

CONSTRAINED PREFERENCE ELICITATION[†]

YARON AZRIELI*, CHRISTOPHER P. CHAMBERS** AND PAUL J. HEALY***

ABSTRACT. A planner wants to elicit information about an agent’s preference relation, but not the entire ordering. Specifically, preferences are grouped into “types,” and the planner only wants to elicit the agent’s type using a (possibly random) incentive mechanism. Assuming beliefs about randomization are subjective, a space of types is elicitable if and only if each type is defined by what the agent would choose from a given list of menus. When beliefs are objective additional type spaces can be elicited, though a convexity condition on types must be satisfied. These results remain unchanged when we consider a setting with multiple agents.

Keywords: Elicitation; incentive compatibility; random mechanisms.

JEL Classification: D8, C7.

I. INTRODUCTION

In many mechanism design and social choice settings the planner collects only partial information about agents’ preferences, rather than their complete ordering over all alternatives. For example, students in the New York City High School match are only asked to report their top 12 schools, even though hundreds are available (Abdulka-dirođlu et al., 2005; Haeringer and Klijn, 2009). In lab experiments, researchers often measure parameters of subjects’ preferences by eliciting a single choice from a convex set, rather than the entire preference ordering (Kroll et al., 1988; Loomes, 1991). A firm surveying customers before the release of a new product only needs a sense of how often their product would be the consumer’s most-preferred; a complete ranking of similar products is unnecessary (Gabor and Granger, 1977). In other settings—for example, surveys on sensitive topics (Warner, 1965)—partial information is elicited to help respect the privacy of respondents. Yet most of the mechanism design literature focuses on direct revelation mechanisms in which entire orderings are revealed, ignoring these practical constraints.

In this paper we study a planner who wants to elicit only partial information about preferences, and we ask what kinds of partial information can be elicited in an incentive

[†]The authors thank seminar audiences at The Technion and Bar Ilan University for useful comments.

*Dept. of Economics, Ohio State University, azrieli.2@osu.edu.

**Dept. of Economics, Georgetown University, cc1950@georgetown.edu.

***Dept. of Economics, Ohio State University, healy.52@osu.edu.

compatible way. For example, suppose there are three alternatives, x , y , and z , and the planner only wants to know which is the agent’s most-preferred. It is easy to incentivize an agent to reveal this information: ask him to pick an item and pay him what he picks. But not all partial information can be incentivized. As an example, suppose the planner is really only interested in learning whether x is the agent’s most-preferred alternative or not, and that learning how he ranks y versus z would somehow impinge on the agent’s privacy. Thus, the planner only wants to know whether the most-preferred item is “ x ” or “ y or z .” We show that there is no way to guarantee that the agent has a strict incentive to reveal his answer truthfully.¹ In general, our goal is to understand what information about preferences *can* be elicited in an incentive compatible way, and what cannot.

We model a planner as having a partition of the space of preferences, and her objective is to identify to which partition element the agent’s true preference belongs. In the first example above there are three partition elements: The first contains those orderings that rank x at the top, the second contains those that rank y at the top, and the third contains those that rank z at the top. In the second example the latter two sets are merged, giving a partition with only two elements. We refer to each element of the partition as a possible *type* of the agent. The entire partition is therefore called a *type space*. We assume strict preferences and study random mechanisms: Agents report their type and are paid a lottery over alternatives. Given this framework, our goal is to characterize those type spaces for which there exists a strictly incentive compatible random mechanism, meaning the agent always has a strictly dominant strategy to report his type truthfully.

Our focus is slightly different from the standard approach in the mechanism design literature. Usually the planner has in mind some objective (such as efficiency or revenue maximization), and incentive compatibility is a constraint on that objective. The restriction to a particular type space might be given as an explicit constraint (e.g., for privacy or simplicity), or may just be a natural feature of the given objective. Here, our goal is more modest: We take the type space as given and ask whether *any* incentive compatible mechanism exists. On the other hand, we require strict incentive compatibility while most of the mechanism design literature only requires a weak version of this constraint.

We begin with a simple sufficiency result. To gain some insight, consider the mechanism from the first example above, where the agent announces his most-preferred element from the set $\{x, y, z\}$ and is paid his announcement. We could construct a more

¹By “guarantee,” we mean regardless of his true preference ordering *and* his risk preferences. For example, if reporting “ x ” pays x and reporting “ y or z ” pays a 50-50 lottery over y and z , then an agent with $y > x > z$ might lie and report “ x ” because of risk aversion. To prevent this possibility we require that truth-telling gives a random payment that stochastically dominates the payment from any misreport.

general type of mechanism where the agent announces his most-preferred element from multiple sets, such as $\{x, y\}$ and $\{y, z\}$; we call each of these sets a “menu.” The mechanism randomly selects *one* of these menus and pays the agent’s announced choice from that menu. Such a mechanism is called a Random Problem Selection (RPS) mechanism, and is known to be strictly incentive compatible (see Azrieli et al., 2018, for example). Notice the type space that this mechanism elicits: Each possible vector of choices from the list of menus defines a type. For example, the type that picks y from both $\{x, y\}$ and $\{y, z\}$ is the set of preferences that rank y at the top. The type that picks x from $\{x, y\}$ and y from $\{y, z\}$ contains the single preference ordering $x > y > z$. And so on. In general, any list of menus will generate a specific type space in this way. We say the resulting type space is *generated by top elements*. Our simple sufficiency result is that every type space generated by top elements is elicitable, because one can simply use the corresponding RPS mechanism.

Our main question then is whether there are other kinds of type spaces that are elicitable. In a Savage setting—where agents do not necessarily have probabilistic beliefs about the likelihood of each item being paid—we show that in fact it *is* necessary: A type space can be elicited if and only if it is generated by top elements.

If, however, subjects’ preferences respect the objective probability distributions given by the mechanism then their beliefs are much more restricted, and so more can be elicited. As an example, suppose there are m alternatives and the planner wants to elicit an agent’s *least* preferred among them. This type space is not generated by top elements, but can be elicited: An agent who announces x as their least-preferred alternative gets paid a lottery in which everything *except* x is paid with equal probability, and x is paid with probability zero. This is incentive compatible because the agent’s ranking of all other alternatives is irrelevant: those are all paid with equal probability, so the only thing that matters is which option is ranked at the bottom. But this construction definitely requires exact indifference between the $m - 1$ unannounced alternatives, which only occurs because they are all paid with the exact same probability. In the Savage framework we cannot control the agent’s beliefs, and thus cannot guarantee this indifference. In that setting, no mechanism can reliably elicit the agent’s least-preferred alternative.

Clearly, any type space that can be elicited in the acts framework can also be elicited with objective lotteries. Extending the example above, we can show that any type space defined by the top *set* of alternatives from menus (rather than the top single alternative) can also be elicited. This can be done by paying everything in the top set with equal probability.² We then demonstrate with an example that even more type spaces

²In the example above the mechanism elicited the top $m - 1$ alternatives.

can be elicited with lotteries. While we do not have a general characterization with lotteries, we show that a certain convexity condition must be satisfied for a type space to be elicitable. If we restrict attention to type spaces that only provide information about which alternatives occupy which rankings—we call these *positional* type spaces—then this convexity condition becomes sufficient. We therefore achieve a complete characterization of the set of elicitable positional type spaces.

We also explore the structure of the set of elicitable type spaces. In both frameworks the set forms a lattice: if two type spaces (again, modeled as partitions) are elicitable, then so is their join.³ The finest type space—in which every ordering is its own type—is elicitable and finer than any other elicitable type space. And the coarsest type space—in which all orderings belong to a single type—is elicitable and coarser than any other elicitable type space.⁴ Thus, the information that can be elicited ranges from “everything” to “nothing,” and is related by the join operator. Furthermore, we show that in the acts framework there is a set of “basic” type spaces in the lattice from which all other non-trivial type spaces in the lattice can be constructed. These are the type spaces generated by the top element from a single menu. Any type space generated by top elements from multiple menus is the join of these single-menu type spaces.

Finally, we extend our results to the case of multiple agents in a social choice framework where all agents must be paid the same lottery over alternatives. The extension is trivial: If each agent’s type space is elicitable on its own, then all of their type spaces can be elicited jointly. This is done using a mechanism that is a convex combination of the mechanisms that would elicit each type space individually. The resulting mechanism is dominant strategy incentive compatible (DIC); each agent has a strict incentive to reveal their type truthfully, regardless of others’ announcements. This provides a complete characterization: a joint type space can be elicited using a DIC mechanism if and only if each individual type space is elicitable.

