Theoretical Economics 18 (2023), 97–127 1555-7561/20230097

On guarantees, vetoes, and random dictators

Anna Bogomolnaia
Adam Smith Business School, University of Glasgow and CNRS-CES

Ron Holzman
Department of Mathematics, Technion-Israel Institute of Technology

Hervé Moulin
Adam Smith Business School, University of Glasgow

A mechanism guarantees a certain welfare level to its agents, if each of them can secure that level against unanimously adversarial others. How high can such a guarantee be, and what type of mechanism achieves it?

In the n-person probabilistic voting/bargaining model with p deterministic outcomes a guarantee takes the form of a probability distribution over the ranks from 1 to p. If $n \geq p$, the *uniform* lottery is shown to be the only maximal (unimprovable) guarantee. If $n < p$, combining (variants of) the familiar *random dictator* and *voting by veto* mechanisms yields a large family of maximal guarantees: it is exhaustive if $n = 2$ and almost so if $p \leq 2n$.

Voting rules à la Condorcet or Borda, even in probabilistic form, are ruled out by our worst case viewpoint.

Keywords. Worst case, guarantees, voting by veto, random dictator.

JEL classification. D63, D71.

1. Guarantees and mechanisms

Two staples of the collective decision literature, voting by veto and the random dictator mechanism, are commonly used, informally, to simplify negotiations: the committee may agree to dismiss first a number of "obviously bad" outcomes, or to resolve disagreements by flipping a coin. We show formally that, when the goal is to offer the best ex ante protection to individual agents, the two mechanisms and their many combinations stand out in the *worst case* analysis.

Fix an arbitrary collective decision problem by its feasible outcomes (allocation of resources, public decision making, etc.), the domain of individual preferences and the

Anna Bogomolnaia: anna.bogomolnaia@glasgow.ac.uk
Ron Holzman: holzman@technion.ac.il
Hervé Moulin: herve.moulin@glasgow.ac.uk
We are grateful for the comments of seminar participants in the University Paris 1, Lancaster, Oxford, and Ashoka Universities. We thank two anonymous referees for their comments, which helped to substantially improve the exposition. Moulin's research was supported by a Blaise Pascal Chair of the Region Ile-de-France, 2020–2021. Work done during a Holzman's visit at Princeton University.

Available at https://econtheory.org. https://doi.org/10.3982/TE4832

number n of relevant agents. We evaluate a mechanism (game form) solving this problem by the *guarantee* it offers to the participants: this is the welfare level each one can secure in the resulting game no matter what other players know or how they will play their part. The worst case assumption is that their moves are collectively adversarial, so my guarantee is the value of the two-person zero-sum game where I play first against the rest of the world. A higher guarantee is a better default option, it encourages acceptance of and participation in the mechanism.

Question. How high can such a guarantee be, and what type of mechanism achieves it?

Guarantees test the ex ante fairness of a mechanism, from the viewpoint of an agent clueless about other participants or unwilling to engage in risky strategic moves. It is also an ex post test: an agent using a best reply to the other agents' strategies gets at least her guaranteed welfare (because she has a safe strategy achieving that level no matter what), so any Nash equilibrium of the game delivers at least that level of welfare to everyone.

The guarantee approach is far from new to economic theory (Section 2 briefly reviews the literature), but this paper is the first to use it in the probabilistic voting model, interpreted equivalently as a bargaining model. There are finitely many pure (deterministic) outcomes and we can choose a convex compromise between these by running a lottery, or allocating time shares, or dividing a budget; for clarity we use the lottery interpretation. We also maintain a symmetric treatment of agents (anonymity) and of outcomes (neutrality). Therefore, it only takes the number n of agents and p of pure outcomes to define a guarantee: it is a lottery λ over the ranks 1 to p (where rank 1 is the worst and p the best), which is *feasible* in the sense that for any profile of utilities of the agents we can find a lottery over pure outcomes for which every agent's expected utility is at least that from the lottery λ.[1] An equivalent definition for agents endowed with purely ordinal preferences plays a key role throughout; see Definitions 1 and 2 in Section 3.

The worst case viewpoint immediately rules out the usual voting methods à la Condorcet or Borda, whether in deterministic or probabilistic form. When there are three or more agents, if everyone else reports the preference opposite to mine, my worst outcome is selected. We show that it is always feasible to offer a much better guarantee to everyone. As explained below, our approach is relevant for problems involving more outcomes than agents selecting one of them: this points to small committees and direct bargaining, emphatically not to political elections.

In our model, a guarantee λ is maximal (unimprovable) if no other guarantee λ' stochastically dominates λ. Our results are of two types: for some values of (n, p) we describe completely the set of maximal guarantees and simple mechanisms to implement them; for the other pairs (n, p) the structure of this set resists a full characterization, so we only construct and implement a large subset of maximal guarantees. In both cases, the guarantees at the center of our analysis, and corresponding mechanisms, are built

[1] That is the utility of her worst outcome with probability λ_1, of her next worst outcome with probability λ_2, and so on.

from the three ingredients voting by veto, random priority rules, and uniformly random choice.

The simplest guarantee of all is the *uniform guarantee*, denoted $\mathrm{uni}(p)$, choosing each outcome with probability $\frac{1}{p}$. One way to implement it is to run this lottery and ignore individual preferences entirely. A more interesting way uses $\mathrm{uni}(p)$ as a canonical disagreement option: to any given mechanism Φ, we add a last round where each participant can reject the outcome proposed by Φ and force a uniformly random choice of the final outcome.

The uniform guarantee $\mathrm{uni}(p)$ is maximal for any (n, p). If $n \geq p$, and only then, it is the *only* maximal guarantee: any other guarantee is weakly worse for everyone in all problems, and sometimes strictly (the proof is easy; see Section 4).

Problems with two agents, $n = 2$, are not much harder to crack and the result is more interesting. A maximal guarantee is any lottery symmetric with respect to the median rank (Proposition 2, Section 4); these lotteries cover a polytope of which the vertices are easily described and implemented. To fix ideas suppose $p = 6$, so the polytope is a triangle with $\mathrm{uni}(6) = (\frac{1}{6}, \frac{1}{6}, \frac{1}{6}, \frac{1}{6}, \frac{1}{6}, \frac{1}{6})$ in its center and the following vertices. The *random dictator* guarantee $\mathrm{rd}(2, 6) = (\frac{1}{2}, 0, 0, 0, 0, \frac{1}{2})$ obtains as usual by tossing a fair coin and letting the winner choose a pure outcome. To implement the guarantee $\lambda = (0, 0, \frac{1}{2}, \frac{1}{2}, 0, 0)$, we ask first each agent to *veto* two outcomes,[2] after which the rule picks uniformly at random a nonvetoed outcome. To implement the third vertex $\mu = (0, \frac{1}{2}, 0, 0, \frac{1}{2}, 0)$, we give one veto token to one agent and four tokens to the other, choosing the roles by tossing a fair coin (equivalently each agent gets to veto one outcome, then a random dictator completes the choice).

Our next example illustrates some of the difficulties in tackling problems with more than two agents.

Example: Three agents $n = 3$, six outcomes $p = 6$ Define the *veto guarantee* $\mathrm{vt}(3, 6) = (0, \frac{1}{3}, \frac{1}{3}, \frac{1}{3}, 0, 0)$ (recall the first coordinate is the worst rank), implemented by first giving one veto token to each agent then choosing uniformly at random one of the remaining outcomes: after the vetoing round, my worst case is that the other agents kill my two best outcomes and I kill my worst outcome; hence the rank distribution cannot be worse than $\mathrm{vt}(3, 6)$. Note that distributing one veto token followed by a deterministic choice implements $\lambda = (0, 1, 0, 0, 0, 0)$, dominated by $\mathrm{vt}(3, 6)$.

The "naive" random dictator mechanism implements the guarantee $\lambda^1 = (\frac{2}{3}, 0, 0, 0, 0, \frac{1}{3})$: my worst case is that the two other agents pick my worst outcome. We can do better. Let each agent report (one of) his top outcome(s); if they all agree on a choose a; if they each choose a different outcome, pick one of them with uniform probability; finally, if the choices are a, a, b, we randomize uniformly between a, b *and an arbitrary third outcome* c. This implements the correct *random dictator guarantee* $\mathrm{rd}(3, 6) = (\frac{1}{3}, \frac{1}{3}, 0, 0, 0, \frac{1}{3})$, that dominates λ^1.[3]

[2] Whether these choices are simultaneous or sequential has no impact on the implemented guarantee.

[3] Modify the above mechanism RD as follows: if they all agree on a select the outcome uniformly between a and two other arbitrary outcomes. We still implement the guarantee $\mathrm{rd}(3, 6)$ but the new mechanism itself is dominated in the obvious sense by RD.

It is easy to check directly that uni(6), rd(3, 6), and vt(3, 6) are maximal guarantees. This follows for rd(3, 6) and vt(3, 6) by inspecting respectively the left or right profile of strict ordinal preferences

$$\begin{array}{ll} \prec_1 a\ b\ x\ y\ z\ c & \quad \prec_1 a\ x\ y\ z\ b\ c \\ \prec_2 b\ c\ y\ z\ x\ a & \quad \prec_2 b\ y\ z\ x\ c\ a \\ \prec_3 c\ a\ z\ x\ y\ b & \quad \prec_3 c\ z\ x\ y\ a\ b \end{array}$$

(where agent 1's worst is a and best is c). At the left profile, giving a $\frac{1}{3}$ chance of their best outcome to each agent requires to pick a, b, or c, each with probability $\frac{1}{3}$: then each agent experiences exactly rd(3, 6) over her ranked outcomes, implying that any lottery λ dominating rd(3, 6) is not feasible at this profile. Similarly at the right profile, implementing vt(3, 6) implies zero probability on a, b, c, and at most (hence exactly) $\frac{1}{3}$ on each of x, y, and z. The symmetry of these two arguments is not a coincidence: it is explained by the duality relation connecting vt(3, 6) and rd(3, 6) (Section 5).

What other guarantees are maximal for $n = 3,\ p = 6$? Convex combinations preserve feasibility but not maximality: for instance, an equal chance of the two mechanisms implementing vt(3, 6) and rd(3, 6), respectively, delivers the guarantee $\frac{1}{2}\mathrm{vt}(3,6) + \frac{1}{2}\mathrm{rd}(3,6) = (\frac{1}{6}, \frac{1}{3}, \frac{1}{6}, \frac{1}{6}, 0, \frac{1}{6})$ dominated by uni(6). But lotteries between uni(6) and vt(3, 6), or between uni(6) and rd(3, 6) are maximal. In fact, the two intervals [uni(6), vt(3, 6)] and [uni(6), rd(3, 6)] capture *all* maximal lotteries for $n = 3,\ p = 6$ (Theorem 1, Section 4.3).

The two intervals in the above example illustrate the typical choices faced by a designer poised to maximize individual welfare guarantees. The veto guarantee is a reasonable option when bargaining is about choosing an expensive infrastructure project, or a person to hold a position for life; the random dictator approach makes sense if we are dividing time between several public decisions, like two alternating roman consuls; the uniform guarantee stands out if we value a disagreement outcome revealing no information about individual preferences.

The punch lines We can describe the set of maximal guarantees if $n = 2$ or $n \geq p$, as well as if there are at most two outcomes per person: if $3 \leq n < p \leq 2n$, just like in the example above, the set of maximal guarantees contains the *nonconvex* union of the two intervals of lotteries from the uniform guarantee uni(p) to either the veto vt(n, p) or random dictator rd(n, p) and not much more; see Theorem 1 and Proposition 3 in Section 4.

In the remaining cases $n \geq 3,\ p > 2n$, we do not know the full structure of the set of maximal guarantees but we provide some useful insights. First, this set is a nonconvex finite union of polytopes, all sharing the uniform guarantee as a vertex (Proposition 5). Second, if d is the strict integer part[4] of $\frac{p}{n}$, we can construct 2^d such polytopes by concatenating exactly d elementary rounds of vetoes or random dictator (Theorem 2, Section 6).

[4]The largest integer strictly smaller than ...

We also provide simple mechanisms implementing the maximal guarantees we identify. Critical to their practical application, these mechanisms rely on ordinal preferences only, as do the agents' safe actions when they report which outcome(s) they veto, or which ones they prefer among those still in play.

2. Related literature

The optimal design of a mechanism under the risk averse assumption that other agents are adversarial is discussed by the early literature on implementation in several slightly different formulations: implementation in maximin (Thomson (1978), Dasgupta, Hammond, and Maskin (1979)), prudent (Moulin (1981)) or protective strategies (Barbera and Dutta (1982)). As explained in Section 1, our guarantees are compatible with a wide range of strategic behaviors.

Steinhaus' seminal papers (Steinhaus (1949), see also Dubins and Spanier (1961), Kuhn (1967)) invented the worst case approach for cutting a cake fairly among any number of agents. His simple mechanism generalizes Divide and Choose and guarantees to each agent a *fair share*: one that is worth at least $\frac{1}{n}$ of the whole cake. The main focus of the subsequent literature is envy- free divisions: how to achieve one by simple cuts and queries (Brams and Taylor (1995), Robertson and Webb (1998), Aziz and McKenzie (2016)) and proving its existence under preferences more general than additive utilities (Stromquist (1980), Woodall (1980)). An exception is the recent paper of Bogomolnaia and Moulin (2020) returning to the worst case approach under very general preferences and identifying the MinMaxShare (my best share in the worst partition of the cake I can be offered) as a feasible guarantee, though not a maximal one.

The last decade saw an explosion of research to define and compute a fair allocation of indivisible items, proposing in particular a new definition of the fair share as the MaxMinShare (Budish (2011)): my worst share in the best partition of the objects I can propose. This guarantee may not be feasible (Procaccia and Wang (2014)) but this happens very rarely (Kurokawa, Procaccia, and Wang (2016)); the real concern is that the mechanisms approximating this guarantee are all but simple.

