Theoretical Economics 16 (2021), 507–538 1555-7561/20210507

Constrained preference elicitation

Yaron Azrieli
Department of Economics, The Ohio State University

Christopher P. Chambers
Department of Economics, Georgetown University

Paul J. Healy
Department of Economics, The Ohio State University

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 wants only to elicit the agent's type. We first assume that beliefs about randomization are subjective, and show that a space of types is elicitable if and only if each type is defined by what the agent would choose from some list of menus. If beliefs are objective, then additional type spaces can be elicited, though a convexity condition must be satisfied. These results remain unchanged when we consider a setting with multiple agents.

Keywords. Elicitation, incentive compatibility, random mechanisms.

JEL classification. C7, D8.

1. 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 (Abdulkadiroğlu et al. 2005, Haeringer and Klijn 2009). In lab experiments, researchers often measure parameters of subjects' preferences by eliciting a single choice from a set of alternatives, rather than their entire preference ordering over the set (Kroll et al. 1988, Loomes 1991). A firm surveying customers before the release of a new product needs only 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 or it may simply be too complex to communicate all of the information needed (Hurwicz 1960, 1977, Mount and Reiter 1974, Nisan and Segal 2006, Kos

Yaron Azrieli: azrieli.2@osu.edu
Christopher P. Chambers: cc1950@georgetown.edu
Paul J. Healy: healy.52@osu.edu
The authors thank the editor, anonymous referees, and seminar audiences at The Technion, Bar Ilan University, University of Montreal, Johns Hopkins, and the Economic Science Association 2019 North American Meeting for useful comments.

 https://doi.org/10.3982/TE4208

2012). Yet most of the mechanism design literature focuses on direct revelation mechanisms in which entire orderings are revealed, ignoring these practical considerations.

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 compatible way. For example, suppose there are three alternatives, x, y, and z, and the planner wants to know only which is the agent's most preferred. This can be done simply by asking him to pick an item and paying 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, and that learning how she ranks y versus z would somehow be costly to the planner or impinge on the agent's privacy.[1] Thus, the planner wants to know only whether the most preferred item is "x" or "not x." We show that there is no way to guarantee that the agent has a strict incentive to reveal his answer truthfully.[2] 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, meaning agents report their type and are paid either an act or a lottery over alternatives. Given this framework, our goal is to characterize those type spaces that are *elicitable*, meaning there exists an incentive compatible random mechanism under which the agent always has a strictly dominant strategy to report her 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 type space would typically be given by the problem. For example, in a classic principal–agent setting, the worker's disutility of effort (type) is assumed to be either high or low. The principal wants to incentivize effort to maximize profit, and does so by picking a type-dependent payment scheme (mechanism), but cannot select the first-best mechanism because it is not incentive compatible. Here our goal is different: for any type space, we ask whether *any* incentive compatible mechanism exists. Our planner simply wants to learn about the agent's preferences, so any incentive compatible mechanism will suffice.

[1]We develop the costly-information motivation in Section 3. Privacy could also be maintained by the researcher committing not to look at certain data, though in some settings this may be impossible or may not be trusted by the subjects.

[2]By "guarantee," we mean regardless of her true preference ordering *and* her risk preferences. For example, if reporting "x" pays x and reporting "not x" pays a 50–50 lottery over y and z, then a very risk-averse agent with $y \succ x \succ z$ might lie and report "x" because it gives a riskless payment. To prevent this possibility, we require that truth-telling gives a random payment that stochastically dominates the payment from any lie.

Alternatively, if the planner is trying to maximize an objective, our question is whether even *one* mechanism that satisfies the incentive compatibility constraint exists. We also differ from the standard literature in that we require strict incentive compatibility, as we want to ensure the agent reports truthfully; most of the mechanism design literature requires only the weak version of this constraint.

To illustrate, in Section 3, we describe a planner who wants to run a lab experiment to test whether some axiom (such as independence of irrelevant alternatives (IIA)) holds. We show that the planner is always able to elicit the entire preference relation and from that can check whether the axiom is satisfied. But this entails collecting far more information than is needed. If information acquisition is costly, then the planner may want to elicit a coarser type space. In particular, the planner really wants to elicit only two types: the set of preferences that satisfy the axiom and the set of preferences that do not. For most axioms of interest, however, we can show that this two-element type space is not elicitable.[3] In these cases the planner would need to find a different type space that is "close" to the two-element type space, is not too costly, and is elicitable. The contribution of our paper is to analyze exactly which type spaces satisfy that elicitability requirement.

We begin with a simple sufficiency condition for elicitability. 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 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. We call such a mechanism a random problem selection (RPS) mechanism, and it is known to be strictly incentive compatible (see Azrieli et al., 2018, for example). This mechanism elicits a type space in which 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 $x \succ y \succ z$, and so on. In general, any list of menus generates a specific type space in this way. We say that 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 of preference over acts—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.

When there are objective probabilities, however, more can be elicited. To understand how, consider again an example with three alternatives $\{x, y, z\}$, but now the planner wants to elicit the agent's *least* preferred item. This type space is not generated by top elements and so cannot be elicited using acts. But it can be elicited with objective

[3] Depending on the costs, it may actually be that some other two-element type space would be optimal. Regardless, we still find that the desired type space is not elicitable in most cases of interest; see Section 3 for details.

lotteries using the following mechanism: The agents announce their least preferred element as their type. If they announce x, then they get paid a lottery in which everything *except* x is paid with equal probability, and x is paid with probability 0. In other words, they receive a 50–50 lottery over y and z. To see that this is incentive compatible, consider subjects who truly rank x at the bottom, but instead announce y. By doing so, they reduce their probability of being paid y to zero and increase their probability of being paid x to $1/2$. This is first-order stochastically dominated (as is announcing z), so truth-telling is strictly optimal.

We can extend the example above to show that any type space defined by the top *set* of alternatives from menus (rather than the top single alternative) can also be elicited. Indeed, eliciting the bottom-ranked alternative is equivalent to eliciting the top two alternatives. As above, this is accomplished by paying everything in the top set with equal probability. We then demonstrate with an example that even more type spaces 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. Roughly, convexity here means that types contain only "adjacent" preference orderings, where two orderings are adjacent if they differ only in a single switch of two alternatives. We do achieve a full characterization on the subset of *neutral* type spaces, which treat all alternatives in X symmetrically. Formally, these are type spaces that are invariant to permutations of X. These type spaces answer questions like, "What item do you rank third?," but do *not* answer questions like, "Where in the rankings do you place x?" We show that a neutral type space is elicitable with lotteries if and only if it is convex. We also derive a characterization of elicitable neutral type spaces in the Savage framework, which is more restrictive.

The lotteries framework is arguably more natural than the Savage framework because the experimenter is free to choose the randomization device and, therefore, should be able to induce objective probabilities. We still find value in the Savage framework, however, because our results for lotteries make extensive use of *exactly* equal probabilities. Recall that to incentivize subjects to reveal x as their least-preferred element from $X = \{x, y, z\}$, we must pay them a 50–50 lottery over y and z. If the subjects do not perceive those probabilities to be exactly equal, they may have an incentive to deviate from truth-telling.[4] The Savage framework is robust to such "trembles," as it assumes no probabilities whatsoever. By taking away our ability to use exactly equal probabilities, the Savage framework takes away our ability to elicit any type space that is not generated by top elements. Thus, we view our results in this setting as more robust, which of course comes at the cost of being more restrictive. See Azrieli et al. (2018, 2020) for further discussions of this distinction.

We also identify structure in the set of elicitable type spaces. In both frameworks, the set forms a lattice, meaning that if two type spaces (again, modeled as partitions) are elicitable, then so is their join.[5] The finest type space—in which every ordering is its

[4]Specifically, suppose they believe $\Pr(z) > 0.5$ when x is announced and $\Pr(z) < 0.5$ when y is announced. If $z \succ x \succ y$ and the cardinal utility of z is high compared to x, then they may prefer to announce x to get the slightly higher $\Pr(z)$.

[5]The join is the coarsest type space that is finer than either of the two.