Related Literature

Probably the most closely related works are Lambert et al. (2008) and Lambert (2018). 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

³The join is the coarsest type space that is finer than either of the two.

⁴The finest type space is generated by the top elements from the list of all binary menus. The coarsest type space is generated by the “top element” from any singleton menu $\{x\}$.

class of probability 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 Lambert et al. (2008) (and also Osband, 1985), 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.

Our work is also reminiscent of the proper scoring rule literature (Brier, 1950; Savage, 1971; Schervish, 1989) because we require strict incentive compatibility. Elicitation with weak incentive compatibility can always be achieved via constant payments, but then it may not be reliable to interpret reports as truthful.

More generally, this paper is also related to our previous work (Azrieli et al., 2018) on eliciting choices from list of menus. There the list of menus (and thus, the type space) is given, and the goal is to identify the set of incentive compatible mechanisms that can be applied to any such type space. Here, we ask for which type spaces does there exist an incentive compatible mechanism. Unlike Azrieli et al. (2018), we also study type spaces that are not generated by any list of menus. Finally, Gibbard (1977) is a classic work on eliciting preferences from multiple individuals using random mechanisms; our work extends his by requiring strict incentive compatibility and eliciting only partial information instead of the complete ranking.

II. 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 complete strict orders (complete, reflexive, transitive, and antisymmetric binary relations) on X . Typical elements of O are \geq, \geq', \geq'' , etc. We write $x > y$ when $x \geq y$ and $x \neq y$.⁶ In examples we often write xyz to denote the relation \geq for which $x > y > z$.

Our notion of a type is simply a set of preference orderings that may (or may not) share some common property. A type space is then a collection of possible types.

⁵In 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.

⁶To be clear, “indifference” only occurs between x and itself: If $x \neq y$ then $x \geq y$ indicates a strict ordering of x and y ; the only case where both $x \geq y$ and $y \geq x$ is when $x = y$.

Definition 1. A type $t \subseteq O$ is a set of preference orderings in O . A type space $T = \{t_1, \dots, t_k\}$ is a collection of types that forms a partition of O .⁷

For example, if $X = \{x, y, z\}$ then one possible type is $t_1 = \{xyz, xzy\}$, which is the set of orders whose maximal element in X is x . Similarly define $t_2 = \{yxz, yzx\}$ and $t_3 = \{zxy, zyx\}$. The type space $T = \{t_1, t_2, t_3\}$ forms a partition of O . In this type space an agent's type reveals their most-preferred element from X , and nothing more.

Given T and \succeq , let $t(\succeq)$ be the (unique) type in T that contains \succeq . The finest possible type space is the one for which $t(\succeq) = \{\succeq\}$ for every $\succeq \in O$. This type space reveals the agent's entire preference ordering. Denote the finest type space by $\overline{T} = \{\{\succeq\}_{\succeq \in O}\}$. The coarsest type space is $\underline{T} = \{O\}$; this type space lumps together all orderings into one type and therefore conveys no information about the agent's preferences.

In general, type space T' refines type space T if for every $t' \in T'$ there is some $t \in T$ for which $t' \subseteq t$. Clearly, the finest type space \overline{T} refines every type space, and the coarsest type space \underline{T} is refined by every type space. For any two type spaces T and T' , the join $T \vee T'$ is the least upper bound according to the refinement relation (equivalently, the coarsest common refinement). It is comprised of all non-empty sets of the form $t' \cap t$, where $t \in T$ and $t' \in T'$.

The planner does not know the agent's true ordering \succeq or her true type $t(\succeq)$. Instead, the agent is asked to announce a type $t \in T$ and is paid an element of X based on her announcement. We allow for random payments, and we consider two possible settings: In one the subject views random payments as objective lotteries over X . In the other the subject may have her own subjective beliefs regarding the likelihood of different outcomes of the randomization device, so random payments are modeled as acts. We now define incentive compatibility for each of these two settings.

Incentive Compatibility with Lotteries

Let $\Delta(X)$ be the set of lotteries on X . We use p , q , and r to denote typical 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 . We recall the following standard definition:

Definition 2. 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 $p \succ^* q$).

⁷ $T = \{t_1, \dots, t_k\}$ forms a partition of O if t_1, \dots, t_k are all non-empty, pairwise disjoint, and $\bigcup_{i=1}^k t_i = O$.

Recall that $p \succeq^* q$ implies that p is preferred to q for any expected utility maximizer consistent with \succeq , regardless of their risk preferences. This interpretation will be used extensively in Section IV below when analyzing incentive compatibility.

A T -mechanism with lotteries is a mapping $g : T \rightarrow \Delta(X)$. The interpretation is that the agent announces an element of T and the mechanism outputs a lottery over X . Incentive compatibility is defined with respect to first-order stochastic dominance:

Definition 3. A T -mechanism with lotteries g is *incentive compatible (IC)* if for every $\succeq \in O$ and every $t \neq t(\succeq)$

$$g(t(\succeq)) \succ^* g(t).$$

By using first-order stochastic dominance we require no assumptions on risk preferences. Put another way, types do not reveal agents' risk preferences, and all risk preferences of a given type are considered possible. Notice also that Definition 3 requires that the lottery obtained by truth-telling *strictly* dominates any other obtainable lottery.

Our interest is in understanding for which type spaces an IC mechanism exists; we call such type spaces elicitable.

Definition 4. A type space T is *elicitable with lotteries* if there exists an IC lotteries T -mechanism.

Incentive Compatibility with Acts

Here we model random payments as acts that map the states of some randomizing device into X . If the state space of the randomizing device is Ω , then the set of possible acts is X^Ω . Thus, a T -mechanism with acts is a pair (Ω, f) , where Ω is a finite state space and $f : T \rightarrow X^\Omega$ is the payment function.

Notice that $f(t) \in X^\Omega$ is the act paid under announcement t , and $f(t)(\omega) \in X$ is the final alternative paid in state $\omega \in \Omega$. When discussing a T -mechanism with acts (Ω, f) we often refer to it only by its payment function f , rather than the pair (Ω, f) .

With acts, our notion of incentive compatibility requires that the act paid under the truthful announcement dominates (state-by-state) the act paid under any other announcement. We additionally require strict incentive compatibility, meaning the truthful announcement must lead to a strictly-preferred outcome in at least one state. One possible justification is that the planner believes that every subjective expected utility agent (with full-support beliefs) is possible. But Definition 5 below guarantees that truthfulness is the unique best-response for other classes of preferences over acts as well, including preferences that lack probabilistic sophistication as defined in Machina and Schmeidler (1995). For example, the subject may be uncertainty averse, e.g., she may have maxmin preferences á-la Gilboa and Schmeidler (1989). All that we require is

that her preferences respect statewise dominance, which is true of almost every decision-theoretic model of ambiguity in the literature.

Definition 5. A T -mechanism with acts (Ω, f) is *incentive compatible (IC)* if for every $\succeq \in O$, every $t \neq t(\succeq)$, and every $\omega \in \Omega$,

$$f(t(\succeq))(\omega) \succeq f(t)(\omega),$$

with a strict preference (meaning $f(t(\succeq))(\omega) \neq f(t)(\omega)$) for some $\omega \in \Omega$.

Definition 6. A type space T is *elicitable with acts* if there exists an IC acts T -mechanism.

We begin our analysis in the acts framework, where a complete characterization of elicitable type spaces is reasonably straightforward.

III. ELICITABLE TYPE SPACES WITH ACT PAYMENTS

In this section we characterize type spaces that are elicitable under acts. Throughout the section we omit the qualifier “with acts” when referring to a T -mechanism or to elicibility of a type space.

To state the result we will need one more piece of notation: For every $\succeq \in O$ and subset of alternatives $X' \subseteq X$, denote by $dom_{\succeq}(X')$ the (unique) maximal element in X' according to \succeq . That is $dom_{\succeq}(X') \in X'$ and $dom_{\succeq}(X') \succeq x$ for every $x \in X'$.

To understand our characterization, we first begin with the preliminary observation that the finest type space is elicitable.

Observation 1. For any X the finest type space $\bar{T} = \{\{\succeq\}_{\succeq \in O}\}$ is elicitable.

