The difference between $u(b, \theta) - u(a, \theta)$ and $\bar{u}(b, \theta) - \bar{u}(a, \theta)$ is the constant $(a-1)^2 - (a-1/n)^2$. Thus we obtain $\bar{u}(b, \theta) - \bar{u}(a, \theta) = 2a + 2\theta/n - 1$. In particular, the coordinates of any vector collinear with Δ_a^b must feature constant differences in θ . However, there is no σ for which $\bar{X}(b) - \sigma\bar{X}(a)$ satisfies this property. To see this, note that if $x < 1/n$, $\bar{X}(b) - \sigma\bar{X}(a)$ is constant across all $\theta \notin \{a, b\}$, and if $x \geq 1/n$, then its coordinates for $\theta = a$ and $\theta = b$ are equal. Therefore, there is no σ such that $\bar{X}(b) - \sigma\bar{X}(a)$ is collinear with Δ_a^b , and by Theorem 2, X is not incentivizable.

This result is not specific to the quadratic payoffs in the decision problem; the same conclusion applies if the payoff is replaced with any strictly proper scoring rule (one for which reporting an action close to the expectation of θ is optimal).

Faced with this result, what should the researcher do? One option is to use a difference measure of decision confidence that *is* incentivizable. For instance, according to Corollary 1, the expected regret question of Example 9 is incentivizable in every decision problem.

6. ADJACENCY CYCLES

When the adjacency graph forms a tree, the incentivizable questions are those that are piecewise aligned with u . Cycles in the adjacency graph impose additional restrictions, further limiting which questions are incentivizable.

An *adjacency cycle* is a tuple $C = (a_0, \dots, a_n)$ such that $a_n = a_0$ and actions a_i and a_{i+1} are adjacent for each $i = 0, \dots, n-1$. We say that n is the *length* of the cycle C . We abuse notation slightly by writing $a \in C$ to mean that $a = a_i$ for some i .

An adjacency cycle is *internally independent* if, for some $a \in C$, the set of vectors $\{\Delta_a^b : b \in C \setminus \{a\}\}$ is linearly independent. Let V_C be the linear space spanned by $\{\Delta_a^b : b \in C \setminus \{a\}\}$. The cycle C is internally independent if and only if $\dim V_C = n - 1$, where n is the length of the cycle. Notice that the space V_C and the linear independence of $\{\Delta_a^b : b \in C \setminus \{a\}\}$ do not depend on the choice of a , only on C itself.

Lemma 3. *Let C be an internally independent adjacency cycle. If X is incentivizable, then either X is aligned with u on C , or $\bar{X}(a) \in V_C \setminus \{0\}$ for all $a \in C$.*

Proof. Without loss of generality, let $a_0 \in C$ be such that either $\bar{X}(a_0) = 0$ or $\bar{X}(a_0) \notin V_C$. (If no such a_0 exists, the lemma holds trivially.)

By Lemma 2, for each $i = 1, \dots, n$, we have $\bar{X}(a_i) = \rho_i \Delta_{a_{i-1}}^{a_i} + \sigma_i \bar{X}(a_{i-1})$ for some $\sigma_i \neq 0$ and some ρ_i . Iterating these equations gives

$$\bar{X}(a_0) = \bar{X}(a_n) = \sum_{i=1}^n \Gamma_i \rho_i \Delta_{a_{i-1}}^{a_i} + \Gamma_0 \bar{X}(a_0),$$

where $\Gamma_i = \sigma_n \cdots \sigma_{i+1}$ and $\Gamma_n = 1$. Because $\Delta_{a_0}^{a_1} + \cdots + \Delta_{a_{n-1}}^{a_n} = 0$, we get

$$\sum_{i=1}^{n-1} (\Gamma_i \rho_i - \rho_n) \Delta_{a_{i-1}}^{a_i} + (\Gamma_0 - 1) \bar{X}(a_0) = 0.$$

Since either $\bar{X}(a_0) = 0$ or $\bar{X}(a_0) \notin V_C$, it follows that

$$\sum_{i=1}^{n-1} (\Gamma_i \rho_i - \rho_n) \Delta_{a_{i-1}}^{a_i} = 0.$$

Internal independence implies that $\Gamma_i \rho_i = \rho_n$ for each $i = 1, \dots, n-1$.

If $\rho_n = 0$, then $\rho_i = 0$ for each i since $\Gamma_i \neq 0$. In this case, all $\bar{X}(a_i)$ are collinear, which implies that X is trivially aligned with u on C .

For the case of $\rho_n \neq 0$, first note that, by the same iteration as above, for each k ,

$$\begin{aligned} \bar{X}(a_k) &= \sum_{i=1}^{k-1} \rho_i \frac{\Gamma_i}{\Gamma_k} \Delta_{a_{i-1}}^{a_i} + \frac{\Gamma_0}{\Gamma_k} \bar{X}(a_0) \\ &= \frac{\rho_n}{\Gamma_k} \sum_{i=1}^{k-1} \Delta_{a_{i-1}}^{a_i} + \frac{\Gamma_0}{\Gamma_k} \bar{X}(a_0) \\ &= \frac{\rho_n}{\Gamma_k} (\bar{u}(a_k) - \bar{u}(a_0)) + \frac{\Gamma_0}{\Gamma_k} \bar{X}(a_0). \end{aligned}$$

Letting $\gamma(a_k) = \rho_k = \rho_n / \Gamma_k \neq 0$ and $d = -\bar{u}(a_0) + \rho_n^{-1} \Gamma_0 \bar{X}(a_0)$, we have $\bar{X}(a_k) = \gamma(a_k) (\bar{u}(a_k) + d)$ for all k , and thus X is (non-trivially) aligned with u on C . \square

While Lemma 3 identifies conditions that an incentivizable question must satisfy on a given internally independent cycle, it says nothing about more complicated adjacency graphs. The following lemma shows how alignment with u on multiple cycles or other subsets of actions can, under mild conditions, be combined to obtain alignment with u on their union.

Lemma 4. *Suppose X is aligned with u on sets of actions B and D . If there exist actions $a_0, a_1 \in B \cap D$ such that $\bar{X}(a_0)$ and $\bar{X}(a_1)$ are not collinear, then X is aligned with u on $B \cup D$.*

Proof. Alignment with u on B implies that there exist $\gamma_0^B, \gamma_1^B \in \mathbb{R}$ and $d^B \in \mathbb{R}^\Theta$ such that

$$\gamma_0^B \bar{X}(a_0) - \bar{u}(a_0) = \gamma_1^B \bar{X}(a_1) - \bar{u}(a_1) = d^B.$$

Similarly, for D ,

$$\gamma_0^D \bar{X}(a_0) - \bar{u}(a_0) = \gamma_1^D \bar{X}(a_1) - \bar{u}(a_1) = \mathbf{d}^D.$$

Rearranging these equations gives

$$\begin{aligned} \gamma_0^B \bar{X}(a_0) - \gamma_1^B \bar{X}(a_1) &= \bar{u}(a_0) - \bar{u}(a_1) \\ \text{and } \gamma_0^D \bar{X}(a_0) - \gamma_1^D \bar{X}(a_1) &= \bar{u}(a_0) - \bar{u}(a_1). \end{aligned}$$

Subtracting one equation from the other leads to

$$(\gamma_0^B - \gamma_0^D) \bar{X}(a_0) - (\gamma_1^B - \gamma_1^D) \bar{X}(a_1) = 0.$$

Since $\bar{X}(a_0)$ and $\bar{X}(a_1)$ are linearly independent, it must be that $\gamma_k^B = \gamma_k^D$ for each $k = 0, 1$. This in turn implies that $\mathbf{d}^B = \mathbf{d}^D = \mathbf{d}$.

All that remains is to show that for any other action $a \in B \cup D$, the corresponding parameters γ_a^B and γ_a^D are equal. Since $\bar{X}(a)$ cannot be collinear with both $\bar{X}(a_0)$ and $\bar{X}(a_1)$, we can repeat the argument replacing one of a_0 or a_1 with a to obtain $\gamma_a^B = \gamma_a^D$. \square

6.1. Complete adjacency. In some common decision problems, the adjacency graph is complete. This is the case, for instance, in Example 10 when $s < 1/2$.

Theorem 3. *Suppose the adjacency graph is complete and $|A| \geq 4$. Suppose in addition that for any four actions a, b_0, b_1, b_2 , the set of vectors $\{\Delta_a^{b_i}\}_{i=0,1,2}$ is linearly independent. Then X is incentivizable if and only if it is aligned with \mathbf{u} .*

That alignment with \mathbf{u} is sufficient for incentivizability follows from part (3) of Proposition 1. Necessity follows from Theorem 4 below, which provides a more general abstract condition for arbitrary adjacency graphs under which only questions aligned with \mathbf{u} are incentivizable.

The main idea behind Theorem 3 is that, given any cycle C through an action a , Lemma 3 implies that any incentivizable X is either aligned with \mathbf{u} on C or lies in $V_C \setminus \{0\}$. The linear independence condition can be used to eliminate the latter possibility by considering multiple cycles through a ; more precisely, there must be *some* cycle C on which X is aligned with \mathbf{u} , as the intersection of the sets $V_{C'}$ across cycles C' containing a is $\{0\}$. Varying a and applying Lemma 4 leads to alignment with \mathbf{u} on the entire action set.

Example 13. Consider a decision problem with at least four states in which the DM is asked to guess the state and receives a reward for guessing correctly that may depend on which state is realized. Thus $A = \Theta$ and

$$u(a, \theta) = \begin{cases} r_\theta & \text{if } a = \theta \\ 0 & \text{otherwise,} \end{cases}$$

where $r_\theta > 0$ for all θ . The adjacency graph for this problem is complete: given any two states θ and θ' , the actions θ and θ' are the only optimal actions at the belief that assigns probability $r_{\theta'}/(r_\theta + r_{\theta'})$ to state θ and all of the remaining probability to state θ' .

Suppose the researcher seeks to elicit the DM's belief about whether she correctly guessed the state, which is described by the question

$$X(a, \theta) = \begin{cases} 1 & \text{if } a = \theta \\ 0 & \text{otherwise.} \end{cases}$$

Taking $\gamma(a) = 1/r_a$ and $\kappa(a) = 0$ for each a and $d = 0$, we see that X is aligned with u , and therefore can be incentivized using the BDM construction of Proposition 1.

