

Fair Mixing: the Case of Dichotomous Preferences

HARIS AZIZ, UNSW Sydney and Data61 CSIRO

ANNA BOGOMOLNAIA, University of Glasgow and HSE St Petersburg

HERVÉ MOULIN, University of Glasgow and HSE St Petersburg

We consider a setting in which agents vote to choose a fair mixture of public outcomes. The agents have dichotomous preferences: each outcome is liked or disliked by an agent. We discuss three outstanding voting rules. The *Conditional Utilitarian* rule, a variant of the random dictator, is strategyproof and guarantees to any group of like-minded agents an influence proportional to its size. It is easier to compute and more efficient than the familiar *Random Priority* rule. We show, both formally and by numerical experiments, that its inefficiency is low when the number of agents is low. The efficient *Egalitarian* rule protects individual agents but not coalitions. It is *excludable strategyproof*: an agent does not want to lie if she cannot consume outcomes she claims to dislike. The efficient *Max Nash Product* rule offers the strongest welfare guarantees to coalitions, which can force any outcome with a probability proportional to their size. But it even fails the excludable form of strategyproofness.

CCS Concepts: • **Computing methodologies** → **Multi-agent systems**; *Cooperation and coordination*; • **Applied computing** → *Economics*.

Additional Key Words and Phrases: participatory budgeting, proportional representation, time-sharing, portioning, approval voting, strategyproofness

ACM Reference Format:

Haris Aziz, Anna Bogomolnaia, and Hervé Moulin. 2019. Fair Mixing: the Case of Dichotomous Preferences. In *The 20th ACM conference on Economics and Computation (EC '19)*, June 24–28, 2019, Phoenix, AZ, USA. ACM, New York, NY, USA, 27 pages. <https://doi.org/10.1145/3328526.3329552>

1 INTRODUCTION

In participatory budgeting [20] the stake-holders (citizens, employees of a firm, club members) vote to decide which subset of public projects and in what proportions the community, firm, or club should implement. We discuss a version of this process in the probabilistic voting model [26, 29].

The guiding principle of our analysis is that the selection of a single (deterministic) public outcome is *prima facie* unfair: fairness requires compromise, we must select a *mixture* of several mutually exclusive outcomes. This mixed outcome can be interpreted in a variety of ways, depending on the specific problem. It may be an actual randomization between outcomes, or an allocation of fractions of time each outcome is in place (“time-shares”), or a distribution of a fixed amount of some resource (e.g., money) over these outcomes. Some typical examples follow.

In the participatory budgeting problem [20], the city authority must divide funds or staff between several projects (library, sports center, concert hall) taking into account the citizens’ wishes. The scheduling of one or several weekly club meetings (gym classes, chess club, study group) must accommodate the time constraints reported by the club members. The local public TV must divide

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

EC '19, June 24–28, 2019, Phoenix, AZ, USA

© 2019 Association for Computing Machinery.

ACM ISBN 978-1-4503-6792-9/19/06...\$15.00

<https://doi.org/10.1145/3328526.3329552>

broadcasting time between different languages. In the *fair knapsack* problem (see, for example, [28]), the server schedules jobs of different reported or observed sizes under a capacity constraint, and must pick a (random) serving protocol. In all these examples, fairness requires to give some share of the public resources to everyone: each club member should have access to some meetings; everyone should enjoy at least some TV programs, etc. This contrasts with traditional high-stakes/low-frequency voting contexts, where the first best is to select a single (deterministic) outcome, and randomization over outcomes is only second best.¹

We run into the familiar conflict between protecting minorities and submitting to the will of the majority (as in the discussion of cumulative voting [47]). On the one hand, the larger the support for a public outcome, the bigger should be its share in the final compromise: numbers matter. On the other hand, we must protect minorities with their idiosyncratic preferences for outcomes disliked by the majority. So the club meetings will be more frequent when many members can attend, but nobody will be entirely excluded; the knapsack server will favor short jobs because this increases the number of satisfied customers, but it cannot ignore long jobs entirely; and so on.

We analyze this tradeoff when preferences can be represented in a very simple *dichotomous* form, as is used in approval ballots: each agent likes or dislikes each outcome,² and her utility is simply the total share of her *likes*. Agents in the knapsack problem care only about their expected service time, and in the club example, about the number of meetings they can attend. Though less natural in the public TV and the library funding examples, where they rule out any complementarities between outcomes, dichotomous preferences are still of practical interest because they are easy to elicit.

We discuss the fairness and incentive compatibility properties of three well-known social choice rules.

Our results. The *fair share guarantee* principle is central to the fair division literature since the earliest *cake division* papers [43]. In our model, this is the *Individual Fair Share* (IFS) axiom: each one of the n agents “owns” a $1/n$ -th share of decision power, so she can ensure an outcome she likes at least $1/n$ -th of the time (or with probability at least $1/n$). To capture more subtle ideas that minorities should be protected, and numbers should matter as well, we strengthen IFS to *Unanimous Fair Share* (UFS), giving to any group of like-minded agents an influence proportional to its size: so if 10% of the agents have identical preferences, they should like the outcome at least 10% of the time.

Our starting point is the impossibility result in [16], where our model and the two fairness properties IFS and UFS first appear: *no mixing rule can be efficient (ex-ante), incentive compatible in the prior-free sense of strategyproofness (SP), and meet Unanimous, or even Individual, Fair Share* (Proposition 6, [16]). One of our first results is that the impossibility disappears under the natural analogue for single-peaked preferences: outcomes can be ordered in such a way that the approval set of every agent is an interval. For such structured preferences, the random priority rule satisfies all the conditions. We then introduce *new* fairness and incentives properties and offer instead possibility results even when there is no structure on the dichotomous preferences. In particular, a natural and attractive incentive requirement we consider is *excludable strategyproofness* (EXSP). If public outcomes are non-rival but *excludable*, and we can force agents to consume only those outcomes they claim to like, so it becomes more costly to fake a dislike, and the strategyproofness is correspondingly weakened. A meeting of the club is such an *excludable* public outcome: it is easy

¹It is used to break ties, or to play the role of an absent deterministic Condorcet winner: for instance [5, 19, 33] identify a lottery that, in a certain sense, wins the majority tournament.