This type space can be elicited using the following mechanism: First, let $X_1, X_2, \dots, X_{\binom{m}{2}}$ be the collection of all two-element menus of X (so $X_1 = \{x, y\}$, $X_2 = \{x, z\}$, and so on). When the agent submits a type $t(\succeq) = \{\succeq\}$, the mechanism randomly selects one of these two-element menus and pays the item from that menu that \succeq ranks higher. Formally, there are $\binom{m}{2}$ two-element menus that could be chosen, so the state space of the mechanism is $\Omega = \{\omega_1, \dots, \omega_{\binom{m}{2}}\}$ and the payment function is given by $f(t(\succeq))(\omega_i) = dom_{\succeq}(X_i)$ for each $i \in \{1, \dots, \binom{m}{2}\}$. This is incentive compatible because an agent with true preference \succeq who announces $t(\succeq)$ receives the \succeq -most-preferred alternative from X_i in each state ω_i , while announcing any other $t(\succeq')$ would lead to the \succeq' -most-preferred alternative. The latter is never strictly preferred (under the true ordering \succeq), and in at least one X_i must be strictly less-preferred (since $\succeq \neq \succeq'$).

Observation 1 means that we can always elicit all information about preferences. Given any other type space T , one could simply use the above mechanism to elicit the agent’s entire preference ordering, which would of course reveal $t(\succeq)$ in T . But this

would elicit more information than needed to identify types in T . We require that only the information contained in T be elicited, and nothing more. This may be to protect the privacy of the agent, to keep reports simple (as in the New York City schools example), or perhaps to keep the mechanism itself simple.

The mechanism used to elicit \bar{T} is an example of a *Random Problem Selection (RPS)* mechanism. In general, an RPS mechanism is defined by a list X_1, \dots, X_l of subsets (or, menus) from X . The agent reports her top element in each of the menus, one menu is randomly chosen, and the agent is paid her most-preferred alternative from that menu. Clearly, the amount of information elicited depends on which menus are used. If all two-element subsets are used then the entire preference relation is elicited. If only $X_1 = X$ is used then only the most-preferred alternative from X is elicited. In this way each list of menus X_1, \dots, X_l can be identified with the type space it elicits through the RPS mechanism.

But not every type space corresponds to some list of menus. For example, the type space $T = \{\{xyz, xzy\}, \{yxz, zxy, yzx, zyx\}\}$ does not correspond to any list of menus. Thus, T cannot be elicited using any RPS mechanism. Our characterization shows that, when this is the case, T cannot be elicited using *any* mechanism, i.e., T is not elicitable.

Characterization

We begin by formalizing those type spaces that can be elicited using an RPS mechanism.

Definition 7. A type space T is *generated by top elements* if there are $l \geq 1$ menus $X_1, \dots, X_l \subseteq X$ such that $t(\geq) = t(\geq')$ if and only if, for all $1 \leq i \leq l$, $\text{dom}_{\geq'}(X_i) = \text{dom}_{\geq}(X_i)$. Denote by $\tilde{T}(X_1, \dots, X_l)$ the type space generated by observing the top elements of the menus X_1, \dots, X_l .

We now have our main result of this section.

Proposition 1. A type space T is elicitable if and only if it is generated by top elements.

Proof. (If) Suppose $T = \tilde{T}(X_1, \dots, X_l)$. This can be elicited using an RPS mechanism with menus X_1, \dots, X_l . Specifically, let $\Omega = \{\omega_1, \dots, \omega_l\}$ and for each $t \in T$ choose an arbitrary representative $\geq^t \in t$. Define $f(t)(\omega_i) = \text{dom}_{\geq^t}(X_i)$ for $i = 1, \dots, l$. Note that by assumption the choice of the representative \geq^t does not affect the resulting mechanism.

To see that the above mechanism is IC fix some \geq and some $t \in T$. Then for each i we have

$$f(t(\geq))(\omega_i) = \text{dom}_{\geq^{t(\geq)}}(X_i) = \text{dom}_{\geq}(X_i) \geq \text{dom}_{\geq^t}(X_i) = f(t)(\omega_i).$$

The first equality is by the definition of f . The second follows from the fact that T is generated by top elements, so that \geq and $\geq^{t(\geq)}$ have the same most-preferred element in

every menu. The next relation follows from the definition of dom , and the last equality is again by construction of f . Moreover, if $t \neq t(\geq)$ then there exists i such that $dom_{\geq}(X_i) \neq dom_{\geq t}(X_i)$ which gives a strict preference at ω_i .

(*Only If*) Suppose T is elicitable and let (Ω, f) be an IC T -mechanism. Enumerate the states so that $\Omega = \{\omega_1, \dots, \omega_l\}$ for some positive integer l , and for each $i = 1, \dots, l$ define $X_i = \{f(t)(\omega_i)\}_{t \in T} \subseteq X$. 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 $T = \tilde{T}(X_1, \dots, X_l)$. Suppose that \geq, \geq' are in the same element of T , and fix some $1 \leq i \leq l$. Then by incentive compatibility we have that $f(t(\geq))(\omega_i) \geq f(t)(\omega_i)$ for every $t \in T$, which implies that $f(t(\geq))(\omega_i) = dom_{\geq}(X_i)$. Applying the same argument to \geq' gives $f(t(\geq'))(\omega_i) = dom_{\geq'}(X_i)$. But since $t(\geq) = t(\geq')$ we get $dom_{\geq}(X_i) = dom_{\geq'}(X_i)$. Repeating for each $i = 1, \dots, l$ shows that \geq, \geq' are in the same element of $\tilde{T}(X_1, \dots, X_l)$.

Conversely, suppose that \geq, \geq' are in the same element of $\tilde{T}(X_1, \dots, X_l)$. From the previous paragraph we have $f(t(\geq))(\omega_i) = dom_{\geq}(X_i)$ and $f(t(\geq'))(\omega_i) = dom_{\geq'}(X_i)$ for each i , so $f(t(\geq)) = f(t(\geq'))$. Incentive compatibility now implies that $t(\geq) = t(\geq')$, which concludes the proof. \square

The Lattice of Elicitable Type Spaces

Consider two elicitable type spaces T and T' . Proposition 1 implies that each one of them is generated by observing the top elements of some collection of menus, say $T = \tilde{T}(X_1, \dots, X_l)$ and $T' = \tilde{T}(X'_1, \dots, X'_k)$. But it is immediate to verify that $\tilde{T}(X_1, \dots, X_l) \vee \tilde{T}(X'_1, \dots, X'_k) = \tilde{T}(X_1, \dots, X_l, X'_1, \dots, X'_k)$, and thus that $T \vee T'$ is elicitable as well. We state this fact in the next corollary.

Corollary 1. If T and T' are both elicitable, then so is their join.

It follows that the set of elicitable type spaces equipped with the refinement relation forms a lattice. In this lattice, the meet of T' and T'' is the elicitable type space T such that both T' and T'' refine T , and no elicitable refinement of T is refined by both T' and T'' . This exists because all type spaces refine \underline{T} (the coarsest type space), which is elicitable. The meet, however, does not coincide with the finest common coarsening of the type spaces, because the finest common coarsening may not be elicitable. This is illustrated in the following example:

Example 1. Let $X = \{x, y, z\}$, and $T' = \{\{xyz, xzy\}, \{yxz, yzx\}, \{zxy, zyx\}\}$. This type space corresponds to observing the top element of X . Let $T'' = \{\{xyz, xzy\}, \{zyx, yzx\}, \{zxy\}, \{yxz\}\}$. This type space corresponds to observing the top elements of the menus $\{x, y\}$ and $\{x, z\}$. The finest common coarsening of T' and T'' is $\{\{xyz, xzy\}, \{yzx, yxz, zxy, zyx\}\}$. This type

space is not generated by top elements and hence not elicitable. In fact, the meet of T' and T'' in the lattice of *elicitable* type spaces is the coarsest type space \underline{T} .

The set of type spaces generated by observing the top element of a single, non-trivial menu (given by $\{\tilde{T}(X') : X' \subseteq X, |X'| \geq 2\}$) forms the set of *atoms* of the lattice of elicitable type spaces. Moreover, any elicitable type space can be formed by taking the join of some collection of atoms. This is because $\tilde{T}(X_1, \dots, X_l) = \tilde{T}(X_1) \vee \tilde{T}(X_2) \vee \dots \vee \tilde{T}(X_l)$. Thus, the lattice of elicitable type spaces is *atomistic*.

Finally, imagine that a principal is interested in eliciting a type space T , but T is not elicitable. Of course she could elicit the finest type space \overline{T} , but she wants to minimize the excess information she collects. Is there a unique coarsest elicitable refinement of T that she could elicit instead? The next example demonstrates that, in general, the answer is no.

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

$$\tilde{T}(\{x, y, z\}, \{y, z\}) = \{\{xzy\}, \{yzx, yxz\}, \{xyz\}, \{zxy, zyx\}\}, \text{ and}$$

$$\tilde{T}(\{x, z\}, \{y, z\}) = \{\{xyz, yxz\}, \{yzx\}, \{xzy\}, \{zxy, zyx\}\}.$$

Clearly there is no coarser refinement of T , so there is no unique coarsest elicitable refinement.