Now consider the same question X in a different decision problem where the DM can also receive a smaller reward for a “close” guess. Let $\Theta = \{1, \dots, n\}$ and

$$u(a, \theta) = \begin{cases} r_\theta & \text{if } a = \theta \\ r_\theta/2 & \text{if } |a - \theta| = 1 \\ 0 & \text{otherwise,} \end{cases}$$

where $\min_\theta r_\theta > \max_\theta r_\theta/2 > 0$. The adjacency graph for this problem is again complete. In this case, however, X is not aligned with u . Since the linear independence condition in Theorem 3 holds, X is not incentivizable.

6.2. General necessary conditions. A set of actions $B \subseteq A$ is *cycle-rich* if it contains at least four elements and, for any proper subset $B' \subset B$ with at least three elements, there exists $a \in B \setminus B'$ such that

$$\bigcap \{V_C : C \text{ is internally independent, } a \in C, \text{ and } |C \cap B'| \geq 2\} = \{0\}.$$

The intersection above goes over all internally independent cycles that contain a and at least two elements of B' .

Example 14. Consider a variant of Example 13 in which there is an additional safe action. Thus $A = \Theta \cup \{a_s\}$ with

$$u(a, \theta) = \begin{cases} r_\theta & \text{if } a = \theta \\ s & \text{if } a = a_s \\ 0 & \text{otherwise.} \end{cases}$$

Suppose in addition that $\Theta = \{\theta_0, \theta_1, \theta_2, \theta_3\}$, with $r_\theta = 1/2$ for $\theta = \theta_0, \theta_1$ and $r_\theta = 1$ for $\theta = \theta_2, \theta_3$. Let $s = 3/10$. Then the adjacency graph—which is depicted in Figure 6.1—is incomplete since actions θ_0 and θ_1 are not adjacent, but A is cycle-rich.

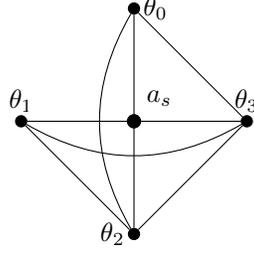


FIGURE 6.1. Adjacency graph for Example 14.

Theorem 4. *Suppose $B \subseteq A$ is cycle-rich. If X is incentivizable, then it is aligned with u on B .*

Proof. We begin with the following observation.

Lemma 5. *If X is non-trivially aligned with u on B , then for any $a, b \in B$, if $\bar{X}(a)$ and $\bar{X}(b)$ are collinear, they are also collinear with Δ_a^b .*

Proof. Let $a, b \in B$ be such that $\bar{X}(a)$ and $\bar{X}(b)$ are collinear and let $\gamma(\cdot) \neq 0$ and d be such that $\bar{X}(a) = \gamma(a)(\bar{u}(a) + d)$ and $\bar{X}(b) = \gamma(b)(\bar{u}(b) + d)$. By the collinearity assumption, there exists $\alpha \neq 0$ such that

$$\bar{u}(a) + d = \alpha(\bar{u}(b) + d).$$

Because $\bar{u}(a) \neq \bar{u}(b)$, it must be that $\alpha \neq 1$. It follows that $d = \frac{1}{1-\alpha}(\alpha\bar{u}(b) - \bar{u}(a))$ and

$$\frac{1}{\gamma(a)}\bar{X}(a) = \bar{u}(a) + \frac{1}{1-\alpha}(\alpha\bar{u}(b) - \bar{u}(a)) = \frac{\alpha}{1-\alpha}(\bar{u}(b) - \bar{u}(a)),$$

as needed. \square

Cycle-richness implies that there is some $a \in B$ and a collection of internally independent cycles \tilde{C} with $a \in \tilde{C} \subseteq B$ for which the intersection of the spaces $V_{\tilde{C}}$ is $\{0\}$. Thus either $\bar{X}(a) = 0$ or $\bar{X}(a) \notin V_C$ for some such cycle C . By Lemma 3, X is aligned with u on C .

Let $B' \subseteq B$ be a subset of B of maximal cardinality on which X is aligned with u . By the above argument, B' has at least three elements. Suppose for contradiction that $B' \neq B$. By the same argument as in the preceding paragraph, cycle-richness implies that there exists $a \in B \setminus B'$ and a cycle C containing a such that $|C \cap B'| \geq 2$ and either $\bar{X}(a) = 0$ or $\bar{X}(a) \notin V_C$. By Lemma 3, X is aligned with u on C .

If there exists a pair of distinct actions $a_0, a_1 \in C \cap B'$ such that $\bar{X}(a_0)$ and $\bar{X}(a_1)$ are not collinear, Lemma 4 implies that X is aligned with u on $C \cup B'$, contradicting the maximality of B' .

From now on, suppose that $a_0, a_1 \in C \cap B'$ are distinct actions such that $\bar{X}(a_0)$ and $\bar{X}(a_1)$ are collinear.

If $\bar{X}(a_0)$ and $\bar{X}(a_1)$ are not collinear with $\Delta_{a_0}^{a_1}$, then Lemma 5 implies that the alignment with \mathbf{u} on C and that on B' must both be trivial, which further implies that, for all $b \in C \cup B'$, $\bar{X}(b)$ is collinear with $\bar{X}(a_0)$. Thus, \mathbf{X} is trivially aligned with \mathbf{u} on $C \cup B'$, contradicting the maximality of B' .

If $\bar{X}(a_0)$ and $\bar{X}(a_1)$ are collinear with $\Delta_{a_0}^{a_1}$, then, by Lemma 2, $\bar{X}(a) \in \text{span}(\Delta_a^{a_0}, \Delta_{a_0}^{a_1}) = V_C$. The choice of cycle C implies that $\bar{X}(a) = 0$. Another application of Lemma 2 shows that $\bar{X}(a_0)$ is collinear with $\Delta_a^{a_0}$. But the latter contradicts collinearity with $\Delta_{a_0}^{a_1}$ due to the independence assumptions. The contradiction finishes the proof of the Theorem. \square

Proof of Theorem 3. It suffices to show that the set of all actions is cycle-rich; the result then follows from Theorem 4.

Take any proper subset $B \subset A$ with at least three actions $b_0, b_1, b_2 \in B$ and let $a \in A \setminus B$. Consider cycles $C_i = B \setminus \{b_i\} \cup \{a\}$. Then, $V_{C_i} = \text{span}\{\Delta_a^{b_j} : j \neq i\}$. The independence assumption implies that $\bigcap_i V_{C_i} = \{0\}$. \square

Theorem 5. *Let $\{A_0, \dots, A_k\}$ be a splitting collection such that, for each l , either A_l is cycle-rich or it contains exactly two elements. If \mathbf{X} is incentivizable, then it is piecewise aligned with \mathbf{u} .*

Proof. If A_l is cycle-rich, then \mathbf{X} is aligned with \mathbf{u} on A_l by Theorem 4. If A_l contains exactly two elements, then \mathbf{X} is aligned with \mathbf{u} on A_l by Lemma 2. \square

7. PRODUCT PROBLEMS

In many experiments, subjects perform a sequence of tasks. These tasks may be identical or they may differ; the subject's payoff is a sum or weighted average payoff of the payoffs in the various tasks, as in the common experimental design in which one task is randomly selected for payment (see Charness, Gneezy, and Halladay (2016) and Azrieli, Chambers, and Healy (2018) and the references therein). In such cases, the researcher may be interested in eliciting beliefs related to the entire sequence of actions chosen by the subject. For example, a student may solve a test with multiple questions—with their payoff being equal to the score—and the researcher may want to ask the subject what she believes about her overall performance on the test. Alternatively, to gauge the impact of learning across repetitions of the same task, the researcher may want to elicit subjects' beliefs about their change performance between the beginning and the end of the experiment.

To formalize this idea, we define a *product problem* (Θ, A, \mathbf{u}) to be a decision problem in which $\Theta = \times_i \Theta_i$, $A = \times_i A_i$, and

$$u(a, \theta) = \sum_i u_i(a_i, \theta_i)$$

for some sets $(\Theta_1, \dots, \Theta_I)$ and (A_1, \dots, A_I) , and some functions $u_i : A_i \times \Theta_i \rightarrow \mathbb{R}$. As noted above, the additive separability of u captures commonly used incentives in which one choice is randomly selected for payment. We refer to each (Θ_i, A_i, u_i) as a *task*. We write a_{-i} for a profile

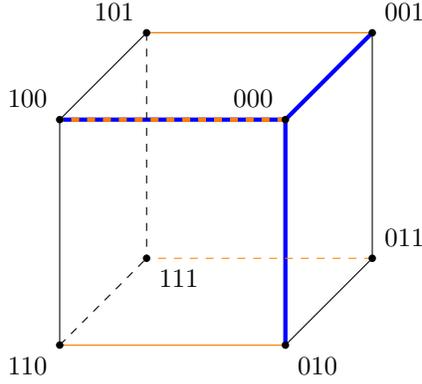


FIGURE 7.1. Adjacency graph for Example 15 with $\Omega = \{0, 1\}$ and $I = 3$.

$(a_1, \dots, a_{i-1}, a_{i+1}, \dots, a_I)$ and $a_i a_{-i}$ for the profile whose i th coordinate is a_i and remaining coordinates are given by a_{-i} .

Example 15. The decision problem is a test consisting of $I \geq 3$ multiple choice questions. The state space and action space are given by $\Theta = A = \Omega^I$, where Ω is a finite set containing at least two elements describing the possible answers to any given question. Coordinate i corresponds to the i th question: θ_i is the correct answer to question i and a_i is the DM's answer to question i . The payoff in the decision problem is the score on the test:

$$u(a, \theta) = \sum_{i=1}^I \mathbb{1}\{\theta_i = a_i\}.$$

The DM has a belief $p \in \Delta(\Theta)$. The optimal choice of action in each task i is the most likely state according to the marginal distribution of p over Θ_i .

Note that we make no assumptions about the correlation among states across tasks: the DM can hold any belief about the joint distribution of $(\theta_1, \dots, \theta_I)$. In particular, the DM need not view the states as independent, nor must there be a fixed state across tasks.