²This is different from the situations where each agent likes, dislikes or does not care about items, which should be described using trichotomous preferences.

to exclude from the meeting those who reported they could not attend; broadcasting via cable TV is similarly excludable, not so via aerial broadcasting.

Three remarkable mixing rules (two of them well known) meet IFS and achieve, loosely speaking, two out of the three goals of efficiency, group fairness (in the sense of UFS or other more demanding properties), and incentive compatibility.

We start with the *Egalitarian* (EGAL) rule, adapting to our model a celebrated principle of distributive justice. Taking the probability that the selected outcome is liked by agent i as her canonical utility, the rule maximizes first the utility level we can guarantee to all agents; among the corresponding mixtures, it maximizes the utility we can guarantee to all agents but one; and so on. It is efficient and satisfies IFS. Therefore it is not strategyproof, by the above-mentioned result. The Egalitarian rule is, however, excludable strategyproof: misreporting one's preferences does not pay, provided an agent is excluded from consuming those public outcomes she reportedly dislikes (Theorem 1). Thus weakening SP to EXSP resolves the impossibility result. But numbers do not matter to the egalitarian rule: it treats a unanimous group of agents exactly as if it contained a single agent, so the UFS property obviously fails. A related problem is that if an agent has a clone (another agent with preferences identical to hers), she can simply stay home and nothing will change. The *Strict Participation* (PART*) axiom takes care of this disenfranchisement problem by insisting that casting her vote is strictly beneficial to each voter. So the EGAL rule is only appealing if we focus on individual guarantees and are comfortable treating a homogenous group as a single person. This makes sense if the club must offer some important training to its members. But in the participatory budgeting or the broadcasting examples, numbers should matter.

The *Conditional Utilitarian* (CUT) rule is a simple variant of the classic "random dictator". Each agent identifies, among the outcomes she likes, those with the largest support from the other agents: then she spreads the probability (time share) of $1/n$ uniformly over these outcomes. So the utilitarian concern is conditional upon guaranteeing one's full utility first. The CUT rule is related to but much simpler than, the *Random Priority* (RP) rule which is discussed in [16] and which averages outcomes of all deterministic priority rules. Both rules are SP, meet PART* and guarantee UFS. It follows from the impossibility result in [16], that they are inefficient. But we show that CUT is strictly more efficient than RP. In numerical simulations (Section 9) and for relatively small values of n , its inefficiency is consistently low.

The third rule that we analyse is the familiar *Max Nash Product* (MNP) rule that picks the mixture that maximizes the product of individual utilities. It is efficient and offers much stronger welfare guarantees to groups than UFS. We introduce two requirements, each one a considerable strengthening of UFS. The *Core Fair Share* (CFS) property has an incentive flavor in the spirit of cumulative voting [30, 42], where a minority can enforce at least a fraction of outcomes by concentrating their votes on those. Any group of agents can pool their shares of decision power and object to the proposed mixture z by enforcing another mixture z' with a probability proportional to the group size. CFS requires that no such objection can benefit all members of the objecting group. Finally, *Average Fair Share* (AFS) applies to any coalition with a common liked outcome: the *average* utility in such a group cannot be smaller than its relative size. In simple examples, AFS limits the set of acceptable efficient mixtures very effectively. Theorem 3 shows that the efficient MNP rule meets PART*, CFS and AFS but even fails EXSP.

Our results suggest several challenging open questions about the impossibility frontiers of our model (see Conclusion). An overview of the properties satisfied by the rules discussed is given in Table 1.

The paper is organized as follows. We discuss related literature in Section 2. The model is specified in Section 3. Section 4 introduces the notions of Individual Fair Share (IFS) and excludable strategyproofness, and presents the efficient Egalitarian rule, which satisfies both. In Section 5 we

	RP	CUT	UTIL	EGAL	MNP
Properties					
Anonymity (ANON) and Neutrality (NEUT)	+	+	+	+	+
EFF (Efficiency)	-	-	+	+	+
MIX-PO (Mix over Pareto optimal outcomes)	+	+	+	+	+
EXSP (Excludable SP)	+	+	+	⊕	⊖
SP (strategyproofness)	+	+	+	-	-
IFS (Individual Fair Share)	+	+	-	+	+
UFS (Unanimous Fair Share)	+	+	-	-	⊕
GFS (Group Fair Share)	+	+	-	-	⊕
AFS (Avg. Fair Share)	⊖	⊖	⊖	⊖	⊕
CFS (Core Fair Share)	⊖	⊖	⊖	⊖	⊕
PART (Participation)	+	+	+	+	⊕
PART* (Strict participation)	+	⊕	-	-	⊕
DEC (Decentralisation)	⊕	⊕	⊖	⊖	⊕

Table 1. Properties satisfied by rules. Key results from this paper are circled.

define fair share notions that go beyond IFS, which the Egalitarian rule fails. Section 6 discusses two strategyproof rules which meet the conditions from Section 5, and compares them. Section 7 discusses the Maximal Nash Product rule. Section 8 is devoted to numerical examples, Section 9 to the Decentralization axiom, and Section 10 concludes. More technical proofs and data on numerical examples are relegated to the appendix.

2 RELATED LITERATURE

Participatory budgeting is an important new aspect of participative democracy, reviewed in [20]. Our model casts this process as a probabilistic voting problem, introduced first by Gibbard [29] as a way to design non-dictatorial strategyproof decision rules. The literature he inspired viewed randomization as a way around the defects of deterministic rules, mostly to allow anonymous and neutral rules, or to circumvent the absence of Condorcet winners [see e.g., 5, 9, 18, 26, 33]. But recent work turns its attention to mixtures of outcomes with time-sharing or compromise in mind [see, e.g., 3, 4, 9, 10, 16, 24].

Our work takes direct inspiration from the original paper of Bogomolnaia et al. [15, 16] who introduced the model of randomised voting under dichotomous preferences. For the same mathematical model, we present several new results about new normative requirements such as participation incentives, weaker forms of strategyproofness, and stronger forms of fairness.

Two of our rules, EGAL and MNP, maximize a familiar social welfare ordering and a classic collective utility function, respectively. The egalitarian rule is the lead mechanism in the related assignment model with dichotomous preferences in [13]. In probabilistic voting, the Egalitarian Simultaneous Reservation rule of Aziz and Stursberg [9] can be seen as an extension of EGAL (see also [11]).