Other early instances of the worst case approach are in production economies (Moulin (1992a, 1992b)) and in the minimal cost spanning tree problem (Hougaard, Moulin, and Osterda (2010)).

The random dictator mechanism is a staple of probabilistic social choice (Gibbard (1977), Sen (2011)). In axiomatic bargaining, it inspires the Raiffa solution (Raiffa (1953)) and the mid-point domination axiom (Sobel (1981), Thomson (1981)) satisfied by both the Nash and Kalai–Smorodinsky solutions.

Voting by veto is another early idea introduced by Mueller (1978) to incentivize agents toward compromising offers: each agent makes one offer, which together with the status quo outcome makes $p = n + 1$ outcomes, after which they take turns to veto one outcome each (in our model the natural status quo is the uniform lottery over outcomes). This procedure is generalized in Moulin (1981). The area monotonic bargaining solution (Anbarci and Bigelow (1994), Anbarci (1993)) is a direct application of voting by veto between two parties, similar to distributing $\lfloor \frac{p-1}{2} \rfloor$ veto tokens to each agent in our model.

A handful of recent papers discuss variants of voting by veto in the classic implementation context; see, for example, De Clippel, Eliaz, and Knight (2014), Barbera and Coelho (2017), Laslier, Nunez, and Sanver (2020). All three papers implement maximal guarantees. Closer to home, Section 4 in Kirneva and Nunez (2021) explains the strategic properties of a veto mechanism implementing arbitrary compositions of our guarantee $vt(n, p)$.

We mention finally the small literature on bargaining with cash compensations and quasilinear utilities (Moulin (1985), Chun (1986)) where only the uniform guarantee is discussed, while our results unveil many more possibilities.

Organization of the paper Basic definitions are in Section 3, including guarantees, maximal or not, and their implementation in two models: first when preferences over outcomes are ordinal and agents compare lotteries by stochastic dominance; second, when they use von Neumann–Morgenstern (vNM) utilities to compare them. The two definitions are equivalent.

In Section 4, we describe and implement the maximal guarantees in the three special cases $n \geq p$, $n = 2$, and $p \leq 2n$. Section 5 introduces two technical tools critical for the proof of our two theorems: a duality operation respecting maximality and pairing voting by veto and random dictator, and the concatenation of fewer than $\frac{p}{n}$ of these building blocks. We derive in Section 6 the geometric structure of the set of maximal guarantees, and its subset obtained by the concatenation just mentioned; we also list a handful of open questions.

Section 7 concludes and the Appendix provides four proofs and illustrates one of our open questions.

3. Guarantees: Definitions

Anonymity and neutrality (symmetric treatment of agents and outcomes, respectively) are hard wired in the definition of a guarantee, which only depends upon the number n of agents and p of deterministic outcomes. It is an element λ of $\Delta(p)$, the simplex of lotteries over the ranks in $[p] = \{1, \dots, p\}$. Here, λ_1 is the probability of the worst rank and λ_p that of the best rank. We give two equivalent definitions of guarantees respectively for the case of agents with ordinal preferences or von Neumann–Morgenstern (vNM) utilities.

The set of deterministic outcomes is A with generic element a, and $\Delta(A)$, with generic element ℓ, is the set of lotteries over A. We keep in mind the alternative interpretations of $\Delta(A)$ as time sharing or division of a budget between the "pure" outcomes in A.

The set of agents is $[n]$, with generic element i. Agent i's ordinal preference over A (a complete, reflexive, and transitive relation) is written $\succsim_i$. Agent i's vNM utility over A is a vector u_i in $\mathbb{R}^A$ and $u_i \cdot \ell = \sum_{a \in A} u_{ia} \ell_a$ is her utility at lottery ℓ.

Notation. For lotteries $\lambda \in \Delta(p)$, we write $[\lambda]_{k_1}^{k_2}$ instead of the sum $\sum_{k_1}^{k_2} \lambda_t$. The symmetric of λ with respect to the middle rank is $\widetilde{\lambda}$: $\widetilde{\lambda}_k = \lambda_{p+1-k}$ for all $k \in [p]$.

The rank-ordered rearrangement (aka order statistics) of $u_i \in \mathbb{R}^A$ is denoted $u_i^* \in \mathbb{R}^p$ and obtained by writing the entries of u_i in nondecreasing order:

$$\forall k \in [p]: \sum_{t=1}^{k} u_{it}^* = \min\left\{\sum_{a\in T} u_{ia} | T \subseteq A, |T| = k\right\}$$

Similarly, given the ordinal preference $\succsim_i$ the rank-ordered rearrangement of the lottery $\ell \in \Delta(A)$ is denoted $\ell^{*i} \in \Delta(p)$ and defined as

$$\forall k \in [p]: \left[\ell^{*i}\right]_1^k = \min\left\{\sum_{a\in T} \ell_a | T \text{ is a } k\text{-tail of } \succsim_i\right\}$$

where a k-tail of $\succsim_i$ is a set of agent i's k worst outcomes (not necessarily unique, due to indifferences): so ℓ_1^{*i} is the weight of agent i's worst outcome (or the smallest weight of i's worst outcomes), and so on.

The stochastic dominance relation (dominance for short) in $\Delta(p)$ plays a central role throughout. We write $\lambda \vdash \mu$ and say that λ dominates μ if we have[5]

$$\left\{\forall k \in [p]: [\lambda]_1^k \leq [\mu]_1^k\right\} \quad \Longleftrightarrow \quad \left\{\forall k \in [p]: [\lambda]_k^p \geq [\mu]_k^p\right\}$$

Definition 1 (ordinal preferences). Given n and p, the lottery $\lambda \in \Delta(p)$ is a guarantee at n, p if for any n-profile π of preferences $\pi = (\succsim_i)_{i=1}^n$ on A there exists a lottery $\ell \in \Delta(A)$ such that $\ell^{*i} \vdash \lambda$ for all $i \in [n]$. Then we say that the lottery ℓ implements λ at profile π.

Definition 2 (vNM utilities). Given n and p, the lottery $\lambda \in \Delta(p)$ is a guarantee at n, p if for any n-profile of utilities $\tau = (u_i)_{i=1}^n$ on A there exists a lottery $\ell \in \Delta(A)$ such that $\ell \cdot u_i \geq \lambda \cdot u_i^*$ for all $i \in [n]$. Then we say that the lottery ℓ implements λ at profile τ.

The ordinal definition is agnostic with respect to the risk attitude of the agents. The cardinal one specifies it completely.

Lemma 1. *Definition 1, Definition 2, and the following property are equivalent:*

$$\textit{for any vNM profile } (u_i)_{i=1}^n: \quad \sum_{i=1}^n u_i = 0 \quad \Longrightarrow \quad \sum_{i=1}^n \lambda \cdot u_i^* \leq 0 \tag{1}$$

We write $\mathcal{G}(n, p)$ for the set of all guarantees at n, p: it is a polytope in $\Delta(p)$.

Property (1) refers to the *sum* of the agents' utilities, rather than to each agent's utility separately. Thus, it naturally corresponds to a variant of the cardinal definition, where utilities are transferable among agents and treated quasilinearly. The lemma shows that a guarantee is feasible if and only if it does not distribute more than the sum of utilities does.

[5]We also use the equivalent definition: $\forall z \in \mathbb{R}^p: \{z_1 \leq z_2 \leq \cdots \leq z_p\} \Longrightarrow \lambda \cdot z \geq \mu \cdot z$.

PROOF. *Definition 1 $\Longrightarrow$ Definition 2*

Fix a lottery $\ell \in \Delta(A)$ and suppose the vNM utility $u_i \in \mathbb{R}^A$ represents the same preference on A as $\succsim_i$. Writing u_i^*, ℓ^{*i} for the corresponding rearrangements, the identity $\ell \cdot u_i = \ell^{*i} \cdot u_i^*$ is easily checked. Fix now λ meeting Definition 1 and an arbitrary profile $(u_i)_{i=1}^n$ of vNM utilities, with associated ordinal preferences $(\succsim_i)_{i=1}^n$. If ℓ implements λ at $(\succsim_i)_{i=1}^n$, the relation $\ell^{*i} \vdash \lambda$ and the identity give $\ell \cdot u_i \geq \lambda \cdot u_i^*$ as desired.

Definition 2 $\Longrightarrow$ property (1)

Fix $(u_i)_{i=1}^n$ such that $\sum_{i=1}^n u_i = 0$ and choose ℓ implementing λ as in Definition 2: the inequalities $\ell \cdot u_i \geq \lambda \cdot u_i^*$ imply (1).

Property (1) $\Longrightarrow$ Definition 1

We check first that (1) implies, for all $(u_i)_{i=1}^n \in (\mathbb{R}^A)^n$, the inequality

$$\sum_{i=1}^n \lambda \cdot u_i^* \leq \max_{a\in A} \sum_{i=1}^n u_{ia} \tag{2}$$

Fix any $(u_i)_{i=1}^n$ and set $z = \max_{a\in A} \sum_{i=1}^n u_{ia}$. Writing $\mathbf{1}$ the vector with all coordinates equal to 1, we pick a profile $(v_i)_{i=1}^n$ such that $u_i \leq v_i$ for all i and $\sum_{i=1}^n v_{ia} = z$ for all a. Applying now (1) to $(w_i)_{i=1}^n$: $w_i = v_i - \frac{z}{n}\mathbf{1}$, we have $\sum_{i=1}^n \lambda \cdot u_i^* \leq \sum_{i=1}^n \lambda \cdot v_i^* \leq z$ and the claim.

Fix now λ meeting property (2) and a preference profile $(\succsim_i)_{i=1}^n$. Call S_i the set of utilities $v_i \in [0, 1]^A$ representing $\succsim_i$ weakly: $a \succsim_i b \Longrightarrow v_{ia} \geq v_{ib}$ for all a, b. Note that S_i is the closure of the set of utilities representing $\succsim_i$ exactly. By property (2) for any profile $(v_i)_{i=1}^n \in \Pi_{i=1}^n S_i$, there exists $a \in A$ such that $\sum_{i=1}^n v_{ia} \geq \sum_{i=1}^n \lambda \cdot v_i^*$, which implies

$$\min_{(v_i)_{i=1}^n \in \Pi_{i=1}^n S_i} \max_{a\in A} \sum_{i=1}^n \left(v_{ia} - \lambda \cdot v_i^*\right) \geq 0$$

The summation is a linear function of the variable $(v_i)_{i=1}^n$ varying in a convex compact, and of a. By the minimax theorem, there exists $\ell \in \Delta(A)$ such that $\sum_{i=1}^n \ell \cdot v_i \geq \sum_{i=1}^n \lambda \cdot v_i^*$ for all $(v_i)_{i=1}^n \in \Pi_{i=1}^n S_i$. Taking $v_i = 0$ for all $i \geq 2$ gives $\ell \cdot v_1 \geq \lambda \cdot v_1^*$ for all $v_1 \in S_1$. Equivalently, $\ell^{*1} \cdot v_1^* \geq \lambda \cdot v_1^*$ for any weakly increasing sequence v_1^* in $[0, 1]^p$: the desired property $\ell^{*1} \vdash \lambda$ follows, and the argument is the same for each $i \geq 2$:

$\mathcal{G}(n, p)$ is a polytope (the bounded intersection of finitely many closed half-spaces)

Fix an ordinal preference profile π. In order for a guarantee $\lambda \in \Delta(p)$ to be implemented by a lottery $\ell \in \Delta(A)$ as in Definition 1, the lottery ℓ has to satisfy a system of linear inequalities of the form $\ell_a \geq \lambda_p$, $\ell_b \geq \lambda_p$, $\ldots$, $\ell_a + \ell_c \geq \lambda_p + \lambda_{p-1}$, $\ell_b + \ell_d \geq \lambda_p + \lambda_{p-1}$, $\ldots$, etc., $\ell_a + \ell_b + \cdots + \ell_z = 1$, $\ell_a \geq 0$, $\ell_b \geq 0, \ldots$ Note that the left-hand side of this linear system is determined by the profile π, and the right-hand side is an affine function of λ. By the Farkas lemma, the existence of a solution ℓ to this system is equivalent to λ lying in the intersection of a finite collection (determined by π) of closed half-spaces. Requiring that λ satisfy this for each of the finitely many possible ordinal profiles π results in a polytope, as claimed. $\square$

If $n = 2$, Lemma 1 gives a simple characterization of guarantees. Check the following identity (recall $\widetilde{\lambda}$ is the symmetric of λ with respect to the middle rank):

$$\forall u \in \mathbb{R}^A: \quad \lambda \cdot (-u)^* = -\widetilde{\lambda} \cdot u^* \tag{3}$$

and use it to rewrite property (1) in the case $n = 2$ as $\lambda \cdot u^* \leq \widetilde{\lambda} \cdot u^*$ for all u. By Footnote 5, the latter is equivalent to $\widetilde{\lambda} \vdash \lambda$, and thus

$$\lambda \in \mathcal{G}(2, p) \quad \Longleftrightarrow \quad [\lambda]_1^k \geq [\lambda]_{p+1-k}^p \quad \text{for all } k = 1, \dots, \left\lfloor \frac{p}{2} \right\rfloor \tag{4}$$

But for $n \geq 3$ it is much harder to discover a set of such inequalities representing $\mathcal{G}(n, p)$, or the set of its extreme points.

The implementation of a guarantee by an abstract mechanism follows in the obvious way from Definitions 1 and 2.