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.[6] Thus, the information that can be elicited ranges from everything to nothing, and is related by the join operator. Interestingly, this shows that what can be elicited is not monotonic in the refinement relation: One might think that coarser information is easier to elicit, but, in fact, the finest type space is elicitable while many coarsenings of it are not. Furthermore, we show that in the acts framework, there is a set of "basic" type spaces in the lattice from which all other nontrivial 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 his/her 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 his/her type truthfully, regardless of others' announcements. Thus, a joint type space can be elicited using a DIC mechanism if and only if each individual type space is elicitable.

Related literature

The papers closest to ours are Lambert et al. (2008) and Lambert (2019). 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 problem is "dual" to theirs (and also to Osband, 1985). Every preference relation can be identified with an equivalence class of probability distributions, meaning a utility index representing the preference can be defined so that all of its values are nonnegative and so that they sum to 1. It is well known that 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. The main distinction is that the "statistics" to be elicited in our situation have the added constraint that they must be identified with a collection of preferences.[7] In particular, Lambert (2019) builds a characterization based on the notion of a power diagram.[8]

Bahel and Sprumont (2020) characterize strategy-proof mechanisms in a Savage setting with many agents. In their framework, agents report their entire (subjective expected utility) preference to the mechanism. Thus, they study only the finest type space.

[6]The 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 $\{x\}$.

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

[8]Though not immediate, it is not too difficult to show, using his techniques, that a similar characterization can hold in our framework. We omit the details.

With a single agent, their result is that a strategy-proof mechanism must pay an act that gives the agent his/her favorite available outcome in every state of nature. This corresponds exactly to using the RPS mechanism, where a list of menus is given, the state of nature is which list is chosen for payment, and the agent indeed receives his/her favorite available outcome in each. Indeed, our results show that other type spaces are elicitable if and only if exactly this mechanism can be used.

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 on eliciting choices from a list of menus (Azrieli et al. 2018). 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 an incentive compatible mechanism exist. 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.

2. Notation and definitions

Let X be a finite set of alternatives, with $|X| = m \geq 2$. Alternatives in X are 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 $\succeq, \succeq', \succeq''$, etc. We write $x \succ y$ when $x \succeq y$ and $x \neq y$.[9] In examples, we often write xyz to denote the relation $\succeq$ for which $x \succ y \succ z$.

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.

Definition 1. A *type* $t \subseteq O$ is a set of preference orderings in O. A *type space* $T = \{t_1, \ldots, t_k\}$ is a collection of types that forms a partition of O.[10]

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 that 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

[9]To be clear, "indifference" occurs only between x and itself: If $x \neq y$, then $x \succeq y$ indicates a strict ordering of x and y; the only case where both $x \succeq y$ and $y \succeq x$ is when $x = y$.

[10]The relationship $T = \{t_1, \ldots, t_k\}$ forms a partition of O if $t_1, \ldots, t_k$ are all nonempty, pairwise disjoint, and $\bigcup_{i=1}^{k} t_i = O$.

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. We often refer to this type space as the *trivial* type space.

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 comprises all nonempty 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. In this setting, agents are paid a lottery in $\Delta(X)$, depending on their announced type. 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 $p \succ^* q$).[11]

Recall that $p \succeq^* q$ implies that p is preferred to q for any expected utility maximizer consistent with $\succeq$, regardless of his/her risk preferences. This interpretation is used extensively in Section 5 when we analyze incentive compatibility.

A T *mechanism with lotteries* is a mapping $g : T \to \Delta(X)$. 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 $t \neq t(\succeq)$,

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

[11]Given that $\succeq$ is a linear order, unless $p = q$, the first-order stochastic relation between p and q is strict.

By using stochastic dominance, we require no assumptions on risk preferences. Put another way, types do not reveal agents' risk preferences, and so 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; this ensures that truth-telling is uniquely optimal. If we had required only weak dominance, then any constant mechanism (whose payments do not depend on t) would be incentive compatible regardless of T.

Our interest is to understand 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 T mechanism with lotteries.

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 $f : T \to X^{\Omega}$ is the payment function. States here do not have objective probabilities; indeed, the decision-maker may have any beliefs over Ω or may not even behave as though they have probability distribution on Ω (Machina and Schmeidler 1995).

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. 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 $\omega \in \Omega$,

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

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 T mechanism with acts.

Remark. IC is the (single-agent) notion of strategy proofness (or dominant strategy equilibrium) imposed state by state. Write $f_\omega(t(\succeq)) = f(t(\succeq))(\omega)$. With an understanding that t here denotes the map carrying a preference to its type, IC is the requirement that $f_\omega \circ t$ is strategy-proof for each ω (with strictness in at least one state).

3. The importance of elicitability: An example

We briefly return to the motivation of our exercise by imagining an experimenter interested in testing whether the subject's preferences satisfy a certain axiom, such as the independence axiom. Let $A \subseteq O$ be the set of preferences that satisfy the axiom. If the experimenter elicits a type space T and finds that $t(\succeq) \subseteq A$, then $\succeq$ is consistent with the axiom. For example, the experimenter could elicit $T_A = \{A, A^c\}$, which directly asks whether A is satisfied, or the experimenter could elicit the finest type space and learn the entire preference. Indeed, the finest type space $\overline{T} = \{\{\succeq\}_{\succeq \in O}\}$ is always elicitable, using either acts or lotteries, and can, therefore, be used to test any axiom.[12]

Eliciting the entire preference may be extremely costly, especially when X is large. It may require many more questions (e.g., more menus of choices) or more detailed reports by subjects; it may be more time-consuming, costing the experimenter's time and requiring higher subject payments. In general, we can imagine an experimenter choosing the type space T that maximizes $V(T) - C(T)$, where $V(T)$ is the value of information gained from the experiment and $C(T)$ is the cost of elicitation. For concreteness, suppose the experimenter starts with a uniform prior over the preference $\succeq$ of the subject, and let $V(T)$ be the probability of the experimenter correctly categorizing the subject as belonging to either A or A^c.[13] Let $C(T)$ be the usual entropy-reduction cost function (Sims 2003, Matějka and McKay 2015, Caplin et al. 2019).

Standard arguments imply that the first-best type space T^*, which maximizes $V(T) - C(T)$, has at most two types. But our Proposition 4 shows that the only two-type type spaces that are elicitable are those defined by a single binary choice. For example, $T_{xy} = \{\{\succeq : x \succeq y\}, \{\succeq : y \succeq x\}\}$ is defined by the choice from the single menu $X_1 = \{x, y\}$ and is, therefore, elicitable. Most axioms, however, are not defined by a single binary choice, and when that is the case, we can prove that T^* is not elicitable.[14]

If the first-best type space T^* is not elicitable then the experimenter must find the second-best type space, which is the T^{**} that maximizes $V(T) - C(T)$ subject to the constraint that T be elicitable. Solving this problem requires knowledge of that constraint set. The aim of our paper is to provide exactly that information.

[12]To elicit $\overline{T}$, first ask the subject to report $\succeq$, randomly draw a pair of alternatives $\{x, y\} \in X$, and then pay the subject the item in $\{x, y\}$ that their reported $\succeq$ ranks higher. This is equivalent to a random problem selection mechanism, which is incentive compatible with acts (Azrieli et al. 2018) and with lotteries (Azrieli et al. 2020).

[13]That is, the experimenter updates his belief given the elicited type t and categorize the preference as belonging to A if and only if the posterior of this event is at least one-half. The ex ante probability of a correct categorization given type space T is $V(T)$.

[14]This is because binary-choice type spaces T_{xy} maximize entropy reduction, so T_A gives both a higher $V(\cdot)$ and a lower $C(\cdot)$ than any T_{xy}. Thus, $T^* \notin \{T_{xy}\}_{xy \in X^2}$. But those are the only elicitable two-element type spaces, so T^* is not elicitable.

4. Elicitable type spaces with act payments

In this section we characterize type spaces that are elicitable under acts. We sometimes omit the qualifier "with acts" when referring to a T mechanism or to elicitability of a type space.

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

Characterization

Recall from the Introduction that the random problem selection (RPS) mechanism gives the subjects a list of menus, has them announce their favorite item from each, and then random selects one of the announced items for payment. In short, our characterization shows that T is elicitable if and only if it can be elicited via an RPS mechanism. Thus, we begin by formalizing what it means for a type space to have this property.