Recent literature emphasizes that the MNP rule is central to the competitive approach of the fair division of private commodities, whether divisible or indivisible [14, 21]. See in particular, the discussion of Moulin [39]. We find here a new application of this rule in the public decision making

context, closer in spirit to Nash’s original bargaining model [40]. Our results are related to those of Fain et al. [24], who also propose the MNP rule for participatory budgeting, reinterpret this rule as a Lindahl equilibrium, and discuss its computational complexity. They allow for more general preferences than ours (in particular, full-fledged vNM utilities), and show the Core Fair Share property (Corollary 1 Section 2.3 in [24]) on larger domain than in statement i) of our Theorem 3. They do not discuss incentives properties or any alternative rule.

The rules CUT and RP are non-welfarist, in that they do not maximize any social welfare ordering. The RP rule is well known (and was discussed by Bogomolnaia et al. [16]), and CUT is a fairly simple twist on the random dictator first introduced by Duddy [23] who noted that it is strategyproof but did not develop its normative appeal.

Fair share is an early design constraint of decision mechanisms: see the mathematical literature on cake cutting [43], and on fair division of microeconomic commodities [38, 44, 45]. The group version of fair share captures the ubiquitous “protection of minorities” principle that is formally related to cooperative stability in standard voting. It is also related to the proportional veto principle [35, 36] and motivates practical twists in the rules such as cumulative voting, especially concerned with the protection of ethnic minorities in political elections [42], or minority stockholders in corporate governance [30, 41, 47]. See also the same concerns for EU enlargement [32].

Our fairness notions are intuitively related to proportional representation axioms in multi-winner voting with dichotomous preferences. The Justified Representation (JR) axiom [7] says: if a large enough group of voters agree in supporting the same candidate, then at least one voter in this group has an approved candidate in the winning committee. The Proportional Justified Representation (PJR) axiom [8] requires that if a group of voters of size ℓk agree to at least like a set of candidates of size ℓ , then the total utility of the group should be at least ℓ . JR can be viewed as an analogue of the IFS axiom adapted to the discrete setting of multi-winner voting with dichotomous preferences. Similarly, PJR can be viewed as a suitable analogue of a strengthening of the UFS axiom adapted to the setting of multi-winner voting with dichotomous preferences.

Strict Participation has been considered in the deterministic voting model, leading mostly to negative results. Our results complement those of Brandl et al. [17] who analyse participation incentives in probabilistic voting.

3 THE MODEL

Let N be a finite set of agents and let A be a finite set of outcomes. A generic agent is $i \in N$, and $n = |N|$. A *pure* public outcome is $a \in A$, and a *mixture* of public outcomes is a vector z , an element of the simplex $\Delta(A)$, interpreted as a lottery over A , or as a vector of time shares (or shares of other types of resources) allocated to the outcomes in A . A utility function (preference) $u_i = (u_{ia})_{a \in A}$ is an element of $\{0, 1\}^{|A|}$. Agents who dislike all outcomes play no role in any of the rules we discuss, thus we exclude them at once: the domain of preferences is $\Omega = \{0, 1\}^{|A|} \setminus \{\mathbf{0}\}$, where $\mathbf{0} = 0^{|A|}$; and $u \in \Omega^{|N|}$ is an instance of utility functions. In the examples, we always represent u as a $|N| \times |A|$ matrix³ filled with 0-s and 1-s, and we use the notation: $u_S = \sum_{i \in S} u_i$ and $u_{SB} = \sum_{i \in S} \sum_{a \in B} u_{ia}$ for $S \subseteq N$ and $B \subseteq A$. A *problem* M is a triple $M = (N, A, u)$ where $u \in \Omega^{|N|}$. Actual utilities (welfare) at a given mixed outcome $z \in \Delta(A)$ are written $U_i = u_i \cdot z$, and the corresponding realized utility profile is written $U = u \cdot z \in [0, 1]^{|N|}$. The set of feasible utility profiles is $\Phi(M) = \{U = u \cdot z \mid z \in \Delta(A)\}$. Given $U \in \Phi(M)$ we set $\varphi^{-1}(U) = \{z \in \Delta(A) \mid U = u \cdot z\}$.

Throughout, efficiency is taken in the ex-ante sense. In problem $M = (N, A, u)$ a feasible utility profile $U \in \Phi(M)$ is *efficient* if there is no profile $U' \in \Phi(M)$ such that $U \leq U'$ and at least one

³We often refer to outcomes as “columns”; for example, when two outcomes are liked by exactly the same set of agents, we speak of “two identical columns”.

inequality $U_i \leq U'_i$ is strict. A mixture $z \in \Delta(A)$ is efficient in M if the profile $u \cdot z$ is efficient. Fix $\varepsilon \in [0, 1]$; the profile $U \in \Phi(M)$ is $(1 - \varepsilon)$ -inefficient if there exists $U' \in \Phi(M)$ such that $U \leq \varepsilon U'$. This inequality reads that, even if we multiply vector U by $\frac{1}{\varepsilon} \geq 1$, it is still Pareto dominated by another feasible vector U' of utilities. Note that a utility profile is more $(1 - \varepsilon)$ -inefficient if ε is smaller.

A rule F picks one $U \in \Phi(M)$ for each problem M ; the mapping f picks the corresponding mixtures: $f(M) = \varphi^{-1}(F(M))$, so that $F(M) = u \cdot f(M)$. We only consider F and f which are anonymous (treat agents symmetrically) and neutral (treat outcomes symmetrically). The rule F is efficient if it selects an efficient profile in every problem. For any n , the rule is $(1 - \varepsilon(n))$ -inefficient if a) there exists a problem M of size n and a profile $U \in \Phi(M)$ such that $F(M)$ is $(1 - \varepsilon(n))$ -inefficient, and b) no smaller number $\varepsilon'(n)$ meets this property.

A rule is “welfarist” by design, in the sense that it does not distinguish between mixtures resulting in the same utility profile. For instance if two outcomes a, b are “clones” in problem M (liked by exactly the same agents), a rule is oblivious to shifting some weight from a to b .