Definition 3. A mechanism Φ on $A = \{a_1, \dots, a_p\}$, $[n] = \{1, \dots, n\}$ is defined by a set of strategies X_i for each $i \in [n]$ and a mapping φ from $X_{[n]}$, the Cartesian product of the sets X_i, to $\Delta(A)$. We say that Φ implements a lottery $\lambda \in \Delta(p)$ if for any $i \in [n]$ and $\succsim_i$ agent i has a strategy $x_i \in X_i$ such that $\varphi(x_i, x_{-i})$ implements λ (according to Definition 1) for all $x_{-i} \in X_{[n] \setminus i}$.

The existence of this safe strategy ensures that, at any utility profile τ each player prefers (at least weakly) each Nash equilibrium outcome of the corresponding Φ-game to the λ outcome.

It is straightforward to see that a lottery $\lambda \in \Delta(p)$ is implementable at every preference profile (according to Definition 1) if and only if it is implementable by some mechanism (according to Definition 3). To check that a certain lottery λ on ranks is a guarantee in the sense of Definitions 1, 2, it is often convenient to describe a simple mechanism implementing it. But nonsimple mechanisms routinely deliver better equilibrium outcomes, from the welfare point of view, than a simple one. For example, the simplest implementation of the uniform guarantee $\mathrm{uni}(p)$ is to pick uniformly at random an outcome in A; but if Ψ is *any* bargaining mechanism, we also implement $\mathrm{uni}(p)$ by (1) playing Ψ, then (2) giving to every agent the option to force a uniform lottery on A (if she prefers this lottery to the outcome of Ψ).

Note that the guarantee λ is anonymous and neutral by construction, but it can be implemented by mechanisms which are neither. For instance, if $n = 2$, $p = 5$, the guarantee $\lambda = (0, 0, 1, 0, 0)$ is implemented by giving two veto tokens to each agent: this can happen sequentially or simultaneously.[6]

From the welfare point of view, the guarantees of interest are those that cannot be improved, the maximal ones.

[6]In the simultaneous version, we can add any mechanism to select among the possibly several nonvetoed outcomes.

Definition 4. The guarantee $\lambda \in \mathcal{G}(n, p)$ is maximal if

$$\forall \mu \in \mathcal{G}(n, p): \quad \mu \vdash \lambda \quad \Longrightarrow \quad \mu = \lambda$$

The set of maximal guarantees is $\mathcal{M}(n, p) \subset \mathcal{G}(n, p)$.

4. Maximal guarantees in three special cases

4.1 *Case 1: $n \geq p$*

The uniform lottery $\mathrm{uni}(p)$ is always feasible and if $n \geq p$ it is the best guarantee. We illustrate the argument for $p = 4$. Pick a profile where the preferences of a subset of four agents have the familiar cyclical pattern $a \prec b \prec c \prec d$, $b \prec c \prec d \prec a$, $c \prec d \prec a \prec b$, and $d \prec a \prec b \prec c$. Suppose that λ is a guarantee implemented by the lottery ℓ at this profile. Considering the first two preferences where a is successively the worst and best outcome gives $\lambda_4 \leq \ell_a \leq \lambda_1$; this is true for all outcomes so by summing up four inequalities we get $\lambda_4 \leq \frac{1}{4} \leq \lambda_1$. Definition 1 again gives $[\lambda]_3^4 \leq \ell_a + \ell_b \leq [\lambda]_1^2$ and the same is true for the pairs b, c, c, d, and d, a: therefore, $[\lambda]_3^4 \leq \frac{1}{2} \leq [\lambda]_1^2$. We conclude $\mathrm{uni}(p) \vdash \lambda$.

Proposition 1. *The uniform guarantee,* $\mathrm{uni}(p)_k = \frac{1}{p}$ *for all* $k \in [p]$*, has the following properties:*

(i) It is maximal for all n, p.

(ii) If $n \geq p$ it dominates every other feasible guarantee: $\mathcal{M}(n, p) = \{\mathrm{uni}(p)\}$.

(iii) If $n \geq 3$ it is a vertex of $\mathcal{G}(n, p)$, hence of $\mathcal{M}(n, p)$, too.

Intuitively, part (i) should be clear: if we start from $\mathrm{uni}(p) = (\frac{1}{p}, \dots, \frac{1}{p})$ and shift any probability mass to the right, we end up assigning higher probability to the top k ranks than to the bottom k ones (for some k). This is not feasible, because one agent's top k outcomes may be another agent's bottom ones. The intuition for part (ii) was illustrated in the $p = 4$ example above: when $n \geq p$ a full cyclical symmetry prevents us from assigning more than $\frac{k}{p}$ to the top k ranks, for any k. Part (iii) is similar, but a bit more technical.

Proof. *Statement* (i). The equality $\mathrm{uni}(p) \cdot u_i^* = \mathrm{uni}(p) \cdot u_i$ for all u_i implies for any profile $(u_i)_{i=1}^n$

$$\sum_{i=1}^{n} u_i = 0 \quad \Longrightarrow \quad \mathrm{uni}(p) \cdot \left(\sum_{i=1}^{n} u_i^*\right) = 0 \tag{5}$$

Suppose some $\mu \in \mathcal{G}(n, p)$ dominates $\mathrm{uni}(p)$ and consider a profile of the form $(u_1, -u_1, 0, \dots, 0)$ where u_1 is arbitrary. Summing up the inequalities $\mu \cdot u_1^* \geq \mathrm{uni}(p) \cdot u_1^*$ and $\mu \cdot (-u_1)^* \geq \mathrm{uni}(p) \cdot (-u_1)^*$ gives $\mu \cdot u_1^* + \mu \cdot (-u_1)^* \geq 0$. Because μ meets property (1) all three inequalities are equalities, and we conclude $\mu = \mathrm{uni}(p)$.

Statement (ii). Assume $n \geq p$ and pick an arbitrary guarantee λ in $\mathcal{G}(n, p)$, and a cyclical permutation σ of A: the latter maps utility u to u^σ: $u_a^\sigma = u_{\sigma(a)}$. We pick any u and consider the profile

$$\left(u, u^\sigma, u^{\sigma^2}, \dots, u^{\sigma^{p-1}}, 0, \dots, 0\right)$$

with $n - p$ null utilities. Clearly, $\sum_{k=0}^{p-1} u^{\sigma^k} = \gamma 1$ for $\gamma = \sum_{a \in A} u_a$, so we can apply property (1) to the profile $(u - \frac{\gamma}{p} 1, u^\sigma - \frac{\gamma}{p} 1, \dots, u^{\sigma^{p-1}} - \frac{\gamma}{p} 1, 0, \dots, 0)$. Together with $(u^{\sigma^k})^* = u^*$ for all k, this gives

$$0 \geq \sum_{k=0}^{p-1} \left(\lambda \cdot \left(u^{\sigma^k}\right)^* - \frac{\gamma}{p}\right) = p\left(\lambda \cdot u^*\right) - \gamma \quad \Longrightarrow \quad \lambda \cdot u^* \leq \frac{\gamma}{p} = \mathrm{uni}(p) \cdot u^*$$

as desired.

Statement (iii). Suppose $\mathrm{uni}(p)$ is a convex combination of two distinct λ^1, λ^2 in $\mathcal{G}(n, p)$. For any profile such that $\sum_{i=1}^n u_i = 0$, property (1) implies $\sum_{i=1}^n \lambda^s \cdot u_i^* \leq 0$ for $s = 1, 2$. Upon writing $\mathrm{uni}(p)$ in (5) as a convex combination of λ^1, λ^2, we see that the two inequalities sum to an equality therefore they are both equalities and each λ^s meets (5). For $n \geq 3$, only $\mathrm{uni}(p)$ does.

We check this claim for a generic lottery λ. Define $\delta_k = [\lambda]_{p+1-k}^p$ for $k \in [p]$ and $\delta_0 = 0$, then pick any three nonnegative integers k, l, m summing to p. Consider a profile of 0, 1 utilities where u_1; u_2; u_3 are equal to 1 precisely on three sets of respective sizes k, l, m partitioning A, while other utilities, if any, are identically zero. Applying (5) to this profile yields $\delta_k + \delta_l + \delta_m = 1$. It is easy to check that this simple integer version of the Cauchy equation implies $\delta_k = \frac{k}{p}$ for all k. □

4.2 *Case 2: $n = 2$*

Assuming $p = 6$ to fix ideas we show that a maximal lottery must be symmetric. Recall from the discussion after Lemma 1 that $\mathcal{G}(2, p)$ is defined by the inequalities (4). We assume that λ is a maximal guarantee, and show that $\lambda_1 = \lambda_6$, $\lambda_2 = \lambda_5$ and $\lambda_3 = \lambda_4$. Indeed, (4) implies that $\lambda_1 \geq \lambda_6$; this inequality cannot be strict, because then $(\lambda_1 - \varepsilon, \lambda_2 + \varepsilon, \lambda_3, \lambda_4, \lambda_5, \lambda_6)$ still satisfies (4) for small enough ε, contradicting the maximality of λ. Next, $\lambda_1 = \lambda_6$ and (4) imply that $\lambda_2 \geq \lambda_5$; again, this cannot be strict, because then $(\lambda_1, \lambda_2 - \varepsilon, \lambda_3 + \varepsilon, \lambda_4, \lambda_5, \lambda_6)$ still satisfies (4). Finally and similarly, we can deduce the remaining equality $\lambda_3 = \lambda_4$.

Proposition 2. *Maximal guarantees for $n = 2$.*

If $n = 2 < p$, the lottery $\lambda \in \Delta(p)$ is a maximal guarantee if and only if it is symmetric:

$$\lambda_k = \lambda_{p+1-k} \quad \text{for } 1 \leq k \leq \left\lfloor \frac{p}{2} \right\rfloor \tag{6}$$

The extreme points (vertices) of the polytope $\mathcal{M}(2, p)$ are the following guarantees λ^t:

$$\lambda_t^t = \lambda_{p+1-t}^t = \frac{1}{2} \quad \text{for } t = 1, \dots, \left\lfloor \frac{p}{2} \right\rfloor; \quad \text{and} \quad \lambda_{\frac{p+1}{2}}^{\frac{p+1}{2}} = 1 \quad \text{if } p \text{ is odd}$$

We see that for $n = 2$ the maximal guarantees are the symmetric ones. Thus, to choose a maximal guarantee we only need to allocate the total probability mass among the rank pairs $\{1, p\}$, $\{2, p-1\}$, etc. The extreme allocations are those that give the entire mass to one such pair. The center of the polytope $\mathcal{M}(2, p)$ is the uniform guarantee $\mathrm{uni}(p)$, contrasting sharply with the case $n \geq 3$, $p > n$ where $\mathrm{uni}(p)$ is a vertex of $\mathcal{M}(n, p)$ by Proposition 1.

Proof. To prove the *if* statement, we fix $\lambda \in \mathcal{G}(2, p)$ and symmetric. Rewrite (6) as $\lambda \cdot u^* = \widetilde{\lambda} \cdot u^*$ for all u^*, which by the identity (3) means $\lambda \cdot u^* = -\lambda \cdot (-u)^*$ for all u^*. The latter is property (5) for $n = 2$, so the maximality of λ follows as in the proof of statement (i) in Proposition 1.

To prove *only if* pick $\lambda \in \mathcal{G}(2, p)$, which means $\widetilde{\lambda} \vdash \lambda$ (i.e., (4)). As the dominance relation is preserved by convex combinations we have $\frac{1}{2}(\widetilde{\lambda} + \lambda) \vdash \lambda$ where $\frac{1}{2}(\widetilde{\lambda} + \lambda)$ is symmetric: thus, λ is dominated if it is not symmetric.

We let the reader check that the extreme points of the polytope defined by (6) are the lotteries with only two ranks in their support (or just the middle rank, when p is odd). $\square$

The mechanisms implementing the vertices of $\mathcal{M}(2, p)$ combine in a simple way the veto and random dictator ideas. Asking one randomly chosen agent to select a pure outcome implements $\lambda^1 = (\frac{1}{2}, 0, \dots, 0, \frac{1}{2})$, which we call the *random dictator* guarantee and denote $\mathrm{rd}(2, p)$ as in Section 1. Giving one veto token to each agent, then selecting one of the remaining outcomes with uniform probability implements the guarantee $\lambda = (0, \frac{1}{p-2}, \dots, \frac{1}{p-2}, 0)$, which we write $\mathrm{vt}(2, p)$. It is maximal, though not a vertex of $\mathcal{M}(2, p)$ except if $p = 3$ or 4.

To implement $\lambda^2 = (0, \frac{1}{2}, 0, \dots, 0, \frac{1}{2}, 0)$, we ask first each agent to veto one outcome, after which we pick a random dictator between the remaining $p - 2$ (or possibly $p - 1$ if the vetoes coincide) outcomes. Similarly, we implement λ^t by giving $t - 1$ veto tokens per person, then choosing a random dictator to pick one of the nonvetoed outcomes.

4.3 *Case 3:* $3 \leq n < p \leq 2n$

We define first the guarantees $\mathrm{vt}(n, p)$ and $\mathrm{rd}(n, p)$, already introduced in the three and two agent cases. First,

$$\mathrm{vt}(n, p) = \left(0, \frac{1}{p-n}, \dots, \frac{1}{p-n}, \overbrace{0, \dots, 0}^{n-1}\right)$$

is implemented by one round of veto (one token per person) followed by the uniformly random choice of a nonvetoed outcome. Next, for

$$\mathrm{rd}(n, p) = \left(\overbrace{\frac{1}{n}, \dots, \frac{1}{n}}^{n-1}, 0, \dots, 0, \frac{1}{n}\right)$$

each agent reports an outcome we deem his best one, then we randomize between n different outcomes containing all the reported best outcomes of the agents.

Both guarantees are maximal: generalizing the proof in the example of Section 1 for $\mathrm{vt}(3,6)$ and $\mathrm{rd}(3,6)$ is straightforward.[7] Moreover, Proposition 4 in the next section shows that, for any n, p such that $n < p$, all lotteries in the intervals $[\mathrm{uni}(p), \mathrm{vt}(n,p)]$ and $[\mathrm{uni}(p), \mathrm{rd}(n,p)]$ are maximal as well. The remarkable fact is that if $3 \le n < p \le 2n$ these two intervals capture most maximal guarantees.