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

In words, if T is generated by top elements, then each type t is defined by what the subject would pick from the menus $X_1, \ldots, X_l$. This type space is easily elicited using an RPS mechanism, which is known to be incentive compatible. The main result of this section is that these are the *only* type spaces that can be elicited with acts.

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

To illustrate, note that the type space $T = \{\{xyz, xzy\}, \{zxy\}, \{yxz\}, \{yzx, zyx\}\}$ can be elicited by asking the subject to choose from the menus $\{x, y\}$ and $\{x, z\}$. Their choice from the first menu $\{x, y\}$ divides the possible preferences into $T_{xy} = \{\{xyz, xzy, zxy\}, \{yxz, yzx, zyx\}\}$ (depending on whether they pick x or y), and their choice from the second menu $\{x, z\}$ divides the preferences into $T_{xz} = \{\{xyz, xzy, yxz\}, \{zyx, yzx, zxy\}\}$. Combining these two pieces of information gives $T_{xy} \vee T_{xz} = T$.

For a non-elicitable type space, recall that $T' = \{\{xyz, xzy\}, \{yxz, yzx, zxy, zyx\}\}$, which is defined by the question, "Is your favorite x or not x?" This cannot be determined from choices over any list of menus. The reason is that any menu would give a partition that splits the "not x" type ($\{yxz, yzx, zxy, zyx\}$) into two.[15] Thus, no list of menus can elicit "not x" without also eliciting additional information.

[15]The menu $\{x, y\}$ would give the split $\{yxz, yzx, zyx\}$ and $\{zxy\}$, the menu $\{x, z\}$ would give $\{yxz\}$ and $\{yzx, zxy, zyx\}$, and the menus $\{y, z\}$ and $\{x, y, z\}$ would both give $\{yxz, yzx\}$ and $\{zxy, zyx\}$. Any combination of menus would further refine these splits.

PROOF OF PROPOSITION 1. (*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 $\succeq^t \in t$. Define $f(t)(\omega_i) = \mathrm{dom}_{\succeq^t}(X_i)$ for $i = 1, \dots, l$. Note that by assumption the choice of the representative $\succeq^t$ does not affect the resulting mechanism.

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

$$f\big(t(\succeq)\big)(\omega_i) = \mathrm{dom}_{\succeq^{t(\succeq)}}(X_i) = \mathrm{dom}_{\succeq}(X_i) \succeq \mathrm{dom}_{\succeq^t}(X_i) = f(t)(\omega_i).$$

The first equality is by the definition of f; the second follows from the fact that T is generated by top elements, so that $\succeq$ and $\succeq^{t(\succeq)}$ 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 $t \neq t(\succeq)$, then there exists i such that $\mathrm{dom}_{\succeq}(X_i) \neq \mathrm{dom}_{\succeq^t}(X_i)$, 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 $\succeq$ and $\succeq'$ are in the same element of T, and fix some $1 \leq i \leq l$. Then by incentive compatibility, we have that $f(t(\succeq))(\omega_i) \succeq f(t)(\omega_i)$ for every $t \in T$, which implies that $f(t(\succeq))(\omega_i) = \mathrm{dom}_{\succeq}(X_i)$. Applying the same argument to $\succeq'$ gives $f(t(\succeq'))(\omega_i) = \mathrm{dom}_{\succeq'}(X_i)$. But since $t(\succeq) = t(\succeq')$, we get $\mathrm{dom}_{\succeq}(X_i) = \mathrm{dom}_{\succeq'}(X_i)$. Repeating for each $i = 1, \dots, l$ shows that $\succeq$ and $\succeq'$ are in the same element of $\tilde{T}(X_1, \dots, X_l)$.

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

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 $T' = \tilde{T}(X'_1, \dots, X'_k)$. Recall that we order type spaces by refinement. So, for two type spaces T and T', $T \vee T'$, their join, is the coarsest common refinement. 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 with acts, then so is their join $T \vee T'$.*

It follows that the set of elicitable type spaces equipped with the refinement relation forms a lattice.[16] In this lattice, the meet of T' and T'' is the elicitable type space T such

[16] Corollary 1 shows that the set of elicitable type spaces equipped with the refinement relation is a $\vee$-sub semilattice of the lattice of all type spaces with the refinement relation. Since the set of elicitable type

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, is 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, nontrivial 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 $\tilde{T}(X_1, \dots, X_l) = \tilde{T}(X_1) \vee \tilde{T}(X_2) \vee \cdots \vee \tilde{T}(X_l)$. 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 two refinements of T,

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

are. Clearly both of these refinements of T are as coarse as possible, 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 $\succeq, \succeq' \in t$,

spaces is finite and since the coarsest possible type space is elicitable (and, therefore, there is a coarsest elicitable type space), we can conclude that each pair of elicitable type spaces has a greatest lower bound (see, e.g., Davey and Priestley (2002), Theorem 2.31, or Birkhoff (1996), Theorem 3, p. 112) and, thus, by definition forms a lattice.

it holds that $\mathrm{dom}_{\succeq}(X') = \mathrm{dom}_{\succeq'}(X')$; in other words this means that $\mathrm{dom}_{\bullet}(X')$ 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 does not perfectly reveal $\succeq$). Let $\tilde{I}(T)$ be the collection of menus that are identified from T.[17]

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

Lemma 1. *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.[18]

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, $\tilde{I}(\overline{T})$ contains every subset of X, but $\overline{T}$ 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 (i) this process of dropping redundant menus always leads to a unique smallest (in terms of inclusion) list of menus that still generates T, and (ii) 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 $X', X_1, \ldots, X_l$ be menus. Say that X' is *surely identified* by $X_1, \ldots, X_l$ if whenever $\succeq$ and $\succeq'$ are such that $\mathrm{dom}_{\succeq}(X_i) = \mathrm{dom}_{\succeq'}(X_i)$ for all $i = 1, \ldots, l$, it also holds that $\mathrm{dom}_{\succeq}(X') = \mathrm{dom}_{\succeq'}(X')$. The menus $X_1, \ldots, X_l$ are *independent* if no X_i is surely identified by $\{X_j\}_{j \neq i}$.

We can now formalize the interval result.

Proposition 2. *For any type space T that is elicitable with acts, there is a unique independent list of menus, denoted $\tilde{B}(T)$, such that $T = \tilde{T}(\tilde{B}(T))$. Furthermore, $T = \tilde{T}(X_1, \ldots, X_l)$ if and only if*

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

The list $\tilde{B}(T)$ is of particular significance as it provides a measure of the amount of information contained in type space T. Specifically, if $\tilde{B}(T) = \{X_1, \ldots, X_l\}$, 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.

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

[18] One can show a bit more than Lemma 1; in particular, for any type space T and set of nonsingleton, nonempty menus M, that

$$\tilde{T}(M) \leq T \quad \Leftrightarrow \quad M \subseteq \tilde{I}(T),$$

where $\leq$ refers to the refinement relation. Thus, $(\tilde{T}, \tilde{I})$ forms a Galois connection in the sense of Davey and Priestley (2002), p. 155). Lemma 1 follows from Lemma 7.26 of that reference, together with Proposition 1. Other simple results also follow. We omit the details.

5. Elicitable type spaces with lottery payments

This section analyzes elicitability in the lotteries framework. Throughout this section, if not otherwise mentioned, a T mechanism refers to lotteries T mechanisms, 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.

Corollary 2. *If T is generated by top elements, then T is elicitable with lotteries.*

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

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

In words, T reveals the least preferred alternative from X, or, equivalently, it reveals the top two alternatives but not their order. Type space T is not generated by top elements. This can easily be checked using Lemma 1: $\tilde{I}(T)$ contains no nontrivial 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

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

that is, g randomly chooses one of the top two ranked alternatives with equal probability. It is immediate to check that g is IC. $\diamondsuit$

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 $\text{dom}^k_{\succeq}(X')$ to be the (unique) set of elements satisfying (i) $\text{dom}^k_{\succeq}(X') \subseteq X'$, (ii) $|\text{dom}^k_{\succeq}(X')| = k$, and (iii) if $x \in \text{dom}^k_{\succeq}(X')$ and $y \in X' \setminus \text{dom}^k_{\succeq}(X')$, then $x \succeq y$. In words, $\text{dom}^k_{\succeq}(X')$ is the set of top k ranked elements in X'. Note that $\text{dom}^1_{\succeq}(X') = \text{dom}_{\succeq}(X')$.

