

Worst Case in Voting and Bargaining

Anna Bogomolnaia*, Ron Holzman[†], Hervé Moulin[‡]

April 14, 2021

Abstract

The guarantee of an anonymous mechanism is the worst case welfare an agent can secure against unanimously adversarial others. How high can such a guarantee be, and what type of mechanism achieves it?

We address the worst case design question in the n -person probabilistic voting/bargaining model with p deterministic outcomes. If $n \geq p$ the uniform lottery is the only maximal (unimprovable) guarantee; there are many more if $p > n$, in particular the ones inspired by the random dictator mechanism and by voting by veto.

If $n = 2$ the maximal set $\mathcal{M}(n, p)$ is a simple polytope where each vertex combines a round of vetoes with one of random dictatorship. For $p > n \geq 3$, writing $d = \lfloor \frac{p-1}{n} \rfloor$, we show that the dual veto and random dictator guarantees, together with the uniform one, are the building blocks of 2^d simplices of dimension d in $\mathcal{M}(n, p)$. Their vertices are guarantees easy to interpret and implement. The set $\mathcal{M}(n, p)$

*University of Glasgow, CNRS-CES, and Higher School of Economics, St Petersburg; anna.bogomolnaia@glasgow.ac.uk

[†]Technion-Israel Institute of Technology. Work done during a visit at Princeton University; holzman@technion.ac.il

[‡]University of Glasgow and Higher School of Economics, St Petersburg; herve.moulin@glasgow.ac.uk

may contain other guarantees as well; what we can say in full generality is that it is a finite union of polytopes, all sharing the uniform guarantee.

Keywords: *worst case, guarantees, voting by veto, random dictator*

JEL classification: D 63, D71

Acknowledgement 1 *We are grateful for the comments of seminar participants in the University Paris 1, Lancaster, Oxford and Ashoka Universities. Moulin's research was supported by a Blaise Pascal Chair of the Region Ile-de-France, 2020-21.*

1 Guarantees and protocols

Worst case analysis is a simple mechanism design question. Fix an arbitrary collective decision problem by its feasible outcomes – allocation of resources, public decision making, etc. –, the domain of individual preferences and the 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 this game form without any prior knowledge of how others will play their part: the worst case assumption is that their moves are collectively adversarial (what I know or believe about their preferences, what I expect about their behaviour is irrelevant). My guarantee is the value of the two-person zero-sum game pitting me against the rest of the world.

Given the decision problem, what guarantees can any mechanism offer, and which mechanisms implement such guarantees? We are particularly interested in the *maximal* guarantees, those that cannot be improved: a higher guarantee is a better default option if I am clueless about other participants or unwilling to engage in risky strategic moves; it encourages acceptance of and participation in the mechanism.

These questions were first addressed by the cake-cutting literature ([32], [12], [16]). Two agents divide a cake over which their utilities are additive

and non atomic. In the Divide and Choose mechanism (D&C for short) they each can guarantee a share worth $1/2$ of the whole cake: the Divider must cut the cake in two parts of equal worth, any other move is at her own risk. This guarantee is not only maximal but also optimal (higher than any other feasible guarantee): when the two agents have identical preferences, their common guarantee cannot be worth more than $1/2$ of the cake.

Contrast D&C with the simple Nash demand game: each agent claims a share of the cake, demands are met if they are compatible, otherwise nobody gets any cake. Against an adversarial player I will not get anything: this mechanism offers no guarantee at all. The appeal of D&C is that I cannot be tricked to accept a share worth less than $1/2$.

Worst case analysis is related to the familiar implementation methodology in mechanism design, but only loosely. We speak of a mechanism *implementing* a certain guarantee – mapping my preferences to a certain welfare level – but we do not postulate that each agent behaves under the worst case assumption, nor do we ask what social choice function will then be realised, as some papers reviewed in section 2 did. Instead the guarantee of a game form is just one of its features, an important one for two reasons:

- 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 the guaranteed welfare to everyone,
- if an agent plays the mechanism repeatedly with changing sets of participants, the safe strategy is always available when she happens to be clueless about the other agents’ behaviour,

Many different mechanisms can implement the same guarantee, as the example below makes clear.¹ We call the whole class of game forms shar-

¹Similarly we can implement the optimal “one half of the whole cake” guarantee by D&C or by any one of Dubins and Spanier’ moving knife procedures ([12])

ing a certain guarantee the *protocol* implementing it. The two concepts of guarantee and protocol and their relation is the object of worst case analysis.

We initiate this approach in the probabilistic voting model, where the protocols we identify can be interpreted as the guidelines for a partially informal bargaining process. There are finitely many pure (deterministic) outcomes and we must choose a convex compromise (probabilistic or otherwise) between these. For tractability, we maintain a symmetric treatment of agents (Anonymity) and of outcomes (Neutrality). We find that, depending on the number n of agents and p of pure outcomes the set of maximal guarantees and their protocols can be either very simple and dull (when $n \geq p$, see below) or dauntingly complex.

A good starting point is the simple case of deterministic voting over p outcomes with ordinal preferences. An anonymous and neutral guarantee is a rank k from 1 to p (where rank 1 is the worst): it is feasible if for any preference profile there is at least one outcome ranked k or above by each voter. Suppose first $n \geq p$: at a profile where each outcome is the worst for some agent, the rank k must be 1, so the guarantee idea has no bite. Now if $n < p$ we can give to each voter the right to veto up to $d = \lfloor \frac{p-1}{n} \rfloor$ outcomes: this is feasible because $nd < p$, so the rank $k = d + 1$ is a feasible guarantee, clearly the best possible one. Worst case analysis' simple advice to a committee smaller than the number of outcomes it chooses from, is to distribute d veto rights to its members, not at all what a standard voting rule à la Condorcet or Borda does. However the corresponding protocol contains, inter alia, the following mechanisms: ask everyone to pick independently d outcomes to veto, then use any voting rule to pick among the remaining free outcomes, often more than $p - nd$ of them.

We allow compromises between the pure outcomes, interpreted as lotteries, time shares, or the division of a budget. Distributing veto tokens is a natural way to achieve a high guarantee, but there are others. The familiar

random dictator² mechanism ([13]) with two voters implements the guarantee putting a 1/2 probability on both my first and worst ranks. And the uniform lottery over all ranks is yet another guarantee implemented by any mechanism where each participant has the right, at some stage of the game that could depend upon the agent, the play of the game etc., to force the decision by flipping a fair coin between all outcomes.

Example: three agents, six outcomes The uniform guarantee $UNI(6)$ is the lottery $\lambda^{uni} = (\frac{1}{6}, \frac{1}{6}, \frac{1}{6}, \frac{1}{6}, \frac{1}{6}, \frac{1}{6})$, where each rank is equally probable.

Distributing one veto token to each agent implements $\lambda^1 = (0, 1, 0, 0, 0, 0)$ (recall the first coordinate is the worst rank), as in the deterministic case. But λ^1 is not maximal: it is improved by making the protocol a bit more precise. After the veto tokens have been used, we can pick one of the remaining outcomes uniformly, or give the option to force this random choice to each agent. Then the rank distribution cannot be worse for anyone than $\lambda^{vt} = (0, \frac{1}{3}, \frac{1}{3}, \frac{1}{3}, 0, 0)$, because my worst case after the vetoing phase is that the two other agents killed my two best outcomes. And λ^{vt} stochastically dominates λ^1 . We will use the notation $VT(3, 6)$ instead of λ^{vt} when it is important to specify n and p .