Product problems share a structure that distinguishes them from the problems we have analyzed so far. Two actions $a, b \in A$ are adjacent only if they differ in exactly one task, that is, if there is some i such that $a_i \neq b_i$ and $a_{-i} = b_{-i}$. Conversely, if a_i and b_i are adjacent in task i , then the product actions $a_i a_{-i}$ and $b_i a_{-i}$ are adjacent for all a_{-i} . The adjacency graph of the product problem is typically neither complete (even if the adjacency graphs in each task are complete) nor a tree.

Figure 7.1 depicts one example. Note that the necessary conditions for incentivizability in Theorems 4 and 5 do not apply due to an absence of internally independent cycles. While the collection of payoff difference vectors associated with edges exiting a single node—such as the blue edges Δ_{000}^{100} , Δ_{000}^{010} , and Δ_{000}^{001} —are linearly independent, no cycle is internally independent because parallel edges

correspond to identical payoff difference vectors. Thus, for instance, the orange edges correspond to $\Delta_{000}^{100} = \Delta_{010}^{110} = \Delta_{001}^{101} = \Delta_{011}^{111}$.

We say that question X depends on task i *trivially* if, for each a_{-i} , the vectors $\bar{X}(a_i a_{-i})$ are collinear for all a_i . The following result provides necessary and sufficient conditions for incentivizability in product problems.

Theorem 6. *Let $(\Theta_i, A_i, u_i)_{i=1}^I$ be a product problem with $I \geq 3$. Suppose that for each i , either A_i contains only two actions, or the adjacency graph for problem (Θ_i, A_i, u_i) is complete and the vectors $\{\Delta_{a_i}^{b_i}, \Delta_{a_i}^{c_i}\}$ are linearly independent for all $a_i, b_i, c_i \in A_i$. Suppose in addition that there are at least three tasks on which X does not depend on trivially. Then X is incentivizable if and only if there exist $v(a), \kappa(a) \in \mathbb{R}$ with $v(a) \neq 0$ for each $a \in A$, $\tau_i \in \mathbb{R}$ for each i , and $d \in \mathbb{R}^\Theta$ such that*

$$(7.1) \quad X(a, \theta) = \kappa(a) + v(a) \left(d(\theta) + \sum_i \tau_i u_i(a_i, \theta_i) \right)$$

for all a and θ .

Theorem 6 requires that the question non-trivially depends on at least three tasks. If instead the question depends non-trivially on only one problem, it is straightforward to adapt the analysis from Section 6.1 to apply here. In the intermediate case in which the question depends non-trivially on exactly two tasks, we do not know whether the conclusion of Theorem 6 holds.

Along the same lines as Theorem 3, we assume that each task has a complete adjacency graph. However, relative to Theorem 3, the other requirements for each task are significantly weaker: there are no restrictions on the number of actions, and we require linear independence only of pairs of payoff difference vectors instead of triples.

The characterization of incentivizable questions in (7.1) is more permissive than that in Theorem 3. If τ_i is constant across i , the expression in (7.1) implies that X is aligned with u . However, in contrast to the case of a single problem with complete adjacency, there are many questions not aligned with u that are also incentivizable (namely, those for which τ_i varies across i). The additional freedom in the product problem results from a smaller number of cycles and a larger number of linear dependencies in the payoffs.

The following example illustrates the added flexibility afforded by (7.1).

Example 16 (Example 15 continued). It is straightforward to verify that in the product problem of Example 15, if $|\Omega| > 2$, the adjacency graph for each problem i is complete and all pairs $\{\Delta_{a_i}^{b_i}, \Delta_{a_i}^{c_i}\}$ are linearly independent. Therefore, Theorem 6 applies to any question that depends non-trivially on at least three tasks.

Suppose the test consists of two parts: the first part comprises questions 1 through I_1 , while the second part comprises questions $I_1 + 1$ through I . Consider the question that asks the DM about the

expected improvement in her average score from the first part of the test to the second:

$$X(a, \theta) = \sum_{i=1}^I \zeta_i \mathbb{1}(a_i = \theta_i) \text{ where } \zeta_i = \begin{cases} -\frac{1}{I_1} & \text{if } i \leq I_1 \\ \frac{1}{I-I_1} & \text{if } i > I_1 \end{cases}$$

This question is of the form given in (7.1), and is therefore incentivizable by Theorem 6. At first glance, this result may be surprising as X seems to create opposing incentives in the two parts of the test. Nevertheless, X can be incentivized using a simple modification of the BDM elicitation mechanism from Proposition 1 that involves adding the payoff from the decision problem:

$$\begin{aligned} V(r, a, \theta) &= \int_0^r X(a, \theta) dx + \int_r^1 x dx - \frac{1}{2} + \sum_i u_i(a_i, \theta_i) \\ &= \sum_i (1 + r\zeta_i) u_i(a_i, \theta_i) - \frac{r^2}{2}. \end{aligned}$$

The two integral terms provide incentives for truthful reporting of r . On their own, since ζ_i is negative for some i , these terms distort the incentives in the original decision problem and change the optimal choice of a . Adding the final sum restores the correct incentives since $1 + r\zeta_i$ is positive for each i .

The next example illustrates the restrictiveness of (7.1).

Example 17 (Example 15 continued). Let $x \in \{1, \dots, I\}$. The researcher would like to elicit the probability the DM assigns to receiving a score of at least x ,¹⁰ which corresponds to the question

$$X(a, \theta) = \begin{cases} 1 & \text{if } \sum_{i=1}^I \mathbb{1}\{\theta_i = a_i\} \geq x \\ 0 & \text{otherwise.} \end{cases}$$

We claim that X is not incentivizable. Since there is no problem on which X depends trivially, Theorem 6 applies. Thus it suffices to show that X is not of the form specified in (7.1). Suppose for contradiction that it is. Notice that there exist at least two distinct scores on the test that are either both below x or both at least x . It follows that, given any a and i , there exist θ and θ' such that $a_i = \theta_i \neq \theta'_i$, $\theta_j = \theta'_j$ for all $j \neq i$, and $X(a, \theta) = X(a, \theta')$. From (7.1), since $v(a) \neq 0$, it follows that $d(\theta) + \tau_i = d(\theta')$. Applying the same argument to the action a' defined by

$$a'_j = \begin{cases} \theta'_i & \text{if } j = i \\ a_j & \text{otherwise,} \end{cases}$$

we obtain $d(\theta') + \tau_i = d(\theta)$. Therefore, $\tau_i = 0$ for all i . Since $X(a, \theta) = X(a, \theta')$, it now follows that $d(\theta) = d(\theta')$. Letting a'' be any action that agrees with θ on exactly x coordinates, including coordinate i , we have $X(a'', \theta) = 1 \neq 0 = X(a'', \theta')$, implying that $d(\theta) \neq d(\theta')$, a contradiction.

¹⁰This question is similar to that of Möbius et al. (2022), who elicit the subject's belief that their score on an incentivized IQ test is above the median among the participants.

7.1. Sketch of proof. The proof of Theorem 6 can be found in Appendix A. Here, we describe the key ideas.

The proof that (7.1) is sufficient for incentivizability is relatively straightforward: the argument directly extends the construction in Example 16.

For the remainder of this section, we assume that question X is incentivizable. By the Adjacency Lemma (Lemma 2), for any pair of adjacent actions a and b , there are coefficients $x(a, b) \neq 0$ and $y(a, b)$ such that $\bar{X}(a) = x(a, b)\bar{X}(b) + y(a, b)\Delta_a^b$. For the purpose of this discussion, we assume that the coefficients $x(a, b)$ and $y(a, b)$ are uniquely defined.

The key step in the proof is to show that each cycle $a_0, \dots, a_n = a_0$ of adjacent actions is *exact*, that is, that

$$x(a_0, a_1)x(a_1, a_2) \cdots x(a_{n-1}, a_n) = 1.$$

We first explain the connection between exactness of all cycles and (7.1), and then explain why all cycles are exact.

Fix an arbitrary action a_0 . The exactness of all cycles implies that we can define

$$v(a) = x(a_0, a_1) \cdots x(a_{n-1}, a_n)$$

for any path $a_0, \dots, a_n = a$ of adjacent actions, and the definition does not depend on the chosen path.¹¹ In particular, for any two adjacent actions, $v(a) = x(a, b)v(b)$. Letting $\bar{X}^*(a) = \frac{1}{v(a)}\bar{X}(a)$ and $y^*(a, b) = \frac{1}{v(a)}y(a, b)$, the Adjacency Lemma implies that, for any two adjacent actions a and b ,

$$(7.2) \quad \bar{X}^*(a) = \bar{X}^*(b) + y^*(a, b)\Delta_a^b.$$

To illustrate how (7.2) implies (7.1), consider Example 15. Let a and b be adjacent actions, and recall that adjacency implies there must be some task i such that $a_{-i} = b_{-i}$. Notice that $\Delta_a^b = \Delta_a^{b_i a_{-i}} = \Delta_{a_i}^{b_i}$, where the last vector refers to the payoff difference $\bar{u}_i(b_i) - \bar{u}_i(a_i)$ in task i . We will argue that the coefficient $y^*(a, b)$ depends only on a_i and b_i , i.e., that $y^*(a, b) = y^*(a_i a'_{-i}, b_i a'_{-i})$ for all a'_{-i} . Indeed, for any two problems i, j , an application of (7.2) to two paths $a, b_i a_{-i}, b_i b_j a_{-ij}$ and $a, b_j a_{-j}, b_i b_j a_{-ij}$ yields

$$\begin{aligned} y^*(a, b_i a_{-i})\Delta_a^{b_i a_{-i}} + y^*(b_i a_{-i}, b_i b_j a_{-ij})\Delta_{b_i a_{-i}}^{b_i b_j a_{-ij}} \\ = y^*(a, b_j a_{-j})\Delta_a^{b_j a_{-j}} + y^*(b_j a_{-j}, b_i b_j a_{-ij})\Delta_{b_j a_{-j}}^{b_i b_j a_{-ij}}. \end{aligned}$$

Since $\Delta_a^{b_i a_{-i}} = \Delta_{b_j a_{-j}}^{b_i b_j a_{-ij}} = \Delta_{a_i}^{b_i}$ and $\Delta_a^{b_j a_{-j}} = \Delta_{b_i a_{-i}}^{b_i b_j a_{-ij}} = \Delta_{a_j}^{b_j}$, this last equation simplifies to

$$(y^*(a, b_i a_{-i}) - y^*(b_j a_{-j}, b_i b_j a_{-ij}))\Delta_{a_i}^{b_i} + (y^*(b_i a_{-i}, b_i b_j a_{-ij}) - y^*(a, b_j a_{-j}))\Delta_{a_j}^{b_j} = 0.$$

¹¹The term “exact” is borrowed from differential geometry, where an exact vector field is a gradient of a scalar function.