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 $\text{dom}^{k_i}_{\succeq}(X_i) = \text{dom}^{k_i}_{\succeq'}(X_i)$ 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 Corollary 2.

PROPOSITION 3. *If T is generated by top sets, then it is elicitable with lotteries.*

PROOF. We sketch the argument here, 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 $\mathrm{dom}^{k_1}_{\succeq^t}(X_1)$, where $\succeq^t$ is any member of t. Then g is clearly IC, so T is elicitable.

Second, it is straightforward to verify that

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

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 $t \cap t' \neq \varnothing$) 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. □

At this point, it is tempting to guess that the converse of Proposition 3 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,

$$\begin{aligned} T = \big\{ &t_0 = \{xyzw, yxzw, xywz, yxwz\}, \{xzyw\}, \{xwyz\}, \{ywxz\}, \{yzxw\}, \\ &\{xzwy, xwzy\}, \{yzwx, ywzx\}, \{zxyw, zyxw\}, \{zywx, zwyx\}, \\ &\{zxwy, zwxy\}, \{wxyz, wyxz\}, \{wyzx, wzyx\}, \{wxzy, wzxy\}\big\}. \end{aligned}$$

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 is not difficult, but perhaps mildly tedious to check that g is IC.

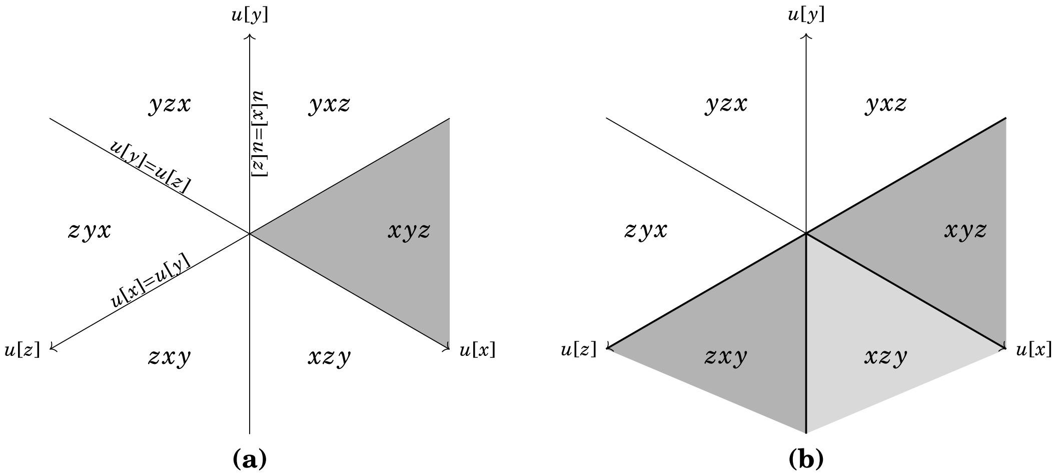

FIGURE 1. The sets of cardinal utility vectors associated with each ordering $\succeq$ (denoted by $U(\succeq)$) for $X = \{x, y, z\}$. (a) The set $U(xyz)$ is shaded. (b) For the nonconvex type $t = \{xyz, zxy\}$, $U(t) = U(xyz) \cup U(zxy)$ is shaded in dark and $\overline{U(\text{Cons}(t))}$ is the union of all three shaded areas.

Second, T is not generated by top sets. The easiest way to show that is by checking that there is no pair consisting of a menu $X' \subseteq X$ and a number $1 \leq k < |X'|$ for which the top k elements in X' are always identified by T.

This example is based on a characterization of convex capacities found in Shapley (1971, p. 14). As we discussed elsewhere, convex capacities are closely related to IC mechanisms (Azrieli et al. 2020). Shapley lists the "extreme rays" of the set of convex capacities on four elements, and this example derives from the first in this list. ◊

Two necessary conditions

We show that a certain convexity condition on a type space is necessary for it to be elicitable. To understand this condition, consider Figure 1, in which we graph the space of cardinal utility vectors over $X = \{x, y, z\}$. The axes are labeled as $u[x]$, $u[y]$, and $u[z]$, and we can view this figure as looking at the origin of $\mathbb{R}^3$ from above, from the point of view of the ray of equal coordinates. Each point $u \in \mathbb{R}^3$ corresponds to an ordering in O. For example, vector $u = (3, 2, 1)$ (meaning $u[x] = 3$, $u[y] = 2$, and $u[y] = 1$) corresponds to the ordering xyz. The interior of the shaded cone in panel (a) represents all utility vectors that correspond to xyz. Define $U(\succeq) = \{u \in \mathbb{R}^X : x \succ y \Leftrightarrow u[x] > u[y]\}$, so 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 on the planes defined by indifference: $u[x] = u[y]$, $u[x] = u[z]$, and $u[y] = u[z]$.[19]

[19]The three planes all intersect along the line of total indifference, defined by $u = (\alpha, \alpha, \alpha)$ for $\alpha \in \mathbb{R}$. The three-dimensional figure is drawn from a vantage point along this line; from this perspective, each plane projects as a line 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 $u + (\alpha, \alpha, \alpha) \in U(\succeq)$. Equivalently, the two-dimensional picture can be thought of as a projection of each $U(\succeq)$ onto the plane defined by $\{u : \sum_x u[x] = 0\}$.

For any type $t \subseteq O$, 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 1 corresponds to the type $t = \{xyz, zxy\}$. Notice that $\overline{U(t)}$ is not a convex set. As a consequence, any type space containing t cannot be elicited, because if lottery $g(t)$ is most preferred for some $u \in U(zxy)$ and some $u' \in U(xyz)$, then $g(t)$ is also most preferred for any $\lambda u + (1-\lambda)u' \in U(xzy)$. Thus, a person with ordering xzy will wrongly report his/her type to be t.

We now present an order-theoretic definition of convexity that corresponds exactly to the geometric notion of convexity given above. Given a set of orderings $t \subseteq O$, let $\sqsupseteq_t = \bigcap_{\succeq \in t} \succeq$ be the maximal relation that all orderings in t agree on; that is, $x \sqsupseteq_t y$ if and only if $x \succeq y$ for all $\succeq \in t$. Note that $\sqsupseteq_t$ is a partial order. For example, for the type t in panel (b) of Figure 1, we have $x \sqsupseteq_t y$, but $\sqsupseteq_t$ does not rank $\{x, z\}$ or $\{y, z\}$. Notice that xzy also agrees with $\sqsupseteq_t$, but $xzy \notin t$; it is in this sense that we say that t is not convex. Formally, let

$$\text{Cons}(t) = \left\{\succeq \in O : (\forall x, y \in X) x \sqsupseteq_t y \Rightarrow x \succeq y\right\}$$

be the set of orderings that are consistent with t.[20] Using this, we can now define convexity.

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

We now show equivalence between convexity of $\overline{U(t)}$ in $\mathbb{R}^X$ and the convexity of t according to Definition 10. If t is convex, then $\overline{U(t)} = \overline{U(\text{Cons}(t))}$. But $\overline{U(\text{Cons}(t))}$ is the intersection of closed half-spaces (each defined by some $x \sqsupseteq_t y$), so $\overline{U(t)}$ must be convex. Conversely, if $\overline{U(t)}$ is convex, then necessarily $\overline{U(t)} = \overline{U(\text{Cons}(t))}$. This implies that $t = \text{Cons}(t)$, so t is convex. This proves the following lemma.

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

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 (2019, pp. 10–11), who attributes this observation to Osband (1985).[21]

Proposition 4. *If T is elicitable with either lotteries or acts, then it is convex.*[22]

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

[20] Equivalently, $\text{Cons}(t) = \{\succeq \in O : \sqsupseteq_t \subseteq \succeq\}$.

[21] Although our result is not technically a corollary of Lambert (2019), it is very similar in spirit since both follow immediately from linearity of the agent's objective function. Here incentive compatibility requires maximality with respect to the stochastic dominance relation, which is equivalent to maximizing expected utility for all risk preferences, and expected utility is linear in lottery payments.

[22] The fact that convexity is necessary with acts follows from the fact that any type space that can be elicited with acts can also be elicited with lotteries.