The random dictator mechanism between our three agents delivers the guarantee $\lambda^2 = (\frac{2}{3}, 0, 0, 0, 0, \frac{1}{3})$: my worst case is that the two other agents pick my worst outcome. Again λ^2 is not maximal, and improved by the following protocol: agents report (one of) their top outcome(s); if they all agree on a we choose a ; if they each choose a different outcome, we pick one of them with uniform probability; but 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 $RD(3, 6) : \lambda^{rd} = (\frac{1}{3}, \frac{1}{3}, 0, 0, 0, \frac{1}{3})$, that stochastically dominates λ^2 .

It is easy to check directly that $UNI(6)$, $VT(3, 6)$ and $RD(3, 6)$ are all

²Each agent has an equal chance to choose the final outcome.

maximal. For instance this follows for λ^{rd} and λ^{vt} by inspecting respectively the left or right profile of strict ordinal preferences

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

(where agent 1's worst is a and best is c). At the left profile, to give a $\frac{1}{3}$ chance of their best outcome to all agents a protocol implementing λ^{rd} must pick a, b or c , each with probability $\frac{1}{3}$: then each agent experiences exactly the distribution λ^{rd} over her ranked outcomes, and no other lottery λ stochastically dominating λ^{rd} is a feasible guarantee at this profile. Similarly at the right profile, implementing λ^{vt} 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: a critical duality relation connects λ^{vt} and λ^{rd} (section 4).

What other guarantees are maximal for $n = 3, p = 6$? Convex combinations preserve feasibility but not maximality: for instance an equal chance of the protocols implementing $VT(3, 6)$ and $RD(3, 6)$ delivers the feasible guarantee $\frac{1}{2}\lambda^{rd} + \frac{1}{2}\lambda^{vt} = (\frac{1}{6}, \frac{1}{3}, \frac{1}{6}, \frac{1}{6}, 0, \frac{1}{6})$ which is dominated by λ^{uni} . But lotteries between $UNI(6)$ and $VT(3, 6)$, or between $UNI(6)$ and $RD(3, 6)$ are in fact maximal. Moreover for this choice of n and p , the maximal guarantees cover exactly the two intervals $[\lambda^{uni}, \lambda^{vt}]$ and $[\lambda^{uni}, \lambda^{rd}]$ (Theorem 1 in section 5).

The choice facing the worst case designer in this example is sharp, and its resolution is context dependent: the veto guarantee is a good fit 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 different activities, or choosing a pair of roman consuls; the uniform guarantee stands out if we value a disagreement outcome revealing no information about individual preferences.

Critical to their practical application, the protocols implementing UNI, VT

and RD above rely on ordinal preferences only, as do the agents' safe action when they report which outcome(s) they veto, or which ones they prefer among those still in play.

The punchline Our results cast a new light on two familiar collective decision mechanisms, random dictator and voting by veto. Together and in combination with the uniform guarantee, they generate all maximal guarantees if $n = 2$, and essentially all of them again if $3 \leq n < p \leq 2n$. In the general case they can be sequentially combined to produce a very large subset of maximal guarantees.

1.1 Contents of the paper

After a review of the literature in section 2 we define in section 3 the concept of guarantee in three related models. In the first one, agents have ordinal preferences over the pure outcomes, and incomplete preferences over lotteries by stochastic dominance. In the second they have von Neuman Morgenstern (vNM) utilities over lotteries. In the third they have quasi-linear utilities over outcomes and money, and lotteries are replaced by cash compensations. A guarantee is a convex combination of the ranks 1 to p where rank 1 is the worst. It is feasible if at each profile of preferences, there is a lottery over pure outcomes, or a pure outcome and a balanced set of cash compensations in the quasi-linear model, that everyone weakly prefers to her guaranteed utility.

Lemma 1 shows that the three definitions are equivalent and that feasible guarantees cover a canonical polytope $\mathcal{G}(n, p)$ in the simplex with p ranked vertices. Its Corollary gives a compact though abstract characterisation of $\mathcal{G}(n, p)$.

Section 4 focuses on the subset $\mathcal{M}(n, p)$ of maximal guarantees, starting with a complete characterisation in two easy cases (section 4.1):

If $n \geq p$ the unique maximal guarantee is $UNI(p)$, dominating every other

feasible guarantee (Proposition 1), so the worst case viewpoint tells us to allow each agent to force this canonical anonymous and neutral disagreement outcome. In every other case there are many more options.

If $n = 2 < p$ a guarantee λ is maximal if and only if it is symmetric with respect to the middle rank (Proposition 2). For instance $\mathcal{M}(2, 6)$ is the convex hull of $\lambda^{rd} = (\frac{1}{2}, 0, 0, 0, 0, \frac{1}{2})$ ($RD(2, 6)$), $\lambda^{mix} = (0, \frac{1}{2}, 0, 0, \frac{1}{2}, 0)$, and $\lambda^{vt} = (0, 0, \frac{1}{2}, \frac{1}{2}, 0, 0)$ (two veto tokens per person). Here is the protocol for λ^{mix} : the agents veto one outcome each, then choose randomly a dictator (equivalently, we randomly give one veto token to one agent and four tokens to the other). Note that $UNI(p)$ is the center of the polytope $\mathcal{M}(2, p)$.

When $3 \leq n < p$, the structure of $\mathcal{M}(n, p)$ is much more complicated. Section 4.2 describes a critical duality property inside $\mathcal{M}(n, p)$, relating $VT(n, p)$ and $RD(n, p)$, while $UNI(p)$ is self-dual: Proposition 3. We define in section 4.3 the large set $\mathcal{C}(n, p)$ of *canonical guarantees*: for three or more agents they are a rich family of vertices of $\mathcal{M}(n, p)$. Their protocols combine up to d successive rounds (recall $d = \lfloor \frac{p-1}{n} \rfloor$) of either veto (one token each) or a (partial) random dictator.

Our first main result Theorem 1 in section 5.1, gives a fairly complete picture of all maximal guarantees with three or more agents and at most twice as many pure outcomes ($p \leq 2n \iff d = 1$). As long as $p \neq 2n - 1$ and $(n, p) \neq (4, 8), (5, 10)$, they cover exactly the two intervals $[UNI(p), VT(n, p)]$ and $[UNI(p), RD(n, p)]$, as in the numerical example above. There are additional maximal guarantees when $p = 2n - 1$ or $(n, p) = (4, 8), (5, 10)$, some of them described after the Theorem (Proposition 4).

In section 5.2 we turn to the general case $3 \leq n < p$ with no restrictions on d . The set $\mathcal{M}(n, p)$ is a union of polytopes (faces of $\mathcal{G}(n, p)$), of which $UNI(p)$ is always a vertex. Theorem 2 uses the canonical guarantees in $\mathcal{C}(n, p)$ to construct 2^d simplices of dimension d , one for each sequence Γ of length d in $\{VT, RD\}$: the vertices of such a simplex are lotteries in $\mathcal{C}(n, p)$ obtained from the d initial subsequences of Γ , plus $UNI(p)$. For instance the sequence

$\Gamma = (VT, RD)$ gives the triangle in $\mathcal{C}(3, 7)$ with vertices $UNI(p)$, $VT(3, 7)$, and $\lambda = (0, \frac{1}{3}, \frac{1}{3}, 0, \frac{1}{3}, 0, 0)$ denoted $VT \otimes RD$; the latter is implemented by a first round of one veto each, followed by $RD(3, q)$ over the remaining q outcomes (four or more). This construction does not cover the entire set $\mathcal{M}(n, p)$ but delivers a large subset built from simple combinations of veto and random dictator steps.