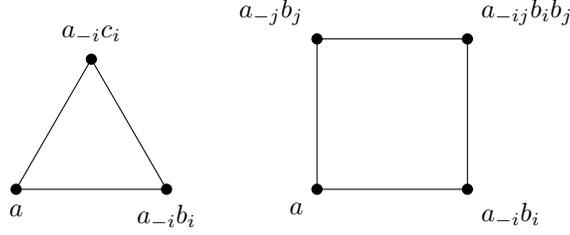


FIGURE 7.2. Two basic types of cycles in product problems.

Both coefficients must be equal to 0 since $\Delta_{a_i}^{b_i}$ and $\Delta_{a_j}^{b_j}$ are linearly independent; in particular,

$$y^*(a, b_i a_{-i}) = y^*(b_j a_{-j}, b_i b_j a_{-ij}).$$

Applying this equation repeatedly to change a_{-i} to any a'_{-i} one component at a time shows that $y^*(a, b_i a_{-i})$ depends only on a_i and b_i .

Now consider Example 15 with $\Omega = \{0, 1\}$. Let $y_i = y^*(0_i a_{-i}, 1_i a_{-i})$ for any a_{-i} . (By the previous observation, y_i is independent of the choice of a_{-i} .) Taking $b = 0$ in (7.2), we obtain

$$\begin{aligned} \bar{X}^*(a) &= \bar{X}^*(0) + \sum_{i \text{ s.t. } a_i=1} y_i \Delta_{0_i}^{1_i} \\ &= \bar{X}^*(0) - \sum_i y_i \bar{u}_i(0_i) + \sum_i y_i \bar{u}_i(a_i). \end{aligned}$$

Taking $d(\theta) = \bar{X}^*(0) - \sum_i y_i \bar{u}_i(0_i)$, we obtain an expression of the form in (7.1).

This argument is not specific to Example 15: it extends directly to all product problems with binary actions in every task. Extending beyond binary actions requires more care and conditions on linear independence of payoffs within tasks. We leave the details to the formal proof.

We now return to the question of why all cycles are exact. The first step is to notice that it is sufficient to establish the exactness of two types of cycles, depicted in Figure 7.2. Cycles of the first type consist of three actions that differ only in choices within the same task. Those of the second type consist of four actions that differ in choices in two tasks. Under the assumptions of Theorem 6, one can show that all cycles can, in a sense, be decomposed into cycles of these two types whose exactness implies exactness of the original cycle.

In Example 15, only cycles of the second type exist. Consider the four-action cycle corresponding to the front face in Figure 7.1. A repeated application of the Adjacency Lemma (Lemma 2) along the cycle leads to

$$\bar{X}(000) = x(000, 100)x(100, 110)x(110, 010)x(010, 000)\bar{X}(000) + s_1 \Delta_{0_1}^{1_1} + s_2 \Delta_{0_2}^{1_2}$$

for some $s_1, s_2 \in \mathbb{R}$. If $\bar{X}(000)$ is not in the subspace spanned by vectors $\Delta_{0_1}^{1_1}$ and $\Delta_{0_2}^{1_2}$, then the product of the x coefficients on the right-hand side must be equal to 1, meaning that the cycle is exact. Otherwise, by analyzing a number of cases, one can show that the cycles corresponding to the other five faces in Figure 7.1 are all exact. These five cycles can be combined in such a way that all edges not belonging to the front face “cancel out,” thereby implying exactness of the original cycle.

8. DISCUSSION

8.1. Three states. We have assumed throughout that there are at least four states. Our necessary conditions for incentivizability rely on this assumption insofar as it ensures that, for any two adjacent actions a and b , intersections of the set of beliefs at which the DM is indifferent between a and b with the level sets of $X(a, \cdot)$ or $X(b, \cdot)$ have dimension at least one.

If there are only two states, the problem becomes trivial as the DM’s belief about the state can be elicited independent of the action, which is sufficient for the researcher to determine the expectation of any question X . With three states, eliciting the entire belief may still be a practical option as it requires asking for only two probabilities.

If the researcher wants to ask for only one number in a problem with three states, although our necessary conditions no longer apply, looking at adjacencies can nonetheless be useful for understanding incentivizability. Suppose the overall payoff is a weighted sum of the payoff from the decision problem and that from a scoring rule applied at the belief elicitation stage. Suppose moreover that the scoring rule depends only on the reported belief and the realized value of $X(a, \theta)$, and not on θ directly.¹² At any belief at which the DM is indifferent between two actions, the value of truthful reporting at the belief elicitation stage must be equal following these two actions. Depending on the structure of the decision problem, following these constant values along cycles of adjacent actions can imply restrictions on X ; see Figure 8.1.

8.2. Independent questions. Our necessary conditions make use of independence assumptions on the payoffs in the decision problem. Similar results can be obtained if one replaces these assumptions with assumptions about independence of \bar{X} across actions. For example, along the lines of Lemma 3, if X is incentivizable and the set of questions $\bar{X}(a)$ is linearly independent for actions a in some cycle C of adjacencies, then one can show that X must be aligned with u on C . Lemma 4 can then be used to obtain necessary conditions on the full set of actions.

8.3. Non-affine questions. We have restricted attention to eliciting beliefs about the expectation of some function $X(a)$. Lambert (2019) studies elicitation of “properties” of beliefs, where a property

¹²This is a natural restriction that we do not impose in our model; doing so would have no effect on our results. When there are three states, we do not know whether this restriction has any bite; we expect that the convexity of the value function would place restrictions on incentivizable questions even if the value can depend on θ directly.

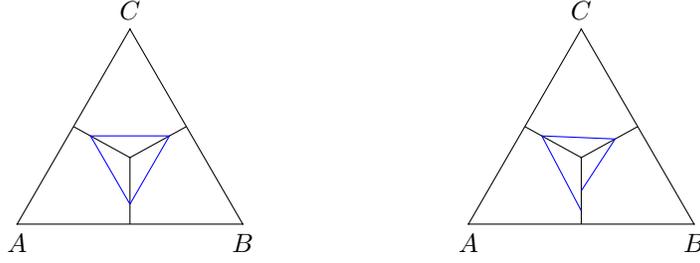


FIGURE 8.1. Problems with three states. Each triangle depicts the simplex of beliefs. The black line segments illustrate the partition of the simplex according to which action is optimal. Blue line segments represent sets of beliefs on which $X(a^*)$ is constant for the optimal action a^* . In the left triangle, because the blue lines form a triangle, adjacency considerations do not rule out incentivizability of X . In the right triangle, X is not incentivizable if the payment for elicitation depends only on the value of X and the reported belief.

corresponds to a partition of the simplex. He characterizes which properties are incentivizable, i.e., for which ones there exists a scoring rule incentivizing truthful reporting when the DM is asked only about the property associated with her belief. A question $X(a)$ in our framework corresponds to a property that partitions the simplex into parallel hyperplanes (unless the question is trivial, in which case there is a single property for the entire simplex). This formulation captures many properties of interest and ensures incentivizability in the absence of an additional decision problem, or if the question is independent of the action choice. There are, however, properties—such as the median of some X —which may be of interest that are incentivizable in Lambert’s context but to which our results do not apply. Nonetheless, we expect our general approach of focusing on adjacencies between actions to be useful for such non-affine questions.

8.4. Multi-dimensional question. In our model, we assume that the researcher can only ask a single question. Here, we show how our methods can be extended to multiple questions. For simplicity, we focus on the case of two questions; the logic extends directly to more than two questions.

We say that questions $X, Y : A \rightarrow \mathbb{R}^\Theta$ are *jointly incentivizable* if there exists an elicitation method $V : \mathbb{R}^2 \times A \times \Theta \rightarrow [0, 1]$ such that, for every $p \in \Delta(\Theta)$,

$$\arg \max_{a, r, s} \mathbb{E}_p V(r, s, a, \theta) \subseteq \left\{ (\mathbb{E}_p X(a; \theta), \mathbb{E}_p Y(a; \theta), a) : a \in \arg \max_{b \in A} \mathbb{E}_p u(b; \cdot) \right\}.$$

If X and Y are both aligned with u , it is straightforward to extend Proposition 1 to show that they are jointly incentivizable.

For necessary conditions, our key result—Lemma 2—extends as follows.

Lemma 6. *Suppose that X and Y are jointly incentivizable. If actions a and b are adjacent, then there are ρ_X, ρ_Y and σ_x^y for $x, y = X, Y$, not all equal to 0, such that*

$$\begin{aligned} \bar{X}(b) &= \rho_X (\bar{u}(b) - \bar{u}(a)) + \sigma_X^X \bar{X}(a) + \sigma_X^Y \bar{Y}(a) \\ \text{and } \bar{Y}(b) &= \rho_Y (\bar{u}(b) - \bar{u}(a)) + \sigma_Y^X \bar{X}(a) + \sigma_Y^Y \bar{Y}(a). \end{aligned}$$

If $\bar{X}(a)$ or $\bar{X}(b)$ is collinear with $\bar{u}(b) - \bar{u}(a)$, then we can take $\rho \neq 0$.

The proof, which we omit, follows the same reasoning as that of Lemma 2. We leave to future research the details of how to use this lemma to identify precise conditions for joint incentivizability.

8.5. Robust scoring rules. We have assumed throughout that the researcher knows the utility function in the decision problem. This assumption is reasonable in many lab experiments, but questionable in other settings—such as field experiments—in which a researcher may want to elicit beliefs. Suppose instead that the researcher has in mind a set of possible utility functions, and requires that questions be incentivizable for every utility function in that set. For questions that (as in our model) are independent of the utility function, this requirement is typically very demanding. For instance, under the conditions of Theorem 3, X must be aligned with every u . Notice, however, that some questions are naturally formulated in a way that depends on the utility function; the expected ex post regret of Example 2 is one such question. Allowing X to depend on u accommodates these questions. Our results can then be applied separately for each utility function in the set considered by the researcher to determine whether a question is incentivizable.

