Coarse Preference Elicitation

Yaron Azrieli, Christopher Chambers, Paul Healy*

May 12, 2019

Abstract

A planner wants to extract information about an agent's preference relation, but not necessarily the entire relation. Formally, a partition of the set of all possible orderings of alternatives (a 'type space') is given, and the planner wants to know to which partition element ('type') the agent's true preference belongs. We say that a type space is elicitable if there exists a mechanism mapping types to (possibly random) outcomes in which the agent strictly prefers truth-telling over lying. In the Savage framework a type space is elicitable if and only if it can be elicited by offering the agent a list of menus and paying one randomly-chosen choice. When the planner can use objective lotteries, more type spaces can be elicited.

Keywords: Elicitation; incentive compatibility

JEL Classification: D8, C7

1 Introduction

In mechanism design and social choice settings, a planner collects information about agents' preferences and uses this information to select a desired outcome. Often the entire preference relation is elicited, but in some settings the planner collects only partial information. For example, students in the New York City High School match are only asked to report their top 12 schools, rather than their complete ranking of hundreds of schools [1]. In that case information is truncated for simplicity. In other cases the planner simply doesn't need the entire ordering to execute her objective. As a very simple example, suppose she wants to give a single agent his most-preferred alternative (meaning, she wants to implement the dictatorial social choice function). Here she only needs to ask the agent for his single most-preferred alternative, not his entire ranking of alternatives.

We refer to the information collected as the agent's *type*.[1] In the single-agent example, the agent's type is simply which alternative he ranks at the top.[2] More formally, if his most-preferred alternative

*Azrieli and Healy: Dept. of Economics, The Ohio State University. Chambers: Dept. of Economics, Georgetown University. Azrieli: azrieli.2@osu.edu. Chambers: cc1950@georgetown.edu. Healy: healy.52@osu.edu.

[1]Note that types here do not refer to hierarchies of beliefs as in the framework of Harsanyi. Rather, they correspond to the private information of the agent that the planner wants to elicit.

[2]We assume strict preferences.

is x, then his type is the set of all preferences that rank x at the top. In the New York example, each of the student's types contains all preference orderings that agree on the top-twelve options. The type is exactly what the planner elicits, nothing more and nothing less. And the set of all possible types forms a partition of the set of preferences: every preference ordering belongs to exactly one type. We therefore refer to this partition as the *type space*.

In this paper we explore what information about preferences—what type spaces—have strictly incentive compatible mechanisms. Formally, we ask for which partitions of preferences does there exist an incentive compatible method of eliciting the agent's type. In the dictator example, paying the agent their announced favorite alternative is clearly incentive compatible. Thus, that particular type space is *elicitable*. Not all type spaces are elicitable, however. Consider three alternatives and the type space defined by the second-best element, meaning the planner wishes to know which of the three objects the agent ranks in the middle. We show that there is no incentive compatible payment mechanism on that type space. Thus, it is not an elicitable type space.

Our model assumes a finite number of alternatives and strict preferences. We allow the planner to use random mechanisms, which may pay an uncertain outcome. And we require that truthfully revealing one's type be *strictly* optimal, regardless of risk preferences. Thus, a mechanism is strictly incentive compatible on a type space if, for every type, announcing their type truthfully pays a gamble that strictly dominates any gamble they could be paid by lying.

Our notion of elicitability requires that no extra information is learned about preferences beyond the agent's type in the given type space. This is crucial, since the finest type space—in which every preference relation is its own type—is elicitable. To do so, the planner simply offers the agent every possible pair of alternatives, asks the agent to pick their favorite from each pair, and then pays them their chosen alternative from one randomly-selected pair.[3] But if the type space is coarse, as in the dictatorial example above, this procedure elicits far more information than wanted by the planner. This may be undesirable because it violates the agent's privacy by asking them to reveal more information than necessary. It is also inefficient, as the amount of information transferred (i.e., the number of questions the agent must answer) is much larger. Motivated by these considerations, we proceed assuming the planner wants only to know the type and nothing else.

We consider two ways of modeling the random payments generated by the mechanism: Lotteries or Savage-style acts. Lotteries represent the case of objective uncertainty, where the probabilities of the various outcomes are set by the planner. Acts, on the other hand, allow for subjective uncertainty. In the acts case the planner cannot control the payment probabilities. This limits their ability to construct incentive compatible mechanisms. Thus, there are type spaces that are elicitable with lotteries (using objective probabilities) but not with acts.

To illustrate, consider the type space based on the agent's *least*-preferred alternative out of m possible alternatives. If the planner can use objective lotteries then this type space is elicitable. To do so, she asks the subject to announce his least-preferred alternative. If the agent reports x, then the

[3]This is a description of an 'indirect mechanism', but clearly the same mapping from types to outcomes can be described as a direct mechanism: The agent reports his preference ordering over all alternatives, a pair of alternatives is randomly drawn, and the higher ranked of the two is awarded to the agent.

planner gives him a lottery that pays every alternative *except* x with probability $1/(m-1)$, and pays x with probability zero. Essentially, the planner asks the agent which one alternative he would like *not* to receive in the lottery. It is easy to see that announcing anything other than the true lowest-ranked alternative gives a (first-order) stochastically-dominated lottery to the agent. This is only true, however, because the probabilities on the $m-1$ other alternatives are exactly equal. In the acts framework the planner would use a physical randomizing device (such as an $m-1$-sided die) but could not be sure that the agent believes each outcome is equally likely. Incentive compatibility would not hold in this case since, regardless of the mapping from the state of the die to alternatives, the Savage act generated by truth-telling does not dominate state-by-state the acts generated by all possible lies.

To understand our characterization in the acts framework, consider the dictator example above. There the type is elicited by offering a single menu of all alternatives and paying the one chosen by the agent. We say that particular type space is *generated by* that single menu. In the case of eliciting the entire preference relation (the finest type space), the agent is offered all binary menus and one is randomly selected for payment. Thus, the finest type space is generated by the list of all binary menus. In general, we prove that a type space is elicitable if and only if it is generated by some list of menus. And the incentive compatible mechanism in that case is simple: offer exactly those menus to the agent and randomly select one for payment. The sufficiency of this condition is a simple extension of our previous work on incentives in experiments [2]; that this condition is also necessity is a new result.

In the lotteries framework we do not have a complete characterization for the general case. Clearly any partition generated by menus can also be elicited using lotteries, since lotteries are more powerful than acts. For necessity, we show that a convexity condition on partitions must be satisfied for a partition to be elicitable.[4] Convexity is also sufficient in the important special case of type spaces that provide information only about the position of alternatives in the ranking.

All of our main analyses focus on the case of a single agent. In the final section we extend these to the case of multiple agents and focus on dominant strategy incentive compatibility and lottery payments. The extension is trivial: If each agent's type space can be elicited individually, then their types can also be jointly elicited. To do so, simply run each individual's incentive compatible mechanism and then take the equally-weighted convex combination of all lotteries generated. The resulting centroid lottery is paid to all agents. Each agent can only influence the resulting lottery to a small degree, but otherwise their incentives are exactly as in the individual case. Thus, regardless of what others announce, each will find truth-telling to strictly dominate any lie.