Section 6 gathers some open questions and potential research directions. Several proofs are gathered in the Appendix, section 7.

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 ([35], [11]), prudent ([20]) or protective strategies ([5]). As explained in section 1 our protocols do not define complete game forms, and our guarantees are compatible with a wide range of strategic behaviours.

Steinhaus' seminal papers ([32], see also [12], [16]) invented the worst case approach for cutting a cake fairly among any number of agents. His simple protocol generalises 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 ([7], [29], [3]) and proving its existence under preferences more general than additive utilities ([33], [36]). An exception is the recent paper ([6]) 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 ([8]): my worst share in the best partition

of the objects I can propose. This guarantee may not be feasible ([27]) but this happens very rarely ([17]); the real concern is that the protocols approximating this guarantee are all but simple.

Other early instances of the worst case approach are in production economies ([21], [22]) and in the minimal cost spanning tree problem ([14]).

The random dictator mechanism is a staple of probabilistic social choice ([13], [30]). In axiomatic bargaining it inspires the Raiffa solution ([28]) and the mid-point domination axiom ([31], [34]) satisfied by both the Nash and Kalai-Smorodinsky solutions.

Voting by veto is another early idea introduced by Mueller ([26]) to incentivise 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 generalised in ([20]). The area monotonic bargaining solution ([2], [1]) 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: ([9]), ([4]),([18]). All three papers implement maximal guarantees. Closer to home section 4 in ([15]) 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 quasi-linear utilities ([19], [10]) where only the uniform guarantee is discussed, while our results unveil many more possibilities.

3 Feasible guarantees

Anonymity and Neutrality (symmetric treatment of agents and outcomes, respectively) are hard wired in the model so a guarantee is well defined once

we fix 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 three equivalent definitions of the same concept of feasibility, after which when we speak of a guarantee we always mean that it is feasible.

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

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 the following three equivalent properties hold

$$\forall k \in [p] : [\lambda]_1^k \leq [\mu]_1^k$$

$$\forall k \in [p] : [\lambda]_k^p \geq [\mu]_k^p$$

$$\forall z \in \mathbb{R}^p : \{z_1 \leq z_2 \leq \dots \leq z_p\} \implies \lambda \cdot z \geq \mu \cdot z$$

The set of deterministic outcomes is A , with generic element a , and $\Delta(A)$, with generic element ℓ , is that 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 . An agent i 's ordinal preference over A (a complete, reflexive and transitive relation) is written \succsim_i . A k -tail of the preference \succsim_i is a subset T of A with cardinality k such that $b \succsim_i a$ whenever $a \in T, b \in A \setminus T$. Indifferences in \succsim_i may produce several k -tails.

Given \succsim_i and $\ell \in \Delta(A)$ the rank-ordered rearrangement of ℓ is the following lottery ℓ^{*i} in $\Delta(p)$

$$\forall k \in [p] : [\ell^{*i}]_1^k = \min \left\{ \sum_{a \in T} \ell_a \mid T \text{ is a } k\text{-tail of } \succsim_i \right\} \quad (1)$$

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 of preferences $\pi = (\succsim_i)_{i=1}^n$ on A there exists a lottery $\ell \in \Delta(A)$ s.t. $\ell^{*i} \vdash \lambda$ for all $i \in [n]$. Then we say that the lottery ℓ implements λ at profile π .

An 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 . We write $u_i^* \in \mathbb{R}^p$ the rank-ordered rearrangement (aka order statistics) of u_i :

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

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 $(u_i)_{i=1}^n$ on A there exists a lottery $\ell \in \Delta(A)$ s.t. $\ell \cdot u_i \geq \lambda \cdot u_i^*$ for all $i \in [n]$.

The ordinal definition is agnostic w.r.t. the risk attitude of the agents. The cardinal one specifies it completely.

In the third model the agents have quasi-linear utilities over the pure outcomes in A : in lieu of randomisation (or any convex combinations) compromises are achieved by cash compensations between the agents. Agent i 's utility is still any $u_i \in \mathbb{R}^A$ and an outcome is a pair (a, t) where $a \in A$ and $t = (t_i)_{i=1}^n$ is a balanced set of transfers between agents, $\sum_1^n t_i = 0$; the corresponding utilities are $u_{ia} + t_i$.

Definition 3 (quasi-linear utilities): Given n and p , the convex combination $\lambda \in \Delta(p)$ is a guarantee at n, p if for any n -profile $(u_i)_{i=1}^n$ of utilities on A , there exists an outcome $a \in A$ and a balanced set of transfers $(t_i)_{i=1}^n$ such that $u_{ia} + t_i \geq \lambda \cdot u_i^*$ for all $i \in [n]$.

Lemma 1 *These three definitions are equivalent.*

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

Proof

Definition 1 \implies Definition 2

The identity $\ell \cdot u_i = \ell^{*i} \cdot u_i^*$ for all $\ell \in \Delta(A)$, $u_i \in \mathbb{R}^A$ is easily checked. Now assume λ meets Definition 1 and fix 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 \implies *Definition 3*

Definition 3 says that λ is a guarantee if and only if for all $(u_i)_{i=1}^n \in (\mathbb{R}^A)^n$ we have:

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

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

$$\sum_{i=1}^n \ell \cdot u_i \leq \max_{a \in A} \sum_{i=1}^n u_{ia}$$

together imply (2).

Definition 3 \implies *Definition 1*. Fix λ as in Definition 3 and a preference profile $(\succsim_i)_{i=1}^n$. Call S_i the set of utilities $v_i \in \mathbb{R}^A$ representing \succsim_i ($a \succsim_i b \iff v_{ia} \geq v_{ib}$ for all a, b) and such that $v_{ia} \in [0, 1]$ for all a . By property (2) for any profile $(v_i)_{i=1}^n \in \prod_{i=1}^n S_i$ there exists $a \in A$ such that $\sum_{i=1}^n v_{ia} \geq \sum_{i=1}^n \lambda \cdot v_i^*$, which implies

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

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 $(v_i)_{i=1}^n \in \prod_{i=1}^n S_i$. 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. Definition 1 implies that the set $\mathcal{G}(n, p)$ is a polytope in $\Delta(p)$ (the convex hull of a finite set, and the intersection of finitely many half-spaces) for any n, p . Indeed feasibility of λ at some fixed ordinal profile

π means that a system of linear inequalities in ℓ of the form $M\ell \geq \theta$ has a solution ℓ , where the matrix M is independent of λ and θ depends affinely on λ : by the Farkas Lemma this is equivalent to finitely many linear inequalities on λ , and there are only finitely many ordinal profiles. ■

Corollary to Lemma 1 *The lottery $\lambda \in \Delta(p)$ is in $\mathcal{G}(n, p)$ if and only if for any n -profile $(u_i)_{i=1}^n$ in $(\mathbb{R}^A)^n$ we have*

$$\sum_{i=1}^n u_i = 0 \implies \sum_{i=1}^n \lambda \cdot u_i^* \leq 0 \quad (3)$$

“Only if” holds because (3) is a special case of the characteristic property (2). For “if” we pick an arbitrary profile $(u_i)_{i=1}^n$ 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 , then we can apply (3) to $(w_i)_{i=1}^n$: $w_i = v_i - \frac{z}{n}\mathbf{1}$. This gives $\sum_{i=1}^n \lambda \cdot v_i^* \leq z$ and the claim.