Theorem 1. *(i) For any n, p such that $3 \le n < p$, we have*

$$\big[\mathrm{uni}(p), \mathrm{vt}(n,p)\big] \cup \big[\mathrm{uni}(p), \mathrm{rd}(n,p)\big] \subset \mathcal{M}(n,p) \tag{7}$$

(ii) This is an equality if $p \le 2n-2$ and if $p = 2n$ except when $(n,p) = (4,8)$ or $(5,10)$.

The hard proof of part (ii) requires the technical tools developed in the next section, and is done in the Appendix A.3.

Our next result shows why additional maximal guarantees appear in the cases excluded by statement (ii) above, and describes the full set $\mathcal{M}(n,p)$ in two such cases. As shown in the proof, the mechanisms used in these cases require different ideas beyond vetoes and random dictators. Moreover, the case $n = 4$, $p = 7$ exhibits extremal maximal guarantees which (unlike all previous examples) are *not* uniform on their support.

Proposition 3. *(i) If $p = 2n-1$ and if $(n,p) = (4,8)$ or $(5,10)$, the inclusion (7) is strict.*

(ii) For $n = 3$, $p = 5$, there are two pairs of maximal guarantees on the boundary of $\Delta(5)$: $\mathrm{vt}(3,5)$, $\mathrm{rd}(3,5)$ and the pair

$$\lambda = \left(\frac{1}{2}, 0, 0, \frac{1}{2}, 0\right); \qquad \lambda^\star = \left(\frac{1}{3}, 0, \frac{1}{3}, \frac{1}{3}, 0\right)$$

The set $\mathcal{M}(3,5)$ is the union of the four intervals joining $\mathrm{uni}(5)$ to these guarantees.

(iii) For $n = 4$, $p = 7$, there are three pairs of maximal guarantees on the boundary of $\Delta(7)$: $\mathrm{vt}(4,7)$, $\mathrm{rd}(4,7)$ and the two pairs

$$\lambda = \left(\frac{1}{2}, 0, 0, 0, \frac{1}{2}, 0, 0\right); \qquad \lambda^\star = \left(\frac{1}{5}, \frac{1}{5}, 0, \frac{1}{5}, \frac{1}{5}, \frac{1}{5}, 0\right)$$

$$\mu = \left(\frac{1}{3}, \frac{1}{9}, \frac{2}{9}, 0, 0, \frac{1}{3}, 0\right); \qquad \mu^\star = \left(\frac{1}{4}, 0, \frac{1}{4}, \frac{1}{4}, \frac{1}{12}, \frac{1}{6}, 0\right)$$

The set $\mathcal{M}(4,7)$ is the union of the six intervals joining $\mathrm{uni}(7)$ to these guarantees.

The duality operator discussed in the next subsection pairs $\mathrm{vt}(n,p)$ with $\mathrm{rd}(n,p)$ as well as λ with $\lambda^\star$ and μ with $\mu^\star$.

[7] For example, for $\mathrm{rd}(n,p)$ pick a profile with the preference $a_1 \prec \cdots \prec a_{n-1} \prec B \prec a_n$ and the $n-1$ others obtained by a cyclical permutation of the a_i-s leaving the block B fixed.

Proof. *Statement* (i). Assume $p = 2n - 1$. At any profile, we can choose a set of $n - 1$ outcomes meeting (containing at least one of) the top two outcomes of each agent. A uniform lottery over these outcomes guarantees to every agent a probability of at least $\frac{1}{n-1}$ for his top two outcomes. Hence, there must be a maximal guarantee that does that, but neither $uni(p)$, $\mathrm{vt}(n, p)$ nor $\mathrm{rd}(n, p)$ does this, hence neither does a convex combination of these.

For $(4, 8)$, one checks easily that we can always choose a triple of outcomes meeting the top three outcomes of each agent in *at most* one element. A uniform lottery over the complement of that triple guarantees to every agent at least $\frac{2}{5}$ for his top three outcomes, and the argument is completed as above. For $(5, 10)$, a simple case check shows that we can choose a triple of outcomes meeting the top three outcomes of each agent. A uniform lottery over them guarantees to every agent at least $\frac{1}{3}$ for his top three outcomes, and the argument is completed as above.

Statements (ii) and (iii). The mechanisms implementing λ and $\lambda^\bigstar$ in each case follow the same logic as above. For λ, we can always pick two outcomes x, y meeting the top two (when $(n, p) = (3, 5)$) or three (when $(n, p) = (4, 7)$) of any agent, then we draw x and y each with probability $\frac{1}{2}$. For $\lambda^\bigstar$, we can always pick two outcomes x, y such that the worst two (when $(n, p) = (3, 5)$) or three (when $(n, p) = (4, 7)$) of any agent contain at least one of them, then we randomize uniformly over the remaining outcomes.

For the rest of the proof, which becomes quite tedious, we give only a brief outline here (details are available upon request from the authors). The mechanisms implementing μ and $\mu^\bigstar$ when $(n, p) = (4, 7)$ require non-uniform lotteries tailored to the configuration of the top two and the bottom two outcomes of each agent. The maximality of λ and $\lambda^\bigstar$ in each case and of μ and $\mu^\bigstar$ when $(n, p) = (4, 7)$ is shown using suitable preference profiles, with a similar argument to that given in Section 1 for rd and vt. These facts, together with Proposition 4 in the next section, imply that the union of the intervals joining $\mathrm{uni}(p)$ to $\mathrm{vt}(n, p)$, $\mathrm{rd}(n, p)$, λ and $\lambda^\bigstar$ (and also to μ and $\mu^\bigstar$ when $(n, p) = (4, 7)$) is contained in $\mathcal{M}(n, p)$ in each case.

To prove that this containment is actually an equality, one considers an arbitrary guarantee $\nu \in \mathcal{G}(n, p)$ and shows that it is dominated by a guarantee in one of those intervals. This part breaks into cases as follows. If $[\nu]_1^k \geq \frac{k}{p}$ for $k = 1, \ldots, p - 1$, then ν is dominated by $\mathrm{uni}(p)$. If $\nu_1 < \frac{1}{p}$ or $\nu_p > \frac{1}{p}$, then the arguments in the corresponding cases in step 3 of the proof of Theorem 1 show that ν is dominated by a guarantee in the interval $[\mathrm{uni}(p), \mathrm{vt}(n, p)]$ or $[\mathrm{uni}(p), \mathrm{rd}(n, p)]$, respectively. When $(n, p) = (3, 5)$, this leaves the cases $\nu_1 + \nu_2 < \frac{2}{5}$ or $\nu_4 + \nu_5 > \frac{2}{5}$, where the implementation constraints for suitable profiles show that ν is dominated by a guarantee in $[\mathrm{uni}(5), \lambda^\bigstar]$ or $[\mathrm{uni}(5), \lambda]$, respectively. When $(n, p) = (4, 7)$ four cases remain: $\nu_1 + \nu_2 < \frac{2}{7}$, $\nu_1 + \nu_2 + \nu_3 < \frac{3}{7}$, $\nu_5 + \nu_6 + \nu_7 > \frac{3}{7}$ or $\nu_6 + \nu_7 > \frac{2}{7}$. They can be handled similarly, showing that ν is dominated by a guarantee in $[\mathrm{uni}(7), \mu^\bigstar]$, $[\mathrm{uni}(7), \lambda^\bigstar]$, $[\mathrm{uni}(7), \lambda]$ or $[\mathrm{uni}(7), \mu]$, respectively. □

Note that Theorem 1 and Proposition 3 together give a full description of maximal guarantees whenever $3 \leq n < p \leq n + 3$.

5. Duality and concatenation

The technically more involved results of this section are key to the proof of Theorem 1 and the general results in Section 6.

5.1 *A duality operation preserving* $\mathcal{M}(n, p)$

Although the results in this subsection apply for all n, p, they are only useful if $3 \le n < p$: we maintain this assumption from now on.

The duality operator is easy to explain if it applies to a lottery λ on the boundary $\partial\Delta(p)$ of the simplex $\Delta(p)$ (at least one coordinate of λ is zero): it goes from λ to its symmetric lottery $\widetilde{\lambda}$, then from $\widetilde{\lambda}$ to the end point in $\Delta(p)$ of the half-line starting at $\widetilde{\lambda}$ toward $\mathrm{uni}(p)$. For instance, it maps $\mathrm{vt}(3, 7) = (0, \frac{1}{4}, \frac{1}{4}, \frac{1}{4}, \frac{1}{4}, 0, 0)$ first to $\mu = (0, 0, \frac{1}{4}, \frac{1}{4}, \frac{1}{4}, \frac{1}{4}, 0)$ then to the intersection of the half-line $\{\mu + \alpha(\mathrm{uni}(p) - \mu) | \alpha \ge 0\}$ with the boundary $\partial\Delta(p)$:

$$\mathrm{vt}(3,7)^{\bigstar} = \mu + \frac{7}{3}\big(\mathrm{uni}(7) - \mu\big) = \left(\frac{1}{3}, \frac{1}{3}, 0, 0, 0, 0, \frac{1}{3}\right) = \mathrm{rd}(3,7)$$

Note that if λ, like $\mathrm{vt}(3, 7)$, is uniform on its nonfull support, the same is true of its dual $\lambda^{\bigstar}$. Moreover, the support of $\lambda^{\bigstar}$ is obtained from the support of λ by reversing the order on $[p]$, then taking the complement. This explains the intuition behind this duality: If a uniform lottery on a proper subset S of A implements the guarantee λ at some preference profile, then the uniform lottery on the complement $A \diagdown S$ implements $\lambda^{\bigstar}$ at the reverse preference profile.

Given a lottery λ in $\Delta(p)$ different from $\mathrm{uni}(p)$, the *radius to* λ is the interval of the half-line from $\mathrm{uni}(p)$ toward λ contained in $\Delta(p)$ (it ends on the boundary $\partial\Delta(p)$), containing all lotteries of the form $\mathrm{uni}(p) + \delta(\lambda - \mathrm{uni}(p))$ for some $\delta \ge 0$. The *antiradius from* $\widetilde{\lambda}$ is the interval in $\Delta(p)$ of the half-line from $\mathrm{uni}(p)$ away from $\widetilde{\lambda}$, that is, the set of all lotteries of the form $\mathrm{uni}(p) + \delta(\mathrm{uni}(p) - \widetilde{\lambda})$ for some $\delta \ge 0$.

If λ is a boundary lottery its dual $\lambda^{\bigstar}$ is the end point of the antiradius from $\widetilde{\lambda}$. Figure 1 illustrates this construction for $p = 3$.

Write the largest coordinate of a lottery as $\lambda_+ = \max_{1 \le k \le p} \lambda_k$, and note that $\lambda_+ > \frac{1}{p}$ because $\lambda \in \partial\Delta(p)$; recalling $\widetilde{\lambda}_k = \lambda_{p+1-k}$, an easy computation gives

$$\begin{aligned} &\lambda^{\bigstar} = (1+\alpha)\,\mathrm{uni}(p) - \alpha\widetilde{\lambda} \\ \Longleftrightarrow \quad &\lambda^{\bigstar}_k = \alpha(\lambda_+ - \widetilde{\lambda}_k) \quad \text{for } 1 \le k \le p, \text{ where } \alpha = \frac{1}{p \cdot \lambda_+ - 1} \end{aligned} \tag{8}$$

Keeping in mind that $\min_{1 \le k \le p} \lambda_k = 0$ it is easy to check the identity $(\lambda^{\bigstar})^{\bigstar} = \lambda$.

For nonboundary lotteries, we extend this definition linearly on the radius to λ

$$\big\{\mu \in \partial\Delta(p) \text{ and } \lambda = \delta\,\mathrm{uni}(p) + (1-\delta)\mu\big\} \quad \Longrightarrow \quad \lambda^{\bigstar} = \delta\,\mathrm{uni}(p) + (1-\delta)\mu^{\bigstar}, \tag{9}$$

so that $\lambda \to \lambda^{\bigstar}$ is a proper duality in $\Delta(p)$. In particular, $\mathrm{vt}(n, p)$ and $\mathrm{rd}(n, p)$ are dual of each other, while $\mathrm{uni}(p)$ is self-dual.

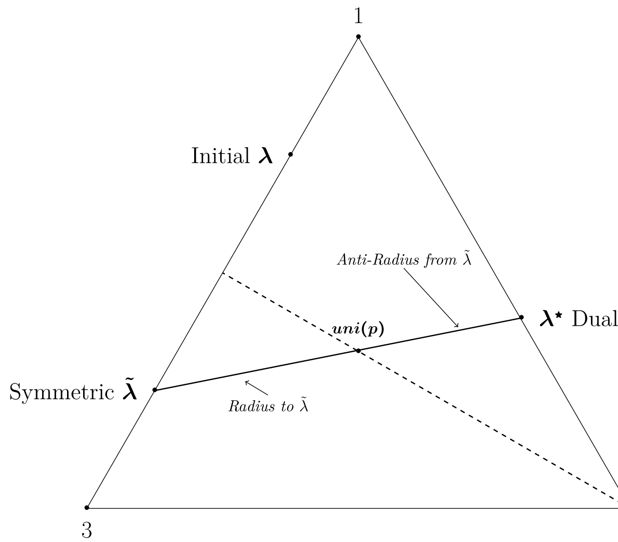

FIGURE 1. The duality operation.