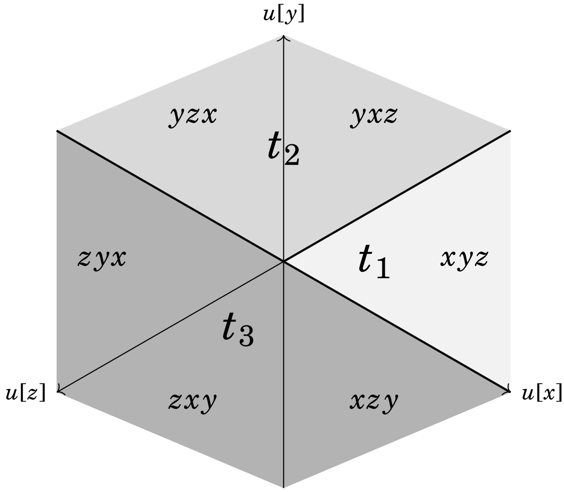

FIGURE 2. The type space from Example 5 is convex but not elicitable.

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

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

This type space is shown in Figure 2. Convexity of T is immediate, as each $\overline{U(t)}$ is convex. We now show, however, that T is not elicitable.

Suppose to the contrary 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 $g(t_1) \cdot u_2 < g(t_2) \cdot u_2$. Letting $M \to +\infty$ 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 elicitability 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 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.[23] In fact, we now demonstrate that this no-cycles condition is also a necessary condition for elicitability, and is a more stringent condition than convexity.

To formalize the no-cycles condition, let $r_{\succeq}(x) = |\{y : y \succeq x\}|$ be the ranking of x in the ordering $\succeq$. We say that $\succeq$ and $\succeq'$ are *adjacent via an x to 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 to y switch if there are $\succeq \in t$ and $\succeq' \in t'$ that are adjacent via an x to y switch.[24] The sets t and t' are said to be *adjacent* if they are adjacent via some x to y switch.

Proposition 5. *Suppose that T is elicitable with lotteries or with acts. 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 to 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 to x switch.*

Remark. An implication of Proposition 5 is that if T is elicitable and $t, t' \in T$ are adjacent via an x to y switch, then any other adjacency between these types must be via a y to x switch. If that were not the case then the cycle (t, t', t) would fail the necessary condition.

We now argue that the no-cycles condition implies convexity. Consider Figure 3, in which convexity fails. Let $t_x = \{zxy, yxz\}$ be the type that ranks x second, and define t_y and t_z analogously. Consider moving along the dashed line from $u \in U(t_y)$ to $u' \in U(t_y)$. It first passes through $U(t_x)$, which gives an x to y switch at the boundary. Then it passes through $U(t_z)$, giving a z to x switch, and finally back to $U(t_y)$, giving a y to z switch. Thus, the cycle (t_y, t_x, t_z, t_y) has an x to y switch, but no z to x switch, and so the type space fails the no-cycles condition. In fact, we can show that *any* failure of convexity generates a failure of the no-cycles condition in exactly this way, giving the following result.

Proposition 6. *If T satisfies the no-cycles condition of Proposition 5, then T also satisfies the convexity condition of Proposition 4. Thus, the no-cycles condition is the more stringent necessary condition.*

---

[23]At each step, the weight on x strictly decreases when x moves down the ranking or does not change when the ranking of x does not change.

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

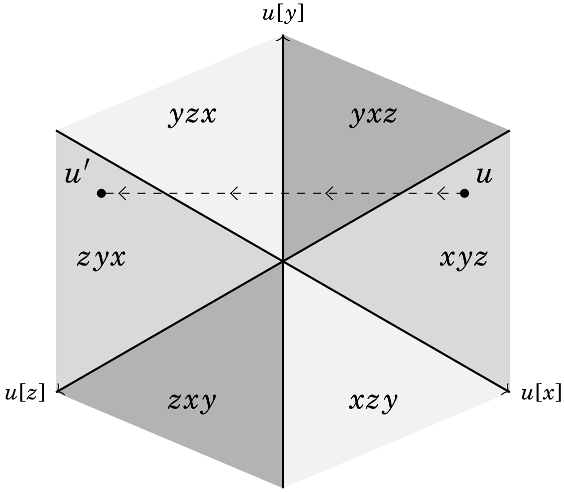

FIGURE 3. The type space defined by the second-ranked alternative is not convex and fails the no-cycles condition.

We do not know whether the no-cycles condition of Proposition 5 is enough to guarantee elicitability in the lotteries framework. While we could not find a counterexample, 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.

The lattice of elicitable type spaces under lotteries

Recall that in the acts framework, the lattice of elicitable type spaces is atomistic, with all nontrivial elicitable type spaces being the join of type spaces of the form $\tilde{T}(X_1)$. Recall also from the proof of Proposition 3 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 type space 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 $\overline{T}$ in which $t(\succeq) = \{\succeq\}$ for all $\succeq \in O$. From Lemma 5 in the proof of Proposition 5, the set of IC $\overline{T}$ mechanisms has the property that if $\succeq$ is adjacent to $\succeq'$ via an x to 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 to 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.[25]

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

Let $\overline{G}$ be the set of mechanisms that come from altering some IC $\overline{T}$ mechanism in this way. Formally, $\overline{G}$ is the set of $\overline{T}$ mechanisms that satisfy the following "weak" incentive compatibility property: If $\succeq$ and $\succeq'$ are adjacent via an x to 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 $\overline{G}$ is defined by finitely many linear inequalities, so it is a polyhedral subset of $\Delta(X)^O$. Let $ext(\overline{G})$ be the extreme points of this set.

For every $g \in \overline{G}$ define the type space T_g it elicits by letting $\succeq$ and $\succeq'$ 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.[26] Moreover, if $g, g' \in \overline{G}$ 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 $\{T_g : g \in ext(\overline{G})\}$. These are the basic elicitable type spaces from which all other elicitable type spaces can be formed.

6. Characterizations for neutral type spaces

In this section, we restrict attention to type spaces that treat all alternatives in X symmetrically, which we refer to as *neutral* type spaces. Neutral type spaces are invariant to the labeling of alternatives. Within this class, we find that convexity is not only necessary, but also sufficient for elicitability with lotteries. In the Savage framework, we achieve a characterization essentially identical to Proposition 1, but reinterpreted for the class of neutral type spaces.

Formally, neutral type spaces are invariant to permutations on X: Given a permutation $\pi : X \to X$ and a preference $\succeq$, define the permuted preference $\succeq_\pi$ such that $x \succeq_\pi y$ if and only if $\pi(x) \succeq \pi(y)$. For each type t, define $\pi(t) = \{\succeq_\pi : \succeq \in t\}$ to be the permuted type, consisting of all permuted preferences from t. Finally, $\pi T = \{\pi(t)\}_{t \in T}$ is the permuted type space. We say T is neutral if it is invariant to any such permutation.

Definition 11. A type space T is *neutral* if for any permutation π of X, we have that $\pi T = T$.

For our characterization theorems, it is useful to define a special subclass of neutral type spaces, called *positional* type spaces. These are type spaces that provide information only about which alternatives occupy each ranking in the subject's preferences. They answer questions like, "What item does the subject rank third?," but do *not* answer questions like, "Where in the rankings does she place x?"

One way to think of positional type spaces is to start with the finest type space $\overline{T}$, which is positional. This type space asks what the subject ranks first, second, and so on, all the way down to their last-ranked alternative. From that, we can generate coarser type spaces by "lumping together" certain rankings. 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 third- and fourth-ranked alternatives. The set $\{xyzw, xywz\}$ would

[26] See the remark following the proof of Proposition 5 for details.

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. These all represent examples of positional type spaces.

To formalize our definition, think of the ranking function $r_{\succeq}(x) = |\{y : y \succeq x\}$ as a bijection from X to $\{1, \dots, m\}$, where $r_{\succeq}(x) = 1$ indicates that x is top-ranked. If $B \subseteq \{1, \dots, m\}$ is a set of "lumped-together" ranking values, 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 $Q = \{B, B', \dots, B''\}$ of $\{1, \dots, m\}$, say that $\succeq$ and $\succeq'$ have the same Q rankings if $r_{\succeq}^{-1}(B) = r_{\succeq'}^{-1}(B)$ for every $B \in Q$. In this way, the partition Q defines a type space, which we call the Q-positional type space.