If $n = 2$ property (3) is easy to interpret, after checking the following identity (recall $\tilde{\lambda}$ is the symmetric of λ w.r.t. the middle rank):

$$\forall u \in \mathbb{R}^A : \lambda \cdot (-u)^* = -\tilde{\lambda} \cdot u^* \quad (4)$$

Property (3) means $\lambda \cdot u^* \leq \tilde{\lambda} \cdot u^*$ for all u , equivalently $\tilde{\lambda} \vdash \lambda$. Therefore

$$\lambda \in \mathcal{G}(2, p) \iff [\lambda]_1^k \geq [\lambda]_{p+1-k}^p \text{ for all } k = 1, \dots, \lfloor \frac{p}{2} \rfloor \quad (5)$$

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.

Remark. We downplay the quasi-linear interpretation because the corresponding protocols, though quite straightforward to construct, are less palatable. They require to report the willingness to pay for different outcomes, and thus essentially rely on the entire cardinal utility profile.

4 Maximal guarantees

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

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

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

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

4.1 Two easy cases: $n \geq p$ and $n = 2$

Proposition 1 *The uniform guarantee $UNI(p)$, $\lambda_k^{uni} = \frac{1}{p}$ 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: $\mathcal{M}(n, p) = \{\lambda^{uni}\}$.*
- iii) If $n \geq 3$ it is a vertex of $\mathcal{G}(n, p)$, hence of $\mathcal{M}(n, p)$ too.*

Proof.

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

$$\sum_{i=1}^n u_i = 0 \implies \lambda^{uni} \cdot \left(\sum_{i=1}^n u_i^* \right) = 0 \quad (6)$$

Suppose some $\mu \in \mathcal{G}(n, p)$ dominates λ^{uni} and consider a profile of the form $(u_1, -u_1, 0, \dots, 0)$ where u_1 is arbitrary. Summing up the inequalities $\mu \cdot u_1^* \geq \lambda^{uni} \cdot u_1^*$, $\mu \cdot (-u_1)^* \geq \lambda^{uni} \cdot (-u_1)^*$ gives $\mu \cdot u_1^* + \mu \cdot (-u_1)^* \geq 0$. Because μ meets property (3) both inequalities are equalities, and we conclude $\mu = \lambda^{uni}$.

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

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

with $n - p$ null utilities. Clearly $\sum_{k=0}^{p-1} u^{\sigma^k} = \gamma \mathbf{1}$ for $\gamma = \sum_{a \in A} u_a$, so we can apply property (3) to the profile $(u - \frac{\gamma}{p} \mathbf{1}, u^\sigma - \frac{\gamma}{p} \mathbf{1}, \dots, u^{\sigma^{p-1}} - \frac{\gamma}{p} \mathbf{1}, 0, \dots, 0)$. Together with $(u^{\sigma^k})^* = u^*$ for all k , this gives

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

as desired.

Statement iii). Suppose λ^{uni} is a convex combination of two distinct λ^1, λ^2 in $\mathcal{G}(n, p)$. For any profile s . t. $\sum_{i=1}^n u_i = 0$, property (3) implies $\sum_{i=1}^n \lambda^s \cdot u_i^* \leq 0$ for $s = 1, 2$. But by (6) the relevant convex combination of these inequalities is an equality, therefore they both are equalities.

For $n \geq 3$, a lottery λ meeting (6) must be λ^{uni} . Indeed, 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 (6) 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 . ■

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:

$$\lambda_k = \lambda_{p+1-k} \text{ for } 1 \leq k \leq \lfloor \frac{p}{2} \rfloor \quad (7)$$

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

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

We see that for $n = 2$ the uniform guarantee $UNI(p)$ is the center of the polytope $\mathcal{M}(2, p)$, contrasting sharply with the case $n \geq 3, p > n$ where $UNI(p)$ is an extreme point of the non convex set $\mathcal{M}(n, p)$: Theorem 2.

Proof. Fix $\lambda \in \mathcal{G}(2, p)$ and symmetric. Rewrite (7) as $\lambda \cdot u^* = \tilde{\lambda} \cdot u^*$ for all u^* , which by the identity (4) means $\lambda \cdot u^* = -\lambda \cdot (-u)^*$ for all u^* . The latter is property (6) for $n = 2$, so the maximality of λ follows as in the above proof of statement *i*) Proposition 1.

To prove the converse statement pick $\lambda \in \mathcal{G}(2, p)$, which means $\tilde{\lambda} \vdash \lambda$ (see property (5) in the previous section). As the dominance relation is preserved by convex combinations we have $\frac{1}{2}(\tilde{\lambda} + \lambda) \vdash \lambda$ where $\frac{1}{2}(\tilde{\lambda} + \lambda)$ is symmetric: thus λ is dominated if it is not symmetric. ■

The protocols implementing the vertices of $\mathcal{M}(2, p)$ combine in a simple way the veto and random dictator ideas explained in section 1. 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 $RD(2, p)$. The guarantee $VT(2, p)$ is $\lambda = (0, \frac{1}{p-2}, \dots, \frac{1}{p-2}, 0)$ implemented by giving one veto token per person, then selecting one of the remaining outcomes with uniform probability (as in the example of section 1). 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 $p - 1)$ outcomes: we write this guarantee as $VT \times RD(2, p)$. And so on: $VT^{t-1} \times RD(2, p)$ is the guarantee λ^t : let each agent veto $t - 1$ pure outcomes, then pick a random dictator for the remaining ones. If p is odd the guarantee $\lambda^{\frac{p+1}{2}}$ denoted $VT^{\frac{p-1}{2}}(2, p)$ simply gives $\frac{p-1}{2}$ veto tokens to each agent.

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

Although the results in this subsection apply for all n, p , they are only useful if $3 \leq n < p$, the only case not covered by Propositions 1 and 2: we maintain this assumption from now on. Besides the uniform rule $UNI(p)$ the two most

basic guarantees come from one round of veto or of random dictator:

$$VT(n, p) \rightarrow (0, \frac{1}{p-n}, \dots, \frac{1}{p-n}, \overbrace{0, \dots, 0}^{n-1})$$

$$RD(n, p) \rightarrow (\overbrace{\frac{1}{n}, \dots, \frac{1}{n}}^{n-1}, 0, \dots, 0, \frac{1}{n})$$

Given a lottery λ in $\Delta(p)$ different from λ^{uni} , the *radius to* λ is the interval of the half-line from λ^{uni} toward λ contained in $\Delta(p)$ (its other end is on the boundary $\partial\Delta(p)$), i. e. all lotteries of the form $\lambda^{uni} + \alpha(\lambda - \lambda^{uni})$ for some $\alpha \geq 0$. The *anti-radius from* $\tilde{\lambda}$ is the interval in $\Delta(p)$ of the half-line from λ^{uni} away from $\tilde{\lambda}$, i. e., the set of all lotteries of the form $\lambda^{uni} + \alpha(\lambda^{uni} - \tilde{\lambda})$ for some $\alpha \geq 0$.