From Menus to Type Spaces and Back Again

Since elicitable type spaces in the acts framework are exactly those that are generated by top elements, it is useful to understand the structure of such type spaces as well as the connection between lists of menus and the type spaces they generate. For a type space T and a menu $X' \subseteq X$, say that X' is identified by T if for every $t \in T$ and every $\succeq, \succeq' \in t$ it holds that $dom_{\succeq}(X') = dom_{\succeq'}(X')$; in other words this means that $dom_{\bullet}(X')$ is a T -measurable function from O to X . Intuitively, if X' is identified by T then knowing the agent's type $t \in T$ will also reveal his most-preferred element in X' (even if it doesn't perfectly reveal \succeq). Let $\tilde{I}(T)$ be the collection of menus that are identified from T .⁸

The following proposition provides a simple way to check whether a given type space is generated by top elements.

Proposition 2. A type space T is generated by top elements if and only if $T = \tilde{T}(\tilde{I}(T))$.

Omitted proofs appear in the appendix.

⁸We exclude singleton and empty menus from $\tilde{I}(T)$.

While $\tilde{I}(T)$ contains all the menus that are identified by T , it is possible that some of them are redundant when generating T . For example, $\tilde{I}(\bar{T})$ contains every subset of X , but \bar{T} can also be generated using only the two-element subsets. In that sense the larger subsets are redundant and can be dropped.

In fact, we can show that (1) this process of dropping redundant menus always leads to a unique smallest (in terms of inclusion) list of menus that still generates T , and (2) any combination of those redundant menus can be kept or dropped and T will still be generated.

To characterize the smallest list we first provide a formal notion of redundancy, and lack thereof.

Definition 8. Let X', X_1, \dots, X_l be menus. Say that X' is *surely identified* by X_1, \dots, X_l if whenever \succeq, \succeq' are such that $\text{dom}_{\succeq}(X_i) = \text{dom}_{\succeq'}(X_i)$ for all $i = 1, \dots, l$, it also holds that $\text{dom}_{\succeq}(X') = \text{dom}_{\succeq'}(X')$. The menus X_1, \dots, X_l are *independent* if no X_i is surely identified by $\{X_j\}_{j \neq i}$.

We can now formalize the interval result.

Proposition 3. For any elicitable type space T there is a unique independent list of menus, denoted $B(\tilde{I}(T))$, such that $T = \tilde{T}(B(\tilde{I}(T)))$. Furthermore, $T = \tilde{T}(X_1, \dots, X_l)$ if and only if

$$B(\tilde{I}(T)) \subseteq \{X_1, \dots, X_l\} \subseteq \tilde{I}(T).$$

The list $B(\tilde{I}(T))$ is of particular significance as it provides a measure of the amount of information contained in type space T . Specifically, if $B(\tilde{I}(T)) = \{X_1, \dots, X_l\}$ then type space T can be thought of as encoding l bits of information, which are the top elements from each of these l menus.

IV. ELICITABLE TYPE SPACES WITH LOTTERY PAYMENTS

This section analyzes elicitability in the lotteries framework. Thus, throughout this section T -mechanism refers to lotteries T -mechanism, and elicitable type space refers to a type space that is elicitable with lotteries.

A Sufficient Condition

We begin with the simple observation that any type space that is elicitable in the acts framework is also elicitable with lotteries.

Proposition 4. If T is generated by top elements then T is elicitable.

The proof is simple: Take the RPS mechanism used in the acts framework to elicit T and turn it into a lottery mechanism by assigning any full-support probability distribution to its state space. The resulting mechanism pays in objective lotteries and is incentive compatible; see the appendix for details.

The following example shows that the set of elicitable type spaces is strictly larger than in the acts framework.

Example 3. Let $X = \{x, y, z\}$ and let $T = \{t_1, t_2, t_3\}$ where

$$t_1 = \{xyz, yxz\}, t_2 = \{xzy, zxy\}, t_3 = \{yzx, zyx\}.$$

In words, T reveals the least-preferred alternative from X . Or, equivalently, it reveals the top-2 alternatives but not their order. Type space T is not generated by top elements. This can easily be checked using Proposition 2 above: $\tilde{I}(T)$ contains no non-trivial menus, so $\tilde{T}(\tilde{I}(T)) = \underline{T}$ (the trivial type space) and hence $\tilde{T}(\tilde{I}(T)) \neq T$. However, consider the mechanism g given by

$$g(t_1) = (x, 0.5; y, 0.5; z, 0), \quad g(t_2) = (x, 0.5; y, 0; z, 0.5), \quad g(t_3) = (x, 0; y, 0.5; z, 0.5),$$

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

Example 3 suggests the following natural generalization of type spaces generated by top elements: Instead of revealing the top element in some list of menus, the type space reveals the set of top- k elements (without their order) for each menu in the list, where k may vary across different menus. To formalize this, for $\succeq \in O$, $X' \subseteq X$ with $|X'| \geq 2$, and $k \in \{1, \dots, |X'| - 1\}$ define $dom_{\succeq}^k(X')$ to be the (unique) set of elements satisfying (1) $dom_{\succeq}^k(X') \subseteq X'$; (2) $|dom_{\succeq}^k(X')| = k$; and (3) if $x \in dom_{\succeq}^k(X')$ and $y \in X' \setminus dom_{\succeq}^k(X')$ then $x \succeq y$. In words, $dom_{\succeq}^k(X')$ is the set of top k ranked elements in X' . Note that $dom_{\succeq}^1(X') = dom_{\succeq}(X')$.

Definition 9. A type space T is *generated by top sets* if there are l , menus $X_1, \dots, X_l \subseteq X$, and numbers k_1, \dots, k_l with $1 \leq k_i \leq |X_i| - 1$ for each i , such that $t(\succeq) = t(\succeq')$ if and only if $dom_{\succeq}^{k_i}(X_i) = dom_{\succeq'}^{k_i}(X_i)$ for every $i = 1, \dots, l$. Denote by $\tilde{T}(X_1, \dots, X_l; k_1, \dots, k_l)$ the type space generated by observing the top k_i elements of the menu X_i , $i = 1, \dots, l$.

Obviously, $\tilde{T}(X_1, \dots, X_l) = \tilde{T}(X_1, \dots, X_l; 1, \dots, 1)$, so if a type space is generated by top elements then in particular it is generated by top sets. We have the following strengthening of Proposition 4.

Proposition 5. If T is generated by top sets then it is elicitable.

Proof. We sketch here the argument, leaving some of the details for the reader to fill out. First, for a single menu $X_1 \subseteq X$ and a number $1 \leq k_1 \leq |X_1| - 1$, consider the type space

$T = \tilde{T}(X_1; k_1)$. Let g be the T -mechanism that for each $t \in T$ pays the uniform lottery over $\text{dom}_{\succeq^t}^{k_1}(X_1)$, where \succeq^t is any member of t . Then g is clearly IC, so T is elicitable.

Second, it is straightforward to verify that

$$\tilde{T}(X_1, \dots, X_l; k_1, \dots, k_l) \vee \tilde{T}(X'_1, \dots, X'_{l'}; k'_1, \dots, k'_{l'}) = \tilde{T}(X_1, \dots, X_l, X'_1, \dots, X'_{l'}; k_1, \dots, k_l, k'_1, \dots, k'_{l'}).$$

Therefore, any type space generated by top sets is equal to the join of type spaces of the form $\tilde{T}(X_1; k_1)$. In other words, the collection of type spaces generated by top sets forms an atomistic lattice (with respect to refinement) whose atoms are the type spaces generated by single menus.

Third, if T and T' are both elicitable, then so is $T \vee T'$. Indeed, if g is an IC T -mechanism and g' is an IC T' -mechanism, then the $T \vee T'$ -mechanism defined by $g^*(t \cap t') := (1/2)g(t) + (1/2)g'(t')$ (whenever $t \cap t' \neq \emptyset$) is IC. That is, the collection of elicitable type spaces also forms a lattice when equipped with the refinement ordering.

The combination of the above three claims proves the proposition. \square

At this point it is tempting to guess that the converse of Proposition 5 is true as well, namely, that every elicitable type space is generated by top sets. The following example demonstrates that this is not the case.