Definition 12. Let Q be a partition of $\{1, \dots, m\}$. The *Q-positional* type space, denoted T_Q, is the type space in which $t(\succeq) = t(\succeq')$ if and only if $\succeq$ and $\succeq'$ 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 $T_Q = \overline{T}$, the finest type space. For $Q = \{\{1, 2, 3\}\}$, $T_Q = \underline{T}$. ◊

Why do we study positional type spaces? Because if a neutral type space is elicitable, then we know it must be convex, and we now show that any convex neutral type space must be positional.[27] This allows us to restrict attention to positional type spaces. Furthermore, convex positional type spaces are identified by a partition Q in which each element must be an interval in $\{1, \dots, m\}$. For example, the type space from Figure 3 is a positional type space generated by $Q = \{\{1, 3\}, \{2\}\}$. The set $\{1, 3\}$ is not an interval, so the type space is not convex.

Lemma 3. *Type space T is neutral and convex if and only if T is a positional type space in which Q consists only of (possibly degenerate) intervals in* $\{1, \dots, m\}$.

Lemma 3 allows us to focus only on interval-based positional type spaces when searching for elicitable neutral type spaces. Our main result in this section is that, in fact, *all* such type spaces are elicitable, giving a complete characterization within the class of neutral type spaces.

Proposition 7. *Let T be a neutral type space. Then the following conditions are equivalent:*

(i) Type space T is elicitable with lotteries.

(ii) Type space T is convex.

[27] Positional type spaces are clearly neutral, but there are neutral type spaces that are not positional. One example is $T = \{\{xyz, yzx, zxy\}, \{zyx, yxz, xzy\}\}$.

(iii) *Type space T is positional with partition Q consisting only of (possibly degenerate) intervals in* $\{1, \dots, m\}$.

The order-theoretic structure of the set of neutral type spaces elicitable with lotteries is then especially simple. Let $\mathcal{Q}$ denote the set of partitions of $\{1, \dots, m\}$ into nonempty intervals. Observe that each $Q \in \mathcal{Q}$ can be identified with the set of "break points" between elements of the partition; in particular, each $Q \in \mathcal{Q}$ can be identified with a subset of $\{1, \dots, m-1\}$. Thus, the cardinality of $\mathcal{Q}$ is 2^{m-1}. We can define the meet and join on $\mathcal{Q}$ in the usual way. For example, if $Q \in \mathcal{Q}$ has break points $\{1, 4, 7\}$ and $Q' \in \mathcal{Q}$ has break points $\{3, 7, 9\}$, then $Q \vee Q'$ (the join) has break points $\{1, 3, 4, 7, 9\}$ and $Q \wedge Q'$ (the meet) has the single break point $\{7\}$. Both of these consist of intervals and are, therefore, elicitable. These observations motivate the following proposition.

Proposition 8. *The set of neutral type spaces elicitable with lotteries is a distributive sublattice of the set of all type spaces, equipped with the refinement relation. It has cardinality* 2^{m-1}.

In particular, the lattice described in Proposition 8 is atomistic, where the atoms are those type spaces defined by partitions with exactly one break point. For example, if Q has break points $\{1, 4, 7\}$, then Q can be represented as the join of the three partitions with break points $\{1\}$, $\{4\}$, and $\{7\}$.

In the Savage framework, convexity is not enough to guarantee sufficiency: elicitability also requires that T be generated by top elements. This puts additional restrictions on Q. To understand those restrictions, note that $Q = \{\{1\}, \{2, \dots, m\}\}$ is generated by top elements from the menu X, which is of size m. Suppose instead we offer all menus of size $m-1$. This would elicit both the top-ranked alternative and the second-ranked alternative, giving $Q = \{\{1\}, \{2\}, \{3, \dots, m\}\}$. In general, if we offer all menus of size $m-k$, then we elicit the positional type space with $Q = \{\{1\}, \dots, \{k+1\}, \{k+2, \dots, m\}\}$. In fact, this completely characterizes the neutral type spaces that can be elicited using acts.

Proposition 9. *Let T be a neutral type space. Then the following conditions are equivalent:*

(i) *Type space T is elicitable with acts.*

(ii) *Type space T is positional and there exists $1 \leq k \leq m$ such that the partition Q that defines T is given by* $Q = \{\{1\}, \{2\}, \dots, \{k-1\}, \{k, \dots, m\}\}$.

Analogous to Proposition 8, the set of neutral lotteries elicitable with acts have an extremely simple structure: they are linearly ordered, can be indexed by k, and there are exactly m of them. Furthermore, elicitation is simple: The kth type space can be elicited by offering all menus of size $m-k+2$ and selecting one randomly for payment.[28]

[28] If $k = 1$, then the empty menu is used.

7. Multiple agents

In this section, we show that our analysis of elicitability can be extended easily 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 ith coordinate of a vector is omitted. The principal wishes to learn the entire vector t, but instead of paying each individual separately, must ultimately select a single alternative for the society of n agents. This corresponds to a typical social choice setting, or an independent private values setting but without a numeraire.[29]

A T mechanism is a mapping $g : T \to \Delta(X)$. 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 13. 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 $t_{-i} \in T_{-i}$,

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

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.

Definition 14. A profile of type spaces $T = (T_1, \dots, T_n)$ is DIC-elicitable if there exists a DIC T mechanism $g : T \to \Delta(X)$.

We now show that elicitability of each T_i on its own is both necessary and sufficient for DIC elicitability 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 10. *The profile of type spaces $T = (T_1, \dots, T_n)$ is DIC-elicitable if and only if each T_i is elicitable with lotteries.*

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 $g_i : T_i \to \Delta(X)$ be an IC T_i mechanism. For every $t \in T$, define

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

[29]Formally, the principal wants to elicit the joint type space that is the Cartesian product of individual type spaces (see Williams, 2008 for an analysis of product structures in mechanism design). The problem of eliciting other partitions of O^n remains a question for future research. Furthermore, we do not assume a distribution over types and, thus, cannot leverage correlation among types in designing our mechanism.

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

A result similar to Proposition 10 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 $g : T \to \Delta(X)$, $i \in N$, 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 $\mathbb{E}_{\mu_{-i}}[g(t_i(\succeq), t_{-i})] \succ_i^* \mathbb{E}_{\mu_{-i}}[g(t_i, t_{-i})]$ for every i, every $\succeq_i$, and every $t_i \neq t_i(\succeq_i)$. 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: Proofs

Proof of Lemma 1

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 $\succeq$ and $\succeq'$ are in the same element of T, then by definition $\text{dom}_{\succeq}(X') = \text{dom}_{\succeq'}(X')$ for every $X' \in \tilde{I}(T)$. This implies that $\succeq$ and $\succeq'$ 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 2

We start with the following lemma, which originally appeared in our earlier work (Azrieli et al. (2020), Lemma 3).

Lemma 4. *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 $\succeq$ and $\succeq'$ be a pair of orders that (i) rank all members of $X \setminus X'$ above X', (ii) rank x and y above all remaining elements of X', and (iii) 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 to the contrary 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, nonsingleton menus, let the set $\text{SI}(\mathcal{X})$ denote the collection of sets surely identified by $\mathcal{X}$. Observe that

(i) $\mathrm{SI}(\mathrm{SI}(\mathcal{X})) = \mathrm{SI}(\mathcal{X})$

(ii) $\mathcal{X} \subseteq \mathrm{SI}(\mathcal{X})$

(iii) $\mathcal{X} \subseteq \mathcal{X}'$ implies $\mathrm{SI}(\mathcal{X}) \subseteq \mathrm{SI}(\mathcal{X}')$.

Therefore, SI forms a *closure operator*. Furthermore, this closure operator has the anti-exchange property, as defined in Edelmean and Jamison (1985). Namely, if $\mathcal{X} = \mathrm{SI}(\mathcal{X})$, $X', X'' \notin \mathcal{X}$, $X' \neq X''$, and $X' \in \mathrm{SI}(\mathcal{X} \cup \{X''\})$, then $X'' \notin \mathrm{SI}(\mathcal{X} \cup \{X'\})$. To see this latter point, observe that it follows from Lemma 4 that there are x and y for which $\{x, y\} \subseteq X'' \subseteq X'$; in particular, $X'' \subseteq X'$. If $X'' \in \mathrm{SI}(\mathcal{X} \cup \{X'\})$, then, similarly, $X' \subseteq X''$, contradicting $X' \neq X''$.