If λ is a boundary lottery its dual λ^\star is the end point of the anti-radius from $\tilde{\lambda}$

$$\lambda^\star = (1 + \alpha)\lambda^{uni} - \alpha\tilde{\lambda} \text{ where } \alpha = \frac{1}{p \cdot \max_{1 \leq k \leq p} \lambda_k - 1} \quad (8)$$

(where $\max_{1 \leq k \leq p} \lambda_k > \frac{1}{p}$ because $\lambda \in \partial\Delta(p)$). Keeping in mind that $\min_{1 \leq k \leq p} \lambda_k = 0$ it is easy to check the identity $(\lambda^\star)^\star = \lambda$. For non boundary lotteries we extend this definition linearly on the radius to λ

$$\{\mu \in \partial\Delta(p) \text{ and } \lambda = \delta\lambda^{uni} + (1 - \delta)\mu\} \implies \lambda^\star = \delta\lambda^{uni} + (1 - \delta)\mu^\star. \quad (9)$$

so that $\lambda \rightarrow \lambda^\star$ is a proper duality in $\Delta(p)$.

We see that $VT(n, p)$ and $RD(n, p)$ are dual of each other, while $UNI(p)$ is self-dual.

Proposition 3

i) If $\lambda \neq \lambda^{uni}$ is a maximal guarantee, the radius to λ and the anti-radius from $\tilde{\lambda}$ (the symmetric of λ w.r.t. the middle rank) are contained in $\mathcal{M}(n, p)$.

ii) The duality operation $\lambda \rightarrow \lambda^\star$ in $\Delta(p)$ preserves maximal lotteries:

$$[\mathcal{M}(n, p)]^\star = \mathcal{M}(n, p)$$

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

Note that the statement in Proposition 3 holds if we replace $\mathcal{M}(n, p)$ by $\mathcal{G}(n, p)$: the radius to a feasible guarantee, and the anti-radius from its symmetric are feasible as well; duality preserves feasibility. This follows at once from the Corollary to Lemma 1, the identity (4) and the definition (8).

The proof for $\mathcal{M}(n, p)$ is much harder: we need a technical result characterising $\mathcal{M}(n, p)$ in $\mathcal{G}(n, p)$ by its position w. r. t. 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 $z \in \mathcal{G}(n, p)^\ominus$ s. t. $\sum_{k=1}^p z_k = 0$, $z_1 < z_2 < \dots < z_p$ 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 the Appendix.

Proof of Proposition 3.

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 $\mu = \lambda^{uni} + \alpha(\lambda - \lambda^{uni})$ in the radius to λ . That μ is feasible as well ($\mu \in \mathcal{G}(n, p)$) is clear by checking property (3) in section 3. For maximality we use $\sum_{k=1}^p z_k = 0$ and $\lambda \cdot z = 0$ to compute $\mu \cdot z = (1 - \alpha)\lambda^{uni} \cdot z + \alpha\lambda \cdot z = 0$ and conclude $\mu \in \mathcal{M}(n, p)$ by Lemma 2 again.

Still fixing $\lambda \in \mathcal{M}(n, p)$ and z , we pick next a lottery $\mu = \lambda^{uni} + \alpha(\lambda^{uni} - \tilde{\lambda})$ in the anti radius from $\tilde{\lambda}$. For feasibility we check property (3) at an arbitrary profile $(u_i)_{i=1}^n$ s. t. $\sum_1^n u_i = 0$. Compute

$$\sum_1^n \mu \cdot u_i^* = (1 + \alpha) \sum_1^n \lambda^{uni} \cdot u_i - \alpha \sum_1^n \tilde{\lambda} \cdot u_i^* = \alpha \sum_1^n \lambda \cdot (-u_i)^* \leq 0$$

where the last equality is the identity (4), and the inequality is from property (3) for λ .

The argument just made shows that for any $\xi \in \mathcal{G}(n, p)$ the lottery $(1 + \alpha)\lambda^{uni} - \alpha\tilde{\xi}$ is feasible as well, in particular

$$0 \geq ((1 + \alpha)\lambda^{uni} - \alpha\tilde{\xi}) \cdot z = -\alpha\tilde{\xi} \cdot z$$

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 (4), 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 : $\mu \cdot w = -\alpha\tilde{\lambda} \cdot w = \alpha\lambda \cdot z = 0$, which proves the maximality of μ .

Statement ii) follows from statement *i)* and the definition of the duality operation. ■

4.3 Canonical guarantees

We write the largest coordinate of a lottery as $\lambda_+ = \max_{1 \leq k \leq p} \lambda_k$. We see from (8) that the dual λ^\star of the boundary lottery λ is

$$\lambda_k^\star = \frac{1}{p\lambda_+ - 1}(\lambda_+ - \tilde{\lambda}_k) \text{ for } 1 \leq k \leq p \quad (10)$$

(where $\tilde{\lambda}_k = \lambda_{p+1-k}$)

Definition 5 *Composition by VT and RD*

For any $\lambda \in \Delta(p)$ the lottery $VT \otimes \lambda \in \Delta(p+n)$ obtains by inserting λ between one zero in rank 1 and $n-1$ zeros after rank $p+1$:

$$VT \otimes \lambda = (0, \lambda, \overbrace{0, \dots, 0}^{n-1}) \quad (11)$$

If $\lambda \in \partial\Delta(p)$ the lottery $RD \otimes \lambda \in \partial\Delta(p+n)$ obtains by filling uniformly $n-1$ ranks before λ and one after as follows:

$$RD \otimes \lambda = \left(\overbrace{\left(\frac{\lambda_+}{n\lambda_+ + 1}, \dots, \frac{\lambda_+}{n\lambda_+ + 1} \right)}^{n-1}, \frac{1}{n\lambda_+ + 1} \cdot \lambda, \frac{\lambda_+}{n\lambda_+ + 1} \right) \quad (12)$$

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

$$RD \otimes \lambda = [VT \otimes \lambda^\star]^\star \quad (13)$$

If $\lambda \in \partial\Delta(p)$ we must check that the two definitions (12) and (13) coincide. Write μ for the boundary lottery on the right-hand side of equation (12): applying (10) and $\mu_+ = \frac{\lambda_+}{n\lambda_+ + 1}$ we get

$$\begin{aligned} \mu^\star &= \frac{1}{(p+n)\frac{\lambda_+}{n\lambda_+ + 1} - 1} \left(\frac{\lambda_+}{n\lambda_+ + 1} \mathbf{1} - \tilde{\mu} \right) \\ &= \frac{1}{p\lambda_+ - 1} (\lambda_+ \mathbf{1} - (\lambda_+, \tilde{\lambda}, \overbrace{\lambda_+, \dots, \lambda_+}^{n-1})) = VT \otimes \lambda^\star \end{aligned}$$

as desired.

Note that Definition 5 implies in particular $VT \otimes UNI(p) = VT(n, n+p)$

and $RD \otimes UNI(p) = RD(n, n + p)$.

Lemma 3

- i) The guarantees $VT(n, p)$ and $RD(n, p)$ are maximal.*
- ii) The composition of guarantees by VT and RD respects their feasibility and maximality. For any $\lambda \in \Delta(p)$*

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

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 characterisation of maximal guarantees; the proof, much easier than that of Lemma 2, is also in the Appendix.

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*

$$\max_{i \in [n]} [\ell^{*i}]_1^k = [\lambda]_1^k \tag{14}$$

Proof of Lemma 3

Statement i) The proof that $VT(n, p)$ is maximal, done in section 1 for $n = 3, p = 6$, is an application of Lemma 4. Its generalisation is straightforward. Then its dual $RD(n, p)$ is maximal by Proposition 3.

Statement ii) 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 conclude that $VT \otimes \lambda$ is feasible.