Example 4. Let $X = \{x, y, z, w\}$. Define T as follows: The four relations \succeq that rank x and y as the top two elements are collected into one type, denoted t_0 . 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. Namely,

$$T = \left\{ t_0 = \{xyzw, yxzw, xywz, yxwz\}, \{xzyw\}, \{xwyz\}, \{ywzx\}, \{yzxw\}, \right. \\ \left. \{xzw y, xwzy\}, \{yzwx, ywzx\}, \{zxyw, zyxw\}, \{zywx, zwyx\}, \right. \\ \left. \{zxwy, zwxy\}, \{wxyz, wyxz\}, \{wyzx, wzyx\}, \{wxzy, wzxy\} \right\}$$

First, we claim that T is elicitable. Indeed, consider the following mechanism g : If t_0 is announced then the output is a random draw between x and y . If any other type 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, T is not generated by top sets. The easiest way to show that is by checking that there are no menu $X' \subseteq X$ and number $1 \leq k < |X'|$ for which the top- k elements in X' are always identified by T .

This example is based on the first extreme ray of the set of convex capacities in Shapley (1971, p.14). The connection between convex capacities and IC mechanisms has been discussed in our previous work Azrieli et al. (2019).

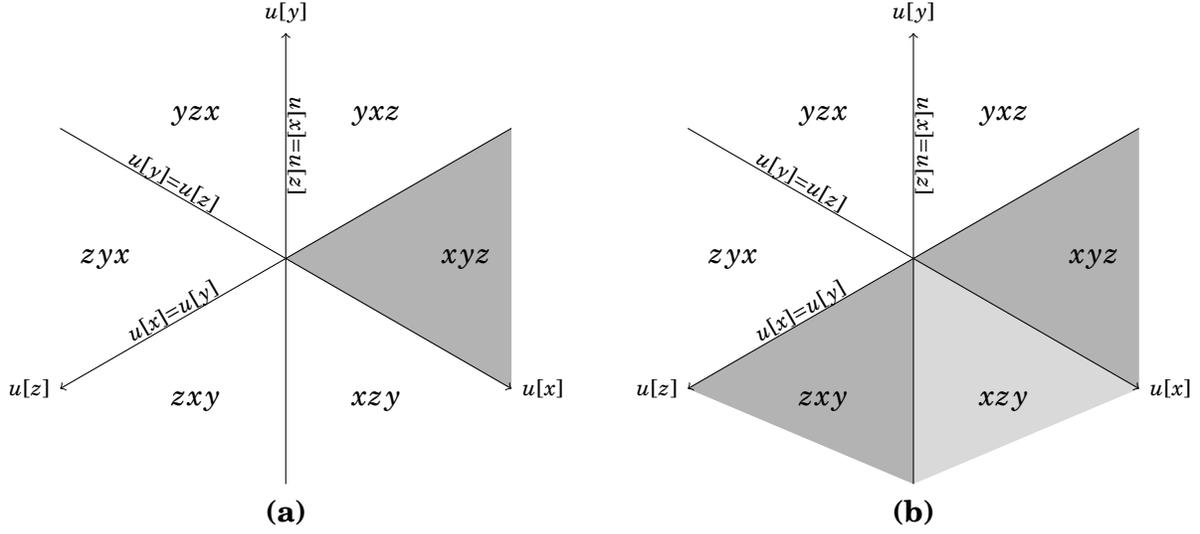


FIGURE I. The sets of cardinal utility vectors associated with each ordering \succeq (denoted by $U(\succeq)$) for $X = \{x, y, z\}$. Panel (a): The set $U(xyz)$ is shaded. Panel (b): For the non-convex type $t = \{xyz, zxy\}$, $U(t) = U(xyz) \cup U(zxy)$ is shaded in dark and $\overline{U(Cons(t))}$ is the union of all three shaded areas.

Necessary Conditions

Given a set of orderings $t \subseteq O$, let $\supseteq_t = \bigcap_{\succeq \in t} \succeq$ be the maximal relation that all orderings in t agree on; that is, $x \supseteq_t y$ if and only if $x \succeq y$ for all $\succeq \in t$. Note that \supseteq_t is a partial order. Let

$$Cons(t) = \{\succeq \in O : (\forall x, y \in X) x \supseteq_t y \Rightarrow x \succeq y\}$$

be the set of orderings that are consistent with t .⁹ In a sense we will make precise, $Cons(t)$ is the convex hull of t . Roughly, if $x \supseteq_t y$ then one of the defining characteristics of t is that all orders in t rank x over y . But t may have “gaps,” in that it may not include all orders that satisfy t ’s defining characteristics. $Cons(t)$ fills in these “gaps” by collecting all such orders.

Definition 10. A set $t \subseteq O$ is *convex* if $t = Cons(t)$. A type space T is convex if every $t \in T$ is convex.

Figure I demonstrates why we call such a type convex. Here we graph all possible cardinal utility vectors over $X = \{x, y, z\}$. Each vector $u \in \mathbb{R}^3$ corresponds to an ordering in O . For example, vector $u = (3, 2, 1)$ (meaning $u[x] = 3$, $u[y] = 2$, and $u[z] = 1$) corresponds to the ordering xyz , since $u[x] > u[y] > u[z]$. The shaded cone in panel (a) represents all utility vectors corresponding to xyz . Define $U(\succeq) = \{u \in \mathbb{R}^X : x > y \Leftrightarrow u[x] > u[y]\}$, so

⁹Equivalently, $Cons(t) = \{\succeq \in O : \supseteq_t \subseteq \succeq\}$.

the shaded region panel (a) is $U(xyz)$. There are six such cones, one for each $\succeq \in O$. The boundaries of these cones are the planes defined by indifference: $u[x] = u[y]$, $u[x] = u[z]$, and $u[y] = u[z]$.¹⁰

For any type (or set) $t \subseteq O$, define $U(t) = \bigcup_{\succeq \in t} U(\succeq)$, and let $\overline{U(t)}$ be its closure. For example, the dark-shaded region in panel (b) of Figure I corresponds to the type $t = \{xyz, zxy\}$. Notice that $\overline{U(t)}$ is not a convex set, and indeed t is not convex according to Definition 10: We have only that $x \succ_t y$, so $Cons(t) = \{xyz, zxy, xzy\} \neq t$. On the other hand, $\overline{U(Cons(t))}$ is convex: it “fills in the gap” in $\overline{U(t)}$.

More generally, if t is convex (as defined above) then $\overline{U(t)} = \overline{U(Cons(t))}$. But $\overline{U(Cons(t))}$ is the intersection of closed half-spaces, which implies that $\overline{U(t)}$ is a convex set. And conversely, if $\overline{U(t)}$ is convex then necessarily $\overline{U(t)} = \overline{U(Cons(t))}$, which implies that $t = Cons(t)$, i.e., that t is a convex type. Thus, we have the following:

Lemma 1. The set $t \subseteq O$ is convex if and only if $\overline{U(t)}$ is convex in \mathbb{R}^X .

Remark. 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 t is convex if and only if for every $\succeq, \succeq' \in t$, if \succeq'' is in a shortest path between \succeq and \succeq' then $\succeq'' \in t$. We omit the proof.

The following proposition shows that convexity is a necessary condition for elicibility in the lotteries framework. Variants of this result in different contexts have been obtained in previous works, see for example Lambert (2018, pp. 10–11) who attributes this observation to Osband (1985).

Proposition 6. If T is elicitable then it is convex.

Unfortunately, convexity of T is not sufficient for elicibility, as shown in Example 5.

Example 5. Let $X = \{x, y, z\}$ and let $T = \{t_1, t_2, t_3\}$, where

$$t_1 = \{xyz\}, \quad t_2 = \{yxz, yzx\}, \quad t_3 = \{zyx, zxy, xzy\}.$$

¹⁰The three planes all intersect along the line of total indifference, defined by $u = (\alpha, \alpha, \alpha)$ for $\alpha \in \mathbb{R}$. The 3-dimensional figure is drawn from a vantage point along this line; from this perspective each plane projects exactly onto one of the three axes. Each cone is actually a three-dimensional cylinder set projecting toward (and away from) the observer: if $u \in U(\succeq)$ then for every $\alpha \in \mathbb{R}$ we have $(u + (\alpha, \alpha, \alpha)) \in U(\succeq)$. Equivalent, the 2-dimensional picture can be thought of as projection of each $U(\succeq)$ onto the plane defined by $\{u : \sum_x u[x] = 0\}$.