APPENDIX A. PROOF OF THEOREM 6

This Appendix is divided into the following subsections. Section A.1 shows that each incentivizable question in the product problem can be decomposed into linearly independent vectors that correspond to different tasks. We use this decomposition together with Lemma 2 to derive restrictions on questions for adjacent actions. Subsections A.2 to A.5 show that all adjacency cycles are exact. Section A.3 develops useful tools, and Sections A.4 and A.5 deal with different classes of cycles. Section A.6 concludes the proof.

A.1. Decomposition. Assume throughout that X is incentivizable.

For each i and each $t_i \in \Theta_i$, let $\mathbf{e}_i(t_i) \in \mathbb{R}^\Theta$ be the vector such that, for each $\theta \in \Theta$, $\mathbf{e}_i(\theta|t_i) = \mathbf{1}\{\theta = t_i\}$. Let $E_i = \text{span}\{\mathbf{e}_i(t_i) : t_i \in \Theta_i\}$ be the linear subspace spanned by such vectors. Notice that we can interpret

$$\Delta_{a_i}^{b_i} = \bar{u}_i(a_i) - \bar{u}_i(b_i) = \sum_{t_i} (\bar{u}_i(t_i|a_i) - \bar{u}_i(t_i|b_i)) \mathbf{e}_i(t_i)$$

for each i , a_i , and b_i as a vector in the subspace E_i . Let E_0 be a complementary space to the sum of E_1 through E_I . Note that the subspaces E_i are linearly independent.

For each a , $\bar{X}(a)$ admits a unique decomposition

$$(A.1) \quad \bar{X}(a) = \sum_{i=0}^I \mathbf{w}_i(a)$$

with $\mathbf{w}_i(a) \in E_i$ for all i . Note that the vectors $\mathbf{w}_0(a), \mathbf{w}_1(a), \dots, \mathbf{w}_I(a)$ are linearly independent.

Take any a, b , and i , such that $a_{-i} = b_{-i}$. By Lemma 2, there are $x(a, b) \neq 0$ and $y(a, b)$ such that $\bar{X}(a) = x(a, b)\bar{X}(b) - y(a, b)\Delta_b^a$. Using the above decomposition, we have

$$0 = [\mathbf{w}_i(a) - x(a, b)\mathbf{w}_i(b) - y(a, b)\Delta_{b_i}^{a_i}] + \sum_{j \neq i} [\mathbf{w}_j(a) - x(a, b)\mathbf{w}_j(b)] + [\mathbf{w}_0(a) - x(a, b)\mathbf{w}_0(b)].$$

The proof of Lemma 2 shows that $x(a, b)$ is uniquely defined if and only if $\bar{X}(a) \neq 0$ and $\bar{X}(a)$ is not collinear with $\Delta_{b_i}^{a_i}$. When this is not the case, we say that the transition (a, b) is *free*. The values of $x(a, b)$ for free transitions are carefully chosen below. Our choice always satisfies $x(a, b)x(b, a) = 1$ (which always holds for non-free transitions). Let $x(a, a) = 1$.

Because all vectors in square brackets form a linearly independent system, all these vectors must be equal to 0:

$$(A.2) \quad \mathbf{w}_i(a) - x(a, b)\mathbf{w}_i(b) = y(a, b)\Delta_{b_i}^{a_i},$$

$$(A.3) \quad \mathbf{w}_j(a) - x(a, b)\mathbf{w}_j(b) = 0 \text{ for each } j \neq i,$$

$$(A.4) \quad \mathbf{w}_0(a) - x(a, b)\mathbf{w}_0(b) = 0.$$

Lemma 7. *For each a , there exist $\gamma_i(a) \neq 0$ and vectors $\mathbf{w}_i^*(a_i)$ for each $i = 1, \dots, I$ and $\gamma_0(a) \neq 0$ and a vector \mathbf{w}_0^* such that*

$$\mathbf{w}_i(a) = \gamma_i(a)\mathbf{w}_i^*(a_i) \text{ for each } i, \text{ and } \mathbf{w}_0(a) = \gamma_0(a)\mathbf{w}_0^*.$$

Proof. For the first claim, fix an action a^* . For each a_i , let $\mathbf{w}_i^*(a_i) = \mathbf{w}_i(a_{-i}^* a_i)$. For each a , fix an arbitrary a path of adjacent actions $a^0 = a_{-i}^* a_i, \dots, a^n = a$ such that for each $l < n$, $a_i^l = a_i$. A repeated application of (A.3) shows that

$$\mathbf{w}_i(a) = x(a^n, a^{n-1}) \cdots x(a^1, a^0)\mathbf{w}_i^*(a_i).$$

Let $\gamma_i(a) = x(a^n, a^{n-1}) \cdots x(a^1, a^0)$.

For the second claim, fix $\mathbf{w}_0^* = \mathbf{w}_0(a^*)$ and repeatedly apply (A.4). □

This result says that, for a fixed a_i , the vectors $\mathbf{w}_i(a_{-i} a_i)$ for a_{-i} are either all collinear or all equal to 0.

Lemma 8. *For each i , one of the following is true:*

- (a) $\mathbf{w}_i^*(a_i) = 0$ for each a_i ,

- (b) *there exists $a_i^0 \in A_i$ such that $w_i^*(a_i^0) = 0$ and $w_i^*(a_i) \neq 0$ for each $a_i \neq a_i^0$,*
(c) *$w_i^*(a_i) \neq 0$ for each a_i .*

If question X does not depend trivially on problem i , then there is an action a_i such that $w_i^(a_i) \neq 0$.*

Proof. Suppose that there are three different actions a_i , b_i , and c_i such that $w^*(a_i) = w^*(b_i) = 0$ and $w^*(c_i) \neq 0$. Equation (A.2) together with Lemma 7 implies that $w_i(c_i)$ is simultaneously collinear with $\Delta_{c_i}^{a_i}$ and $\Delta_{c_i}^{b_i}$. But the latter is impossible given the independence assumption.

For the last claim, suppose that $w^*(a_i) = 0$ for each a_i . Equation (A.2) implies that $y(a, b) = 0$ for each a and b such that $a_{-i} = b_{-i}$. But then, for each a_{-i} , the vectors $\bar{X}(a_{-i}a_i)$ and $\bar{X}(a_{-i}b_i) = x(a_{-i}b_i, a_{-i}a_i)\bar{X}(a_{-i}a_i)$ are collinear. Hence, X depends on problem i trivially. \square

Because X depends non-trivially on at least three problems, Lemma 8 implies that at least three problems satisfy (b) or (c).

A.2. Exact cycles. Say that an adjacency cycle $a^0, \dots, a^n = a^0$ is *exact* if

$$x(a^0, a^1)x(a^1, a^2) \cdots x(a^{n-1}, a^n) = 1.$$

The goal of this subsection as well as subsections A.3 to A.5 is to prove the following result.

Lemma 9. *The values $x(a, b)$ for free transitions (a, b) can be chosen so that (i) $x(a, b)x(b, a) = 1$ for all adjacent a and b , and (ii) every adjacency cycle is exact.*

We begin with the following observation.

Lemma 10. *Suppose that $w_0^* \neq 0$. Then, all adjacency cycles are exact.*

Proof. The result follows from a repeated application of equality (A.4). \square

From now on, we assume that $w_0^* = 0$.

A.3. Tools. The two results in this section develop tools that we use in the subsequent analysis.

The first tool allows us to replace the question of whether a cycle is exact with an analogous question about related cycles. For each path $c = (a^0, \dots, a^{n_c})$ and each transition (a, b) between two adjacent actions, define

$$m_a^b(c) = \#\{l < n_c : a^l = a, a^{l+1} = b\} - \#\{l < n_c : a^l = b, a^{l+1} = a\}.$$

Lemma 11. *Suppose that, for some adjacency cycle c , there exists a collection D of exact adjacency cycles such that for each adjacent pair a, b , $m_a^b(c) = \sum_{d \in D} m_a^b(d)$. Then cycle c is exact.*

Proof. Let \prec be an arbitrary strict order on the set of actions A . Then,

$$\begin{aligned} \prod_{l < n_c} x(a^{c,l}, a^{c,l+1}) &= \prod_{a \prec b} (x(a, b))^{m_a^b(c)} = \prod_{a \prec b} (x(a, b))^{\sum_{d \in D} m_a^b(d)} \\ &= \prod_{d \in D} \prod_{a \prec b} (x(a, b))^{m_a^b(d)} = \prod_{d \in D} \prod_{l < n_d} x(a^{d,l}, a^{d,l+1}) = 1, \end{aligned}$$

where the last equality comes from the fact that cycles in D are exact. \square

The second result shows that it is enough to consider a particular category of “small” cycles.

Lemma 12. *Each adjacency cycle is exact if and only if the following cycles are exact:*

- (1) (a, b, a) for any two adjacent actions a and b ,
- (2) (a, b, c, a) for any three actions a , b , and c such that $a_{-i} = b_{-i} = c_{-i}$,
- (3) $(a, a_{-i}b_i, a_{-ij}b_ib_j, a_{-j}b_j, a)$ for any action a , any $i \neq j$, and any b_i and b_j .

Proof. Take any “large” cycle of adjacent actions $a = a^0, \dots, a^n = a$ and let i_l be such that, for each $l < n$, $a_{-i_l}^l = a_{-i_l}^{l+1}$. For future reference, notice that if the action in some problem i is ever changed, then it must be changed at least twice: if $i_l = i$ for some l , then there is some $l' \neq l$ such that $i_{l'} = i$. We use the “small” cycles to re-order and reduce the “large” cycle without changing the value of the product of associated the x terms:

- if $i_l > i_{l+1}$ for $l < n - 1$ we use the “small” cycle of type (3) to switch the order of the two problems, i.e., replace the cycle fragment $\dots, a^l, a^{l+1}, a^{l+2}, \dots$, where $a^{l+1} = a_{-i_l}^l a_{i_l}^{l+2}$, with $\dots, a^l, a_{-i_{l+1}}^l a_{i_{l+1}}^{l+2}, a^{l+2}, \dots$;
- if $i_l = i_{l+1}$, including $i_{n-1} = i_0$, we use either type (1) or type (2) to reduce the “large” cycle, i.e, replace the cycle fragment $\dots, a^l, a^{l+1}, a^{l+2}, \dots$ with $\dots, a^l, a^{l+2}, \dots$ in the case of a type (2) cycle or with \dots, a^l, \dots in the case of a type (1) cycle.