The efficient pure outcomes in A are easy to recognize: a is efficient *if and only if* there is no b such that the set of agents liking b is strict superset of the set of agents liking a . We call such outcomes *undominated*. Consider the following example.

Example 3.1.

	A	a	b	c	d	e	
N							
1		0	0	0	1	1	
2		0	0	1	1	0	
3		1	1	0	0	0	(1)
4		1	0	1	0	0	
5		0	1	0	1	1	

Outcome e is dominated by d , and the four other outcomes are undominated. However, convex combinations of undominated outcomes may well be inefficient. In the example, any mixture $z = (z_a, z_b, z_c, z_d, z_e)$ such that z_b, z_c are both positive, say $z_b, z_c \geq \alpha > 0$, can be improved by redistributing the weight α to a and to d . That is, z is Pareto inferior to the mixture $z' = (z_a + \alpha, z_b - \alpha, z_c - \alpha, z_d + \alpha, z_e)$.

Of special interest are those problems where any mixture of undominated pure outcomes is efficient: in the probabilistic interpretation of our model this means that ex post efficiency implies ex ante efficiency. Indeed the four rules we discuss below mix only undominated outcomes, so in such problems their efficiency is guaranteed.

Our first (minor) result presents two examples where it is the case. Here, the set of outcomes liked by an agent is called her *approval set*.

Lemma 1

- i) *There exists an instance of preferences for which a mixture of undominated outcomes is dominated, if and only if $|A| > 3$ and $|N| > 4$*
- ii) *If A can be ordered in such a way that the approval set of every agent is an interval, then any mixture of undominated (pure) outcomes is efficient.*

PROOF. “Only if” part of statement i) is proven by Duddy [23]; “If” part follows from Example 3.1 (once we eliminate e), which has the smallest sizes of A and N for which a combination of undominated outcomes is inefficient. We provide the argument for statement ii).

Fix a problem M as in statement ii). If some outcomes are “clones” (liked by exactly the same set of agents), a class of clones is an interval as well, and it is clearly enough to prove the statement for

the “decloned” problem where each interval of clones has shrunk to a single outcome. Thus we can assume that our problem has no clones.

Let A^* denote the subset of undominated pure outcomes. We fix a mixture z with support in A^* ($z \in \Delta(A^*)$) and assume some other mixture $y \in \Delta(A^*)$ makes everyone weakly better off than z : we will show $y = z$, which implies the statement.

We keep in mind that for any two a, b in A^* there is some agent i who likes a but not b , because a and b are not clones. Write the ordered set A^* as $\{1, \dots, K\}$ and apply this remark to the first two agents: some agent i likes 1 but not 2, hence i likes only 1 and $u_i \cdot z \leq u_i \cdot y$ implies $z_1 \leq y_1$. Some agent j likes 2 but not 3, hence j likes 1, 2 or just 2, so $u_j \cdot z \leq u_j \cdot y$ is either $z_{12} \leq y_{12}$ or $z_2 \leq y_2$ and either way we deduce $z_{12} \leq y_{12}$. Similarly there is some k who likes 3 but not 4, so $u_k \cdot z \leq u_k \cdot y$ means that at least one of z_3, z_{23} , and z_{123} increases weakly and inequality $z_{123} \leq y_{123}$ follows in each case. An induction argument gives

$$z_{12\dots k} \leq y_{12\dots k} \text{ for all } k, 1 \leq k \leq K.$$

The symmetric argument starting from outcome K gives

$$z_{k(k+1)\dots K} \leq y_{k(k+1)\dots K} \text{ for all } k, 1 \leq k \leq K$$

and the desired conclusion $y = z$ follows. \square

Since RP mixes over Pareto optimal outcomes [16], it is efficient if the agent preferences satisfy condition (ii). Hence, under the condition (ii), the impossibility result from [16] disappears.

4 EXCLUDABLE STRATEGYPROOFNESS AND THE EGALITARIAN RULE

We start with the familiar prior-free incentive compatibility requirement that misreporting one’s preferences is never profitable if no agent can coordinate this move with other agents.⁴ Notation: upon replacing in the profile u the coordinate u_i by another $u'_i \in \Omega$, the resulting profile is $(u|u'_i)$.

$$\textbf{Strategyproofness (SP): } u_i \cdot f(M) \geq \max_{z' \in f(N, A, (u|u'_i))} u_i \cdot z' \text{ for all } M, i \text{ and } u'_i.$$

The simplest strategyproof rule adapts approval voting to our model: it selects only those outcomes liked by the largest number of agents. Write $\Phi^p(M)$ for the set of utility profiles implemented by pure outcomes in A : $\Phi^p(M) = \{U \in [0, 1]^N \mid \exists a \in A \forall i \in N, U_i = u_{ia}\}$. We use $avg(Y)$ to denote the average of the elements from the set Y of utility profiles, i.e. $avg(Y) = \frac{1}{|Y|} \sum_{y \in Y} y$. The utilitarian rule is defined as follows.

$$\textbf{Utilitarian rule (UTIL): } F^{ut}(M) = avg\{\arg \max_{U \in \Phi^p(M)} U_N\}.$$

UTIL admits a linear-time algorithm. It finds all the different utility vectors which result from pure outcomes liked by the largest number of agents, and returns the uniform lottery on this set of utility profiles. Note that this rule deliberately treats a problem with two identical columns exactly as the reduced problem where only one column remains.

The careful reader can check that this defines a rule as defined before (welfare-wise singleton, anonymous and neutral), one that is efficient and strategyproof. However, UTIL ignores minority opinions entirely so it fails to address the normative concerns described in the introduction.

If an agent gets a fair $1/n$ -th share of total decision power, she will use it on an outcome she likes. We take the following lower bound on individual welfare as the first test that mixing is fair:

$$\textbf{Individual Fair Share (IFS): } U = F(M) \implies U_i \geq \frac{1}{n} \text{ for all } M \text{ and all } i.$$

⁴Propositions 2 and 3 in [16] show that in our model group versions of SP are not compatible with efficiency, even in the ex post sense.