Now, by Edelmean and Jamison (1985, Theorem 2.1), for any $\mathcal{X}$, there is a unique minimal collection of menus $B(\mathcal{X})$ such that $\mathrm{SI}(\mathcal{X}) = \mathrm{SI}(B(\mathcal{X}))$. Since $\tilde{T}(X_1, \dots, X_l) = \tilde{T}(X'_1, \dots, X'_k)$ if and only if $\mathrm{SI}(\{X_1, \dots, X_l\}) = \mathrm{SI}(\{X'_1, \dots, X'_k\})$, it follows that if T is elicitable, then the set of collections of menus that generate T is the interval $[B(\tilde{I}(T)), \tilde{I}(T)]$. Denoting $\tilde{B}(T) = B(\tilde{I}(T))$ gives the result of the proposition. 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 $\mathrm{SI}(B(\mathcal{X}) \setminus \{X'\}) = \mathrm{SI}(B(\mathcal{X}))$, contradicting the minimality of $B(\mathcal{X})$.

Proof of Corollary 2

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

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

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 $x \in X$,

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

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 Corollary 2 directly. Suppose $T = \tilde{T}(X_1, \dots, X_l)$. Let λ be a full-support distribution on $\{1, \dots, l\}$ and define $g(t)(x) = \lambda(\{1 \leq i \leq l : \mathrm{dom}_{\succeq^t}(X_i) = x\})$, where $\succeq^t$ is an arbitrary choice from t. It is not hard to check that g is IC.

Proof of Proposition 4

Let g be an IC T mechanism and let $t \in T$. We 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 suffices 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.[30]

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 a nonempty 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.

Proof of Proposition 5

The key to the proof is the following lemma.

Lemma 5. *If g is an IC T mechanism and if $t, t' \in T$ are adjacent via an x to 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 to 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 $\bar{u}(z) = 0$ and $\bar{u}(w) = 1$ 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 $g(t)$ and $g(t')$ 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 is ranked above x according to $\succeq'$). ☐

[30]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 proposition now easily follows. Indeed, let g be an IC mechanism and suppose $\{t_1, \dots, t_k\} \subseteq T$ satisfies the assumption of the proposition. Then by Lemma 5 we have that $g(t_1)(x) > g(t_2)(x)$. Suppose to the contrary that there is no $1 < i \leq k$ and z such that t_i and t_{i+1} are adjacent via a z to x switch. Then it follows again from Lemma 5 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 5 says that a necessary condition for a T mechanism g to be IC is that if t and t' are adjacent via an x to 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 that 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 5 is not only necessary for g to be IC, it is also sufficient.

Proof of Proposition 6

We prove the contrapositive by showing that if $t \in T$ is a nonconvex type, then T contains a cycle. Since t is not convex, there are two utility vectors $u, u' \in U(t)$ such that the interval $[u, u'] = \{\lambda u' + (1-\lambda)u : \lambda \in [0, 1]\}$ intersects the set $U(t')$ for some $t' \neq t$. By perturbing u and u' if needed, we may assume that at most two alternatives are tied at any point in the interval $[u, u']$. Let $(t_1 = t, t_2, \dots, t_k, t_1 = t)$ be the (ordered) collection of types that the interval $[u, u']$ intersects as λ goes from 0 to 1. We now argue that the above cycle contains a violation of the condition in Proposition 5.[31]

Consider a directed graph where the set of nodes is X and a directed edge (x, y) exists if and only if there is $1 \leq i \leq k$ such that the transition from t_i to $t_{(i+1) \bmod k}$ occurs via an x to y switch (i.e., x is ranked just above y in t_i but just below y in $t_{(i+1) \bmod k}$). Now, if the condition of Proposition 5 is satisfied, then any alternative that has an incoming edge also has an outgoing edge. Thus, the no-cycles condition on the type space implies that this directed graph of alternatives *does* have a cycle of the form $(x_1, x_2, \dots, x_l, x_1)$.

For any two adjacent alternatives in this cycle, say x_j and $x_{j+1 \bmod l}$, since we have an x_j to $x_{j+1 \bmod l}$ switch, we must have that $u(x_j) > u(x_{j+1 \bmod l})$. (This follows because $(1-\lambda)[u(x_j) - u(x_{j+1 \bmod l})] + \lambda[u'(x_j) - u'(x_{j+1 \bmod l})]$ is linear in λ, so there cannot be both an x_j to $x_{j+1 \bmod l}$ switch *and* a $x_{j+1 \bmod l}$ to x_j switch in the cycle.) But then we have $u(x_1) > u(x_2) > \cdots > u(x_l) > u(x_1)$, a contradiction.

Proof of Lemma 3

Suppose T is neutral and convex. Let $\bar{t} \in T$. Because of convexity, there is a partial order $\sqsupseteq_t$ for which $\succeq \in t$ if and only if $\succeq$ extends $\sqsupseteq_t$. Suppose there are $x, y \in X$ for which x and y are unranked according to $\sqsupseteq_t$. We claim that there is $\succeq \in t$ for which $x \succeq y$ and for which

[31] To get a violation of the condition as written in the proposition, one may need to "start the cycle" at one of the other types along the way.

there is no z with $x \succ z \succ y$ (i.e., x covers y). The argument is a standard analogue of the finite version of Szpilrajn's theorem (see Davey and Priestley (2002) Exercise 1.29(ii)), so we only sketch it here.

Define $\succeq_1$ to be the transitive closure of $\sqsupseteq_t \cup\{(x, y)\}$ and observe that $\succeq_1$ remains a partial order extending $\sqsupseteq_t$ for which x covers y. Now supposing that $\succeq_n$ is a partial order extending $\sqsupseteq_t$ for which x covers y, we define $\succeq_{n+1}$ by picking any pair (z, w) that is unranked. If neither z nor w is ranked, let $\succeq_{n+1}$ denote the transitive closure of $\succeq_n \cup\{(z, w)\}$. Suppose without loss that $w \notin \{x, y\}$: If $z = x$, define $\succeq_{n+1}$ to be the transitive closure of $\succeq_n \cup\{(w, x)\}$; if $z = y$, define $\succeq_{n+1}$ to be the transitive closure of $\succeq_n \cup\{(y, w)\}$. Continue inductively and use the fact that X is finite to establish that for large enough n, $\succeq_n$ will be complete (and, hence, a linear order).

Observe next that by taking the permutation $\pi(x) = y$, $\pi(y) = x$, and $\pi(z) = z$ for $z \notin \{x, y\}$, it follows that $\succeq_\pi \in t$ (since $\succeq_\pi$ also extends $\sqsupseteq_t$).

This fact together with neutrality implies that $\pi(t) = t$. Now suppose that $z \succeq x$ for all $\succeq \in t$. Then, by using the permutation π in the previous paragraph, we know that $z \succeq_\pi y$ for all $\succeq \in t$; but the previous paragraph established that $\{\succeq_\pi : \succeq \in t\} = t$.

We have shown that for any t and any $\sqsupseteq_t$ defining t, if x and y are incomparable, then $z \sqsupseteq_t x$ if and only if $z \sqsupseteq_t y$. Similarly, we can show that $x \sqsupseteq_t z$ if and only if $y \sqsupseteq_t z$.

So fix $\bar{t} \in T$. We can write $x \parallel_{\bar{t}} y$ when x and y are unranked according to $\sqsupseteq_{\bar{t}}$. The preceding establishes that $\succeq^*_{\bar{t}} = \sqsupseteq_{\bar{t}} \cup \parallel_{\bar{t}}$ is a weak order, and that $\succeq \in \bar{t}$ if and only if $x \succ^*_{\bar{t}} y$ implies $x \succ y$. Let us write the equivalence classes of $\succeq^*_{\bar{t}}$ in descending order as $B_1, \ldots, B_k$, where $k \geq 1$. Now define $Q \equiv \{\{1, \ldots, |B_1|\}, \ldots, \{\sum_{l=1}^{k-1} B_l| + 1, \ldots, \sum_{l=1}^{k} |B_l|\}\}$. In particular, all $\succeq \in \bar{t}$ have the same Q ranking, and it is obvious that Q consists of intervals.