Consider a process in which one of the above operations is applied until it cannot be applied anymore. Because the operations either reduce the size of the cycle or they re-order problems in an increasing direction, the process never reverts and it will eventually stop. If the process stops at a single-element cycle a , then, because $x(a, a) = 1$, the original cycle must be exact.

Otherwise, the process stops with a non-trivial cycle $a = a^0, \dots, a^m = a$ for some $2 \leq m \leq n$. Then it must be that $i_l < i_{l+1} + 1$ for each $l < m$. But this contradicts the earlier observation that, in an adjacency cycle, if i appears at least once, it must appear at least twice. \square

A.4. Cycles without free transitions. Here, we consider the exactness of small cycles without free transitions. We refer to i from the definition of type (2) cycles as the *relevant problem* for this cycle; similarly, we refer to i and j as the relevant problems for type (3) cycles. We say that a type (2) or type (3) cycle $(a^0, \dots, a^n = a^0)$ is *grounded* if there exists some k such that k is not a relevant problem and $w_k^*(a_k^0) \neq 0$.

Lemma 13. *Any grounded type (2) or type (3) cycle is exact.*

Proof. Suppose that $w_k^*(a_k^0) \neq 0$ for some irrelevant problem k . Then $w_k(a^i) \neq 0$ for each action a^i in the cycle. A repeated application of (A.3) shows that for each $l \leq n$,

$$w_k(a^l) = w_k(a)x(a^0, a^1) \cdots x(a^{l-1}, a^l).$$

The result follows from the fact that $a^n = a^0$. \square

Lemma 14. *Suppose that for some a and $i \neq j$ and any b_i and b_j , $(a, a_{-i}b_i, a_{-ij}b_ib_j, a_{-j}b_j, a)$ is a type (3) “small” cycle such that either (i) $w_i^*(a_i) \neq 0$ and $w_i^*(a_i)$ is not collinear with $\Delta_{a_i}^{b_i}$, or (ii) $w_j^*(a_j) \neq 0$ and $w_j^*(a_j)$ is not collinear with $\Delta_{a_j}^{b_j}$. Then the cycle is exact.*

Proof. Using A.2 and A.3, we get

$$\begin{aligned} & x(a_{-i}b_i, a_{-ij}b_ib_j) (x(a, a_{-i}b_i)w_i(a) + x(a, a_{-i}b_i)\Delta_{a_i}^{b_i}) \\ &= x(a_{-i}b_i, a_{-ij}b_ib_j)w_i(a_{-i}b_i) \\ &= w_i(a_{-ij}b_ib_j) \\ &= x(a_{-j}b_j, a_{-ij}b_ib_j)w_i(a_{-j}b_j) + y(a_{-j}b_j, a_{-ij}b_ib_j)\Delta_{a_i}^{b_i} \\ &= x(a_{-j}b_j, a_{-ij}b_ib_j)x(a, a_{-j}b_j)w_i(a) + y(a_{-j}b_j, a_{-ij}b_ib_j)\Delta_{a_i}^{b_i}. \end{aligned}$$

Suppose without loss of generality that $w_i^*(a_i) \neq 0$ and $w_i^*(a_i)$ is not collinear with $\Delta_{a_i}^{b_i}$, which implies that $w_i(a) \neq 0$ and $w_i(a)$ is not collinear with $\Delta_{a_i}^{b_i}$. The first and the last line of the above sequence of equalities yield

$$x(a_{-i}b_i, a_{-ij}b_ib_j)x(a, a_{-i}b_i) = x(a_{-j}b_j, a_{-ij}b_ib_j)x(a, a_{-j}b_j).$$

Hence, the cycle is exact. \square

Lemma 15. *Suppose that a type (2) or type (3) “small” cycle is such that $w_i^*(a_i) \neq 0$ for each action a in the cycle and each relevant problem i . Then it is exact.*

Proof. If the cycle is grounded, the result follows from Lemma 13. Accordingly, suppose henceforth that the cycle is not grounded. Then, because of Lemma 8, and due to the assumption that at least three different problems are nontrivial, there exist non-relevant k and action $b_k \neq a_k$ such that $w_k^*(b_k) \neq 0$.

Suppose that the original cycle $(a, a_{-i}b_i, a_{-ij}b_ib_j, a_{-j}b_j, a)$ has type (3). The cycle corresponds to the orange face on the left-side of Figure A.1. Consider the type (3) cycles that are associated with all five remaining faces of the cube:

- $a, a_{-i}b_i, a_{-ik}b_ib_k, a_{-k}b_k, a$;
- $a_{-i}b_i, a_{-ij}b_ib_j, a_{-ijk}b_ib_jb_k, a_{-ik}b_ib_k, a_{-i}b_i$;

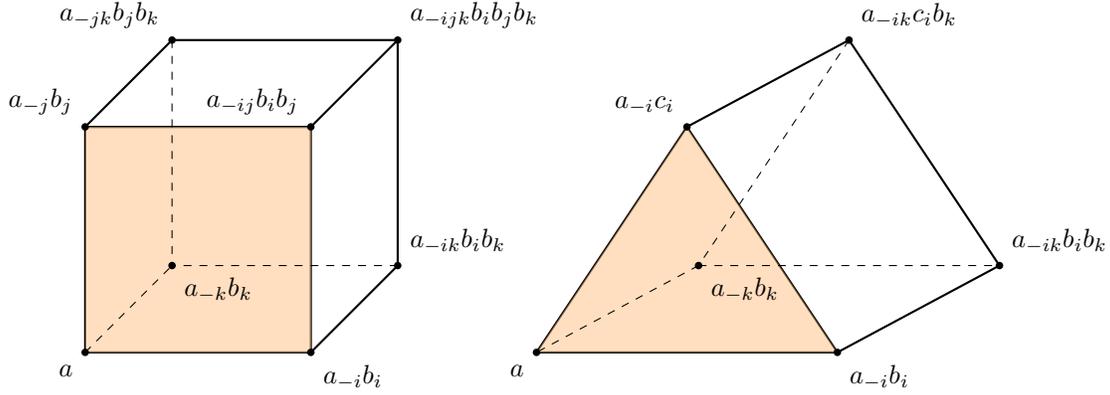


FIGURE A.1. Small cycles of type (3) and type (2)

- $a_{-j}b_j, a_{-j}b_jb_k, a_{-i}b_ib_jb_k, a_{-i}b_ib_j, a_{-j}b_j$;
- $a, a_{-k}b_k, a_{-j}b_jb_k, a_{-j}b_j, a$;
- $a_{-k}b_k, a_{-i}b_ib_k, a_{-i}b_ib_jb_k, a_{-j}b_jb_k, a_{-k}b_k$.

The first cycle corresponds to the bottom face, the second to the right face, the third to the top face, the fourth to the left face, and the fifth to the back face. All five cycles are grounded, and hence exact by Lemma 13. Moreover, the conditions of Lemma 11 are satisfied. Therefore, the original cycle is exact as well.

An analogous argument applies to type (2) cycles (see the right panel of Figure A.1). \square

Lemma 16. *If a type (2) or (3) cycle has no free transitions, then it is exact.*

Proof. By Lemmas 13 and 15, it is enough assume that the cycle is not grounded and $w_i^*(a_i) = 0$ for some relevant i and action a in the cycle. It follows that $w_k^*(a_i) \neq 0$ if and only if $k = j$, where j is the other relevant problem of the cycle.

In such a case, if the cycle were of type (2), all transitions to action a would be free.

Suppose the cycle is of type (3). Let b be the action in the cycle such that $a_{-j} = b_{-j}$. Then

$$\bar{X}(a) = \gamma_j(a)w_j^*(a_j)$$

and $\bar{X}(b) = \gamma_j(b)w_j^*(b_j) = x(a, b)\gamma_j(a)w_j^*(a_j) + y(a, b)\Delta_{a_j}^{b_j}$.

Because the transition (a, b) is not free, it must be that $w_j^*(a_j) \neq 0$ and $w_j^*(a_j)$ is not collinear with $\Delta_{a_j}^{b_j}$. Lemma 14 therefore implies that the cycle is exact. \square

A.5. Cycles with free transitions. Next, we consider cycles with free transitions.

Notice first that, if a transition between adjacent actions a and b such that $a_{-i} = b_{-i}$ is free, then it must be that $w_j^*(a_j) = 0$ for each $j \neq i$. Indeed, if, for some $\alpha \neq 0$,

$$\bar{X}(a) = \alpha \bar{X}(b) = x(a, b) \bar{X}(b) + y(a, b) \Delta_{a_i}^{b_i},$$

then $\bar{X}(b)$, and hence $\bar{X}(a)$, must be collinear with $\Delta_{a_i}^{b_i}$. Because of the linear independence of $\Delta_{a_i}^{b_i}$ and $w_j(a)$ for each $j \neq i$, it must be that $w_j(a) = 0$, and hence $w_j^*(a_j) = 0$. We refer to this property as the *test* for freeness of the transition (which provides necessary conditions). It follows that, if $j \neq i$ is non-trivial, then $a_j = a_j^0$.

In what follows, we consider two cases:

- I: There exists a single non-trivial i such that $w_i^*(a_i) \neq 0$ for all a_i . Assume without loss of generality that $i = I$. In this case, all free transitions must be between adjacent actions a and b such that $a_{-I} = b_{-I}$. Moreover, it must be that $w_I^*(a_I)$ and $w_I^*(b_I)$ are collinear with $\Delta_{a_I}^{b_I}$.

Assume without loss of generality that problem 1 is non-trivial and fix action $b_1 \neq a_1$ so that $w^*(b_1) \neq 0$. Let

$$x(a, b) = x(a, a_{-1}b_1)x(a_{-1}b_1, a_{-1I}b_1b_I)x(a_{-1I}b_1b_I, b).$$

The above definition implies that the cycle $(a, a_{-1}b_1, a_{-1I}b_1b_I, b, a)$ is exact. This cycle corresponds to the red face(s) in Figure A.5. The orange edge corresponds to the free transition. Notice that each red face cycle contains only one free transition. This is because of the test: all other transitions of the cycle either keep fixed the action in problem I or the action b_1 , and those actions are associated with non-zero $w^*(\cdot)$ vectors.