1.1 Related Literature

Probably the most closely related work is that of Lambert [8, 9]. These authors are specifically concerned with which types of statistics of probability distributions can be elicited in a strictly incentive compatible way, when the individual in question has risk neutral subjective expected utility preferences. Roughly, our work on eliciting preferences using lotteries can be viewed geometrically as a special case of this framework. Namely, every preference relation can be identified with an equivalence class of probability

[4] We also show that a particular kind of cycles of 'neighboring' types cannot exist when a partition is elicitable.

distributions: a utility index representing the preference can be defined so that all of its values are nonnegative and so that they sum to one. It is well-known a lottery p first order stochastically dominates a lottery q for a given preference relation if and only if every utility index consistent with the relation ascribes a higher expected utility to p than to q. Therefore, we have a "dual" problem to the one studied by [8] (and also [11]), which is mathematically identical. The main distinction is that the "statistics" to be elicited in our situation are those which satisfy an additional measurability constraint: each region to be elicited can be identified with a collection of preferences.[5] In particular, any results from these papers relating to the elicitation of finite statistics hold here.

More generally, our work is also related to our previous work ([2]) on eliciting multiple choices. Here, the focus is not on the choices themselves, but rather the set of properties of preferences which are elicited by such mechanisms. Finally, [6] is a classic work on eliciting preferences from multiple individuals using stochastic payment rules.

2 Notation and definitions

Let X be a finite set of alternatives, with $|X| = m \geq 2$. Alternatives in X will be denoted by x, y, z, w etc. Let O be the set of all strict orders (complete, transitive, antisymmetric binary relations) on X. Typical elements of O are $\succeq, \succeq', \succeq''$. We write $x \succ y$ when $x \succeq y$ and $x \neq y$. Subsets of O are typically denoted by $A, A', A_1, A_2, \ldots$ etc. A collection of subsets $P = \{A_1, \ldots, A_k\}$ is a partition of O if these sets are non-empty, pairwise disjoint, and $\bigcup_{i=1}^{k} A_i = O$. Given a partition P and $\succeq \in O$, let $P(\succeq)$ be the (unique) element of P that contains $\succeq$.

For two partitions P and P', we say that P' refines P if for every $\pi' \in P'$ there is some $\pi \in P$ for which $\pi' \subseteq \pi$. The join $P \vee P'$ is the least upper bound according to the refinement relation, and consists of all nonempty sets of the form $\pi \cap \pi'$, where $(\pi, \pi') \in P \times P'$.

2.1 Lotteries

Let $\Delta(X)$ be the set of lotteries on X. We use p, q, r etc. to denote elements of $\Delta(X)$. If $p \in \Delta(X)$ then $p(x)$ is the probability assigned to the element $x \in X$ by the lottery p.

A lottery p First-Order Stochastically Dominates a lottery q relative to $\succeq$ (denoted $p \succeq^* q$) if for every $x \in X$

$$\sum_{\{y : y \succeq x\}} p(y) \geq \sum_{\{y : y \succeq x\}} q(y).$$

If there is a strict inequality for at least one x then p strictly dominates q relative to $\succeq$ (denoted $p \succ^* q$). It is worth pointing out the well-known fact that $p \succeq^* q$ can be interpreted as the lottery p yielding (weakly) higher expected utility than the lottery q for every utility function $u : X \to \mathbb{R}$ consistent with $\succeq$ (i.e., for every u satisfying $u[x] > u[y]$ whenever $x \succ y$). This interpretation will be used extensively in Section 4 below when analyzing incentive compatibility.

[5] In the context of probabilities, this would simply mean that any two probabilities inducing the same ranking on atoms must be included in the same region.

Given a partition P of O, a P-adaptable mechanism (in the lotteries framework) is a mapping $g : P \to \Delta(X)$. The interpretation is that the agent announces an element of P and the mechanism outputs a lottery over X. A mechanism is IC if the lottery obtained by announcing the element of P that contains the true preference dominates any other feasible lottery.[6] Formally,

Definition 1. A P-adaptable mechanism g is IC if for every $\succeq \in O$ and every $A \neq P(\succeq)$

$$g(P(\succeq)) \succ^* g(A).$$

Our interest is in understanding which partitions P admit an IC P-adaptable mechanism, hence the following definition.

Definition 2. A partition P is elicitable (in the lotteries framework) if there exists an IC P-adaptable mechanism $g : P \to \Delta(X)$.

2.2 Acts

In the acts framework we do not restrict the agent to view the randomization device of the mechanism as generating objective lotteries. Rather, the agent may have his own subjective beliefs about the likelihood of different realizations, or he may even have beliefs that cannot be represented by an additive probability measure. From the point of view of the agent the output of the mechanism is an act, i.e. a mapping from the state space of the randomization device to X. By varying his announcement the agent controls which act he receives.

Formally, given a partition P of O, a P-adaptable mechanism (in the acts framework) is a pair (Ω, f), where Ω is a finite state space (of the randomization device) and f is a mapping $f : P \to X^\Omega$.

Definition 3. A P-adaptable mechanism (Ω, f) is IC if for every $\succeq \in O$ and every $A \neq P(\succeq)$

$$f(P(\succeq))(\omega) \succeq f(A)(\omega),$$

with a strict preference for some $\omega \in \Omega$.

Thus, incentive compatibility requires that the act obtained by truth-telling (reporting $P(\succeq)$ when the preference is $\succeq$) weakly dominates state-by-state any other feasible act, with a strict preference in at least one state. One way to interpret this is that every subjective expected utility agent with full-support beliefs strictly prefers truth-telling over lying. But Definition 3 guarantees that truthfulness is the unique best-response for other classes of preferences over acts as well (e.g., essentially any model in the Knightian uncertainty literature).

Definition 4. A partition P is elicitable (in the acts framework) if there exists an IC P-adaptable mechanism (Ω, f).

[6] Our notion of incentive compatibility is not completely standard as it requires strict rather than weak dominance. In particular, constant mechanisms are not incentive compatible according to our definition.

3 Incentive compatibility with acts

3.1 Characterization

For every $\succeq \in O$ and a subset of alternatives $X' \subseteq X$, denote by $dom_{\succeq}(X')$ the maximal element in X' according to $\succeq$. That is $dom_{\succeq}(X') \in X'$ and $dom_{\succeq}(X') \succeq x$ for every $x \in X'$.

Given a finite collection of subsets of alternatives ('menus') $X_1, \ldots, X_l \subseteq X$, say that $\succeq$ cannot be distinguished from $\succeq'$ based on $\{X_1, \ldots, X_l\}$ if $dom_{\succeq}(X_i) = dom_{\succeq'}(X_i)$ for every $i = 1, \ldots, l$. Clearly this defines an equivalence relation on O. Let $\tilde{P}(X_1, \ldots, X_l)$ be the partition of O into the equivalence classes of this equivalence relation. Intuitively, $\tilde{P}(X_1, \ldots, X_l)$ contains the information revealed by observing the agent chooses his favorite element from each one of the menus $X_1, \ldots, X_l$.

Definition 5. A partition P is generated by menus if there are l and $X_1, \ldots, X_l \subseteq X$ such that $P = \tilde{P}(X_1, \ldots, X_l)$.

Proposition 1. A partition P is elicitable in the acts framework if and only if it is generated by menus.