If now $\lambda \in \mathcal{M}(n, p)$, we fix an index $k \in [p - 1]$ and an (n, p) -profile π ensuring property (14) as in the premises of Lemma 4. We construct the

following $(n, p + n)$ profile θ

$$\begin{array}{ccccccc}
\prec_1 & a_1 & \overbrace{\pi}^p & a_2 & \cdots & a_n & \\
\cdots & \cdots & \pi & \cdots & & \cdots & \\
\prec_n & a_n & \pi & a_1 & \cdots & a_{n-1} &
\end{array} \tag{15}$$

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 (14) 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 feasible, resp. maximal if λ is follows from the duality relation (13) and the fact that duality respects maximality and feasibility (Proposition 3). Here is for completeness the protocol implementing $RD \otimes \lambda$ if λ is a boundary feasible lottery: 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 6 *Canonical guarantees*

Fix $n, p, 3 \leq n < p$, s. t. $d = \lfloor \frac{p-1}{n} \rfloor$ and $p = dn + q$ for some $q = 1, \dots, n$. Each sequence $\Gamma = (\Gamma^t)_{t=1}^h$ in $\{VT, RD\}$ of length $h, h \leq d$, defines a canonical guarantee $\Gamma^1 \otimes \Gamma^2 \otimes \dots \otimes \Gamma^h$ by iterating the composition operation, i.e.,

$$\Gamma^1 \otimes \Gamma^2 \otimes \dots \otimes \Gamma^h = \Gamma^1 \otimes (\Gamma^2 \otimes (\dots \otimes \Gamma^h) \dots)$$

where Γ^h acts on $(d-h+1)n+q$ outcomes, $\Gamma^{h-1} \otimes \Gamma^h$ acts on $(d-h+2)n+q$ outcomes, etc. We write their set as $\mathcal{C}(n, p)$, of cardinality $2^{d+1} - 2$.

By Lemma 3 and the fact that the composition by each Γ^t adds n outcomes to the previous ones, all canonical lotteries are maximal. By duality (13), canonical lotteries come in dual pairs: exchanging VT and RD in each

term of the sequence Γ produces the dual lottery.

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

We give some examples.

If $d = 1$ ($p \leq 2n$) $VT(n, p)$ and $RD(n, p)$ are the only canonical guarantees.

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

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

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

A simple protocol for the former gives h veto tokens to each agent, then randomises 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 randomise 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 $\mathcal{C}(3, 7)$:

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

The protocol 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 $VT(3, 4)$ among the remaining outcomes.

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

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

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

5 General results with three or more agents

5.1 Maximal guarantees for $3 \leq n < p \leq 2n$

If $d = 1$ we have only two canonical guarantees $VT(n, p)$ and $RD(n, p)$ and by Proposition 3 any convex combination of $UNI(p)$ with one of these two is also maximal. It turns out that, for the most part, this exhausts all maximal guarantees.

Theorem 1

i) For any n, p s. t. $3 \leq n < p$ let $\lambda^{vt}, \lambda^{rd}, \lambda^{uni}$ be the guarantees from $VT(n, p), RD(n, p)$ and $UNI(p)$. Then

$$[\lambda^{uni}, \lambda^{vt}] \cup [\lambda^{uni}, \lambda^{rd}] \subset \mathcal{M}(n, p) \tag{16}$$

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

The proof is in the Appendix.

Our next result explains 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.

Proposition 4

- i) If $p = 2n - 1$ and if $(n, p) = (4, 8)$ or $(5, 10)$, the inclusion (16) is strict.
ii) For $n = 3, p = 5$ there are two dual pairs of maximal guarantees on the boundary of $\Delta(5)$: $VT(3, 5), RD(3, 5)$ and the pair

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

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

- iii) For $n = 4, p = 7$ there are three dual pairs of maximal guarantees on the boundary of $\Delta(7)$: $VT(4, 7), RD(4, 7)$ and the two pairs

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

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

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

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 $UNI(p), VT(n, p)$ nor $RD(n, p)$ 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 protocols implementing λ and λ^\star 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 λ^\star 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 randomise uniformly over the remaining outcomes.

We omit for brevity the tedious arguments, available upon request from the authors, showing that these guarantees, as well as μ and μ^\star are maximal, and generate the entire sets $\mathcal{M}(3, 5)$ and $\mathcal{M}(4, 7)$. ■

Note that, in particular, Theorem 1 and Proposition 4 give a full description of maximal guarantees whenever $3 \leq n < p \leq n + 3$.

5.2 Maximal guarantees for $3 \leq n < p$

For higher values of $d = \lfloor \frac{p-1}{n} \rfloor$ we know only a few general facts about the structure of $\mathcal{M}(n, p)$. 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 $UNI(p)$ 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 s. t. $3 \leq n < p$, $d = \lfloor \frac{p-1}{n} \rfloor$. For each sequence Γ of length d in $\{VT, RD\}$, the canonical guarantees from*

the d initial subsequences³ 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.

The simplest example not covered in Theorem 1 is $n = 3, p = 7$, 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 $UNI(7)$ to each of the following two dual non canonical guarantees:

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

In general, we keep in mind that many more guarantees than the ones described in Theorem 2 are maximal. Pick any non canonical guarantee λ in $\mathcal{M}(n, p) \cap \partial\Delta(p)$, for instance those described in Proposition 4 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 non canonical maximal guarantees in $\mathcal{M}(n, p + hn) \cap \partial\Delta(p + hn)$.

³I. e., the guarantees $\Gamma^1, \Gamma^1 \otimes \Gamma^2, \Gamma^1 \otimes \Gamma^2 \otimes \Gamma^3$, etc..

6 Concluding comments

The set $\mathcal{M}(n, p)$ remains simple if $n = 2$ and/or $p \leq 2n$ (Proposition 2 and Theorem 1), but its combinatorial/geometric structure becomes complicated, perhaps severely, as $\frac{p}{n}$ increases while $n \geq 3$. Questions that remain open for further investigation include:

- can we get new maximal guarantees by other convex combinations of canonical guarantees than those described in Theorem 2? We conjecture the answer is No.
- what is the maximal dimension of a simplicial component of $\mathcal{M}(n, p)$? We conjecture it is $d = \lfloor \frac{p-1}{n} \rfloor$.
- can we evaluate the number of such components?

The early voting by veto literature stresses the guarantees it offers to coalitions of like-minded voters ([26]). We could similarly define, given n and p , a guarantee for each size of a coalition, and try to link our modeling approach to the design of voting rules where the strategic formation of coalitions promotes stability ([23], [24],[25]).

Even for the individual guarantees discussed here, it may be possible to connect a maximal guarantee with the "best" game form(s) to implement it, where best may refer to simplicity, or to strategic or normative properties.

7 Appendix

7.1 Proof of Lemma 2

We 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 :

$$W = \{z = \sum_{i=1}^n u_i^* \mid \text{for some } U = (u_i)_{i=1}^n \text{ s.t. } \sum_{i=1}^n u_i = 0\} \quad (17)$$

By its characteristic property (3) $\mathcal{G}(n, p)$ 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 $\overleftrightarrow{W} = \mathcal{G}(n, p)^\ominus \cap \{z \mid \sum_{k=1}^p z_k = 0\}$.