Now let T_Q be given; we claim that $T = T_Q$. Let $t \in T$. Suppose that $\succeq, \succeq' \in t$. Let π be a permutation for which $\succeq_\pi \in \bar{t}$. Then by neutrality, $\succeq'_\pi \in \bar{t}$, so that they must have the same Q ranking. Conversely if $\succeq$ and $\succeq'$ have the same Q ranking, then let π be any permutation for which $\succeq_\pi \in \bar{t}$. Since $\succeq'_\pi$ has the same Q ranking as $\succeq_\pi$, we know that they are each in $\bar{t}$, which implies that $\succeq$ and $\succeq'$ are in the same element $t \in T$.

The other direction is clear: If T_Q is positional, then it is neutral, and if Q is defined by intervals, then T_Q is convex.

Proof of Proposition 7

(i)$\Longrightarrow$(ii). This follows from Proposition 4.

(ii)$\Longrightarrow$(iii). Follows directly from Lemma 3.

(iii)$\Longrightarrow$(i). 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 $i < j$. Then

$$T_Q = \tilde{T}\left(X, \ldots, X; |B_1|, |B_1| + |B_2|, \ldots, \sum_{1 \leq i \leq K-1} |B_i|\right),$$

that is, T_Q is generated by top sets. By Proposition 3 we are done.

Proof of Proposition 9

(i)$\Longrightarrow$(ii). Since T is elicitable with acts, by Proposition 1, it is generated by top elements. By Corollary 2, T is elicitable with lottery payments. By Proposition 4, T is convex, and, therefore, since T is neutral, it follows from Lemma 3 that T is positional with $T = T_Q$ for some partition Q.

By Proposition 1, if T_Q is elicitable, then it is generated by top elements. Since T_Q is positional, it follows that if a certain menu X' is identified by T_Q, then so is any other menu X'' for which $|X''| = |X'|$. Notice that if all menus of a certain size k are identified by some type space, then so are all menus of size larger than $k > 1$ (recall Lemma 4). It follows that either T_Q is trivial or it is generated by the top elements of a collection of menus of the form $\{X' \subseteq X : |X'| = k\}$ for some $2 \leq k \leq m$. This proves that T_Q is as in part (ii) of the proposition.

(ii)$\Longrightarrow$(i). The type space T_Q with $Q = \{\{1\}, \{2\}, \dots, \{k-1\}, \{k, \dots, m\}\}$ is generated by the top elements of the collection of menus $\{X' \subseteq X : |X'| = m - k + 2\}$. By Proposition 1, T_Q is elicitable.

References

Abdulkadiroğlu, Atila, Parag A. Pathak, and Alvin E. Roth (2005), "The New York city high school match." *American Economic Review Papers and Proceedings*, 95, 364–367. [507]

Azrieli, Yaron, Christopher P. Chambers, and Paul J. Healy (2018), "Incentives in experiments: A theoretical analysis." *Journal of Political Economy*, 126, 1472–1503. [509, 510, 512, 515]

Azrieli, Yaron, Christopher P. Chambers, and Paul J. Healy (2020), "Incentives in experiments with objective lotteries." *Experimental Economics*, 23, 1–29. [510, 515, 522, 531]

Bahel, Eric and Yves Sprumont (2020), "Strategyproof choice of social acts." *American Economic Review*, 110, 596–627. [511]

Birkhoff, Garrett (1996), "Lattice theory", volume 25 of Colloquium Publications. American Mathematical Society. [518]

Brier, Glenn W. (1950), "Verification of forecasts expressed in terms of probability." *Monthly weather review*, 78, 1–3. [512]

Caplin, Andrew, Mark Dean, and John Leahy (2019), "Rational inattention, optimal consideration sets, and stochastic choice." *Review of Economic Studies*, 86, 1061–1094. [515]

Carroll, Gabriel (2012), "When are local incentive constraints sufficient?" *Econometrica*, 80, 661–686. [527, 534]

Davey, B. A. and H. A. Priestley (2002), *Introduction to Lattices and Order*. Cambridge University Press. [518, 519, 535]

Edelmean, Paul H. and Robert E. Jamison (1985), "The theory of convex geometries." *Geometriae Dedicata*, 19, 247–270. [532]

Gabor, Andre and Clive William John Granger (1977), *Pricing: Principles and Practices*. Heinemann Educational, London. [507]

Gibbard, Allan (1977), "Manipulation of schemes that mix voting with chance." *Econometrica*, 45, 665–681. [512]

Gilboa, Itzhak and David Schmeidler (1989), "Maxmin expected utility with a non-unique prior." *Journal of Mathematical Economics*, 18, 141–153. [514]

Haeringer, Guillaume and Flip Klijn (2009), "Constrained school choice." *Journal of Economic Theory*, 144, 1921–1947. [507]

Hurwicz, Leonid (1960), "Optimality and informational efficiency in resource allocation processes." In *Mathematical Methods in the Social Sciences: Proceedings of the First Stanford Symposium* (Kenneth J. Arrow, Samuel Karlin, and Patrick Suppes, eds.), volume IV of Stanford Mathematical Studies in the Social Sciences. Stanford University Press. [507]

Hurwicz, Leonid (1977), "On the dimensional requirements of informally decentralized Pareto-satisfactory processes." In *Studies in Resource Allocation Processes* (Kenneth J. Arrow and Leonid Hurwicz, eds.), 413–424, Cambridge University Press, New York. [507]

Kos, Nenad (2012), "Communication and efficiency in auctions." *Games and Economic Behavior*, 75, 233–249. [507, 508]

Kroll, Yoram, Haim Levy, and Amnon Rapoport (1988), "Experimental tests of the separation theorem and the capital asset pricing model." *American Economic Review*, 78, 500–519. [507]

Lambert, Nicolas (2019), "Elicitation and evaluation of statistical forecasts." Unpublished paper, Stanford Graduate School of Business, Stanford University. [511, 523]

Lambert, Nicolas S., David M. Pennock, and Yoav Shoham (2008), "Eliciting properties of probability distributions." In *EC'08 Proceedings of the 9th ACM Conference on Electronic Commerce*, 129–138, Chicago, Illinois. [511]

Loomes, Graham (1991), "Evidence of a new violation of the independence axiom." *Journal of Risk and Uncertainty*, 4, 91–108. [507]

Machina, Mark J. and David Schmeidler (1995), "Bayes without Bernoulli: Simple conditions for probabilistically sophisticated choice." *Journal of Economic Theory*, 67, 106–128. [514]

Matějka, Filip and Alisdair McKay (2015), "Rational inattention to discrete choices: A new foundation for the multinomial logit model." *American Economic Review*, 105, 272–298. [515]

Mount, Kenneth and Stanley Reiter (1974), "The informational size of message spaces." *Journal of Economic Theory*, 8, 161–192. [507]

Nisan, Noam and Ilya Segal (2006), "The communication requirements of efficient allocations and supporting prices." *Journal of Economic Theory*, 129, 192–224. [507]

Osband, Kent H. (1985), *Providing Incentives for Better Forecasting*. Ph.D. thesis, University of California, Berkeley. [511, 523]

Savage, Leonard (1971), "Elicitation of personal probabilities and expectations." *Journal of the American Statistical Association*, 66, 783–801. [512]

Schervish, Mark J. (1989), "A general method for comparing probability assessors." *Annals of Statistics*, 17, 1856–1879. [512]

Shapley, Lloyd (1971), "Cores of convex games." *International Journal of Game Theory*, 1, 11–26. [522]

Sims, Christopher A. (2003), "Implications of rational inattention." *Journal of Monetary Economics*, 50, 665–690. [515]

Warner, Stanley L. (1965), "Randomized response: A survey technique for eliminating evasive answer bias." *Journal of the American Statistical Association*, 60, 63–69. [507]

Williams, Steven R. (2008), *Communication in Mechanism Design: A Differential Approach*. Cambridge University Press. [530]

Co-editor Simon Board handled this manuscript.

Manuscript received 20 March, 2020; final version accepted 13 August, 2020; available online 18 August, 2020.