Proof. (*If*) Suppose $P = \tilde{P}(X_1, \ldots, X_l)$. Let $\Omega = \{\omega_1, \ldots, \omega_l\}$. For each $A \in P$ choose an arbitrary representative $\succeq^A \in A$, and define $f(A)(\omega_i) = dom_{\succeq^A}(X_i)$ for $i = 1, \ldots, l$. Note that by assumption the choice of the representative $\succeq^A$ does not affect the resulting mechanism.

To see that the above mechanism is IC fix some $\succeq$ and some $A \in P$. Then for each i we have

$$f(P(\succeq))(\omega_i) = dom_{\succeq^{P(\succeq)}}(X_i) = dom_{\succeq}(X_i) \succeq dom_{\succeq^A}(X_i) = f(A)(\omega_i),$$

where the first equality is by the definition of f, the second follows from $\succeq \in P(\succeq)$, the next relation follows from the definition of dom, and the last equality is again by construction of f. Moreover, if $A \neq P(\succeq)$ then there exists i such that $dom_{\succeq}(X_i) \neq dom_{\succeq^A}(X_i)$ which gives a strict preference at ω_i.

(*Only If*) Suppose P is elicitable and let (Ω, f) be an IC P-adaptable mechanism. Enumerate the states so that $\Omega = \{\omega_1, \ldots, \omega_l\}$ for some positive integer l, and for each $i = 1, \ldots l$ define $X_i = \{f(A)(\omega_i)\}_{A \in P} \subseteq X$. In words, X_i is the set of all possible alternatives that can be chosen at state ω_i as the agent varies his announcement.

We now show that $P = \tilde{P}(X_1, \ldots, X_l)$. Suppose that $\succeq, \succeq'$ are in the same element of P, and fix some $1 \leq i \leq l$. Then by incentive compatibility we have that $f(P(\succeq))(\omega_i) \succeq f(A)(\omega_i)$ for every $A \in P$, which implies that $f(P(\succeq))(\omega_i) = dom_{\succeq}(X_i)$. Applying the same argument to $\succeq'$ gives $f(P(\succeq'))(\omega_i) = dom'_{\succeq}(X_i)$. But since $P(\succeq) = P(\succeq')$ we get $dom_{\succeq}(X_i) = dom'_{\succeq}(X_i)$. Repeating for each $i = 1, \ldots, l$ shows that $\succeq, \succeq'$ are in the same element of $\tilde{P}(X_1, \ldots, X_l)$.

Conversely, suppose that $\succeq, \succeq'$ are in the same element of $\tilde{P}(X_1, \ldots, X_l)$. From the previous paragraph we have $f(P(\succeq))(\omega_i) = dom_{\succeq}(X_i)$ and $f(P(\succeq'))(\omega_i) = dom'_{\succeq}(X_i)$ for each i, so $f(P(\succeq))(\omega_i) = f(P(\succeq'))(\omega_i)$ for each i. Incentive compatibility now implies that $P(\succeq) = P(\succeq')$, which concludes the proof. $\square$

3.2 The lattice of elicitable partitions

Consider two elicitable partitions under acts P and P'. Proposition 1 implies that each one of them is generated by some collection of menus, say $P = \tilde{P}(X_1, \ldots, X_l)$ and $P' = \tilde{P}(X'_1, \ldots, X'_k)$. But it is immediate to verify that $\tilde{P}(X_1, \ldots, X_l) \vee \tilde{P}(X'_1, \ldots, X'_k) = \tilde{P}(X_1, \ldots, X_l, X'_1, \ldots, X'_k)$, and thus that $P \vee P'$ is elicitable as well. We state this fact in the next corollary.

Corollary 1. If P and P' are both elicitable in the acts framework, then so is their join.

It follows that the set of elicitable partitions equipped with the refinement relation forms a lattice. Furthermore, this lattice is *atomistic*, meaning that every elicitable partition is the join of some collection of the atoms in the lattice. Indeed, if $X' \subseteq X$ is a menu (with $|X'| \geq 2$) then the partition $\tilde{P}(X')$ generated by X' only is an atom – there is no coarser elicitable partition other than the trivial one; and any elicitable partition is the join of some collection of these basic partitions.

It is worth pointing out that the meet in the lattice of elicitable partitions does not coincide with the finest common coarsening. This is illustrated in the following example.

Example 1. Let $X = \{x, y, z\}$, and[7] $P = \{\{xyz, xzy\}, \{yxz, yzx\}, \{zxy, zyx\}\}$. This partition is generated by the menu X. Let $P' = \{\{xyz, xzy\}, \{zyx, yzx\}, \{zxy\}, \{yxz\}\}$. This partition is generated by the menus $\{x, y\}$ and $\{x, z\}$. The finest common coarsening of P and P' is $\{\{xyz, xzy\}, \{yzx, yxz, zxy, zyx\}\}$. This partition is not generated by menus and hence not elicitable. In fact, the meet of P and P' in the lattice of elicitable partitions is the trivial partition.

We end this subsection with the following observation. Suppose that a principal is interested in eliciting a partition P, but P is not elicitable. A natural solution would be to elicit the coarsest elicitable refinement of P, since this would provide the principle with the required information while revealing the minimal amount of redundant information about the agent's preferences. Unfortunately, the next example demonstrates that there need not exist such a coarsest elicitable refinement.

Example 2. Let $X = \{x, y, z\}$ and $P = \{\{xyz, yxz, yzx\}, \{xzy\}, \{zxy, zyx\}\}$. Then P is not elicitable, but the following two refinements of P are:

- $P_1 = \{\{xyz\}, \{xzy\}, \{yzx, yxz\}, \{zxy, zyx\}\}$, generated by the menus $\{x, y, z\}$ and $\{y, z\}$.
- $P_2 = \{\{xyz, yxz\}, \{yzx\}, \{xzy\}, \{zxy, zyx\}\}$, generated by the menus $\{y, z\}$ and $\{x, z\}$.

Since P is the finest common coarsening of P_1 and P_2, it follows that there is no elicitable refinement of P that is coarser than all other elicitable refinements.

3.3 From menus to partitions and back

Since elicitable partitions in the acts framework are exactly those that are generated by menus, it is useful to understand the structure of such partitions as well as the connection between lists of menus

[7]The notation xyz refers to the ordering that ranks x first, y second, and z third. Other orderings are denoted in an analogous way.

and the partitions they generate. For a partition P and a menu $X' \subseteq X$, say that X' is identified by P if for every $A \in P$ and every $\succeq, \succeq' \in A$ it holds that $dom_{\succeq}(X') = dom_{\succeq'}(X')$; in other words this means that $dom_{\bullet}(X')$ is a P-measurable function from O to X. Intuitively, if X' is identified by P then knowing which element of P contains the true ordering is enough to pin down the top-ranked element in X'. Let $\tilde{I}(P)$ be the collection of menus that are identified from P.

The following proposition provides a simple way to check whether a given partition is generated by menus (and hence elicitable under acts).

Proposition 2. A partition P is generated by menus if and only if $P = \tilde{P}(\tilde{I}(P))$.

Proof. If $P = \tilde{P}(\tilde{I}(P))$ then clearly P is generated by menus (by the menus $\tilde{I}(P)$).

To prove the converse note first that, for every partition P, if $\succeq, \succeq'$ are in the same element of P then by definition $dom_{\succeq}(X') = dom_{\succeq'}(X')$ for every $X' \in \tilde{I}(P)$. This implies that $\succeq, \succeq'$ are also in the same element $\tilde{P}(\tilde{I}(P))$. In other words, P is always (weakly) finer than $\tilde{P}(\tilde{I}(P))$.