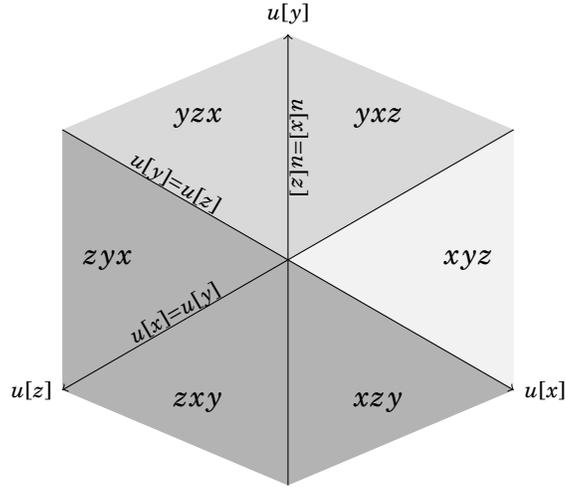


FIGURE II. The type space from Example 5. Type t_1 is shaded off-white, t_2 is light gray, and t_3 is dark gray.

This type space is shown in Figure II. Convexity of T is immediate, as each $\overline{U(t)}$ is convex. We now show, however, that T is not elicitable.

Suppose by contradiction that g is an IC T -mechanism. First we compare t_1 to t_2 by comparing the IC conditions for the two orders that are adjacent to each other in the figure: $xyz \in t_1$ and $yxz \in t_2$. We first show that, because z does not change position between these orders, we must have $g(t_1)(z) = g(t_2)(z)$. To prove this, fix $M > 0$ and take $u_1 = (2, 1, -M) \in U(xyz) = U(t_1)$ and $u_2 = (1, 2, -M) \in U(yxz) \subset U(t_2)$. IC requires $g(t_1) \cdot u_1 > g(t_2) \cdot u_1$ and $g(t_1) \cdot u_2 < g(t_2) \cdot u_2$. Taking M to $+\infty$ gives $g(t_1)(z) = g(t_2)(z)$. Given that, we show that $g(t_1)(x) > g(t_2)(x)$. This is because the first IC inequality reduces to

$$2g(t_1)(x) + 1g(t_1)(y) > 2g(t_2)(x) + 1g(t_2)(y),$$

which, when combined with $g(t_1)(x) + g(t_1)(y) = g(t_2)(x) + g(t_2)(y)$, gives $g(t_1)(x) > g(t_2)(x)$. (The second IC constraint similarly gives $g(t_2)(y) > g(t_1)(y)$.)

Now move from t_2 to t_3 by comparing adjacent orders $yzx \in t_2$ to $zyx \in t_3$. Since x does not change position between these, we can again use a limiting argument (with $u_2 = (-M, 2, 1) \in U(t_2)$ and $u_3 = (-M, 1, 2) \in U(t_3)$) to get that $g(t_2)(x) = g(t_3)(x)$. Finally, move from t_3 back to t_1 by comparing adjacent orders $xzy \in t_3$ and $xyz \in t_1$. Again, x does not change position, so we have $g(t_3)(x) = g(t_1)(x)$.

Through this cycle we have achieved a contradiction: $g(t_1)(x) > g(t_2)(x)$, yet $g(t_2)(x) = g(t_3)(x) = g(t_1)(x)$. Thus, IC is impossible on this convex type space.

The reason that elicibility fails in Example 5 is because we can find a cycle of adjacent types (t_1, t_2, t_3, t_1) where x moves down the ranking in the first step (when crossing

the boundary from t_1 to t_2), but then in no subsequent step does x move up. In general, if there are k types that form a cycle (t_1, \dots, t_k, t_1) in which x moves down from t_1 to t_2 but then never moves up, we'll have

$$g(t_1)(x) > g(t_2)(x) \geq g(t_3)(x) \cdots \geq g(t_k)(x) \geq g(t_1)(x),$$

a contradiction.¹¹

We now provide the definitions necessary to formalize this condition. Recall that $r_{\succeq}(x) = |\{y : y \succeq x\}|$ is the ranking of x in the ordering \succeq , and 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 $z \neq x, y$. Geometrically, this means that $U(\succeq)$ and $U(\succeq')$ share a boundary on the $u[x] = u[y]$ hyperplane. Similarly, say that the sets t and $t' \neq t$ are adjacent via an x - y switch if there are $\succeq \in t$ and $\succeq' \in t'$ that are adjacent via an x - y switch.¹² The sets t and t' are said to be *adjacent* if they are adjacent via some x - y switch.

Proposition 7. Suppose that T is elicitable. For any cycle of adjacent types $(t_1, t_2, \dots, t_k, t_1)$, if t_1 and t_2 are adjacent via an x - y switch then there exist some $1 < i \leq k$ and z such that t_i and $t_{(i+1) \bmod k}$ are adjacent via a z - x switch.

Remark. An implication of Proposition 7 is that if T is elicitable and $t, t' \in T$ are adjacent via an x - y switch, then any other adjacency between these types must be via a y - x switch. If that were not the case then the cycle (t, t', t) would fail the necessary condition.

We do not know whether convexity together with the no-cycles condition of Proposition 7 is enough to guarantee elicibility 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 type spaces in the lotteries framework is therefore still an open question.

Characterization for Positional Type Spaces

We now restrict attention to type spaces that treat all the alternatives symmetrically, and only contain information about which alternatives occupy which positions in the ranking. We refer to these as *positional* type spaces. One example is when the agent's type only identifies his top-ranked alternative from X . Another is the type space in

¹¹At each step the weight on x will strictly decrease when x moves down the ranking or will not change when the ranking of x does not change.

¹²The order of terms matters: If t and t' are adjacent via an x - y switch then t' and t are adjacent via a y - x switch. Also, types t and t' can have multiple adjacencies, but this definition only requires that one of them be via an x - y switch.

which each type reveals the top k rankings ($1 \leq k < m$) from X , but contains no information about the ordering of lower-ranked alternatives. Alternatively, types may reveal the k lowest rankings, but not the relative ordering of the first $m - k$ alternatives.

More generally, we can think about defining types based on what they rank 1st, 2nd, 3rd, and so on, but also we may “lump together” certain rankings that the types don’t distinguish. This is represented by a partition of the numbers $\{1, \dots, m\}$. For example if $m = 4$ then the partition $\{\{1\}, \{2\}, \{3, 4\}\}$ indicates that types are identified by what is ranked first and second, but types do not distinguish between the 3rd- and 4th-ranked alternatives. For example, $\{xyzw, xywz\}$ would be one type in that type space. The partition $\{\{1\}, \{2, \dots, m\}\}$ indicates that types identify only their most-preferred item in X . The partition $\{\{1, \dots, m - 1\}, \{m\}\}$ indicates that types are only identified by their least-preferred item in X . And so on.

To formalize this, think of the ranking function $r_{\geq}(\cdot)$ as a bijection from X to $\{1, \dots, m\}$. If $B \subseteq \{1, \dots, m\}$ then $r_{\geq}^{-1}(B) = \{x \in X : r_{\geq}(x) \in B\}$ is the set of alternatives whose ranking according to \geq is in B . Given a partition Q of $\{1, \dots, m\}$, say that \geq, \ge' have the same Q -rankings if $r_{\geq}^{-1}(B) = r_{\ge'}^{-1}(B)$ for every $B \in Q$. In this way the partition Q defines a type space, which we call the Q -positional type space.

Definition 11. Let Q be a partition of $\{1, \dots, m\}$. The Q -positional type space, denoted T_Q , is the type space in which $\geq, \ge' \in t$ if and only if \geq and \ge' have the same Q -ranking. A type space T is positional if it is Q -positional for some Q .

Example 6. Suppose $X = \{x, y, z\}$. For $Q = \{\{1, 2\}, \{3\}\}$ the type space T_Q is the type space of Example 3, that is, $T_Q = \{\{xyz, yxz\}, \{xzy, zxy\}, \{yzx, zyx\}\}$. For $Q = \{\{1, 3\}, \{2\}\}$ we have $T_Q = \{\{xyz, zyx\}, \{yxz, zxy\}, \{xzy, yzx\}\}$. For $Q = \{\{1\}, \{2\}, \{3\}\}$ we have that $T_Q = \bar{T}$, the finest type space. For $Q = \{\{1, 2, 3\}\}$, $T_Q = \underline{T}$.

We now show that, when we restrict attention to positional type spaces, convexity is not only necessary but also sufficient for a type space to be elicitable. Furthermore, convexity puts structure on the partition of rankings Q : Convex type spaces can only group together ranks that are adjacent. In other words, each element of Q must be an interval. For example, the type space generated by $Q = \{\{1, 3\}, \{2\}\}$ identifies types by their second-ranked alternative; Figure III shows that this is definitely not convex, and therefore not elicitable. In fact, having Q consist entirely of intervals is equivalent to convexity, giving our complete characterization of elicitable positional type spaces:

Proposition 8. Let T be a positional type space. Then the following conditions are equivalent:

- (1) T is elicitable.
- (2) T is convex.