There are three types of cycles that contain free transitions, which are depicted as orange faces in Figure A.5.

The top left panel corresponds to the cycle $(a, a_{-i}b_i, a_{-iI}b_i b_I, a_{-I}b_I, a)$ when problem i is non-trivial. In this case, $b_i \neq a_i^0$, and, by Lemma 8, $w_i^*(b_i) \neq 0$. The test implies that none of the other transitions in the orange cycle are free: either action b_i or action a_I is fixed. Analogously, an application of the test shows that none of the other cycles (the uncolored faces) is free: one of the actions a_I , b_i , or b_1 is fixed. Proceeding as in the proof of Lemma 15, we see that this cycle is exact.

The top right panel corresponds to a situation when problem i is trivial. In this case, the transition $(a_{-i}b_i, a_{-iI}b_i b_I)$ is free and $x(a_{-i}b_i, a_{-iI}b_i b_I)$ can be chosen to make the cycle on the bottom face exact. None of the other transitions are free. Because the red face and the uncolored cycles are exact, the above argument implies that the orange face cycle is exact as well.

The bottom panel corresponds to the orange cycle $(a, a_{-I}b_I, a_{-I}c_I, a)$. The other transitions of the orange cycle are not free (otherwise, $w^*(a_I)$ would be collinear with $\Delta_{a_I}^{c_I}$, which

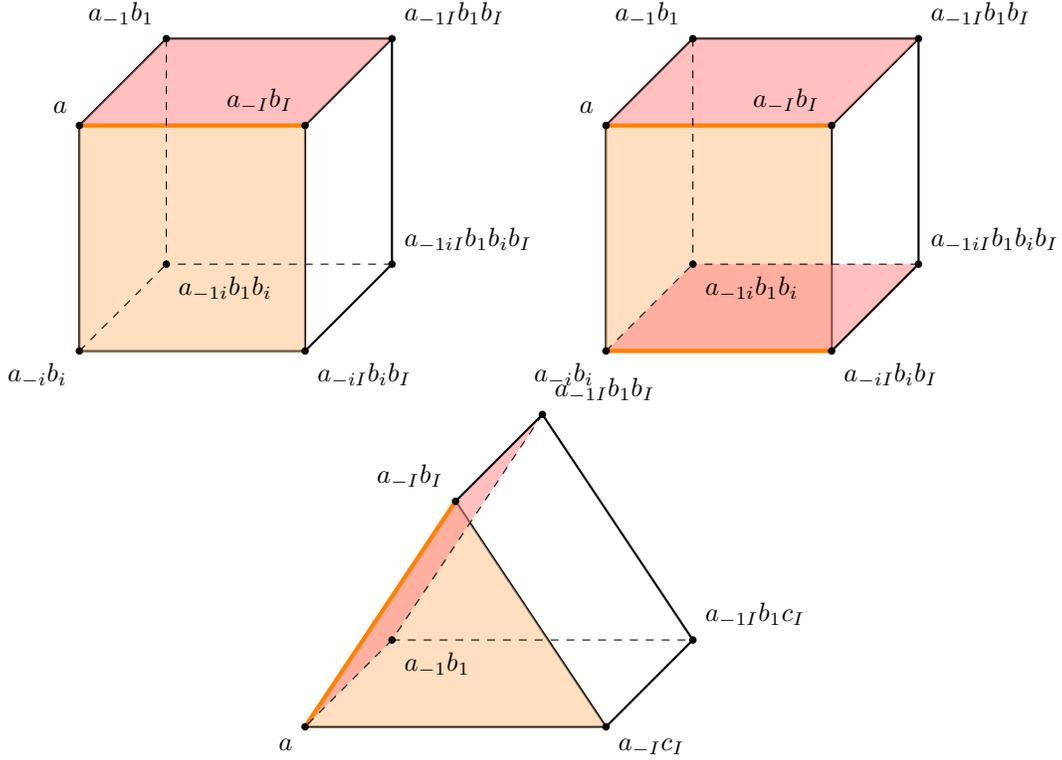


FIGURE A.2. Cycles with free transitions in case I.

would violate linear independence of the latter vector with $\Delta_{a_I}^{b_I}$). All other transitions fix one of the actions: a_I , b_I , c_I , or b_1 . Hence, due to the test, none of the remaining transitions are free. The claim follows from the same reasoning as in Lemma 15.

- II: For all non-trivial i , there exists a unique a_i^0 such that $w_i^*(a_i^0) = 0$. Let a^0 be the product problem action that consists of actions a_i^0 . Assume without loss of generality that problem $i = 1$ is non-trivial. Fix action $a_1^* \neq a_1^0$.

In this case, a transition is free if and only if it takes the form $(a_{-i}^0 a_i, a^0)$ for some i . Indeed, the above observation implies that if transition (a, b) is free, then $a_{-i} = a_{-i}^0$. Furthermore, if neither $a = a^0$ nor $b = a^0$, then both $w_i^*(a_i)$ and $w_i^*(b_i)$ must be collinear with $\Delta_{a_i}^{b_i}$. But, together with the linear independence assumption, this implies that $w_i^*(a_i)$ is not collinear with $\Delta_{a_i}^{a_i^0}$, which contradicts the fact that $x(a, a^0)\gamma_i(a)w_i^*(a_i) + y(a, a^0)\Delta_{a_i}^{a_i^0} = 0$.

For each $b_1 \neq a_1^*$, each $i \neq 1$, and each b_i , let

$$x(a_{-1}^0 a_1^*, a^0) = 1,$$

$$x(a_{-1}^0 b_1, a^0) = x(a_{-1}^0 b_1, a_{-1}^0 a_1^*),$$

$$\text{and } x(a_{-i}^0 b_i, a^0) = x(a_{-i}^0 b_i, a_{-1i}^0 a_1^* b_i) x(a_{-1i}^0 a_1^* b_i, a_{-1}^0 a_1^*).$$

By the above definition, the cycles $(a^0, a_{-1}^0 b_1, a_{-1}^0 a_1^*, a^0)$ and $(a^0, a_{-i}^0 b_i, a_{-1i}^0 a_1^* b_i, a_{-1}^0 a_1^*, a^0)$ are exact. These cycles correspond to the red faces in Figure A.5.

There are three types of cycles that contain free transitions in this case other than the cycles listed above.

First, consider a cycle $(a^0, a_{-i}^0 b_i, a_{-ij}^0 b_i b_j, a_{-j}^0 b_j, a^0)$ for $i \neq 1$. This cycle is depicted in orange on the left panel of Figure A.5. The two cycles depicted in red are exact due to the choice of the x coefficients. Finally, all of the other cycles are exact because they do not contain free transitions.

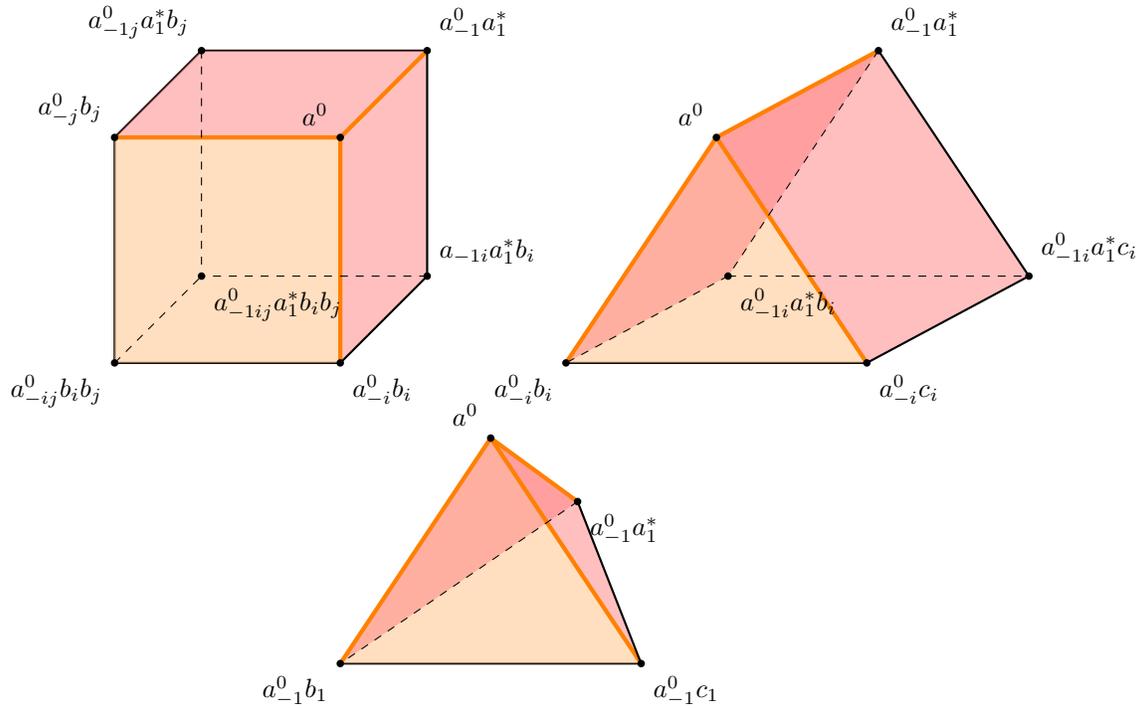


FIGURE A.3. Cycles with free transitions in case II.

Second, consider a cycle $(a^0, a_{-i}^0 b_i, a_{-i}^0 c_i, a^0)$ for $i \neq 1$. This cycle corresponds to front wall depicted in orange on the right panel of Figure A.5. The two cycles depicted in red are exact

due to the choice of the x coefficients. All of the other cycles (corresponding to the back and bottom walls) are exact because they do not contain free transitions.

Third, consider a cycle $(a^0, a_{-1}^0 b_1, a_{-1}^0 c_1, a^0)$. This cycle corresponds to the front wall depicted in orange on the bottom panel of Figure A.5. The two cycles depicted in red (corresponding to the left and right walls) are exact due to the choice of the x coefficients. The remaining cycle (corresponding to the bottom wall) is exact because it does not contain any free transitions.