Now, suppose that P is generated by menus, say $P = \tilde{P}(X_1, \ldots, X_l)$. Then clearly $\{X_1, \ldots, X_l\} \subseteq \tilde{I}(P)$. But adding more menus can only make the resulting partition finer, so $\tilde{P}(\tilde{I}(P))$ is (weakly) finer than $\tilde{P}(X_1, \ldots, X_l) = P$. This completes the proof. □

While $\tilde{I}(P)$ contains all the menus that are identified by P, it is possible that smaller (in the sense of inclusion) collections of menus are already sufficient to generate P. This will happen when the information revealed by the choices from the smaller collection of menus is the same as the information from the larger one. We illustrate this point with the following example.

Example 3. Let $X = \{x, y, z\}$ and let P be the finest partition (each element of P is a singleton). Then clearly $\tilde{I}(P)$ contains all the menus, in particular the grand menu X. Consider the collection of menus $\mathcal{X} = \{\{x, y\}, \{x, z\}, \{y, z\}\}$. Observing the top-ranked element in each of these menus would always reveal the entire ranking, so that $\tilde{P}(\mathcal{X}) = P$ even though $X \notin \mathcal{X}$. In other words, observing the choice from X is redundant if one has already observed the choices from the menus in $\mathcal{X}$.

Our last result in this section shows that the set of collections of menus that generate a given partition is an interval according to $\subseteq$, and characterizes the minimal element in this interval.[8] This minimal element is of particular significance as it corresponds to the smallest number of observations needed to generate a given partition. We will need the following definition.

Definition 6. 1. Let $X', X_1, \ldots, X_l$ be menus. Say that X' is surely identified by $X_1, \ldots, X_l$ if whenever $\succeq, \succeq'$ are such that $dom_{\succeq}(X_i) = dom_{\succeq'}(X_i)$ for all $i = 1, \ldots, l$ it also holds that $dom_{\succeq}(X') = dom_{\succeq'}(X')$.
2. The menus $X_1, \ldots, X_l$ are independent if no X_i is surely identified by $\{X_j\}_{j \neq i}$.

Proposition 3. For any P elicitable in the acts framework, the set of collection of menus generating P is an interval according to $\subseteq$. Furthermore, the minimal such menu is the unique independent collection of menus generating P.

Proof. We start with the following lemma, which originally appeared in our earlier work [3, Lemma 3].

[8] As explained above, the maximal element of the interval is the collection $\tilde{I}(P)$.

Lemma 1. Suppose $|X'| \geq 2$. Then X' is surely identified by $\{X_1, \ldots, X_l\}$ if and only if for every $x, y \in X'$, there is $1 \leq i \leq l$ such that $\{x, y\} \subseteq X_i \subseteq X'$.

Proof. Suppose X' is surely identified by $X_1, \ldots, X_l$, but there are $x, y \in X'$ for which for all i, either $\{x, y\} \subseteq X_i$ is false or $X_i \subseteq X'$ is false. Let $\succeq, \succeq'$ be a pair of orders which (1) rank all members of $X \setminus X'$ above X'; (2) rank x, y above all remaining elements of X'; and (3) differ only in their ranking of x and y, say $x \succ y$ and $y \succ' x$. Observe then that $\text{dom}_{\succeq}(X_i) = \text{dom}_{\succeq'}(X_i)$ for all i, but $\text{dom}_{\succeq}(X') = x \neq y = \text{dom}_{\succeq'}(X')$, a contradiction.

Conversely, suppose that the condition in the lemma holds, and that $\text{dom}_{\succeq}(X_i) = \text{dom}_{\succeq'}(X_i)$ for all i. Suppose by means of contradiction that $x = \text{dom}_{\succeq}(X') \neq \text{dom}_{\succeq'}(X') = y$, so that $x \succ y$ and $y \succ' x$. Let i be such that $\{x, y\} \subset X_i \subseteq X'$. Then $x = \text{dom}_{\succeq}(X')$ implies $x = \text{dom}_{\succeq}(X_i)$, and $y = \text{dom}_{\succeq'}(X')$ implies $y = \text{dom}_{\succeq'}(X_i)$, a contradiction. □

Moving on to the proof of the proposition, for any collection $\mathcal{X}$ of nonempty, non-singleton menus, let the set $\text{SI}(\mathcal{X})$ denote the collection of sets surely identified by $\mathcal{X}$. Observe that:

1. $\text{SI}(\text{SI}(\mathcal{X})) = \text{SI}(\mathcal{X})$
2. $\mathcal{X} \subseteq \text{SI}(\mathcal{X})$
3. $\mathcal{X} \subseteq \mathcal{X}'$ implies $\text{SI}(\mathcal{X}) \subseteq \text{SI}(\mathcal{X}')$.

Therefore, SI forms a *closure operator*. Furthermore, this closure operator has the anti-exchange property, as defined in [5]. Namely, if $\mathcal{X} = \text{SI}(\mathcal{X})$, $X', X'' \notin \mathcal{X}$, $X' \neq X''$, and $X' \in \text{SI}(\mathcal{X} \cup \{X''\})$, then $X'' \notin \text{SI}(\mathcal{X} \cup \{X'\})$. To see this latter point, observe that it follows from Lemma 1 that there are x, y for which $\{x, y\} \subseteq X'' \subseteq X'$; in particular, $X'' \subseteq X'$. If $X'' \in \text{SI}(\mathcal{X} \cup \{X'\})$, then similarly, $X' \subseteq X''$, contradicting $X' \neq X''$.

Now, by [5, Theorem 2.1], for any $\mathcal{X}$ there is a unique minimal collection of menus $B(\mathcal{X})$ such that $\text{SI}(\mathcal{X}) = \text{SI}(B(\mathcal{X}))$. Since $\tilde{P}(X_1, \ldots, X_l) = \tilde{P}(X'_1, \ldots, X'_k)$ iff $\text{SI}(\{X_1, \ldots, X_l\}) = \text{SI}(\{X'_1, \ldots, X'_k\})$, it follows that if P is elicitable then the set of collection of menus that generate P is the interval $[B(\tilde{I}(P)), \tilde{I}(P)]$. Finally, the independence of a collection $B(\mathcal{X})$ (for some $\mathcal{X}$) immediately follows: If $X' \in B(\mathcal{X})$ is surely identified by the other menus in $B(\mathcal{X})$ then $\text{SI}(B(\mathcal{X}) \setminus \{X'\}) = \text{SI}(B(\mathcal{X}))$, contradicting the minimality of $B(\mathcal{X})$. □

4 Incentive compatibility with lotteries

4.1 A sufficient condition

Proposition 4. If P is generated by menus then P is elicitable in the lotteries framework.

Proof. If P is generated by menus then it follows from Proposition 1 that there is an IC P-adaptable mechanism (Ω, f) in the acts framework. Let μ be a full-support probability distribution on Ω, and define the lotteries mechanism g by

$$g(A)(x) = \mu\Big(\{\omega \in \Omega \ : \ f(A)(\omega) = x\}\Big)$$

for any $A \in P$ and $x \in X$. In words, $g(A)$ is the distribution of the X-valued random variable $f(A)$ when the state-space Ω is endowed with the measure μ.