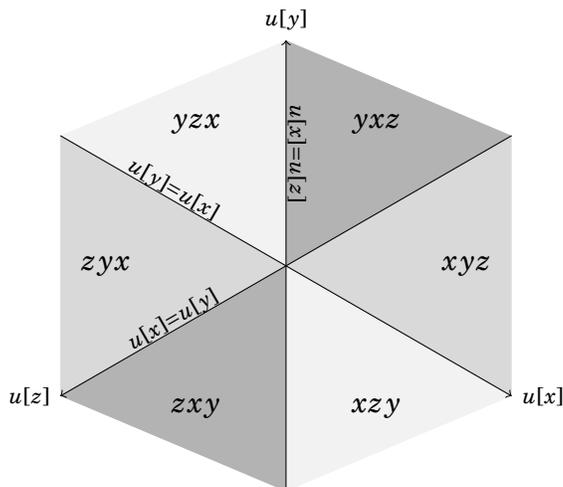


FIGURE III. The non-convex type positional space T_Q generated from $Q = \{\{1, 3\}, \{2\}\}$.

(3) Every element in the partition Q that defines T is a (possibly degenerate) interval in $\{1, \dots, m\}$.

The Lattice of Elicitable Type Spaces under Lotteries

Recall that in the acts framework the lattice of elicitable type spaces is atomistic, with all non-trivial elicitable type spaces being the join of type spaces of the form $\tilde{T}(X_1)$. Recall also from the proof of Proposition 5 that the set of elicitable type spaces under lotteries also forms a lattice. We do not know if the lattice of elicitable type spaces under lotteries is atomistic; however, we now argue that every elicitable lottery is the join of a collection of basic type spaces, though we do not have a good description of these basic spaces.

Consider the finest type space \bar{T} in which $t(\succeq) = \{\succeq\}$ for all $\succeq \in O$. From Lemma 3 in the proof of Proposition 7 the set of IC \bar{T} -mechanisms has the property that if \succeq is adjacent to \succeq' via an x - y switch then $g(\{\succeq\})(x) > g(\{\succeq'\})(x)$, $g(\{\succeq\})(y) < g(\{\succeq'\})(y)$, and $g(\{\succeq\})(z) = g(\{\succeq'\})(z)$ for all $z \neq x, y$. Now take one such g and two orders \succeq and \succeq' that are adjacent via an x - y switch. Alter g to g' , where $g \equiv g'$ except $g'(\{\succeq\}) = g'(\{\succeq'\})$. In other words, g' no longer distinguishes between these two adjacent orders; the strict incentive compatibility condition between them has been weakened to an equality. The type space that g' would elicit is therefore the type space $T_{g'} = \{\{\succeq, \succeq'\}, \{\succeq''\}_{\succeq'' \neq \succeq, \succeq'}\}$, which is the finest type space but with \succeq and \succeq' combined.¹³

¹³Technically g' is not a $T_{g'}$ -mechanism because its domain is \bar{T} instead of $T_{g'}$; however, g' is measurable in $T_{g'}$, so the difference is irrelevant.

Let \bar{G} be the set of mechanisms that come from altering some IC \bar{T} -mechanism in this way. Formally, \bar{G} is the set of \bar{T} -mechanisms satisfying the following “weak” incentive compatibility 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 $z \neq x, y$. The set \bar{G} is defined by finitely many linear inequalities, so it a polyhedral subset of $\Delta(X)^O$. Let $ext(\bar{G})$ be the extreme points of this set.

For every $g \in \bar{G}$ define the type space T_g it elicits by letting \succeq, \succeq' be in the same element of T_g if and only if $g(\succeq) = g(\succeq')$. Clearly we can view g as a T_g -mechanism, and by Carroll (2012) the local incentive constraints for adjacent types also ensure global incentive compatibility of g .¹⁴ Moreover, if $g, g' \in \bar{G}$ and $0 < \alpha < 1$ then $T_{\alpha g + (1-\alpha)g'} = T_g \vee T_{g'}$. Therefore, any elicitable type space is the join of type spaces from the set $\{T_g : g \in ext(\bar{G})\}$. It is these sets that form the basic elicitable type spaces from which all other elicitable type spaces can be formed.

V. MULTIPLE AGENTS

In this section we show that our analysis of elicibility 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 $N = \{1, \dots, n\}$ be the set of agents. For each $i \in N$ a type space T_i of O is given, and we let $T = (T_1, \dots, T_n)$ denote the profile of type spaces. We use t_i for a typical element of T_i , and $t = (t_1, \dots, t_n)$ for a profile of such elements. As usual, a subscript $-i$ indicates that the i th coordinate of a vector is omitted.

A T -mechanism is a mapping $g : T \rightarrow \Delta(X)$. Thus, for every $t = (t_1, \dots, t_n) \in T$ the lottery $g(t) \in \Delta(X)$ is the output of the mechanism when each agent $i \in N$ announces that his preference is in t_i . All agents receive the same lottery $g(t)$, as in a standard social choice setting.

Definition 12. A T -mechanism g is *dominant-strategy IC (DIC)* if for every $i \in N$, every $\succeq_i \in O$, every $t_i \in T_i$ with $t_i \neq t_i(\succeq_i)$, and every $t_{-i} \in T_{-i}$

$$g(t_i(\succeq), t_{-i}) >_i^* g(t_i, t_{-i}).$$

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.

¹⁴See the remark following the proof of Proposition 7 for details.

Definition 13. A profile of type spaces $T = (T_1, \dots, T_n)$ is DIC-elicitable if there exists a DIC T -mechanism $g : T \rightarrow \Delta(X)$.

We now show that elicibility of each T_i on its own is both necessary and sufficient for DIC-elicibility of $T = (T_1, \dots, T_n)$; indeed the T -mechanism that DIC-elicits T is simply the unweighted average of the mechanisms that elicit each T_i .

Proposition 9. The profile of type spaces $T = (T_1, \dots, T_n)$ is DIC-elicitable if and only if T_i is elicitable for each $i \in N$.

Necessity is clear: if one of the T_i 's is not elicitable, then T is not DIC-elicitable (just fix an arbitrary t_{-i}). For sufficiency, suppose that every T_i is elicitable and let $g_i : T_i \rightarrow \Delta(X)$ be an IC T_i -mechanism. For every $t \in T$ define

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

Since $g_i(t_i(\geq_i)) >^*_i g_i(t_i)$, it follows that $g(t_i(\geq_i), t_{-i}) >^*_i g(t_i, t_{-i})$ for any t_{-i} . Thus, g is a DIC T -mechanism.

A similar result to Proposition 9 holds if one replaces the notion of DIC by Bayesian incentive compatibility. Namely, let μ be a full-support product distribution over $\times_{i \in N} T_i$. Given a T -mechanism $g : T \rightarrow \Delta(X)$, $i \in N$, and $t_i \in T_i$, let $\mathbb{E}_{\mu_{-i}}[g(t_i, t_{-i})]$ be the expectation of $g(t_i, t_{-i})$ when t_{-i} is distributed according to the marginal of μ on T_{-i} . Say that g is Bayesian IC (BIC) if $\mathbb{E}_{\mu_{-i}}[g(t_i(\geq), t_{-i})] >^* \mathbb{E}_{\mu_{-i}}[g(t_i, t_{-i})]$ for every i , every \geq_i , and every $t_i \neq t_i(\geq_i)$. Finally, say that T is BIC-elicitable under μ if there exists a BIC T -mechanism g . It is not hard to show that T is BIC-elicitable under μ if and only if each of the T_i 's is elicitable. We note that the assumption that μ is a product measure is important for this result.

APPENDIX A. PROOFS

Proof of Proposition 2

If $T = \tilde{T}(\tilde{I}(T))$ then clearly T is generated by top elements (of the menus $\tilde{I}(T)$).

To prove the converse note first that, for every type space T , if \geq, \geq' are in the same element of T then by definition $dom_{\geq}(X') = dom_{\geq'}(X')$ for every $X' \in \tilde{I}(T)$. This implies that \geq, \geq' are also in the same element $\tilde{T}(\tilde{I}(T))$. In other words, T is always (weakly) finer than $\tilde{T}(\tilde{I}(T))$.

Now, suppose that T is generated by top elements, say $T = \tilde{T}(X_1, \dots, X_l)$. Then clearly $\{X_1, \dots, X_l\} \subseteq \tilde{I}(T)$. But adding more menus can only make the resulting type space finer, so $\tilde{T}(\tilde{I}(T))$ is (weakly) finer than $\tilde{T}(X_1, \dots, X_l) = T$. This completes the proof.