The main result of Bogomolnaia et al. [16] is that a rule cannot be together efficient, strategyproof, and satisfy IFS. Our first result is that this impossibility disappears if we weaken SP as explained below. To motivate this weakening, we adapt to our model the standard idea of equalizing individual utilities while respecting efficiency.

The lexicographic ordering in $[0, 1]^{(1, \dots, n)}$ maximizes the first coordinate, and when this is not decisive, the second one, and so on. For a utility profile $U \in [0, 1]^N$ the vector $U^* \in [0, 1]^{(1, \dots, n)}$ is obtained by rearranging its coordinates increasingly. Then the leximin ordering $>_{leximin}$ compares U^1 and U^2 in $[0, 1]^N$ exactly as the lexicographic ordering compares U^{1*} and U^{2*} in $[0, 1]^{(1, \dots, n)}$.

$$\text{Egalitarian rule (EGAL): } F^{eg}(M) = \arg \max_{U \in \Phi(M)} >_{leximin} .$$

This maximization yields a unique and efficient utility profile (see e. g., Lemma 1.1 in [37]). Anonymity and neutrality are clear. To check Individual Fair Share, pick for each agent i a pure outcome a_i she likes, and observe that the uniform average of the a_i -s ensures utility at least $1/n$ to each agent: therefore the egalitarian profile U^{eg} must have $U_1^{eg*} \geq 1/n$. For EGAL, the outcome can be computed in polynomial-time by solving at most $n + 1$ linear programs each with $|A|$ variables [9].

Here is the simplest problem where the rule EGAL is vulnerable to a misreport of preferences:

	A	a	b	c		A	a	b	c
	N				→	N			
true profile $u =$	1	1	1	0		1	1	$\tilde{0}$	0
	2	0	1	0		2	0	1	0
	3	0	0	1		3	0	0	1

At the true profile u outcome a is dominated and EGAL mixes b and c , $z = (0, \frac{1}{2}, \frac{1}{2})$. After the misreport by agent 1, outcome a no longer appears dominated and EGAL mixes equally the three outcomes, $\tilde{z} = (\frac{1}{3}, \frac{1}{3}, \frac{1}{3})$. Agent 1's utility raises from $1/2$ at z to $2/3$ at \tilde{z} , *because she can enjoy outcome b despite pretending not to*. The latter is avoidable if the public outcomes are excludable: based on reported preferences, the mechanism excludes agents from consuming outcomes they claim to dislike. Recall the discussion of this possibility in the examples of Section 1.

We use the notation $u_i \wedge u'_i$ for the coordinate-wise minimum of the two utility functions. For our setting, for any outcome a , $u_{ia} \wedge u'_{ia} = u_{ia} \cdot u'_{ia}$. The following incentives property captures the resulting weaker incentive compatibility requirement. Below, we interpret u_i as the true utility and u'_i as a contemplated misreport.

Excludable Strategyproofness (EXSP)

$$u_i \cdot f(M) \geq \max_{z' \in f(N, A, (u^i u'_i))} (u_i \wedge u'_i) \cdot z' \text{ for all } M, i \text{ and } u'_i.$$

To make the definition more explicit, we identify the true utility u_i by its *approval set* $L_i = \{a \in A \mid u_{ia} = 1\}$. Assume agent i 's misreport is L'_i . Let $L_i^0 = L_i \cap L'_i$ be the approval set of $u_i \wedge u'_i$. Define also $L_i^- = L_i \setminus L_i^0$ and $L_i^+ = L'_i \setminus L_i^0$, so that agent i pretends to like L_i^+ and to dislike L_i^- .

In this notation, EXSP reads:

$$z_{L_i^0 \cup L_i^-} \geq z'_{L_i^0} \text{ for all } z \in f(M), z' \in f(N, A, (u^i u'_i)).$$

It is useful to decompose EXSP in two properties. Both assume excludable setting (so an agent cannot consume goods she claims to dislike). In the first one agent i misreports only by inflating her approval set ($L_i = L_i^0$):

$$\text{SP}^+: u_i \cdot f(M) \geq \max_{z' \in f(N, A, (u^i u'_i))} u_i \cdot z' \text{ for all } M, i \text{ and } u'_i \text{ s.t. } u_i \leq u'_i.$$

Note that in this case $u_i \wedge u'_i = u_i$. In the second one, only by decreasing this set ($L_i^+ = \emptyset$):

$$\text{SP}^-: u_i \cdot f(M) \geq \max_{z' \in f(N, A, (u^i | u'_i))} u'_i \cdot z' \text{ for all } M, i \text{ and } u'_i \text{ s.t. } u'_i \leq u_i.$$

Note that in this case $u_i \wedge u'_i = u'_i$.

EXSP (“in excludable setting, no manipulation u'_i is profitable under given u_i ”) equals the combination of SP^+ and SP^- . This is clear by applying first SP^- under true u_i and manipulation $u_i \wedge u'_i$, and then SP^+ under true $u_i \wedge u'_i$ and manipulation u'_i (remember that an agent cannot consume goods she pretends to dislike).

However, SP not just implies SP^- and SP^+ but also SP^* :

$$\text{SP}^*: u_i \cdot f(M) \geq \max_{z' \in f(N, A, (u^i | u'_i))} u_i \cdot z' \text{ for all } M, i \text{ and } u'_i \text{ s.t. } u'_i \leq u_i.$$

This is the requirement that agent cannot manipulate by pretending she does not like some outcomes (shrinking her approval set), provided that she still can consume outcomes she pretends to dislike. The example above shows that EGAL violates SP^* .

Theorem 1 *The Egalitarian rule is efficient, excludable strategyproof, and guarantees Individual Fair Share.*

Proof

Efficiency and IFS are immediate, and are left to the reader. We will prove EXP.

Preliminary notation and remarks:

Let $Q \subseteq N$ and $U \in [0, 1]^{|Q|}$, $U = (U_1, \dots, U_{|Q|})$. We define $U^* = (U^{*1}, \hat{a}\check{A}\check{e}, U^{*r})$ as the vector of *distinct* coordinates U^{*k} of U arranged increasingly; so U^* may be of lower dimension than Q (i.e., $r \leq |Q|$), and $U^{*1} < \dots < U^{*r}$.