We fix now a maximal guarantee λ and define the sub-cone Z of \overleftrightarrow{W} :

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

This cone is convex, and every element of Z satisfies $z_1 \leq z_2 \leq \dots \leq z_p$, because these inequalities hold in W . To prove that Z contains some z s. t. $z_1 < z_2 < \dots < z_p$, 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 $z \in W$ s. t. $z_k < z_{k+1}$ 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

$$\sum_{i=1}^n u_i = 0 \text{ and } u_{1k}^* = 0, u_{1,k+1}^* = 1 \quad (18)$$

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 s. t. 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 , s.t. $\lambda \cdot u_i^{s*} = \lambda^i \cdot u_i^s$ 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

$$\sum_{i=1}^n x_i^a = 0 \text{ in each column } a$$

$x_{1a} = 0, x_{1b} = 1$ where a and b are the outcomes ranked k and $k+1$ by each u_1^s

Note that Q is non empty 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

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

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

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 s.t. 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 :

$$\left\{ \sum_{i=1}^n u_i = 0 \text{ and } u_{1k}^* < u_{1,k+1}^* \right\} \implies \sum_{i=1}^n \lambda \cdot u_i^* \leq -\varepsilon (u_{1,k+1}^* - u_{1k}^*) \quad (19)$$

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 (19) we can replace coordinate 1 by any coordinate i . Therefore we have

$$\sum_{i=1}^n u_i = 0 \implies \sum_{i=1}^n \lambda \cdot u_i^* \leq -\frac{\varepsilon}{n} \sum_{i=1}^n (u_{i,k+1}^* - u_{ik}^*) \quad (20)$$

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 (20) 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 $(\bar{u}_i^j)_{i=1}^n$ as in (17). Note that if $k > 1$ then $\bar{u}_{i_0,k-1}^{j_0*} < \bar{u}_{i_0,k}^{j_0*}$ for some i_0, j_0 . Pick such i_0 and j_0 (or arbitrary ones if $k = 1$). Let $a \in A$ be s. t. $\bar{u}_{i_0,k}^{j_0*} = \bar{u}_{i_0,a}^{j_0}$. For some small $\varepsilon > 0$, modify $(\bar{u}_i^{j_0})_{i=1}^n$ to $(u_i)_{i=1}^n$ by letting $u_{i_0,a} = \bar{u}_{i_0,a}^{j_0} - \varepsilon$, $u_{i_1,a} = \bar{u}_{i_1,a}^{j_0} + \varepsilon$ 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 $\sum_{i=1}^n \lambda \cdot u_i^* \geq \sum_{i=1}^n \lambda \cdot \bar{u}_i^{j_0*} = \lambda \cdot z^{j_0} = 0$. 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. ■

7.2 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 (14): 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 $\varepsilon > 0$ s. t. at any profile π there is some

lottery ℓ implementing λ at π and such that

$$\max_{i \in [n]} [\ell^{*i}]_1^k = [\lambda]_1^k - \varepsilon \quad (21)$$

We must 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 (21) implements λ' as well. Suppose next $\lambda_k = 0$. Then we have for all i

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

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 (21). ■

7.3 Proof of Theorem 1

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 $p \leq 2n-2$ and $2 \leq k \leq \lfloor \frac{p}{2} \rfloor$. If k divides p the lemma is obvious (take a partition of $[p]$). Suppose $p = tk + r$ where $1 \leq r \leq k-1$. 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 $\gamma_1 = \dots = \gamma_{t-1} = 1, \gamma_t = \frac{r}{k}, \gamma_{t+1} = \dots = \gamma_{t+k} = \frac{1}{k}$. These weights make S_1, \dots, S_{t+k} a balanced family, and it remains to check that $t+k \leq n$.

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

Assume next $p = 2n$ and $2 \leq k \leq n$. When k divides p a partition works, so we may assume that $3 \leq k \leq n - 1$ and thus $n \geq 4$. We further exclude the exceptional cases $n = 4, 5$ and assume $n \geq 6$. If $k \leq n - 2$ we still have $\frac{p}{k} + k \leq n + 1$ 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 $P_j, j \neq i$, 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 $P_j, j \neq i$, 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 $\frac{k-1}{2} + 4 \leq n$. ■

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

By duality, it suffices to show this for $2 \leq k \leq \lfloor \frac{p}{2} \rfloor$. 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 preferences in which $\{a_j : j \in S_i\}$ is the k -tail of the preferences of agent $i, i = 1, \dots, m$. Let ℓ be a lottery that implements λ at this profile. Then $1 = \sum_{a \in A} \ell_a = \sum_{i=1}^m \gamma_i \sum_{j \in S_i} \ell_{a_j} \leq \sum_{i=1}^m \gamma_i [\lambda]_1^k = \frac{p}{k} [\lambda]_1^k$, implying the desired inequality.

Step 3. We know from Proposition 3 and the maximality of $\lambda^{vt}, \lambda^{rd}$ that $\mathcal{M}(n, p)$ contains the union of the two intervals in the statement. Conversely we fix $\lambda \in \mathcal{G}(n, p)$, where (n, p) are as in Lemma 5, and show that it is dominated by a guarantee in those two intervals. We distinguish three cases.

Case 1. $\lambda_p \geq \frac{1}{p}$. Set $\lambda_p = x$ and keep in mind that feasibility implies

$x \leq \frac{1}{n}$. We will show that λ is dominated (weakly) by the guarantee $\mu \in [\lambda^{uni}, \lambda^{rd}]$ s. t. $\mu_p = x$: that is $\mu_k = x$ for $1 \leq k \leq n - 1$ and $\mu_k = y$ for $n \leq k \leq p - 1$, 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 $\ell_a \geq x$ for each a implying $[\lambda]_1^k \geq kx$ for $1 \leq k \leq n - 1$; moreover $\lambda_p = x$ by assumption. It remains to show that $[\lambda]_{p-r}^p \leq x + ry$ for $1 \leq r \leq q - 1$. 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):

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

Case 2. $\lambda_1 \leq \frac{1}{p}$. Set $\lambda_1 = x$ and $p = n + q$. We show similarly that λ is dominated (weakly) by the guarantee $\mu \in [\lambda^{uni}, \lambda^{vt}]$ s. t. $\mu_1 = x$: that is $\mu_k = x$ for $p - n + 2 \leq k \leq p$ and $\mu_k = y$ for $2 \leq k \leq q + 1$, 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 mimicks that in case 1 by showing first that a lottery implementing λ at this profile has $[\lambda]_{p-k+1}^p \leq kx$ for $1 \leq k \leq n - 1$, then

focusing attention on the first $q + 1$ ranks to show $[\lambda]_1^{r+1} \geq x + ry$ for $1 \leq r \leq q - 1$. 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 λ^{uni} . ■

7.4 Proof of Theorem 2

We fix $1 \leq q \leq n$ 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 $[UNI(p), VT(n, p)]$, $[UNI(p), RD(n, p)]$ are in $\mathcal{M}(n, p)$.

Fix $d \geq 2$ and consider a sequence $\Gamma \in \{VT, RD\}^d$ starting with $\Gamma^1 = VT$. By its definition (11) the composition by VT commutes with convex combinations of $\Gamma^2, \Gamma^2 \otimes \Gamma^3, \dots$. Using the notation $VEX[\cdot]$ for such combinations we have

$$\begin{aligned} VEX[VT, VT \otimes \Gamma^2, \dots, VT \otimes \Gamma^2 \otimes \dots \otimes \Gamma^d] &= \quad (22) \\ &= VT \otimes VEX[UNI, \Gamma^2, \Gamma^2 \otimes \Gamma^3, \dots, \Gamma^2 \otimes \dots \otimes \Gamma^d] \end{aligned}$$

where by the inductive assumption the second convex combination of canonical guarantees in $\mathcal{C}(n, p - n)$ and of $UNI(p - n)$ is a maximal guarantee. By Lemma 3 so is the left-hand convex combination λ , and by Proposition 3 so is a convex combination of $UNI(p)$ 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 (22) where RD replaces VT can only be true if the two sides use different convex combinations.