Now, fix $\succeq$ and $A \neq P(\succeq)$. Since (Ω, f) is IC we have that $f(P(\succeq))(\omega) \succeq f(A)(\omega)$ for all ω and that $f(P(\succeq)) \neq f(A)$. Thus, for every $x \in X$

$$\{\omega \in \Omega \ : \ f(P(\succeq))(\omega) \succeq x\} \supseteq \{\omega \in \Omega \ : \ f(A)(\omega) \succeq x\},$$

with strict inclusion for at least one x. Since μ has full support it follows that $g(P(\succeq)) \succ^* g(A)$, and we are done. □

Remark 1. It is also possible to prove Proposition 4 directly. Suppose $P = \tilde{P}(X_1, \ldots, X_l)$. Let λ be a full-support distribution on $\{1, \ldots, l\}$ and define $g(A)(x) = \lambda\left(\left\{1 \leq i \leq l \ : \ dom_{\succeq^A}(X_i) = x\right\}\right)$, where $\succeq^A$ is an arbitrary choice from A. It is not hard to check that g is IC.

The following example shows that, in contrast to the acts framework, P may be elicitable even if it is not generated by menus. The intuition for this is that in the lotteries framework we make more restrictions than in the acts framework on the preferences that the agent may have over uncertain prospects, so incentive compatibility is easier to satisfy.

Example 4. Let $X = \{x, y, z\}$ and let $P = \{A_1, A_2, A_3\}$ where

$$A_1 = \{xyz, yxz\}, \quad A_2 = \{xzy, zxy\}, \quad A_3 = \{yzx, zyx\}.$$

In words, P reveals the top 2 alternatives but not their order. Then P is not generated by menus as can be easily checked using Proposition 2 above: $\tilde{I}(P)$ contains no non-trivial menus, so $\tilde{P}(\tilde{I}(P))$ is the trivial partition and hence $\tilde{P}(\tilde{I}(P)) \neq P$. However, consider the mechanism g given by

$$g(A_1) = (x, 0.5; y, 0.5; z, 0), \quad g(A_2) = (x, 0.5; y, 0; z, 0.5), \quad g(A_3) = (x, 0; y, 0.5; z, 0.5),$$

that is, g chooses randomly one of the top-2 ranked alternatives. It is immediate to check that g is IC.

4.2 Necessary conditions

Given a set of orderings $A \in O$, let $\sqsupseteq_A = \bigcap_{\succeq \in A} \succeq$ be the maximal relation that all orderings in A agree on; that is, $x \sqsupseteq_A y$ if and only if $x \succeq y$ for all $\succeq \in A$. Note that $\sqsupseteq_A$ is a partial order. Say that $\succeq$ is consistent with A if $\sqsupseteq_A \subseteq \succeq$, i.e., if for all $x, y \in X$, $x \sqsupseteq_A y$ implies $x \succeq y$. Denote by $Cons(A)$ the set of ordering consistent with A.

Definition 7. A set $A \subseteq O$ is convex if $A = Cons(A)$. A partition P is convex if every $A \in P$ is convex.

To understand why we refer to the above property as convexity, recall that $u \in \mathbb{R}^X$ is consistent with an ordering $\succeq$ when $u[x] > u[y]$ iff $x \succ y$. Let $U(\succeq) \subseteq \mathbb{R}^X$ be the set of utility vectors consistent with $\succeq$. The closure $\overline{U(\succeq)}$ of $U(\succeq)$ is the set of utility vectors u satisfying $u[x] \geq u[y]$ whenever $x \succeq y$. Finally, for $A \subseteq O$ we denote $U(A) = \bigcup_{\succeq \in A} U(\succeq)$ and $\overline{U(A)} = \bigcup_{\succeq \in A} \overline{U(\succeq)}$.

Lemma 2. A set $A \subseteq O$ is convex if and only if $\overline{U(A)}$ is convex in $\mathbb{R}^X$.

Proof. It is not hard to check that for any set A

$$\overline{U(Cons(A))} = \bigcap_{\{(x,y)\ :\ x \sqsupseteq_A y\}} \{u\ :\ u[x] \geq u[y]\}.$$

Since the set on the right-hand side is convex, it follows that if A is convex (i.e., $A = Cons(A)$) then $\overline{U(A)}$ is convex. On the other hand, if $\overline{U(A)}$ is convex then it is an intersection of closed half-spaces of the form $\{u\ :\ u[x] \geq u[y]\}$. This implies that $\overline{U(A)} = \overline{U(Cons(A))}$, and hence that $A = Cons(A)$. □

Remark 2. Convexity of a set of orderings can also be characterized using the notion of convexity of a set of nodes in a graph. For $\succeq \in O$ and $x \in X$ let $r_\succeq(x) = |\{y :\ \ y \succeq x\}|$ be the ranking of x in the ordering $\succeq$. Say that two orderings $\succeq$ and $\succeq'$ are adjacent if there are $x, y \in X$ such that $r_\succeq(x) = r_{\succeq'}(y) = r_\succeq(y) - 1 = r_{\succeq'}(x) - 1$ and $r_\succeq(z) = r_{\succeq'}(z)$ for all $z \neq x, y$, that is if $\succeq$ and $\succeq'$ differ only by a single switch of neighboring elements. Consider an undirected graph G where the set of vertices is O and the set of edges is the set of adjacent orderings. Then A is convex if and only if for every $\succeq, \succeq' \in A$, if $\succeq''$ is in a shortest path between $\succeq$ and $\succeq'$ then $\succeq'' \in A$. We omit the proof.

The following proposition shows that convexity is a necessary condition for elicitability in the lotteries framework. Variants of this result in different contexts have been obtained in previous works, see for example Lambert [9, pages 10-11] who attributes this observation to Osband [10].

Proposition 5. If P is elicitable in the lotteries framework then it is convex.

Proof. Let g be an IC P-adaptable mechanism, and let $A \in P$. We will show that

$$\overline{U(A)} = \bigcap_{A' \in P} \{u\ :\ \langle g(A), u\rangle \geq \langle g(A'), u\rangle\},$$

where $\langle \cdot, \cdot \rangle$ is the standard scalar product in $\mathbb{R}^X$. As the set on the right-hand side is clearly convex, this will suffice to prove the proposition.

Suppose first that u is in the set on the left-hand side. Then there is $\succeq \in A$ such that $u \in \overline{U(\succeq)}$. Incentive compatibility of g then implies that $\langle g(A), u\rangle \geq \langle g(A'), u\rangle$ for every $A' \in P$, so u is in the right-hand side as well.[9]

Conversely, suppose that u is in the right-hand side. Then in every open neighborhood of u there is u' for which $\langle g(A), u'\rangle > \langle g(A'), u'\rangle$ holds for all $A' \neq A$ (here we use the fact that the right-hand side is a polyhedral set with non-empty interior). Incentive compatibility of g implies that $u' \notin \overline{U(A^c)}$, so we must have $u' \in \overline{U(A)}$. Since this set is closed we get that $u \in \overline{U(A)}$ as well. □

[9]Recall that lottery p strictly dominates lottery q relative to $\succeq$ if and only if $\langle p, u\rangle > \langle q, u\rangle$ for every $u \in U(\succeq)$.

The following example demonstrates that convexity is not sufficient for elicitability.