We check that among maximal guarantees, $\mathrm{uni}(p)$ is the only self-dual one. By the discussion in Section 3, a feasible guarantee λ is dominated by $\widetilde{\lambda}$. Fix now $\lambda \in \mathcal{G}(n, p) \cap \partial\Delta(p)$ and self-dual: by (8), the interval $[\lambda, \widetilde{\lambda}]$ contains $\mathrm{uni}(p)$, implying $\widetilde{\lambda} \vdash \mathrm{uni}(p) \vdash \lambda$, a contradiction because $\mathrm{uni}(p)$ is maximal. Definition (9) concludes the argument for λ self-dual but not in $\partial\Delta(p)$.

PROPOSITION 4. *(i) If $\lambda \neq \mathrm{uni}(p)$ is a maximal guarantee, the radius to λ and the anti-radius from $\widetilde{\lambda}$ (the symmetric of λ with respect to the middle rank) are contained in $\mathcal{M}(n, p)$.*

(ii) The duality operation $\lambda \to \lambda^{\bigstar}$ in $\Delta(p)$ preserves maximal lotteries:

$$\left[\mathcal{M}(n, p)\right]^{\bigstar} = \mathcal{M}(n, p)$$

For the proof, we need a technical result characterizing $\mathcal{M}(n, p)$ in $\mathcal{G}(n, p)$ by its position with respect to the polar cone of $\mathcal{G}(n, p)$. Notation: we write $G^{\ominus}$ for the polar cone of $G \subset \mathbb{R}^p$: $G^{\ominus} = \{z \in \mathbb{R}^p | \forall y \in G : z \cdot y \leq 0\}$.

LEMMA 2. *The guarantee $\lambda \in \mathcal{G}(n, p)$ is maximal if and only if there exists a vector $z \in \mathcal{G}(n, p)^{\ominus}$ such that $\sum_{k=1}^{p} z_k = 0$, $z_1 < z_2 < \cdots < z_p$ and $\lambda \cdot z = 0$.*

PROOF OF "IF.". Fix λ in $\mathcal{G}(n, p)$ and z in $\mathcal{G}(n, p)^{\ominus}$ as in the statement, and suppose λ is dominated by μ. As the coordinates of z increase strictly, $\mu \vdash \lambda$ and $\mu \neq \lambda$ imply $\lambda \cdot z < \mu \cdot z$. Now feasibility of μ and $z \in \mathcal{G}(n, p)^{\ominus}$ give $\mu \cdot z \leq 0$. This contradicts the assumption $\lambda \cdot z = 0$. □

Note that the condition $\sum_{k=1}^{p} z_k = 0$ was not used therefore Lemma 2 remains valid without this condition. But the condition makes the "only if" part stronger. The long proof of this direction is given in Appendix A.1.

Proof of Proposition 4. Observe first that the statements in Proposition 4 hold if we replace $\mathcal{M}(n, p)$ by $\mathcal{G}(n, p)$. This follows easily from the characteristic property (1), the identity (3), and the definition (8).

Statement (i). We fix $\lambda \in \mathcal{M}(n, p)$ and $z \in \mathcal{G}(n, p)^{\ominus}$ as in Lemma 2. Consider first a lottery $\mu = \text{uni}(p) + \delta(\lambda - \text{uni}(p))$ in the radius to λ: it is in $\mathcal{G}(n, p)$ as well by the observation above. For maximality, we use $\sum_{k=1}^{p} z_k = 0$ and $\lambda \cdot z = 0$ to compute $\mu \cdot z = (1-\delta)\,\text{uni}(p) \cdot z + \delta\lambda \cdot z = 0$ and conclude $\mu \in \mathcal{M}(n, p)$ by Lemma 2 again.

Still fixing $\lambda \in \mathcal{M}(n, p)$ and z, we pick a lottery $\mu = \text{uni}(p) + \delta(\text{uni}(p) - \widetilde{\lambda})$ in the antiradius from $\widetilde{\lambda}$; we know that μ is in $\mathcal{G}(n, p)$. Now for any $\xi \in \mathcal{G}(n, p)$ the observation implies that $(1+\delta)\,\text{uni}(p) - \delta\widetilde{\xi}$ is in $\mathcal{G}(n, p)$, in particular,

$$0 \geq \big((1+\delta)\,\text{uni}(p) - \delta\widetilde{\xi}\big) \cdot z = -\delta\widetilde{\xi} \cdot z$$

where the equality uses $\sum_{k=1}^{p} z_k = 0$. Writing

$$w = (-z_p, -z_{p-1}, \dots, -z_1)$$

and using the identity (3), we conclude that $\xi \cdot w \leq 0$. Thus, w is in $\mathcal{G}(n, p)^{\ominus}$, too, and it satisfies the requirements in Lemma 2 with respect to μ: $\mu \cdot w = -\delta\widetilde{\lambda} \cdot w = \delta\lambda \cdot z = 0$, which proves the maximality of μ.

Statement (ii) follows from statement (i) and the definition of the duality operation. □

5.2 *The operators VT and RD*

We construct a rich family of maximal guarantees by successive compositions of two operators VT and RD mapping a guarantee λ in $\mathcal{G}(n, p)$ to one in $\mathcal{G}(n, p+n)$.

Fixing $\lambda \in \mathcal{G}(n, p)$, we implement $VT \otimes \lambda$ as follows: ask agents to report their worst outcome, eliminate n outcomes containing all the reported ones, then implement λ over the remaining p outcomes. The latter are ranked weakly higher than $2, \dots, p+1$ for each agent, so we see that $VT \otimes \lambda$ is a bona fide guarantee; and that $VT \otimes \lambda$ obtains by inserting λ between one zero in rank 1 and $n-1$ zeros after rank $p+1$.

The implementation of $RD \otimes \lambda$ is similar, but only if λ is a boundary lottery in $\mathcal{G}(n, p)$. Agents report their best outcome, then we pick n outcomes containing all reports; with probability $\frac{n\lambda_+}{n\lambda_+ + 1}$, we choose one of those uniformly, and with probability $\frac{1}{n\lambda_+ + 1}$ we implement λ among the remaining p outcomes.

Definition 5. Fix a lottery $\lambda \in \Delta(p)$.

We set $VT \otimes \lambda = (0, \lambda, 0, \dots, 0) \in \Delta(p+n)$ with $n-1$ zeros after and one before λ.

The lottery $RD \otimes \lambda \in \Delta(p+n)$ is given by

$$RD \otimes \lambda = \big[VT \otimes \lambda^{\bigstar}\big]^{\bigstar} \tag{10}$$

If $\lambda \in \partial\Delta(p)$, we obtain $RD \otimes \lambda$ by filling uniformly $n-1$ ranks before λ and one after as follows:

$$RD \otimes \lambda = \Bigg(\overbrace{\frac{\lambda_+}{n\lambda_+ + 1}, \dots, \frac{\lambda_+}{n\lambda_+ + 1}}^{n-1}, \frac{1}{n\lambda_+ + 1} \cdot \lambda, \frac{\lambda_+}{n\lambda_+ + 1}\Bigg) \tag{11}$$

If $\lambda \in \partial\Delta(p)$, we must check that the two definitions (11) and (10) coincide. Write μ for the boundary lottery on the right-hand side of equation (11): applying (8) and $\mu_+ = \frac{\lambda_+}{n\lambda_+ + 1}$ we get

$$\mu^\star = \frac{1}{(p+n)\dfrac{\lambda_+}{n\lambda_+ + 1} - 1}\left(\frac{\lambda_+}{n\lambda_+ + 1}\mathbf{1} - \widetilde{\mu}\right)$$

$$= \frac{1}{p\lambda_+ - 1}\big(\lambda_+\mathbf{1} - (\lambda_+, \widetilde{\lambda}, \overbrace{\lambda_+, \ldots, \lambda_+}^{n-1})\big) = VT \otimes \lambda^\star$$

as desired.

The definition implies $VT \otimes \text{uni}(p) = \text{vt}(n, p+n)$ and $RD \otimes \text{uni}(p) = \text{rd}(n, p+n)$; we give many more examples in the next subsection.

Lemma 3. *The composition of guarantees by VT and RD respects their feasibility and maximality. For any $\lambda \in \Delta(p)$,*

$$\lambda \in \mathcal{M}(n, p) \quad \Longrightarrow \quad VT \otimes \lambda, RD \otimes \lambda \in \mathcal{M}(n, p+n)$$

and the same statement holds by replacing $\mathcal{M}(n, p)$ by $\mathcal{G}(n, p)$ and $\mathcal{M}(n, p+n)$ by $\mathcal{G}(n, p+n)$.

For the proof, we need a second characterization of maximal guarantees; the proof, much easier than that of Lemma 2, is also in the Appendix A.2.

Lemma 4. *The guarantee $\lambda \in \mathcal{G}(n, p)$ is maximal if and only if for all $k \in [p-1]$ there exists a preference profile π such that, for any lottery ℓ implementing λ at π (Definition 1) we have*

$$\max_{i \in [n]}\big[\ell^{*i}\big]_1^k = [\lambda]_1^k \tag{12}$$

Proof. Proof of Lemma 3 The implementation argument above shows that $VT \otimes \lambda$ is in $\mathcal{G}(n, p)$ if λ is. If now $\lambda \in \mathcal{M}(n, p)$, we fix an index $k \in [p-1]$ and an (n, p)-profile π ensuring property (12) as in the premises of Lemma 4. We construct the following $(n, p+n)$ profile θ:

$$\begin{array}{llclcl} \prec_1 & a_1 & \overbrace{\pi}^{p} & a_2 & \cdots & a_n \\ \cdots & \cdots & \pi & \cdots & & \cdots \\ \prec_n & a_n & \pi & a_1 & \cdots & a_{n-1} \end{array} \tag{13}$$

where the initial profile π on p outcomes occupies the ranks 2 to $p+1$, while the preferences over the n other outcomes are cyclical. If a lottery ℓ implements $VT \otimes \lambda$ at θ, it can put no weight on any a_i outcome because $(VT \otimes \lambda)_1 = 0$, therefore, the restriction of ℓ to the outcomes of π implements λ at π, so property (12) holds for ranks 2 to $p+1$ as well as for the first one and the last $n-1$ ones.

That $RD \otimes \lambda$ is in $\mathcal{G}(n, p)$ respectively in $\mathcal{M}(n, p)$, if λ is follows from the same property for VT, using the duality relation (10) and the fact that duality respects feasibility and maximality (Proposition 4). □

5.3 *Canonical guarantees*

Starting from the uniform guarantee, its composition by an arbitrary sequence of the operators VT and RD generates a large family of maximal guarantees: we call them *canonical*, because the veto and random dictator ideas stand out in the characterization results for $n = 2$ and $p \leq 2n$ (Section 4).

Definition 6 (Canonical guarantees). Fix $n, p, 3 \leq n < p$, such that $d = \lfloor \frac{p-1}{n} \rfloor$ and $p = dn + q$ for some $q = 1, \ldots, n$. Each sequence $\Gamma = (\Gamma^t)_{t=1}^h$ in $\{VT, RD\}$ of length h, $h \leq d$, defines a canonical guarantee in $\mathcal{M}(n, p)$ by composition of these operators starting from $\text{uni}((d-h)n + q)$, that is,

$$\boldsymbol{\Gamma}^1 \otimes \boldsymbol{\Gamma}^2 \otimes \cdots \otimes \boldsymbol{\Gamma}^h \otimes \text{uni}\big((d-h)n + q\big) = \Gamma^1 \otimes (\Gamma^2 \otimes \big(\cdots \otimes \big(\Gamma^h \otimes \text{uni}\big((d-h)n + q\big)\big) \cdots\big)$$

where $\Gamma^t \otimes \Gamma^{t+1} \otimes \cdots \otimes \Gamma^h \otimes \text{uni}((d-h)n + q)$ is in $\mathcal{M}(n, (d-t+1)n + q)$ for all $t \in [h]$. We write their set as $\mathcal{C}(n, p)$, of cardinality $2^{d+1} - 2$.

By Lemma 3, all canonical guarantees are maximal, because the composition by each Γ^t adds n outcomes to the previous ones, they are in $\mathcal{M}(n, p)$. By duality (10), canonical guarantees come in dual pairs: exchanging VT and RD in each term of the sequence Γ produces the dual guarantee.

An important observation is that each $\lambda \in \mathcal{C}(n, p)$ is uniform on its support therefore determined by this nonfull support. This implies that it is a vertex of $\mathcal{G}(n, p)$ (the proof mimics that of statement (iii) in Proposition 1); hence, also a vertex of $\mathcal{M}(n, p)$.

If $d = 1$ $(p \leq 2n)$ $\text{vt}(n, p)$ and $\text{rd}(n, p)$ are the only canonical guarantees. We give some examples where $d \geq 2$, writing for brevity a canonical guarantee as $\Gamma^1 \otimes \Gamma^2 \otimes \cdots \otimes \Gamma^h$ without the initial uniform lottery.

Constant sequences: the composition of h veto steps, or of h random dictator steps, gives dual guarantees of a similar shape: their support is at the extreme ranks or in the center:

$$\overbrace{VT \otimes \cdots \otimes VT}^{h} = \Big(\overbrace{0, \ldots, 0}^{h}, \frac{1}{p - nh}, \ldots, \frac{1}{p - nh}, \overbrace{0, \ldots, 0}^{(n-1)h}\Big)$$

$$\overbrace{RD \otimes \cdots \otimes RD}^{h} = \Big(\overbrace{\frac{1}{nh}, \ldots, \frac{1}{nh}}^{(n-1)h}, 0, \ldots, 0, \overbrace{\frac{1}{nh}, \ldots, \frac{1}{nh}}^{h}\Big)$$

A simple mechanism for the former gives h veto tokens to each agent, then randomizes uniformly between the remaining outcomes, even if there are more than $p - nh$ of those (which will only improve the guaranteed welfare). To implement the latter, we elicit from each agent her h top outcomes, then randomize uniformly between any nh

outcomes containing all reported tops, adding arbitrary outcomes if the reported ones are fewer than nh. The last instruction is important: ignoring it could result in giving too much weight to someone's worst outcomes (as illustrated in the example of Section 1).