Fix a problem $M = (N, A, u)$. For any $Q \subseteq N$ and convex compact $C \subseteq \Delta(A)$ the linear projection on Q of the set of feasible utility profiles $\Phi(C) = \{U = u \cdot z \mid z \in C\}$ is convex and compact, so it admits a unique leximin optimal element that we write $F^{eg}(Q, C, u) \in [0, 1]^{|Q|}$. This extends the domain of the mapping F^{eg} to the cases where the set of feasible mixed outcomes is restricted to C . Note that we abuse notation by keeping u instead of its restriction to $Q \times A$.

Recall the procedure defining $U = F^{eg}(N, C, u)$. Start with $U^{*1} = \max_{z \in C} \min_{j \in N} \{u_j \cdot z\}$. Write N^1 for the set of agents achieving this minimum,⁵ $P^1 = N \setminus N^1$, and $C^1 = \{z \in C \mid u_j \cdot z = U^{*1} \text{ for all } j \in N^1\}$. We stop if $N^1 = N$, otherwise we set $U^{*2} = \max_{z \in C^1} \min_{j \in P^1} \{u_j \cdot z\}$. We let N^2 be the set of agents achieving U^{*2} , $P^2 = N \setminus (N^1 \cup N^2)$, and C^2 the subset of C^1 achieving U^{*2} in N^2 ; we stop if $P^2 = \emptyset$, otherwise we set $U^{*3} = \max_{z \in C^2} \min_{j \in P^2} \{u_j \cdot z\}$, and so on. We end up with a partition $N = \cup_{k=1}^K N^k$ such that U_i equals U^{*k} whenever $i \in N^k$.

Now, to prove EXSP it is enough to show separately SP^- and SP^+ . Fix an arbitrary $M = (N, A, u)$. An agent who likes all outcomes, $u_i = 1$, cannot benefit from any misreport; pick now $i \in N$ such that $u_{ia} = 0$ for at least one a , and a profile \tilde{u} identical to u for all $j \in N \setminus i$ and such that $u_i \not\leq \tilde{u}_i$ (so at least one 0 in u_i is changed to a 1). Let $U = F^{eg}(N, A, u)$ and $\tilde{U} = F^{eg}(N, A, \tilde{u})$ be implemented respectively by some lotteries z and \tilde{z} . We prove successively:

$$\tilde{U}_i \geq U_i \tag{2}$$

$$U_i = u_i \cdot z \geq u_i \cdot \tilde{z} \tag{3}$$

The first inequality implies SP^- (when i with true \tilde{U}_i reports U_i), the second gives SP^+ (when i with true U_i reports \tilde{U}_i).

⁵The sets N_k are uniquely defined, by convexity of $\Phi(C)$.

We clearly have $\tilde{U} \succeq_{\text{leximin}} U$, in particular $\tilde{U}^{*1} \geq U^{*1}$: this proves (2) if $U_i = U^{*1}$. Assume for the rest of the proof $U_i = U^{*\ell}$ where $\ell \geq 2$. We check first $\tilde{U}^{*1} = U^{*1}$. If $\tilde{U}^{*1} > U^{*1}$ we pick $\varepsilon \in (0, 1]$, and note that the mixture $z' = \varepsilon \tilde{z} + (1 - \varepsilon)z$ ensures $u_j \cdot z' > U^{*1}$ for all $j \in N \setminus i$. Indeed, $u_i \cdot \tilde{z} \geq \tilde{U}^{*1} > U^{*1}$ and $u_j \cdot z \geq U^{*1}$. For ε small enough we also have $u_i \cdot z' > U^{*1}$ because $u_i \cdot z > U^{*1}$. This contradicts the definition of U^{*1} .

Set $N^1 = \{j \mid U_j = U^{*1}\}$ and $\tilde{N}^1 = \{j \mid \tilde{U}_j = U^{*1}\}$. We use a similar argument to show next $N^1 \subseteq \tilde{N}^1$. If $j \in N^1$ and $u_j \cdot \tilde{z} > U^{*1}$, then for any $\varepsilon \in (0, 1]$ the mixture $z' = \varepsilon \tilde{z} + (1 - \varepsilon)z$ gives $u_k \cdot z' \geq U^{*1}$ for all $k \in N \setminus \{j, i\}$ and $u_j \cdot z' > U^{*1}$; for ε small enough we also have $u_i \cdot z' > U^{*1}$ (because $U_i = U^{*\ell} > U^{*1}$) and then z' guarantees exactly U^{*1} to a smaller set of agents than z , and strictly more to all others. This implies that $u \cdot z'$ leximin-dominates $u \cdot z$, a contradiction.

Similarly, the strict inclusion $N^1 \subsetneq \tilde{N}^1$ would imply that the vector \tilde{U} is strictly leximin-dominated by U , which we saw is not true.

So far we have shown that the maxi-minimization of feasible utilities – the first step in the algorithm defining the leximin solution – gives at u and \tilde{u} identical values U^{*1} and \tilde{U}^{*1} , and identical sets N^1 and \tilde{N}^1 . Now the second step of the algorithms, delivering U^{*2} , \tilde{U}^{*2} , and N^2 , \tilde{N}^2 , is the same maxi-minimization problem applied in both cases to $C^1 = \{z \in \Delta(A) \mid u_j \cdot z = U^{*1} \text{ for all } j \in N^1\}$ and $P^1 = N \setminus N^1$. Mimicking the above proof, we deduce that, if $U_i = U^{*2}$ then $\tilde{U}_i \geq U_i$, and if $U_i = U^{*\ell}$ for some $\ell \geq 3$, then $U^{*2} = \tilde{U}^{*2}$, $N^2 = \tilde{N}^2$. The induction argument establishing $U^{*k} = \tilde{U}^{*k}$, $N^k = \tilde{N}^k$ up to $k = \ell - 1$, and finally, (2) is now clear.