Proceeding as in the proof of Lemma 15, we see that each of the considered cycles is exact.

Together with Lemma 16, this concludes the proof of Lemma 9.

A.6. Conclusion of the proof of Theorem 6. Recall that $w_i(a) = \gamma_i(a)w_i^*(a_i)$. The next result delivers additional information about the function $\gamma_i(\cdot)$.

Lemma 17. *There exist $\gamma : A \rightarrow \mathbb{R}$, $\gamma_i^* : A_i \rightarrow \mathbb{R}$, and $\gamma_0^* \in \mathbb{R}$ such that, for each a , $\gamma_i(a) = \gamma(a)\gamma_i^*(a_i)$ for each $i = 1, \dots, I$, and $\gamma_0(a) = \gamma(a)\gamma_0^*$. In addition, $x(a, b) = \frac{\gamma(a)}{\gamma(b)}$ for any two adjacent actions a and b .*

Proof. Fix an action a^* . For each a , find a path of adjacent actions $a^* = a^0, \dots, a^m = a$. Define

$$\gamma(a) = x(a^m, a^{m-1}) \cdots x(a^1, a^0).$$

Lemma 9 implies that $\gamma(a)$ is well defined in that its definition does not depend on the choice of path from a^* to a . Moreover, for any two adjacent actions a and b , if $a^* = a^0, \dots, a^m = a$ is an adjacency path from a^* to a , then $a^* = a^0, \dots, a^m, b$ is an adjacency path from a^* to b , and

$$\gamma(b) = x(b, a)x(a^m, a^{m-1}) \cdots x(a^1, a^0) = x(b, a)\gamma(a).$$

Let $\gamma_i^*(a_i) = \gamma_i(a_{-i}^* a_i) / \gamma(a_{-i}^* a_i)$. The claim for $i > 0$ follows from the fact that, for each a , if $w_i^*(a_i) \neq 0$, then

$$w_i(a) = \frac{\gamma_i(a)}{\gamma_i(a_{-i}^* a_i)} w_i(a_{-i}^* a_i) = \frac{\gamma(a)}{\gamma(a_{-i}^* a_i)} w_i(a_{-i}^* a_i).$$

A similar argument establishes the claim for $i = 0$ (see also the proof of Lemma 7). \square

Lemma 18. *There exist $y_i^* \in \mathbb{R}$ and $d_i \in E_i$ such that*

$$\gamma_i^*(a_i)w_i^*(a_i) = y_i^* \bar{u}_i(a_i) + d_i \text{ for any } a_i.$$

Proof. By (A.2), for any actions a and b such that $a_{-i} = b_{-i}$,

$$\begin{aligned}
 (A.5) \quad \gamma_i^*(a_i)\mathbf{w}_i^*(a_i) - \gamma_i^*(b_i)\mathbf{w}_i^*(b_i) &= \frac{1}{\gamma(a)}\mathbf{w}_i(a) - \frac{1}{\gamma(b)}\mathbf{w}_i(b) \\
 &= \frac{1}{\gamma(a)}(\mathbf{w}_i(a) - x(a, b)\mathbf{w}_i(b)) \\
 &= \frac{1}{\gamma(a)}y(a, b)\Delta_{b_i}^{a_i}.
 \end{aligned}$$

Because the left-hand side and $\Delta_{b_i}^{a_i}$ do not depend on a_{-i} , neither does $y(a, b)/\gamma(a)$. Let $y^*(a_i, b_i) = y(a, b)/\gamma(a)$.

If problem i has only two actions, it is easy to see that $y_i^*(a_i, b_i) = y_i^*(b_i, a_i) =: y_i^*$. The claim follows.

If problem i has at least three actions, take a , b , and c such that $a_{-i} = b_{-i} = c_{-i}$ and a_i , b_i , and c_i are distinct. Applying the above equation to pairs (a, b) , (b, c) , and (c, a) yields

$$\begin{aligned}
 y_i^*(a_i, c_i)\Delta_{b_i}^{a_i} + y_i^*(a_i, c_i)\Delta_{c_i}^{b_i} &= y_i^*(a_i, c_i)\Delta_{c_i}^{a_i} \\
 &= \gamma_i^*(a_i)\mathbf{w}_i^*(a_i) - \gamma_i^*(c_i)\mathbf{w}_i^*(c_i) \\
 &= \gamma_i^*(a_i)\mathbf{w}_i^*(a_i) - \gamma_i^*(b_i)\mathbf{w}_i^*(b_i) + \gamma_i^*(b_i)\mathbf{w}_i^*(b_i) - \gamma_i^*(c_i)\mathbf{w}_i^*(c_i) \\
 &= y_i^*(a_i, b_i)\Delta_{b_i}^{a_i} + y_i^*(b_i, c_i)\Delta_{c_i}^{b_i}.
 \end{aligned}$$

The independence assumption implies that $y_i^*(a_i, b_i) = y_i^*(b_i, c_i)$. Because the claim holds for arbitrary and distinct actions, there must be y_i^* such that for all a_i and b_i , we have $y_i^*(a_i, b_i) = y_i^*$.

Finally, fix a_i^* and take $\mathbf{d}_i = \gamma_i^*(a_i)\mathbf{w}_i^*(a_i) - y_i^*\bar{\mathbf{u}}_i(a_i^*)$. The claim follows from equation (A.5). \square

Substituting the observations from the two lemmas back into (A.1), we obtain

$$\begin{aligned}
 \bar{X}(a) &= \sum_{i=0}^n \mathbf{w}_i(a) \\
 &= \gamma(a) \left(\sum_{i=1}^n \gamma_i^*(a_i)\mathbf{w}_i^*(a_i) + \gamma_0\mathbf{w}_0^* \right) \\
 &= \gamma(a) \left(\sum_{i=1}^n y_i^*\bar{\mathbf{u}}_i(a_i) + \left[\sum_{i=1}^n \mathbf{d}_i + \gamma_0\mathbf{w}_0^* \right] \right).
 \end{aligned}$$

Let \mathbf{d} be the vector in the square brackets. The result follows.

A.7. Converse. We have shown that (7.1) is necessary for incentivizability. All that remains is to show that it is sufficient.

Notice that if $\tau_i > 0$ for all i , the product problem is equivalent (in terms of optimal choices) to a problem with payoffs

$$u(a, \theta) = \sum_i \tau_i u_i(a_i, \theta_i).$$

In this latter problem, any X satisfying (7.1) is aligned with u , and is therefore incentivizable. This in turn implies that X is incentivizable in the original problem.

If $\tau_i \leq 0$ for some i , note that, by Proposition 1, to show that X of the form described in (7.1) is incentivizable, it suffices to show that

$$X(a, \theta) = d(\theta) + \sum_i \tau_i u_i(a_i, \theta_i)$$

is, where we may assume $|\tau_i| < 1$ for all i . Letting

$$\begin{aligned} V(r, a, \theta) &= \int_0^r X(a, \theta) dx + \int_r^1 x dx - \frac{1}{2} + \sum_i u_i(a_i, \theta_i) \\ &= rd(\theta) + \sum_i (1 + r\tau_i) u_i(a_i, \theta_i) - \frac{r^2}{2}, \end{aligned}$$

the optimal choice of a is the same as in the original problem since $1 + r\tau_i > 0$ for all i and r , and the optimal choice of r is $\mathbb{E}_p[X(a, \theta)]$, as needed.

REFERENCES

- Amelio, Andrea (2022). “Cognitive uncertainty and overconfidence”. ECONtribute Discussion Paper No. 173.
- Arts, Sara, Qiyang Ong, and Jianying Qiu (2024). “Measuring decision confidence”. In: *Experimental Economics*, pp. 1–22.
- Azrieli, Yaron, Christopher P Chambers, and Paul J. Healy (2018). “Incentives in experiments: A theoretical analysis”. In: *Journal of Political Economy* 126.4, pp. 1472–1503.
- Blanco, Mariana et al. (2010). “Belief elicitation in experiments: is there a hedging problem?” In: *Experimental economics* 13, pp. 412–438.
- Butler, David J and Graham C Loomes (2007). “Imprecision as an account of the preference reversal phenomenon”. In: *American Economic Review* 97.1, pp. 277–297.
- Chambers, Christopher P and Nicolas Lambert (2021). “Dynamic belief elicitation”. In: *Econometrica* 89.1, pp. 375–414.
- Charness, Gary, Uri Gneezy, and Brianna Halladay (2016). “Experimental methods: Pay one or pay all”. In: *Journal of Economic Behavior & Organization* 131, pp. 141–150.
- Charness, Gary, Uri Gneezy, and Vlastimil Rasochoa (2021). “Experimental methods: Eliciting beliefs”. In: *Journal of Economic Behavior & Organization* 189, pp. 234–256.
- Danz, David, Lise Vesterlund, and Alistair J Wilson (2022). “Belief elicitation and behavioral incentive compatibility”. In: *American Economic Review* 112.9, pp. 2851–2883.
- Enke, Benjamin and Thomas Graeber (2023). “Cognitive uncertainty”. In: *Quarterly Journal of Economics* 138.4, pp. 2021–2067.

- Haaland, Ingar, Christopher Roth, and Johannes Wohlfart (2023). “Designing information provision experiments”. In: *Journal of Economic Literature* 61.1, pp. 3–40.
- Healy, Paul J. and Greg Leo (2024). “Belief elicitation: A user’s guide”. To appear in the Handbook of Experimental Methodology, ed. L. Yariv and E. Snowberg.
- Hossain, Tanjim and Ryo Okui (2013). “The binarized scoring rule”. In: *Review of Economic Studies* 80.3, pp. 984–1001.
- Hu, En Hua (2023). “Confidence in Inference”. Working paper.
- 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”. In: *Proceedings of the 9th ACM Conference on Electronic Commerce*, pp. 129–138.
- Möbius, Markus M. et al. (2022). “Managing self-confidence: theory and experimental evidence”. In: *Management Science* 68.11, pp. 7793–7817.
- Schlag, Karl H, James Tremewan, and Joël J Van der Weele (2015). “A penny for your thoughts: A survey of methods for eliciting beliefs”. In: *Experimental Economics* 18, pp. 457–490.
- Xiang, Yang et al. (2021). “Confidence and central tendency in perceptual judgment”. In: *Attention, Perception, & Psychophysics* 83, pp. 3024–3034.