For $d = 2$, we have six canonical guarantees, four from the constant sequences and a dual pair from (VT, RD) and (RD, VT). For instance, in $\mathcal{C}(3, 7)$:

$$VT \otimes RD = \left(0, \frac{1}{3}, \frac{1}{3}, 0, \frac{1}{3}, 0, 0\right); \qquad RD \otimes VT = \left(\frac{1}{4}, \frac{1}{4}, 0, \frac{1}{4}, 0, 0, \frac{1}{4}\right)$$

The mechanism for $RD \otimes VT$ selects three outcomes containing the top ones of each agent; then with probability $3/4$ it picks one of those uniformly, and with probability $1/4$ plays $\mathrm{vt}(3, 4)$ among the remaining outcomes.

Our final example is in $\mathcal{C}(3, 11)$ where $d = 3$ and we have three pairs of nonconstant sequences of length three, for instance:

$$(RD, VT, VT) \to \lambda = \left(\frac{1}{5}, \frac{1}{5}, 0, 0, \frac{1}{5}, \frac{1}{5}, 0, 0, 0, 0, \frac{1}{5}\right)$$
$$(RD, VT, RD) \to \lambda = \left(\frac{1}{6}, \frac{1}{6}, 0, \frac{1}{6}, \frac{1}{6}, 0, 0, \frac{1}{6}, 0, 0, \frac{1}{6}\right)$$

6. General results and open problems

Theorem 1 in Section 4 tells much about the structure of $\mathcal{M}(n, p)$ in the case $d = 1$.

For general values of $d = \lfloor \frac{p-1}{n} \rfloor$, we know only a few general facts. Lemma 2 in Proposition 3 provides our best clue. For any $z \in \mathcal{G}(n, p)^{\ominus}$ such that $\mathcal{G}(n, p)$ intersects the hyperplane $H = \{y | z \cdot y = 0\}$, the intersection $H \cap \mathcal{G}(n, p)$ is a face of $\mathcal{G}(n, p)$, in particular a polytope. The lemma tells us that such a face defined by a vector z with increasing coordinates is a subset of $\mathcal{M}(n, p)$, and that all maximal guarantees obtain for some z. This proves the following.

Proposition 5. *For $3 \leq n < p$, the set $\mathcal{M}(n, p)$ is a finite union of faces of the polytope $\mathcal{G}(n, p)$, each having* $\mathrm{uni}(p)$ *as a vertex.*

Our second main result identifies a large subset of $\mathcal{M}(n, p)$ constructed from the canonical guarantees.

Theorem 2. *Fix n, p such that $3 \leq n < p$, $d = \lfloor \frac{p-1}{n} \rfloor$.*

For each sequence Γ of length d in $\{VT, RD\}$, the canonical guarantees from the d initial subsequences[8] *of Γ, plus the uniform guarantee, are the vertices of a simplex of dimension d contained in $\mathcal{M}(n, p)$.*

The proof is in the Appendix A.4.

Theorem 2 describes 2^d components of $\mathcal{M}(n, p)$ (faces of $\mathcal{G}(n, p)$). To see that this may not exhaust $\mathcal{M}(n, p)$, we take the simplest example not covered in Theorem 1:

[8]That is, the guarantees Γ^1, $\Gamma^1 \otimes \Gamma^2$, $\Gamma^1 \otimes \Gamma^2 \otimes \Gamma^3$, etc.

$n=3$, $p=7$, so $d=2$. Theorem 2 describes four triangles of maximal guarantees coming in dual pairs. The uniform lottery is always a vertex and the other two vertices are canonical guarantees:

sequence	vertex 1	vertex 2
VT, VT	$(0, \frac{1}{4}, \frac{1}{4}, \frac{1}{4}, \frac{1}{4}, 0, 0)$	$(0, 0, 1, 0, 0, 0, 0)$
RD, RD	$(\frac{1}{3}, \frac{1}{3}, 0, 0, 0, 0, \frac{1}{3})$	$(\frac{1}{6}, \frac{1}{6}, \frac{1}{6}, \frac{1}{6}, 0, \frac{1}{6}, \frac{1}{6})$
VT, RD	$(0, \frac{1}{4}, \frac{1}{4}, \frac{1}{4}, \frac{1}{4}, 0, 0)$	$(0, \frac{1}{3}, \frac{1}{3}, 0, \frac{1}{3}, 0, 0)$
RD, VT	$(\frac{1}{3}, \frac{1}{3}, 0, 0, 0, 0, \frac{1}{3})$	$(\frac{1}{4}, \frac{1}{4}, 0, \frac{1}{4}, 0, 0, \frac{1}{4})$

where the dual pairs are the top two and the bottom two rows. In addition to these four triangles, the maximal set $\mathcal{M}(3, 7)$ also contains two intervals, joining $\mathrm{uni}(7)$ to each of the following two dual noncanonical guarantees:

$$\lambda = \left(\frac{1}{3}, 0, 0, \frac{1}{3}, \frac{1}{3}, 0, 0\right); \qquad \lambda^{\star} = \left(\frac{1}{4}, \frac{1}{4}, 0, 0, \frac{1}{4}, \frac{1}{4}, 0\right).$$

In general, we keep in mind that many more guarantees than the ones described in Theorem 2 are maximal. Pick any noncanonical guarantee λ in $\mathcal{M}(n, p) \cap \partial\Delta(p)$, for instance, those described in Proposition 3 or in the previous paragraph: by Lemma 3, successive compositions of λ with VT and/or RD generate, for any $h \geq 1$, 2^h noncanonical maximal guarantees in $\mathcal{M}(n, p+hn) \cap \partial\Delta(p+hn)$.

Conjectures and open questions We conjecture that no convex combinations of canonical guarantees other than those described in Theorem 2 (i.e., corresponding to nested sequences in $\{VT, RD\}$) can produce a maximal lottery.

We conjecture that the maximal dimension of a simplicial component of $\mathcal{M}(n, p)$ is $d = \lfloor \frac{p-1}{n} \rfloor$.

We do not know how to evaluate the number of such components, which we know is at least 2^d.

Finally, we give an example in Appendix A.5 of a three person mechanism implementing several noncomparable guarantees: it offers a two item menu of guarantees among which the agents can choose.[9] We do not know how long a menu of guarantees can be, given n and p.

7. Concluding comments

Our results provide a formal vindication of the role of Voting by Veto (Veto) and Random Dictator (RDict) as ex ante protections of each individual decision-maker's welfare. But much depends on the comparison of n (the number of agents) and p (the number of outcomes).

Our message is especially clear for bargaining situations between only two agents, $n=2$: then the vertices of the set of maximal guarantees simply combine a round of Veto (possibly with multiple veto tokens) followed by one of RDict (Proposition 2).

[9]We thank an anonymous referee for suggesting that this may be the case.

For general problems, the situation is more nuanced. If we need to protect at least as many individuals as there are outcomes ($n \geq p$), the Uniform (Unif) guarantee drawing blindly an outcome is the only acceptable guarantee (Proposition 1). Veto and RDict play no role if the electorate is large.

By contrast, if there are at most twice as many outcomes than agents $n < p \leq 2n$, essentially all maximal guarantees obtained by a convex combination of Unif and Veto, or of Unif and RDict (Theorem 1).

Finally, if we have significantly more outcomes than agents ($p > 2n$ and $n \geq 3$) we produce many more (exponentially more in $d = \lfloor \frac{p-1}{n} \rfloor$) maximal guarantees by combining up to d elementary Veto or RDict blocks; some of their convex combinations are maximal as well (Theorem 2). But the mechanisms implementing them become increasingly opaque. A practical application of our *canonical* guarantees would involve at most one round of Veto with up to $d-1$ tokens per person followed by RDict on the nonvetoed outcomes, or the dual mechanism starting by RDict.

Appendix: Four proofs and one mechanism

A.1 *Proof of Lemma 2*

We must prove the *only if* statement: for any $\lambda \in \mathcal{M}(n, p)$, we can find a vector z as in Lemma 2. Consider the following cone W in $\mathbb{R}^p$:

$$W = \left\{ z = \sum_{i=1}^{n} u_i^* \;\middle|\; \text{for some } U = (u_i)_{i=1}^n \text{ such that } \sum_{i=1}^{n} u_i = 0 \right\} \tag{14}$$

By its characteristic property (1), $\mathcal{G}(n, p)$ is the intersection of $W^{\ominus}$ with $\Delta(p)$ therefore $\mathcal{G}(n, p)^{\ominus}$ is the Minkowski sum of $\overleftrightarrow{W}$ and $\mathbb{R}^p_-$, where $\overleftrightarrow{W}$ is the convex hull of W. Moreover, the identity $\sum_{(i,k)\in[n]\times[p]} u_{ik}^* = \sum_{(i,a)\in[n]\times A} u_{ia}$ implies $\sum_{k=1}^p z_k = 0$ in W therefore $\overleftrightarrow{W} = \mathcal{G}(n, p)^{\ominus} \cap \{z | \sum_{k=1}^p z_k = 0\}$.

We fix now a maximal guarantee λ and define the subcone Z of $\overleftrightarrow{W}$:

$$Z = \left\{ z \in \mathcal{G}(n, p)^{\ominus} \;\middle|\; \sum_{k=1}^{p} z_k = 0 \text{ and } \lambda \cdot z = 0 \right\}$$

This cone is convex, and every element of Z satisfies $z_1 \leq z_2 \leq \cdots \leq z_p$, because these inequalities hold in W. To prove that Z contains some z such that $z_1 < z_2 < \cdots < z_p$, we choose in Z one $\widehat{z}$ in which the number of equalities between consecutive coordinates of $\widehat{z}$ is as small as possible. If there is no equality, we are done. Otherwise, assume that the first equality is $\widehat{z}_k = \widehat{z}_{k+1}$. We will show the existence of some $z \in W$ such that $z_k < z_{k+1}$ and $\lambda \cdot z = 0$; this leads to a contradiction because $\widehat{z} + z \in Z$ has fewer equalities than $\widehat{z}$. Consider two cases.

Case 1 $\lambda_k > 0$

We proceed by contradiction and assume that if $z \in W$ and $z_k < z_{k+1}$, then $\lambda \cdot z < 0$.

Call Π the set of profiles $U = (u_i)_{i=1}^n$ such that

$$\sum_{i=1}^n u_i = 0 \quad \text{and} \quad u_{1k}^* = 0, \qquad u_{1,k+1}^* = 1 \tag{15}$$

The corresponding vector $z = \sum_{i=1}^n u_i^*$ is in W therefore $\sum_{i=1}^n \lambda \cdot u_i^* < 0$ for all $U \in \Pi$. We show next, again by contradiction, that the supremum of $\sum_{i=1}^n \lambda \cdot u_i^*$ over Π cannot be zero.

If it is, there is a sequence U^s in Π such that the sequence $\sum_{i=1}^n \lambda \cdot u_i^{s*}$ converges to zero. By taking subsequences, we can make sure that for each i, the way each u_i^s orders the outcomes in A does not depend on s (but depends on (i)). Then for each i there is a lottery λ^i on A, its coordinates a permutation of those of λ, such that $\lambda \cdot u_i^{s*} = \lambda^i \cdot u_i^s$ for all s.

Consider the polyhedron Q of $n \times p$ matrices $X = [x_i^a]_{i\in[n],a\in A}$ defined by three sets of conditions:

In each row i, the entries are ordered the same way as in every u_i^s
$\sum_{i=1}^n x_i^a = 0$ in each column a,
$x_{1a} = 0$, $x_{1b} = 1$ where a and b are the outcomes ranked k and $k+1$ by each u_1^s

Note that Q is nonempty because it contains each matrix U^s.

By construction, each X in Q defines a profile in Π and $\lambda \cdot x_i^* = \lambda^i \cdot x_i$ for all i. Therefore, we have

$$\sum_{i=1}^n \lambda^i \cdot x_i < 0 \quad \text{for all } X \in Q$$

$$\lim_{s\to\infty} \sum_{i=1}^n \lambda^i \cdot u_i^s = 0 \quad \text{for the sequence } U^s \text{ in } Q$$

This is impossible: if the closed polyhedron Q is disjoint from the hyperplane H : $\sum_{i=1}^n \lambda^i \cdot x_i = 0$, it cannot contain points arbitrarily close to H.

Thus, there is some positive ε such that for any profile U in Π we have $\sum_{i=1}^n \lambda \cdot u_i^* < -\varepsilon$, and we can now conclude the proof in Case 1. These inequalities imply for any profile U:

$$\left\{\sum_{i=1}^n u_i = 0 \text{ and } u_{1k}^* < u_{1,k+1}^*\right\} \implies \sum_{i=1}^n \lambda \cdot u_i^* \le -\varepsilon\big(u_{1,k+1}^* - u_{1k}^*\big) \tag{16}$$

Indeed if $u_{1,k+1}^* - u_{1k}^* = 1$ the profile $(u_1 - u_{1k}^*\mathbf{1}, u_2 + u_{1k}^*\mathbf{1}, u_3, \dots, u_n)$ is in Π, and rescaling our profile by $\frac{1}{u_{1,k+1}^* - u_{1k}^*}$ implies the claim.

Note that in (16) we can replace coordinate 1 by any coordinate i. Therefore, we have

$$\sum_{i=1}^n u_i = 0 \implies \sum_{i=1}^n \lambda \cdot u_i^* \le -\frac{\varepsilon}{n} \sum_{i=1}^n \big(u_{i,k+1}^* - u_{ik}^*\big) \tag{17}$$

Because $\lambda_k > 0$, the lottery μ obtained from λ by shifting $\frac{\varepsilon}{n}$ or λ_k, whichever is less, from λ_k to λ_{k+1} dominates λ, and property (17) implies it is feasible.