Example 5. Let $X = \{x, y, z\}$ and let $P = \{A_1, A_2, A_3\}$ where

$$A_1 = \{xyz, xzy, zxy\}, \quad A_2 = \{yxz, yzx\}, \quad A_3 = \{zyx\}.$$

Convexity of P is immediate (A_1 is the set of all orderings in which x is ranked above y, A_2 is the set of rankings in which y is ranked above both x and z, and A_3 is a singleton).

We show that P is not elicitable. Suppose by contradiction that g is an IC P-adaptable mechanism. Consider the vectors $u_1 \in U(A_1)$, $u_2 \in U(A_2)$ defined by $(u_1[x] = 2,\ u_1[y] = 1,\ u_1[z] = -M)$ and $(u_2[x] = 1,\ u_2[y] = 2,\ u_2[z] = -M)$, where $M > 0$ is large. Then incentive compatibility implies $\langle g(A_1) - g(A_2), u_1 \rangle > 0$ and $\langle g(A_1) - g(A_2), u_2 \rangle < 0$. Taking M to $+\infty$ gives $g(A_1)(z) = g(A_2)(z)$. A similar argument (with M replacing $-M$) shows that $g(A_1)(z) = g(A_3)(z)$. Combining these two equalities gives $g(A_2)(z) = g(A_3)(z)$.

Now, consider $u_3 \in U(A_3)$ given by $(u_3[x] = 1, u_3[y] = 2, u_3[z] = M)$ for large $M > 0$. From incentive compatibility $\langle g(A_2) - g(A_3), u_3 \rangle < 0$ and $\langle g(A_2) - g(A_3), u_2 \rangle > 0$, but this is inconsistent with $g(A_2)(z) = g(A_3)(z)$.

Roughly speaking, the reason that elicitability fails in Example 5 is that the cycle of sets (A_2, A_1, A_3, A_2) is inconsistent with incentive compatibility: The probability of z must be the same in $g(A_2)$ and $g(A_1)$, and must also be the same in $g(A_1)$ and $g(A_3)$, but at the same time must be higher in $g(A_3)$ than in $g(A_2)$. We now generalize this observation to obtain another necessary condition for elicitability.

Recall that $r_{\succeq}(x) = |\{y : \ y \succeq x\}|$ is the ranking of x in the ordering $\succeq$. Say that $\succeq$ and $\succeq'$ are adjacent via an $x - y$ switch if $r_{\succeq}(x) = r_{\succeq'}(y) = r_{\succeq}(y) - 1 = r_{\succeq'}(x) - 1$ and $r_{\succeq}(z) = r_{\succeq'}(z)$ for all $z \neq x, y$. Similarly, say that the sets A and A' ($A \neq A'$) are adjacent via an $x - y$ switch if there are $\succeq \in A$ and $\succeq' \in A'$ that are adjacent via an $x - y$ switch; the sets A and A' are adjacent if they are adjacent via an $x - y$ for some x, y.

Proposition 6. Suppose that P is elicitable. If $\{A_1, \ldots, A_k\} \subseteq P$ are such that A_i is adjacent to A_{i+1} for each $i \in \{1, \ldots, k\}$ (where addition is read modulo k), and if A_1 is adjacent to A_2 via an $x - y$ switch, then there exists $1 < i \leq k$ and z such that A_i and A_{i+1} are adjacent via a $z - x$ switch.

Proof. The key to the proof is the following lemma.

Lemma 3. If g is an IC P-adaptable mechanism, and if $A, A' \in P$ are adjacent via an $x - y$ switch, then $g(A)(x) - g(A)(y) = g(A')(y) - g(A')(x) > 0$, and $g(A)(z) = g(A')(z)$ for all $z \notin \{x, y\}$.

Proof. Let $\succeq \in A$ and $\succeq' \in A'$ be adjacent via an $x - y$ switch, that is $r_{\succeq}(x) = r_{\succeq'}(y) = r_{\succeq}(y) - 1 = r_{\succeq'}(x) - 1$ and $r_{\succeq}(z) = r_{\succeq'}(z)$ for all $z \neq x, y$.

We first show that $g(A)(z) = g(A')(z)$ for all z with $r_{\succeq}(z) < r_{\succeq}(x)$, i.e., for all z ranked above x and y (assuming such z exists). The proof proceeds by induction on $r_{\succeq}(z)$. For $r_{\succeq}(z) = 1$, consider the utility vector $\bar{u}$ with $\bar{u}(z) = 1$ and $\bar{u}(w) = 0$ for all $w \neq z$. Then $\bar{u}$ is both a limit point of

$U(\succeq)$ and a limit point of $U(\succeq')$. Any $u \in U(\succeq)$ has $\langle u, g(A)\rangle > \langle u, g(A')\rangle$ and any $u' \in U(\succeq')$ has $\langle u', g(A)\rangle < \langle u', g(A')\rangle$, from which we conclude that $\langle \bar{u}, g(A)\rangle = \langle \bar{u}, g(A')\rangle$ must be satisfied. But this is the same as $g(A)(z) = g(A')(z)$.

Now, consider z with $r_{\succeq}(z) < r_{\succeq}(x)$ and suppose that $g(A)(w) = g(A')(w)$ for all w for which $r_{\succeq}(w) < r_{\succeq}(z)$. Let $\bar{u}$ be given by $\bar{u}(w) = 1$ for all w with $r_{\succeq}(w) \le r_{\succeq}(z)$ and $\bar{u}(w) = 0$ otherwise. Observe again that $\bar{u}$ is both a limit point of $U(\succeq)$ and a limit point of $U(\succeq')$. Conclude that $\langle \bar{u}, g(A)\rangle = \langle \bar{u}, g(A')\rangle$, so by the induction hypothesis it follows that $g(A)(z) = g(A')(z)$.

A symmetric argument establishes the result when $r_{\succeq}(z) > r_{\succeq}(y)$ (e.g., for $r_{\succeq}(z) = m$, use $\bar{u}(z) = 0$, $\bar{u}(w) = 1$ for $w \neq z$, and proceed by induction). Finally, since $g(A)(z) = g(A')(z)$ for all $z \neq x, y$, and since both $g(A), g(A')$ are lotteries, we must have $g(A)(x) - g(A)(y) = g(A')(y) - g(A')(x)$. The fact that these differences are positive immediately follows from incentive compatibility of g (recall that x is ranked above y according to $\succeq$ and y above x according to $\succeq'$). □

The proposition now easily follows. Indeed, let g be an IC mechanism and suppose $\{A_1, \ldots, A_k\} \subseteq P$ satisfy the assumption of the proposition. Then by Lemma 3 we have that $g(A_1)(x) > g(A_2)(x)$. Suppose by means of contradiction that there is no $1 < i \le k$ and z such that A_i and A_{i+1} are adjacent via a $z - x$ switch. Then it follows again from Lemma 3 that $g(A_i)(x) \ge g_(A_{i+1})(x)$, whereby $g(A_2)(x) \ge g(A_1)(x)$, a contradiction. □

Remark 3. Lemma 3 says that a necessary condition for a P-adaptable mechanism g to be IC is that if A and A' are adjacent via an $x - y$ switch, then the lotteries $g(A)$ and $g(A')$ are identical except that some mass is shifted from x to y. This is a local incentive constraint which guarantees that an agent with true preference in A has no incentive to announce A', and vice versa. Carroll [4, Proposition 2] shows that in a class of models that includes ours, if a mechanism satisfies all the local incentive constraints then it is globally incentive compatible. He works with the standard notion of weak incentive compatibility, but the result goes through with our strict notion. Thus, the condition in Lemma 3 is not only necessary for g to be IC, it is also sufficient.