Proof of Proposition 3

We start with the following lemma, which originally appeared in our earlier work Azrieli et al. (2019, Lemma 3).

Lemma 2. Suppose $|X'| \geq 2$. Then X' is surely identified by $\{X_1, \dots, X_l\}$ 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, \dots, 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 $\{x, y\} \subseteq X_i \subseteq X'$. 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. \square

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 Edelman and Jamison (1985). 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 2 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 Edelman and Jamison (1985, 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 $\tilde{T}(X_1, \dots, X_l) = \tilde{T}(X'_1, \dots, X'_k)$ iff $\text{SI}(\{X_1, \dots, X_l\}) = \text{SI}(\{X'_1, \dots, X'_k\})$, it follows that if T is elicitable then the set of collection of menus that generate T is the interval $[B(\tilde{I}(T)), \tilde{I}(T)]$. 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 $B(\mathcal{X})$.

Proof of Proposition 4

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

$$g(t)(x) = \mu\left(\{\omega \in \Omega : f(t)(\omega) = x\}\right)$$

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

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

$$\{\omega \in \Omega : f(t(\succeq))(\omega) \succeq x\} \supsetneq \{\omega \in \Omega : f(t)(\omega) \succeq x\},$$

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

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

Proof of Proposition 6

Let g be an IC T -mechanism, and let $t \in T$. We will show that

$$\overline{U(t)} = \bigcap_{t' \in T} \{u : \langle g(t), u \rangle \geq \langle g(t'), u \rangle\},$$

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 $\succeq \in t$ such that $u \in \overline{U(\succeq)}$. Incentive compatibility of g then implies that $\langle g(t), u \rangle \geq \langle g(t'), u \rangle$ for every $t' \in T$, so u is in the right-hand side as well.¹⁵

Conversely, suppose that u is in the right-hand side. Then in every open neighborhood of u there is u' for which $\langle g(t), u' \rangle > \langle g(t'), u' \rangle$ holds for all $t' \neq t$ (here we use the fact that the right-hand side is a polyhedral set with non-empty interior). Incentive compatibility of g implies that $u' \notin \overline{U(t^c)}$, so we must have $u' \in \overline{U(t)}$. Since this set is closed we get that $u \in \overline{U(t)}$ as well.

¹⁵Recall that lottery p strictly dominates lottery q relative to \succeq if and only if $\langle p, u \rangle > \langle q, u \rangle$ for every $u \in U(\succeq)$.

Proof of Proposition 7

The key to the proof is the following lemma.

Lemma 3. If g is an IC T -mechanism, and if $t, t' \in T$ are adjacent via an x - y switch, then $g(t)(x) - g(t)(y) = g(t')(y) - g(t')(x) > 0$, and $g(t)(z) = g(t')(z)$ for all $z \notin \{x, y\}$.

Proof. Let $\succeq \in t$ and $\succeq' \in t'$ 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 $g(t)(z) = g(t')(z)$ 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 $\langle u, g(t) \rangle > \langle u, g(t') \rangle$ and any $u' \in U(\succeq')$ has $\langle u', g(t) \rangle < \langle u', g(t') \rangle$, from which we conclude that $\langle \bar{u}, g(t) \rangle = \langle \bar{u}, g(t') \rangle$ must be satisfied. But this is the same as $g(t)(z) = g(t')(z)$.

Now, consider z with $r_{\succeq}(z) < r_{\succeq}(x)$ and suppose that $g(t)(w) = g(t')(w)$ 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 $r_{\succeq}(w) \leq r_{\succeq}(z)$ 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 $\langle \bar{u}, g(t) \rangle = \langle \bar{u}, g(t') \rangle$, so by the induction hypothesis it follows that $g(t)(z) = g(t')(z)$.

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 $g(t)(z) = g(t')(z)$ for all $z \neq x, y$, and since both $g(t), g(t')$ are lotteries, we must have $g(t)(x) - g(t)(y) = g(t')(y) - g(t')(x)$. 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'). \square

The proposition now easily follows. Indeed, let g be an IC mechanism and suppose $\{t_1, \dots, t_k\} \subseteq T$ satisfy the assumption of the proposition. Then by Lemma 3 we have that $g(t_1)(x) > g(t_2)(x)$. Suppose by means of contradiction that there is no $1 < i \leq k$ and z such that t_i and t_{i+1} are adjacent via a z - x switch. Then it follows again from Lemma 3 that $g(t_i)(x) \geq g(t_{i+1})(x)$, whereby $g(t_2)(x) \geq g(t_1)(x)$, a contradiction.

Remark. Lemma 3 says that a necessary condition for a T -mechanism g to be IC is that if t and t' are adjacent via an x - y switch, then the lotteries $g(t)$ and $g(t')$ 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 t has no incentive to announce t' , and vice versa. Carroll (2012, 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.

Proof of Proposition 8

(1) \implies (2): This follows from Proposition 6.

(2) \implies (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 $i, k \in B \in \mathcal{Q}$ 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 $T_{\mathcal{Q}}$, say t .

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 $\succeq'' \notin t$. This proves that $\overline{U(t)}$ is not convex, so $T_{\mathcal{Q}}$ is not a convex type space.

(3) \implies (1): Suppose $\mathcal{Q} = \{B_1, \dots, B_K\}$ where each B_k is an interval in $\{1, \dots, m\}$. 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 $i < j$. Then

$$T_{\mathcal{Q}} = \tilde{T}(X, \dots, X; |B_1|, |B_1| + |B_2|, \dots, \sum_{1 \leq i \leq K-1} |B_i|),$$

that is $T_{\mathcal{Q}}$ is generated by top sets. By Proposition 5 we are done.

REFERENCES

- Abdulkadiroğlu, A., Pathak, P. A., Roth, A. E., 2005. The new york city high school match. *American Economic Review* 95, 364–367.
- Azrieli, Y., Chambers, C. P., Healy, P. J., 2018. Incentives in experiments: A theoretical analysis. *Journal of Political Economy* 126 (4), 1472–1503.
- Azrieli, Y., Chambers, C. P., Healy, P. J., 2019. Incentives in experiments with objective lotteries. *Experimental Economics* forthcoming.
- Brier, G. W., 1950. Verification of forecasts expressed in terms of probability. *Monthly Weather Review* 78, 1–3.
- Carroll, G., 2012. When are local incentive constraints sufficient? *Econometrica* 80, 661–686.

- Edelmean, P. H., Jamison, R. E., 1985. The theory of convex geometries. *Geometriae Dedicata* 19 (3), 247–270.
- Gabor, A., Granger, C. W. J., 1977. *Pricing: Principles and Practices*. Heinemann Educational, London.
- Gibbard, A., 1977. Manipulation of schemes that mix voting with chance. *Econometrica*, 665–681.
- Gilboa, I., Schmeidler, D., 1989. Maxmin expected utility with a non-unique prior 18 (2), 141–153.
- Haeringer, G., Klijn, F., 2009. Constrained school choice. *Journal of Economic Theory* 144 (5), 1921–1947.
- Kroll, Y., Levy, H., Rapoport, A., 1988. Experimental tests of the separation theorem and the capital asset pricing model. *American Economic Review* 78 (3), 500–519.
- Lambert, N., 2018. Elicitation and evaluation of statistical forecasts, Stanford University working paper.
- Lambert, N. S., Pennock, D. M., Shoham, Y., July 2008. Eliciting properties of probability distributions. In: *EC '08 Proceedings of the 9th ACM conference on electronic commerce*. Chicago, IL, USA, pp. 129–138.
- Loomes, G., 1991. Evidence of a new violation of the independence axiom. *Journal of Risk and Uncertainty* 4, 91–108.
- Machina, M. J., Schmeidler, D., 1995. Bayes without bernoulli: Simple conditions for probabilistically sophisticated choice. *Journal of Economic Theory* 67 (1), 106–128.
- Osband, K. H., 1985. Providing incentives for better forecasting. Ph.D. thesis, University of California, Berkeley.
- Savage, L. J., 1971. Elicitation of personal probabilities and expectations. *Journal of the American Statistical Association* 66, 783–801.
- Schervish, M. J., 1989. A general method for comparing probability assessors. *The Annals of Statistics* 17 (4), 1856–1879.
- Shapley, L., 1971. Cores of convex games. *International Journal of Game Theory* 1, 11–26.
- Warner, S. L., 1965. Randomized response: A survey technique for eliminating evasive answer bias. *Journal of the American Statistical Association* 60 (309), 63–69.