Case 2 $\lambda_k = 0$

In this case, because $\widehat{z} \in \overleftrightarrow{W}$ it is a sum of m elements $z^j \in W$, $j \in [m]$, each z^j defined by n utilities $(\overline{u}_i^j)_{i=1}^n$ as in (14). Note that if $k > 1$ then $\overline{u}_{i_0,k-1}^{j_0 *} < \overline{u}_{i_0,k}^{j_0 *}$ for some i_0, j_0. Pick such i_0 and j_0 (or arbitrary ones if $k = 1$). Let $a \in A$ be such that $\overline{u}_{i_0,k}^{j_0 *} = \overline{u}_{i_0,a}^{j_0}$. For some small $\varepsilon > 0$, modify $(\overline{u}_i^{j_0})_{i=1}^n$ to $(u_i)_{i=1}^n$ by letting $u_{i_0,a} = \overline{u}_{i_0,a}^{j_0} - \varepsilon$, $u_{i_1,a} = \overline{u}_{i_1,a}^{j_0} + \varepsilon$ for some $i_1 \neq i_0$, and leaving all other utilities unchanged. Because $\lambda_k = 0$ and by our choice of i_0, j_0, for small enough ε we have $\sum_{i=1}^n \lambda \cdot u_i^* \geq \sum_{i=1}^n \lambda \cdot \overline{u}_i^{j_0 *} = \lambda \cdot z^{j_0} = 0$. As λ is feasible, this must be an equality and, therefore, $z = \sum_{i=1}^n u_i^* \in Z$ and satisfies $z_k < z_{k+1}$ by construction.

A.2 *Proof of Lemma 4*

Statement If. Pick two guarantees λ, μ in $\mathcal{G}(n, p)$, such that λ meets the property above while $\mu \vdash \lambda$. Pick $k \in [p-1]$ and a profile π as in the statement. Choose a lottery ℓ implementing μ at π and an agent i reaching the maximum in (12): we have $[\ell^{*i}]_1^k \leq [\mu]_1^k \leq [\lambda]_1^k$ and $[\ell^{*i}]_1^k = [\lambda]_1^k$. As k was arbitrary in $[p-1]$, we conclude $\mu = \lambda$ therefore λ is maximal.

Statement Only If. Suppose now that $\lambda \in \mathcal{G}(n, p)$ fails the property in the lemma: there is some k and some $\varepsilon > 0$ such that at any profile π there is some lottery ℓ implementing λ at π and such that

$$\max_{i\in[n]} \left[\ell^{*i}\right]_1^k = [\lambda]_1^k - \varepsilon \tag{18}$$

We show that λ is not maximal. Suppose first $\lambda_k > 0$ and construct λ' dominating λ by shifting a weight δ, smaller than ε and λ_k, from λ_k to λ_{k+1} (and no other change). The lottery λ' is still in $\mathcal{G}(n, p)$: at a profile π the lottery ℓ implementing λ and meeting (18) implements λ' as well. Suppose next $\lambda_k = 0$. Then we have for all i,

$$\left[\ell^{*i}\right]_1^{k-1} \leq \left[\ell^{*i}\right]_1^k \leq [\lambda]_1^k - \varepsilon = [\lambda]_1^{k-1} - \varepsilon$$

so that if λ_{k-1} is positive we can apply the previous argument. If $\lambda_{k-1} = 0$ again, we repeat this observation until we find some positive λ_t, $t \leq k-2$, whose existence is assured by (18).

A.3 *Proof of Theorem 1*

We prove part (ii): if $3 \leq n < p$ and $p \leq 2n-2$, or $p = 2n$ but $n \neq 4, 5$, then any maximal guarantee lies in one of the intervals $[\mathrm{uni}(p), \mathrm{vt}(n, p)]$, $[\mathrm{uni}(p), \mathrm{rd}(n, p)]$.

Step 1. Recall the following notion from the Shapley–Bondareva theorem. A family $S_1, \dots, S_m$ of subsets of $[p]$ is *balanced* if there exist positive weights $\gamma_1, \dots, \gamma_m$ such that $\sum_{i:j\in S_i} \gamma_i = 1$ for every $j \in [p]$.

Lemma 5. *Assume that $p \leq 2n-2$, or $p = 2n$ but $n \neq 4, 5$ and let $2 \leq k \leq \lfloor \frac{p}{2} \rfloor$. Then there exists a balanced family $S_1, \dots, S_m$ of subsets of $[p]$ of size k each, such that $m \leq n$.*

Assume first that $p \le 2n-2$ and $2 \le k \le \lfloor \frac{p}{2} \rfloor$. If k divides p, the lemma is obvious (take a partition of $[p]$). Suppose $p = tk + r$ where $1 \le r \le k-1$. Let $S_i = \{(i-1)k + 1, \dots, ik\}$ for $i = 1, \dots, t$. Also, let $S_i = C_i \cup \{tk+1, \dots, p\}$ for $i = t+1, \dots, t+k$, where the sets C_i are of size $k-r$ and form the k cyclic intervals in a cyclic arrangement of S_t. Let $\gamma_1 = \cdots = \gamma_{t-1} = 1$, $\gamma_t = \frac{r}{k}$, $\gamma_{t+1} = \cdots = \gamma_{t+k} = \frac{1}{k}$. These weights make $S_1, \dots, S_{t+k}$ a balanced family, and it remains to check that $t + k \le n$.

We have $t + k < \frac{p}{k} + k \le \max\{\frac{p}{x} + x : x \in [2, \frac{p}{2}]\} = \frac{p+4}{2}$. If $p \le 2n-2$, this gives $t + k < n+1$ as desired.

Assume next $p = 2n$ and $2 \le k \le n$. When k divides p a partition works, so we may assume that $3 \le k \le n-1$, and thus $n \ge 4$. We further exclude the exceptional cases $n = 4, 5$ and assume $n \ge 6$. If $k \le n-2$, we still have $\frac{p}{k} + k \le n+1$ as in the original proof. Thus, we may assume that $k = n-1$. We provide two variants of the construction of the balanced family, depending on parity.

Case 1. $k = n-1$ is even. Partition $[p] = [2k+2]$ into $S, P_1, \dots, P_{\frac{k}{2}+1}$ where $|S| = k$ and the other sets are pairs. Take S with weight 1, and for each P_i, the union of all P_j, $j \ne i$, with weight $\frac{2}{k}$. This gives a balanced family of size $\frac{k}{2} + 2 < n$.

Case 2. $k = n-1$ is odd. Partition $[p] = [2k+2]$ into $S, T, P_1, \dots, P_{\frac{k-1}{2}}$ where $|S| = k$, $|T| = 3$ and the other sets are pairs. Take S with weight 1, for each P_i take the union of T and all P_j, $j \ne i$, with weight $\frac{2}{k}$, and for each element a of T take the union of $\{a\}$ and all the P_i with weight $\frac{1}{k}$. This gives a balanced family of size $\frac{k-1}{2} + 4 \le n$.

Step 2. Assume (n, p) are as in Lemma 5 and let $2 \le k \le p-2$. Then for any $\lambda \in \mathcal{G}(n, p)$ we have $[\lambda]_1^k \ge \frac{k}{p}$.

By duality, it suffices to show this for $2 \le k \le \lfloor \frac{p}{2} \rfloor$. Let $S_1, \dots, S_m$ with weights $\gamma_1, \dots, \gamma_m$ be a balanced family as in the lemma. Consider a profile of strict preferences in which $\{a_j : j \in S_i\}$ is the set of the k worst outcomes of agent i, $i = 1, \dots, m$. Let ℓ be a lottery that implements λ at this profile. Then $1 = \sum_{a \in A} \ell_a = \sum_{i=1}^m \gamma_i \sum_{j \in S_i} \ell_{a_j} \le \sum_{i=1}^m \gamma_i [\lambda]_1^k = \frac{p}{k} [\lambda]_1^k$, implying the desired inequality.

Step 3. For (n, p) as in Lemma 5, we fix $\lambda \in \mathcal{G}(n, p)$ and show that it is dominated by a guarantee in $[\mathrm{uni}(p), \mathrm{vt}(n, p)] \cup [\mathrm{uni}(p), \mathrm{rd}(n, p)]$. This will establish part (ii) of Theorem 1. We distinguish three cases.

Case 1. $\lambda_p \ge \frac{1}{p}$. Set $\lambda_p = x$ and keep in mind that feasibility implies $x \le \frac{1}{n}$. We will show that λ is dominated (weakly) by the guarantee $\mu \in [\mathrm{uni}(p), \mathrm{rd}(n, p)]$ such that $\mu_p = x$, that is, $\mu_k = x$ for $1 \le k \le n-1$ and $\mu_k = y$ for $n \le k \le p-1$, with $nx + (p-n)y = 1$.

Set $p = n + q$ and partition A as $\{a_1, \dots, a_n\} \cup \{b_1, \dots, b_q\}$ then consider a profile of preferences where for everyone:

the a-s occupy the ranks 1 to $n-1$ and p and each a appears exactly once in rank p;

the b-s occupy the ranks n to $p-1$ and the pattern of the b-s is cyclical for the first q agents.

Pick a lottery ℓ implementing λ at this profile. Then $\ell_a \ge x$ for each a implying $[\lambda]_1^k \ge kx$ for $1 \le k \le n-1$; moreover, $\lambda_p = x$ by assumption. It remains to show that $[\lambda]_{p-r}^p \le x + ry$ for $1 \le r \le q-1$. Indeed by summing the implementation constraints for the top

$r+1$ outcomes of the first q agents, we get (denoting the top outcome of agent i by a_i):

$$\begin{aligned} q[\lambda]_{p-r}^{p} &\le \sum_{i=1}^{q} \ell_{a_i} + r\sum_{i=1}^{q} \ell_{b_i} = \left(\sum_{i=1}^{q} \ell_{a_i} + \sum_{i=1}^{q} \ell_{b_i}\right) + (r-1)\sum_{i=1}^{q} \ell_{b_i} \\ &\le \bigl(1-(n-q)x\bigr) + (r-1)(1-nx) = q(x+ry) \end{aligned}$$

Case 2. $\lambda_1 \le \frac{1}{p}$. Set $\lambda_1 = x$ and $p = n+q$. We show similarly that λ is dominated (weakly) by the guarantee $\mu \in [\text{uni}(p), \text{vt}(n, p)]$ such that $\mu_1 = x$: that is, $\mu_k = x$ for $p-n+2 \le k \le p$ and $\mu_k = y$ for $2 \le k \le q+1$, with $nx + qy = 1$.

We consider a profile of preferences over the outcomes in $\{a_1, \dots, a_n\} \cup \{b_1, \dots, b_q\}$ where:

the a-s occupy the ranks 1 and $p-n+2$ to p and each a appears exactly once in rank 1;

the b-s occupy the ranks 2 to $q+1$ and the pattern of the b-s is cyclical for the first q agents.

Then the proof mimics that in case 1 by showing first that a lottery implementing λ at this profile has $[\lambda]_{p-k+1}^{p} \le kx$ for $1 \le k \le n-1$, then focusing attention on the first $q+1$ ranks to show $[\lambda]_1^{r+1} \ge x + ry$ for $1 \le r \le q-1$. We omit the details.

Case 3. $\lambda_p < \frac{1}{p} < \lambda_1$. Combining these inequalities with those in step 2 we see that λ is strictly dominated by $\text{uni}(p)$.

A.4 *Proof of Theorem 2*

We fix $1 \le q \le n$ such that $p = dn + q$ and prove the statement by induction on d. It is clear for $d = 1$ as $\{VT\}$ and $\{RD\}$ are the only two sequences and the intervals $[\text{uni}(p), \text{vt}(n, p)]$, $[\text{uni}(p), \text{rd}(n, p)]$ are in $\mathcal{M}(n, p)$.

Fix $d \ge 2$ and consider a sequence $\Gamma \in \{VT, RD\}^d$ starting with $\Gamma^1 = VT$. By Definition 5, the composition by VT commutes with convex combinations of $\Gamma^2, \Gamma^2 \otimes \Gamma^3, \dots$. Using the notation $\text{Vex}[\cdot]$ for such combinations, and the simplified notation VT instead of $VT \otimes uni(p-n)$, etc., we have

$$\begin{aligned} &\text{Vex}\bigl[VT, VT \otimes \Gamma^2, \dots, VT \otimes \Gamma^2 \otimes \cdots \otimes \Gamma^d\bigr] \\ &\quad = VT \otimes \text{Vex}\bigl[\text{uni}(p-n), \Gamma^2, \Gamma^2 \otimes \Gamma^3, \dots, \Gamma^2 \otimes \cdots \otimes \Gamma^d\bigr] \end{aligned} \tag{19}$$

where by the inductive assumption the second convex combination of canonical guarantees in $\mathcal{C}(n, p-n)$ and of $\text{uni}(p-n)$ is a maximal guarantee. By Lemma 3, so is the left-hand convex combination, call it λ, and by Proposition 4 so is a convex combination of $\text{uni}(p)$ and λ.

The proof of the inductive step for a sequence starting from RD is more involved, because RD does not commute with convex combinations, even of boundary lotteries; therefore, property (19) where RD replaces VT can only be true if the two sides use different convex combinations.