Remark 4. An implication of Proposition 6 is that if P is elicitable and $A, A' \in P$ are adjacent via an $x - y$ switch, then these two sets can't be adjacent via any other switch. Indeed, this corresponds to the case of a cycle of length 2 containing only A and A'.

Remark 5. We do not know whether convexity together with the no-cycles condition of Proposition 6 is enough to guarantee elicitability in the lotteries framework. While we could not find a counter example, a problem may arise if several cycles of sets (each of which is not violating the condition) interact in a way that prevents a single mechanism to work for all of them simultaneously. Characterizing elicitable partitions in the lotteries framework is therefore still an open question.

4.3 Characterization for positional partitions

We now restrict attention to partitions that treat all the alternatives symmetrically, and only contain information about positions in the ranking. One example is when the agent only announces the ranking of his top k alternatives, with $1 \leq k < m$. Another example is when the agent announces his k lowest ranked alternatives (say, not including their order) to indicate what he views as unacceptable. We show below that in this class of partitions convexity is not only necessary but also sufficient for elicitability.

To formalize this, let us think of the ranking function $r_{\succeq}(\cdot)$ as a bijection from X to $\{1, \ldots, m\}$. If $B \subseteq \{1, \ldots, m\}$ then $r_{\succeq}^{-1}(B) = \{x \in X \ : \ r_{\succeq}(x) \in B\}$ is the set of alternatives whose ranking according to $\succeq$ is in B. Given a partition Q of $\{1, \ldots, m\}$, say that $\succeq, \succeq'$ have the same Q-rankings if $r_{\succeq}^{-1}(B) = r_{\succeq'}^{-1}(B)$ for every $B \in Q$.

Definition 8. Let Q be a partition of $\{1, \ldots, m\}$. The Q-positional partition, denoted P_Q, is the partition of O into the equivalence classes of the equivalence relation 'having the same Q-ranking'. A partition P is positional if it is Q-positional for some Q.

Example 6. Suppose $X = \{x, y, z\}$. For $Q = \{\{1, 2\}, \{3\}\}$ the partition P_Q is the partition of Example 4, that is, $P_Q = \{\{xyz, yxz\}, \{xzy, zxy\}, \{yzx, zyx\}\}$. For $Q = \{\{1, 3\}, \{2\}\}$ we have $P_Q = \{\{xyz, zyx\}, \{yxz, zxy\}, \{xzy, yzx\}\}$. For $Q = \{\{1\}, \{2\}, \{3\}\}$ the partition P_Q is the finest partition (a unique ordering in each element). For $Q = \{\{1, 2, 3\}\}$ the partition P_Q is the coarsest partition, i.e., where all orderings are in a single element.

Proposition 7. Let P be a positional partition. Then the following conditions are equivalent:
(1) P is elicitable in the lotteries framework.
(2) P is convex.
(3) Every element in the partition Q that defines P is a (possibly degenerate) interval in $\{1, \ldots, m\}$.

Proof. (1) $\Longrightarrow$ (2): This follows from Proposition 5.

(2) $\Longrightarrow$ (3): We show that if (3) is violated then (2) is wrong as well. Suppose that $1 \leq i < j < k \leq m$ are such that $i, k \in B \in Q$ but $j \notin B$. Let $\succeq$ be an ordering such that $r_{\succeq}(x) = i$, $r_{\succeq}(y) = j$, and $r_{\succeq}(z) = k$ for some three elements $x, y, z \in X$. Let $\succeq'$ be another ordering that is identical to $\succeq$ everywhere except that the rankings of x and z are switched; that is, $r_{\succeq'}(x) = k$, $r_{\succeq'}(z) = i$, and $r_{\succeq}(w) = r_{\succeq'}(w)$ for all $w \neq x, z$. Then by definition $\succeq$ and $\succeq'$ are in the same element of P_Q, say A.

Now, let $u \in U(\succeq)$ be such that $u[x] = 3$, $u[y] = 1$, and $u[z] = 0$. Existence of such u is obvious. Let u' be identical to u except that $u'[x] = 0$ and $u'[z] = 3$. Note that $u' \in U(\succeq')$. Consider $u'' = \frac{1+\epsilon}{3} u' + \frac{2-\epsilon}{3} u$, where $\epsilon > 0$. We have $u''[x] = 2 - \epsilon$, $u''[z] = 1 + \epsilon$, and u'' is identical to u (and to u') otherwise. Let ϵ be sufficiently small such that no two elements of u'' are identical and such that no element of u'' is between 1 and $1+\epsilon$. Call $\succeq''$ to the ordering induced by u''. Then $r_{\succeq''}(z) = r_{\succeq}(y) = j$, which implies that $z \notin r_{\succeq''}^{-1}(B)$. Thus, $r_{\succeq''}^{-1}(B) \neq r_{\succeq}^{-1}(B)$, so $\succeq'' \notin A$. This proves that $\overline{U(A)}$ is not convex, so P_Q is not a convex partition.

(3) $\Longrightarrow$ (1): Suppose $Q = \{B_1, \ldots, B_K\}$ where each B_k is an interval in $\{1, \ldots, m\}$. Without loss assume that the B_k's are ordered, so that if $i \in B_k, j \in B_{k'}$ and $k < k'$ then $i < j$. Let $\delta_1 > \delta_2 > \ldots > \delta_K$ be non-negative numbers satisfying $\sum_{k=1}^{K} |B_k| \delta_k = 1$. For each $A \in P_Q$ choose an arbitrary ordering $\succeq^A \in A$. Define the mechanism g by $g(A)(x) = \delta_k$ if $r_{\succeq^A}(x) \in B_k$. Note that by definition of the partition P_Q the mechanism g does not depend on the choice of the representatives $\succeq^A$. Also, g is IC since misreporting shifts probability from higher to lower ranked alternatives (recall Example 4). $\square$

4.4 Lattice structure

Proposition 8. If P and P' are elicitable in the lotteries framework, then so is their join.

Proof. Let g be an IC P-adaptable mechanism, and let g' be an IC P'-adaptable mechanism. We define a $(P \vee P')$-adaptable mechanism g^* as follows: For any pair $(A, A') \in P \times P'$ with $A \cap A' \neq \varnothing$ let $g^*(A \cap A') = (1/2)g(A) + (1/2)g'(A')$. It is straightforward to verify that g^* is IC, so that $P \vee P'$ is elicitable. □

It follows that the collection of elicitable partitions in the lotteries framework forms a lattice when equipped with the refinement partial ordering. As demonstrated by example 4, this lattice is distinct (strict superset) from the lattice of partitions elicitable in the acts framework.

Recall that in the acts framework the lattice of elicitable partitions is atomistic, and that there is a simple description of its atoms. We do not know if the lattice of elicitable partitions under lotteries is atomistic. However, we now argue that in a certain sense every elicitable lottery is the join of a collection of 'basic' elicitable partitions.