To prove (3) we compare the profiles $u \cdot z$ and $u \cdot \tilde{z}$. We just saw that they coincide on $N \setminus P^{\ell-1} = \cup_{k=1}^{\ell-1} N^k$, and that if a mixture guarantees utility U^{*k} to all agents in N^k for $k = 1, \dots, \ell - 1$, it cannot guarantee (at u) more than $U^{*\ell}$ to all agents in $P^{\ell-1}$: z and \tilde{z} are two such lotteries, so if $u_i \cdot \tilde{z} > u_i \cdot z = U^{*\ell}$, there is some $j \in P^{\ell-1}$ for whom $u_j \cdot \tilde{z} < U^{*\ell}$. But $\tilde{U}^{*\ell} \geq U^{*\ell}$ (because \tilde{U} weakly leximin-dominates U) and $\tilde{u}_j \cdot \tilde{z} = u_j \cdot \tilde{z} \geq \tilde{U}^{*\ell}$, thus we reach a contradiction. ■

5 STRICT PARTICIPATION AND UNANIMOUS FAIR SHARE

A striking feature of the Egalitarian rule is *Clone Invariance*: if at least one voter who shares agent's i preferences does vote, adding her own vote will not change the resulting mixture. Indeed, for any vector V , let \tilde{V} be obtained from it by adding an $(n + 1)$ -th coordinate repeating V_i . Fixing an agent i , the leximin ordering compares two utility profiles U and U' in the same way as \tilde{U} and \tilde{U}' . Thus the Egalitarian rule is oblivious to the size of support for a particular preference, an unpalatable feature in all the examples discussed in the introduction.

We now define two requirements capturing, each in a different way, the concern that numbers should matter. The first one is an incentive property.

Given a problem M and agent i , define $M(-i) = (N \setminus i, A, u_{-i})$ and

$$U_i(-i) = \max_{z \in f(M(-i))} u_i \cdot z.$$

In $M(-i)$, agent i does not vote, so the requirement for the outcome to be a singleton utility-wise does not apply, and she can have different utility from different outcomes in $f(M(-i))$. Thus, $U_i(-i)$ is the best she can hope for if she does not vote.

Participation (PART): $F_i(M) \geq U_i(-i)$ for all M and i .

The violation of Participation is commonly called the No Show Paradox [27]: a voter is better off abstaining from going to the polls. In the context of participatory budgeting, we want more: everyone should have a strict incentive to show up. Otherwise, many agents may stay home or put

a blank ballot, and the result of the vote will not give an accurate picture of the opinion profile.

Strict Participation (PART*):

$$F_i(M) \geq U_i(-i) \text{ and } \{U_i(-i) < 1 \implies F_i(M) > U_i(-i)\} \text{ for all } M \text{ and } i.$$

Under dichotomous preferences that we consider, strong SD-participation and SD-participation as studied by Brandl et al. [17] coincide with PART and very strong SD-participation coincides with PART*. A consequence of PART* is *Clone Responsiveness*: An agent is strictly better off if one or more agents with preferences identical to hers cast their vote. Thus the Egalitarian rule violates PART*, although it satisfies PART.⁶

The second axiom, in the spirit of cumulative voting (which gives minorities control over part of joint outcome by letting them concentrate vote on a few issues), allows groups of agents with identical preferences to pool their respective shares of decision power. This leads to the following strengthening of IFS, where we set again $U = F(M)$:

Unanimous Fair Share (UFS) :

$$\text{for all } S \subseteq N: \{u_i = u_j \text{ for all } i, j \in S\} \implies U_i \geq \frac{|S|}{n} \text{ for all } i \in S.$$

In the statement of UFS the unanimous group S can be a minority or a majority. However, unanimous preferences are much more likely in small than large groups, so this property will be more relevant in practice to minorities.

All three rules discussed in the next two sections meet Strict Participation and Unanimous Fair Share. Thus they cannot be both efficient and strategyproof. We start with two strategyproof rules.

6 INCENTIVE COMPATIBILITY AND FAIRNESS; THE CONDITIONAL UTILITARIAN RULE

We introduce two rules adapting to our model the familiar random dictator mechanism (see [29]). The difficulty is the treatment of indifferences: if an agent can dictate the outcome for a $1/n$ -th share of the time, what should she choose inside her approval set?

The first rule, introduced by Duddy [23], allocates to each agent $\frac{1}{n}$ -th of the decision power. She then uses it to pick the outcomes from her approval set which maximise utilitarian benefit of others. Recall that $\Phi^p(M)$ is the set of utility profiles implemented by *pure* outcomes. Consider the set $\Phi^p(M; i) = \{U \in \Phi^p(M) \mid U_i = 1\}$ of all the utility profiles corresponding to the approval set of agent i . Each agent spreads her share $\frac{1}{n}$ -th of the decision power equally between the profiles in $\Phi^p(M; i)$ with maximal support:

$$\text{Conditional Utilitarian (CUT) rule: } F^{cut}(M) = \frac{1}{n} \sum_{i \in N} \text{avg}\{U \mid U \in \arg \max_{U' \in \Phi^p(M; i)} U'_N\}.$$

Note that the CUT utility profile will not change if we remove “duplicated” pure outcomes. That is, without loss of generality, we can assume that no two pure outcomes are approved by the same set of agents. In this case, effectively, under CUT an agent spreads her decision power uniformly over outcomes in her approval set with the largest utilitarian support. We say that she “loads” those outcomes. CUT admits a linear-time algorithm to compute the outcome. In the approval set of each agent, we simply need to identify those liked by the largest number of other agents.

Remark 1. Our definition of the domain Ω allows for agents who like all outcomes, $u_i = 1^A$. The presence of such agents is of no consequence for the rules UTIL, EGAL, RP and MNP, but it does impact

⁶To show that EGAL satisfies PART, assume agent 1 does not vote. Define $U^* = \arg \max_{U \in \Phi(M)} \succ_{leximin} ; \bar{U}_{-1} = \arg \max_{U \in \Phi(M(-1))} \succ_{leximin}$ and $\bar{U}_1 = U_1(-1)$. Suppose $\bar{U}_1 > U_1^*$, then we have $\bar{U} = (\bar{U}_1, \bar{U}_{-1}) \succeq_{leximin} (\bar{U}_1, U^*)$ and $(\bar{U}_1, U^*) \succ_{leximin} U^*$, contradiction to maximality of U^* .

the mixture selected by the CUT rule, as such agents put their weight on the utilitarian outcomes (those with the largest support). Suppose we choose to exclude those agents in the definition of the CUT rule. This new rule will share the incentives and fairness properties discussed below.