Observe first that if the lottery λ is maximal at n, $p-n$ and in $\partial\Delta(p-n)$, then any $\mu \in \mathrm{Vex}[\mathrm{rd}(n,p), RD \otimes \lambda]$ is in $\mathcal{M}(n,p) \cap \partial\Delta(p)$ as well. That μ is on the boundary is clear. By (11), μ takes the form

$$\mu = \Big(\overbrace{\frac{\alpha}{n}, \dots, \frac{\alpha}{n}}^{n-1}, (1-\alpha)\lambda, \frac{\alpha}{n}\Big)$$

Consider a profile similar to (13) in the proof of Lemma 3, where by maximality of λ we choose π ensuring property (12) in Lemma 4:

$$\begin{array}{ccccccc} \prec_1 & a_1 & \cdots & a_{n-1} & \overbrace{\pi}^{p-n} & a_n \\ \cdots & \cdots & & \cdots & \pi & \cdots \\ \prec_n & a_n & \cdots & a_{n-2} & \pi & a_{n-1} \end{array}$$

If the lottery ℓ implements μ at this profile, we have $\ell_{a_i} = \frac{\alpha}{n}$ therefore its weight on the remaining $p-n$ outcomes in π is $(1-\alpha)$ and the claim follows by Lemma 4 again.

We fix now an arbitrary convex combination

$$\Lambda = \sum_{j=2}^{d} \alpha_j RD \otimes \Gamma^2 \otimes \cdots \otimes \Gamma^j$$

in $\mathcal{G}(n,p)$ and claim that it takes the form $RD \otimes \lambda$ where λ is some other convex combination

$$\lambda = \sum_{j=2}^{d} \beta_j \Gamma^2 \otimes \cdots \otimes \Gamma^j.$$

This claim allows us to complete the induction step as follows. By the induction hypothesis, λ is in $\mathcal{M}(n, p-n)$, and it is easy to see (and explained in detail below) that it is on the boundary. By what we just observed, any $\mathrm{Vex}[\mathrm{rd}(n,p), RD \otimes \lambda]$ is also maximal; by the claim, this means that any convex combination of the guarantees corresponding to the initial subsequences of Γ starting with RD is maximal. Finally, Proposition 4 handles the addition of the uniform guarantee.

Proof. Proof of the claim Recall that canonical guarantees are uniform on their support, which we now describe for the canonical guarantees in our sequence. We partition the ranks $1, \dots, p$ into subsets $S^1, \dots, S^{d+1}$ each of size n except for the last one of size q. The set S^1 is the support of $\mathrm{rd}(n,p)$ (the ranks 1 to $n-1$ and p). If $\Gamma^2 = RD$, then S^2 has the ranks n to $2n-2$ and $p-1$, and the support of $RD \otimes \Gamma^2$ is $S^1 \cup S^2$. If $\Gamma^2 = VT$, then S^2 has the rank n and those from $p-n+1$ to $p-1$, and the support of $RD \otimes \Gamma^2$ is $S^1 \cup S^3 \cup \cdots \cup S^{d+1}$ (the complement of S^2). Continuing in this fashion, each Γ^j defines a new set S^j that is added to its support if $\Gamma^j = RD$, while if $\Gamma^j = VT$ we add $S^{j+1} \cup \cdots \cup S^{d+1}$ to the support. We keep track of this construction by entering a one for sets in the support and a zero for those outside it: with the notation $\varepsilon \in \{0, 1\}$ and $\varepsilon' = 1 - \varepsilon$, our sequence in $\mathcal{C}(n, p)$ is described by a table as follows:

$$
\begin{array}{lccccccc}
 & S^1 & S^2 & S^3 & S^4 & \cdots & S^d & S^{d+1} \\
RD \otimes \Gamma^2 & 1 & \varepsilon_2 & \varepsilon_2' & \varepsilon_2' & \cdots & \varepsilon_2' & \varepsilon_2' \\
RD \otimes \Gamma^2 \otimes \Gamma^3 & 1 & \varepsilon_2 & \varepsilon_3 & \varepsilon_3' & \cdots & \varepsilon_3' & \varepsilon_3' \\
\cdots & \cdots & \cdots & \cdots & \cdots & \cdots & \cdots & \cdots \\
RD \otimes \Gamma^2 \otimes \cdots \otimes \Gamma^d & 1 & \varepsilon_2 & \varepsilon_3 & \varepsilon_4 & \cdots & \varepsilon_d & \varepsilon_d'
\end{array}
$$

where $\varepsilon_j = 1$ if $\Gamma^j = RD$, $\varepsilon_j = 0$ if $\Gamma^j = VT$.

Defining $\Theta_k = n\sum_{j=2}^k \varepsilon_j + (p - kn)\varepsilon_k'$, we see in the table that Θ_k is the size of the support of $\Gamma^2 \otimes \cdots \otimes \Gamma^k$, while that of $RD \otimes \Gamma^2 \otimes \cdots \otimes \Gamma^k$ has cardinality $\Theta_k + n$. On its support $RD \otimes \Gamma^2 \otimes \cdots \otimes \Gamma^k$ is worth $\frac{1}{\Theta_k + n}$ while $\Gamma^2 \otimes \cdots \otimes \Gamma^k$ is $\frac{1}{\Theta_k}$ on its own support.

Clearly, but critically, there is a column with only zeroes. This holds if $\varepsilon_2 = 0$ ($\Gamma^2 = VT$), or if $\varepsilon_2 = 1$ but $\varepsilon_3 = 0$, etc., until, if $\varepsilon_j = 1$ for all j, the last column is null. A symmetric argument shows that in addition to the first column, there is another column full of ones. The first remark implies that Λ and λ are respectively in $\partial\Delta(p)$ and $\partial\Delta(p - n)$; the second that the maximal coordinate of λ is $\lambda_+ = \sum_{j=2}^d \frac{\beta_j}{\Theta_j}$. Now we select the coefficients β_j such that

$$\frac{1}{n\lambda_+ + 1}\frac{\beta_j}{\Theta_j} = \frac{\alpha_j}{\Theta_j + n} \quad \text{for all } j = 2, \ldots, d, \quad \text{and} \quad \sum_{j=2}^d \beta_j = 1$$

Check that β is well-defined because summing the first $d - 1$ equalities above implies

$$\frac{n\lambda_+}{n\lambda_+ + 1} = \sum_{j=2}^d \frac{n}{\Theta_j + n}\alpha_j < 1,$$

which determines λ_+. After rearranging the equation above as

$$\frac{1}{n\lambda_+ + 1} = \sum_{j=2}^d \frac{\Theta_j}{\Theta_j + n}\alpha_j,$$

the last equality $\sum_{j=2}^d \beta_j = 1$ follows.

We check finally the equality $\Lambda = RD \otimes \lambda$ for this choice of β. Because $\lambda \in \partial\Delta(p - n)$ the lottery $RD \otimes \lambda$ is given by (11); in particular, it is constant on each set S^k, just like Λ. We see in the table that $RD \otimes \lambda$ equals $\frac{\lambda_+}{n\lambda_+ + 1}$ in S^1, while Λ is worth $\sum_{j=2}^d \frac{\alpha_j}{\Theta_j + n}$ so they coincide. Each entry in another column S^k at row j adds $\varepsilon\frac{1}{n\lambda_+ + 1}\frac{\beta_j}{\Theta_j}$ to $RD \otimes \lambda$ and $\varepsilon\frac{\alpha_j}{\Theta_j + n}$ to Λ, where ε is the coefficient of that particular entry, so the desired equality follows. $\square$

A.5 *A mechanism implementing several guarantees*

Let $n = 3$, $p = 5$, and $\lambda = (\frac{1}{2}, 0, 0, \frac{1}{2}, 0)$ and $\mu = (\frac{1}{4}, \frac{1}{4}, \frac{1}{2}, 0, 0)$. Note that λ is maximal while μ is a guarantee undominated by λ. We construct a mechanism implementing both λ and μ.

Each agent i can either name a pair L_i of outcomes that he likes, or a single outcome d_i that he dislikes. Then the mechanism selects a set S of outcomes and performs a uniform lottery over its members. In cases 1–3 below, $|S| = 2$, while in case 4, $|S| = 4$.

Case 1: All three agents name pairs they like. Then we take S to be a pair that meets each of the three named pairs.

Case 2: All three agents name outcomes they dislike. Then we take S to be a pair avoiding the three named outcomes.

Case 3: Two agents name pairs they like, say L_1 and L_2, while the third agent names an outcome d_3 he dislikes. Then we take S to be a pair meeting L_1 and L_2 while avoiding d_3.

Case 4: One agent names a pair he likes, say L_1, while the other two agents name outcomes they dislike. Then we take S to be a set of four outcomes containing L_1.

It is easy to check that an agent naming his top two outcomes as the ones he likes guarantees himself at least λ, while an agent naming his bottom outcome as the one he dislikes guarantees himself at least μ.

References

Anbarci, Nejat (1993), "Noncooperative foundations of the area monotonic solution." *The Quarterly Journal of Economics*, 108, 245–258. [101]

Anbarci, Nejat and John P. Bigelow (1994), "The area monotonic solution to the cooperative bargaining problem." *Math. Soc. Sciences*, 28, 133–142. [101]

Aziz, Haris and Simon Mackenzie (2016), "A discrete and bounded envy-free cake cutting mechanism for any number of agents." In *IEEE 57th Annual Symposium on Foundations of Computer Science (FOCS)*, New Brunswick, NJ, USA, 9–11 Oct. 2016. [101]

Barbera, Salvador and Danilo Coelho (2017), "Balancing the power to appoint officers." *Games and Economic Behavior*, 101, 189–203. [102]

Barbera, Salvador and Bhaskar Dutta (1982), "Implementability via protective equilibria." *Journal of Mathematical Economics*, 10, 49–65. [101]

Bogomolnaia, Anna and Hervé Moulin (2020) Maths Operations Research. arXiv:1911.10009v3[econ.TH]. (Guarantees in fair division: general or monotone preferences, forthcoming, Sep 2020). [101]

Brams, Steven J. and Alan D. Taylor (1995), "An envy-free cake division protocol." *American Math. Monthly*, 102, 9–18. [101]

Budish, Eric (2011), "The combinatorial assignment problem: Approximate competitive equilibrium from equal incomes." *J. Polit. Econ.*, 119, 1061–1103. [101]

Chun, Youngsub (1986), "The solidarity axiom for quasi-linear social choice problems." *Social Choice and Welfare*, 3, 297–310. [102]

Dasgupta, Partha, Peter Hammond, and Eric Maskin (1979), "The implementation of social choice rules: Some general results on incentive compatibility." *The Review of Economic Studies*, 46, 185–216. [101]

De Clippel, Geoffroy, Kfir Eliaz, and Brian Knight (2014), "On the selection of arbitrators." *American Economic Review*, 104, 3434–3458. [102]

Dubins, Lester E. and Edwin H. Spanier (1961), "How to cut a cake fairly." *The American Mathematical Monthly*, 68, 1–17. [101]

Gibbard, Allan (1977), "Manipulation of schemes that mix voting with chance." *Econometrica*, 45, 665–681. [101]

Hougaard, Jens Leth Hervé Moulin, and Lars Peter Osterda (2010), "Decentralized pricing in minimum cost spanning trees." *Economic Theory*, 44, 293–306. [101]

Kirneva, Margarita and Matias Nunez (2021). On coordination in selection committees. [102]

Kuhn, Harold W. (1967), "On games of fair division." In *Essays in Mathematical Economics in Honour of Oskar Morgenstern*, 29–37, Princeton University Press. [101]

Kurokawa, David, Ariel D. Procaccia, and Junxing Wang (2016), "When can the maximin share guarantee be guaranteed?" In *Proceedings of the 30th AAAI Conference on Artificial Intelligence (AAAI)*, 523–529. [101]

Laslier, Jean Francois, Matias Nunez, and Remzi Sanver (2020), "A solution to the two-person implementation problem." halshs-02173504v3. [102]

Moulin, Hervé (1981), "Prudence versus sophistication in voting strategy." *Journal of Economic Theory*, 24, 398–412. [101]

Moulin, Hervé (1985), "Egalitarianism and utilitarianism in quasi-linear bargaining." *Econometrica*, 53, 49–68. [102]

Moulin, Hervé (1992a), "All sorry to disagree: A general principle for the provision of non-rival goods." *Scandinavian Journal of Economics*, 94, 37–51. [101]

Moulin, Hervé (1992b), "Welfare bounds in the cooperative production problem." *Games and Economic Behavior*, 4, 373–401. [101]

Mueller, Dennis C. (1978), "Voting by veto." *Journal of Public Economics*, 10, 57–75. [101]

Procaccia, Ariel D. and Junxing Wang (2014), "Fair enough: Guaranteeing approximate maximin shares." In *Proceedings of the 14th ACM Conference on Economics and Computation EC'14*, 675–692. [101]

Raiffa, Howard (1953), "Arbitration schemes for generalized two-person games." In *Contributions to the Theory of Games II* (Harold W. Kuhn and Albert W. Tucker, eds.). Annals of Mathematics Studies. Princeton University Press. [101]

Robertson, Jack and William Webb (1998), *Cake Cutting Algorithms: Be Fair if You Can*. A. K. Peters. [101]

Sen, Arunava (2011), "The Gibbard random dictatorship theorem: A generalization and a new proof." *SERIEs*, 2, 515–527. [101]

Sobel, Joel (1981), "Distortion of utilities and the bargaining problem." *Econometrica*, 49, 597–619. [101]

Steinhaus, Hugo (1949), "Sur la division pragmatique." *Econometrica (supplement)*, 17, 315–319. [101]

Stromquist, Walter (1980), "How to cut a cake fairly." *Amer. Math. Monthly*, 87, 640–644. [101]

Thomson, William (1978), "Maximin strategies and elicitation of preferences." University of Minnesota. Report. [101]

Thomson, William (1981), "A class of solutions to bargaining problem." *Journal of Economic Theory*, 25, 431–441. [101]

Woodall, Douglas R. (1980), "Dividing a cake fairly." *Journal of Mathematical Analysis and Applications*, 78, 233–247. [101]

Co-editor Simon Board handled this manuscript.

Manuscript received 14 April, 2021; final version accepted 30 March, 2022; available online 1 April, 2022.