Observe first that if the boundary lottery λ is maximal, $\lambda \in \mathcal{M}(n, p - n) \cap \partial\Delta(p - n)$, then any $\mu = VEX[RD(n, p), RD \otimes \lambda]$ is in $\mathcal{M}(n, p) \cap \partial\Delta(p)$

as well. That μ is on the boundary is clear. By (12) μ takes the form

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

Consider a profile similar to (15) in the proof of Lemma 3, where by maximality of λ we choose π ensuring property (14) in Lemma 4:

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

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 $VEX[RD(n, p), RD \otimes \lambda]$ 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 3 handles the addition of the uniform guarantee.

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 $RD(n, p)$ (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 $S^1 \cup S^3 \cup \dots \cup S^{d+1}$ (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 $S^{j+1} \cup \dots \cup S^{d+1}$ 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

	S^1	S^2	S^3	S^4	\dots	S^d	S^{d+1}
$RD \otimes \Gamma^2$	1	ε_2	ε'_2	ε'_2	\dots	ε'_2	ε'_2
$RD \otimes \Gamma^2 \otimes \Gamma^3$	1	ε_2	ε_3	ε'_3	\dots	ε'_3	ε'_3
\dots	\dots	\dots	\dots	\dots	\dots	\dots	\dots
$RD \otimes \Gamma^2 \otimes \dots \otimes \Gamma^d$	1	ε_2	ε_3	ε_4	\dots	ε_d	ε'_d

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 $\Gamma^2 \otimes \dots \otimes \Gamma^k$, while that of $RD \otimes \Gamma^2 \otimes \dots \otimes \Gamma^k$ has cardinality $\Theta_k + n$. On its support $RD \otimes \Gamma^2 \otimes \dots \otimes \Gamma^k$ is worth $\frac{1}{\Theta_k + n}$ while $\Gamma^2 \otimes \dots \otimes \Gamma^k$ 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

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

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 (12): 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. ■

References

- [1] Anbarci N. 1993. Noncooperative foundations of the area monotonic solution, *The Quarterly Journal of Economics*, 108, 1, 245-258
- [2] Anbarci N, Bigelow JP. 1994. The area monotonic solution to the cooperative bargaining problem, *Math.Soc.Sciences*, 28, 2, 133-142
- [3] Aziz A, McKenzie S. 2016. A Discrete and Bounded Envy-free Cake Cutting Protocol for Any Number of Agents, *IEEE 57th Annual Sym-*

posium on Foundations of Computer Science (FOCS), New Brunswick, NJ, USA, 9-11 Oct. 2016

- [4] Barbera S, Coelho D. 2017. Balancing the power to appoint officers, *Games and Economic Behavior*, 101, 189-203
- [5] Barbera S, Dutta B. 1982. Implementability via protective equilibria, *Journal of Mathematical Economics*, 10, 1, 49-65
- [6] Bogomolnaia A, Moulin H. 2020. Guarantees in fair division: general or monotone preferences, arXiv:1911.10009v3[econ.TH]16 Sep 2020
- [7] Brams SJ, Taylor AD. 1995. An envy-free cake division protocol, *American Math. Monthly*, 102(1), 9-18
- [8] Budish E. 2011. The combinatorial assignment problem: Approximate competitive equilibrium from equal incomes. *J. Polit. Econ.* 119, 6, 1061–103
- [9] De Clippel G, Eliaz K, Knight B. 2014. On the selection of arbitrators, *American Economic Review*, 104, 11, 3434-58
- [10] Chun Y. 1986. The solidarity axiom for quasi-linear social choice problems, *Social Choice and Welfare*, 3, 297-310
- [11] Dasgupta P, Hammond P, Maskin E. 1979. The Implementation of Social Choice Rules: Some General Results on Incentive Compatibility, *The Review of Economic Studies*, 46, 2, 185-216
- [12] Dubins LE, Spanier EH. 1961. How to cut a cake fairly, *The American Mathematical Monthly*, 68, 1, 1-17
- [13] Gibbard A. 1977. Manipulation of schemes that mix voting with chance, *Econometrica*, 45, 3, 665-681

- [14] Hougaard J. L., Moulin H., Osterdal L. P. 2010. Decentralized pricing in minimum cost spanning trees, *Economic Theory*, 44, 2, 293-306
- [15] Kirneva M, Nunez M. 2021. On coordination in selection committees
- [16] Kuhn, H.1967. On games of fair division, *Essays in Mathematical Economics in Honour of Oskar Morgenstern*, Princeton University Press, pp. 29–37
- [17] Kurokawa D, Procaccia AD, Wang J. 2016. When can the maximin share guarantee be guaranteed?, *Proceedings of the 30th AAAI Conference on Artificial Intelligence (AAAI)*,523-9
- [18] Laslier JF, Nunez M, Sanver MR. 2020. A solution to the two-person implementation problem, halshs-02173504v3
- [19] Moulin H. 1985. Egalitarianism and utilitarianism in quasi-linear bargaining, *Econometrica*, 53, 1, 49-68
- [20] Moulin H. 1981. Prudence versus sophistication in voting strategy, *Journal of Economic Theory*, 24, 398-412
- [21] Moulin H. 1992. All Sorry to Disagree: a General Principle for the Provision of Non-rival Goods, *Scandinavian Journal of Economics*, 94, 37–51
- [22] Moulin H. 1992. Welfare Bounds in the Cooperative Production Problem, *Games and Economic Behavior*, 4, 373–401
- [23] Moulin H. 1981. The Proportional Veto Principle, *Review of Economic Studies*, 48, 407-416
- [24] Moulin H. 1983. *The Strategy of Social Choice*, Advanced Textbooks in Economics, North-Holland
- [25] Peleg B. 1984. *Game theoretic analysis of voting in committees*, Cambridge University Press.

- [26] Mueller D. 1978. Voting by veto. *Journal of Public Economics*, 10, 1, 57-75
- [27] Procaccia AD, Wang J. 2014. Fair Enough: Guaranteeing Approximate Maximin Shares, *Proceedings of the 14th ACM Conference on Economics and Computation EC'14*, 675–692
- [28] Raiffa H. 1953. Arbitration schemes for generalized two-person games, *Annals of Mathematics Studies*, in: H.W. Kuhn and A.W. Tucker, eds., *Contributions to the theory of games II*, Princeton U. P.
- [29] Robertson JM, Webb WA.1998. *Cake Cutting Algorithms: Be Fair If You Can*, A. K. Peters
- [30] Sen A. 2011. The Gibbard random dictatorship theorem: a generalization and a new proof. *SERIEs* 2, 515–527
- [31] Sobel J. 1981. Distortion of utilities and the bargaining problem, *Econometrica*, 49, 3, 597-619
- [32] Steinhaus H. 1949. Sur la division pragmatique *Econometrica* (supplement), 17, 315-319
- [33] Stromquist W. 1980. How to cut a cake fairly, *Amer. Math. Monthly* 87, 640-644
- [34] Thomson W. 1981, A class of solutions to bargaining problem, *Journal of Economic Theory* 25, 431-441
- [35] Thomson W. 1978 Maximin strategies and elicitation of preferences, mimeo, University of Minnesota
- [36] Woodall DR. 1980. Dividing a cake fairly, *Journal of Mathematical Analysis and Applications*, 78(1):233-247