Consider the set G of all mappings $g : O \to \Delta(X)$ satisfying the following property:
If $\succeq$ and $\succeq'$ are adjacent via an $x - y$ switch then $g(\succeq)(x) \geq g(\succeq')(x)$, $g(\succeq)(y) \leq g(\succeq')(y)$, and $g(\succeq)(z) = g(\succeq')(z)$ for all $z \neq x, y$.

For every $g \in G$ define the partition P_g of O by letting $\succeq, \succeq'$ be in the same element of P_g if and only if $g(\succeq) = g(\succeq')$. Then clearly we can view g as a P_g-adaptable mechanism, and by Remark 3 g is IC. Moreover, if $g, g' \in G$ and $0 < \alpha < 1$ then $P_{\alpha g + (1-\alpha)g'} = P_g \vee P_{g'}$. Therefore, any elicitable partition is the join of partitions from the set $\{P_g \ : \ g \textit{ is an extreme point of } G\}$. Every atom of the lattice corresponds to some extreme point of G, but we don't know whether every extreme point corresponds to an atom.

It is natural to ask whether positional partitions as defined in the previous subsection play a special role in the lattice. More precisely, for every menu $X' \subseteq X$ and a number $k = 1, \ldots, |X'| - 1$ consider the partition that reveals the k top-ranked elements from X' (without their ordering). Every such partition is elicitable (recall Proposition 7). Moreover, it is not hard to see that each of these partitions is an atom of the lattice. One may suspect that these are all the atoms, but unfortunately this is not true as the following example demonstrates.

Example 7. Let $X = \{x, y, z, w\}$. Define P as follows: The four relations $\succeq$ that rank x and y as the top two elements are collected into one region A. The remaining orders are partitioned so that two orders are equivalent if and only if the first-ranked and last-ranked elements are the same in both.

First, we claim that P is elicitable with lotteries. Indeed, consider the following mechanism g: If A is announced then the output is a random draw between x and y. If any other set is announced then the top-ranked element is selected with probability 0.5, and each of the two "middle" elements is selected with probability 0.25. It's not hard to check that g is IC.

Second, P is not equal to the join of some collection of positional partitions. In fact, P is an atom. We do not provide the details here, but remark that this can be deduced from the analysis of the extreme rays of the set of convex capacities in [12] (see the first extreme point listed on page 14).

5 Multiple agents

In this section we show that our analysis of elicitability can be extended straightforwardly to multi-agent setups. To make the point we focus here on the case of lotteries, but it should be clear that the results apply (with the necessary changes) to the acts framework.

Let $N = \{1, \ldots, n\}$ be the set of agents. For each $i \in N$ a partition P_i of O is given, and we let $P = (P_1, \ldots, P_n)$ denote the profile of partitions. We use A_i for a typical element of P_i, and $A = (A_1, \ldots, A_n)$ for a profile of such elements. As usual, a subscript $-i$ indicates that the ith coordinate of a vector is omitted.

A P-adaptable mechanism is a mapping $g : P \to \Delta(X)$. Thus, for every $A = (A_1, \ldots, A_n) \in P$ the lottery $g(A) \in \Delta(X)$ is the output of the mechanism when each agent $i \in N$ announces that his preference is in A_i.

Definition 9. A P-adaptable mechanism g is Dominant-strategy IC (DIC) if for every $i \in N$, every $\succeq_i \in O$, every $A_i \in P_i$ with $A_i \neq P_i(\succeq_i)$, and every $A_{-i} \in P_{-i}$

$$g(P_i(\succeq), A_{-i}) \succ_i^* g(A_i, A_{-i}).$$

Notice that the above definition corresponds to the standard notion of a dominant-strategy mechanism, where truthfully reporting one's type is optimal regardless of other agents' reports. However, as in the previous sections, we require strict domination.

Definition 10. A profile of partitions $P = (P_1, \ldots, P_n)$ is DIC-elicitable if there exists a DIC P-adaptable mechanism $g : P \to \Delta(X)$.

Proposition 9. The profile of partitions $P = (P_1, \ldots, P_n)$ is DIC-elicitable if and only if P_i is elicitable for each $i \in N$.

Proof. If one of the P_i's is not elicitable, then clearly P is not DIC-elicitable (just fix an arbitrary A_{-i}). In the other direction, suppose that every P_i is elicitable and let $g_i : P_i \to \Delta(X)$ be a P_i-adaptable IC mechanism. For every $A \in P$ define

$$g(A) = \frac{1}{n} \sum_{i=1}^{n} g_i(A_i).$$

Since $g_i(P_i(\succeq_i)) \succ_i^* g_i(A_i)$, it follows that $g(P_i(\succeq_i), A_{-i}) \succ_i^* g(A_i, A_{-i})$ for any A_{-i}, and the result follows. □

Remark 6. A similar result to Proposition 9 holds if one replaces the notion of DIC by Bayes incentive compatibility. Namely, let μ be a full-support product distribution over $\times_{i \in N} P_i$. Given a P-adaptable mechanism $g : P \to \Delta(X)$, $i \in N$, and $A_i \in P_i$, let $\mathbb{E}_{\mu_{-i}}[g(A_i, A_{-i})]$ be the expectation of $g(A_i, A_{-i})$ when A_{-i} is distributed according to the marginal of μ on P_{-i}. Say that g is Bayesian IC (BIC) if $\mathbb{E}_{\mu_{-i}}[g(P_i(\succeq), A_{-i})] \succ^* \mathbb{E}_{\mu_{-i}}[g(A_i, A_{-i})]$ for every i, every $\succeq$, and every $A_i \neq P_i(\succeq)$. Finally, say that P is BIC-elicitable under μ if there exists a BIC P-adaptable mechanism g. It is not hard to show that

P is BIC-elicitable under μ if and only if each of the P_i's is elicitable. We note that the assumption that μ is a product measure is important for this result.

References

[1] A. Abdulkadiroğlu, P.A. Pathak, and A.E. Roth (2005), The New York city high school match, *American Economic Review* **95**, 364-367.

[2] Y. Azrieli, C. Chambers, and P.J. Healy (2018), Incentives in experiments: A theoretical analysis, *Journal of Political Economy* **126**, 1472-1503.

[3] Y. Azrieli, C. Chambers, and P.J. Healy (2019), Incentives in experiments with objective lotteries, *Experimental Economics* forthcoming.

[4] G. Carroll (2012), When are local incentive constraints sufficient?, *Econometrica* **80**, 661-686.

[5] P.H. Edelman and R.E. Jamison (1985), The theory of convex geometries, *Geometriae Dedicata* **19**, 247-270.

[6] A. Gibbard (1977), Manipulation of schemes that mix voting with chance, *Econometrica* **45**, 665-681.

[7] T. Ichiishi (1981), Super-modularity: Applications to convex games and to the greedy algorithm for LP, *Journal of Economic Theory* **25**, 283-286.

[8] N.S. Lambert, D.M. Pennock, and Y. Shoham (2008), Eliciting properties of probability distributions, *ACM EC* 129-130.

[9] N.S. Lambert (2018), Elicitation and evaluation of statistical forecasts, working paper.

[10] K. Osband (1985), Providing incentives for better cost forecasting, Ph.D. thesis, University of California, Berkeley.

[11] K. Osband (1989), Optimal forecasting incentives, *Journal of Political Economy* **97**, 1091-1112.

[12] L. Shapley (1971), Cores of convex games, *International Journal of Game Theory* **1**, 11-26.