The next rule uses a familiar hierarchical rule to resolve indifferences, that plays a critical role in probabilistic voting [6], as well as for assigning indivisible private goods ([1], [12]). Let $\Theta(N)$ be the set of strict orderings σ of N . For any $\sigma \in \Theta(N)$ the σ -Priority rule F^σ guarantees full utility to agent $\sigma(1)$; next to agent $\sigma(2)$ as well if 1 and 2 like a common outcome, else $\sigma(2)$ is deemed irrelevant; next to agent $\sigma(3)$ if she likes an outcome that all relevant agents before her like, else she is irrelevant; and so on.

Random Priority rule (RP)

$$F^{rp}(M) = \frac{1}{n!} \sum_{\sigma \in \Theta(N)} F^\sigma(M) \text{ where } F^\sigma(M) = \arg \max_{U \in \Phi(M)} \succ_{lexico(\sigma)}.$$

If mixtures in $\Delta(A)$ represent lotteries, the RP rule picks an ordering σ with uniform probability and computes U^σ . But in other interpretations, time shares or the distribution of other resources, this simple implementation is not available. We retain nevertheless the intuitive probabilistic terminology. The RP outcome is #P-complete to compute even under dichotomous preferences [6]. Therefore unless P=NP, it is unlikely that there exists an efficient algorithm for computing the exact RP outcome. For RP, it is even open whether there exists an FPRAS (Fully Polynomial-time Approximation Scheme) for computing the outcome shares/probabilities.

After checking that both rules are incentive compatible and fair, we compare them from the efficiency angle, and recap our discussion in Theorem 2. RP is strategyproof [12] and satisfies PART* [17]. For UFS, it is enough to observe that a member of coalition S is first in σ with probability $\frac{|S|}{n}$. We check now that CUT meets the same three properties. UFS is clear. We now show that CUT satisfies SP. Consider an agent i who has a like set $L_i = L_i^0 \cup L_i^-$ where L_i^- is possibly empty. Suppose she now reports her like set as $L_i' = L_i^0 \cup L_i^+$ where $L_i^+ \subseteq A \setminus L_i$. Agent i 's own contribution of $1/n$ to L_i can only decrease if she reports L_i' . We now focus on the effect of i reporting L_i' on to any other agent $j \neq i$. Note that j 's contribution to L_i either remains the same or decreases. Therefore agent i gets no benefit by reporting L_i' .

Next, we verify that CUT satisfies PART*. Fix a problem M , an agent i , and for every $j \in N \setminus i$ let B_j be the set of outcomes agent j loads in problem $M(-i)$. Set $N^+ = \{j \in N \setminus i \mid B_j \cap L_i \neq \emptyset\}$ and $N^- = N \setminus (N^+ \cup i)$. Before participating agent i 's utility was

$$\frac{1}{n-1} \sum_{j \in N^+} \lambda_j \text{ where } \lambda_j = \frac{|B_j \cap L_i|}{|B_j|}.$$

After i shows up every j in N^+ loads only $B_j \cap L_i$, and agents in N^- may give some of their load to L_i therefore i 's utility is at least $\frac{1}{n}(1 + |N^+|)$. The inequality

$$\frac{1}{n-1} \sum_{j \in N^+} \lambda_j \leq \frac{|N^+|}{n-1} \leq \frac{1}{n}(1 + |N^+|).$$

proves PART. And both inequalities are equalities if and only if each $\lambda_j = 1$ and $|N^+| = n-1 \Leftrightarrow N^+ = N \setminus i$; the latter implies that i 's utility is already 1 in $M(-i)$.

Example (3.1) in Section 3 shows that both RP and CUT are inefficient. Under the CUT rule agents 1, 2 and 5 load only d , while agent 3 spreads her load between a and b , and agent 4 between a and c , resulting in the mixture $z^{cut} = (\frac{1}{5}, \frac{1}{10}, \frac{1}{10}, \frac{3}{5}, 0)$. Under RP we get $z^{rp} = (\frac{1}{5}, \frac{1}{6}, \frac{1}{6}, \frac{7}{15}, 0)$; for instance b is selected in two cases only: if 3 is first, and 5 comes before 4 (probability $\frac{1}{10}$), or 5 is

first and 3 is first among 1, 2, 3 (probability $\frac{1}{15}$). As noted at the end of Section 3, shifting the weight of b and c to a and d is a Pareto improvement. Clearly, then, z^{RP} is more inefficient than z^{cut} .

In our next example, with $n = 6$ and $|A| = 5$,

	A	a	b	c	d	e	
N							
1		1	0	0	1	0	
2		1	0	0	0	1	
3		0	1	0	1	0	(4)
4		0	1	0	0	1	
5		0	0	1	1	0	
6		0	0	1	0	1	

the CUT rule selects the efficient mixture $z^{cut} = (0, 0, 0, \frac{1}{2}, \frac{1}{2})$ and $U_i^{cut} = 0.5$ for all i , while RP picks $z^{RP} = (\frac{1}{9}, \frac{1}{9}, \frac{1}{9}, \frac{1}{3}, \frac{1}{3})$ and $U_i^{RP} = 0.44$ for all i : thus z^{cut} is strictly Pareto superior to z^{RP} . The reverse situation cannot happen: the RP mixture never Pareto dominates the CUT one. This follows because in all problems, total utility under RP is at most that under CUT: $U_N^{RP} \leq U_N^{cut}$. Indeed U^{cut} is the uniform average of profiles $U(i)$ maximizing total utility in $\Phi^P(M; i)$, and for each ordering σ where i is first, the corresponding profile U^σ is in $\Phi^P(M; i)$ as well.⁷

We finally prove that *the CUT rule is efficient more often than RP*: whenever RP picks an efficient mixture, so does CUT. Observe first that both rules only give weight to undominated pure outcomes. In the case of RP every such outcome a has a positive weight, because it is selected whenever the set of agents who like a has the highest priority. Thus the support of the RP mixture is exactly the set of all undominated columns. Therefore RP selects an efficient mixture if and only if all mixtures with support in this set are efficient as well.⁸ The claim follows because the CUT rule is also a combination of undominated columns.

The next theorem reinforces the advantage of CUT over RP in terms of efficiency.

Theorem 2

- i) Both rules CUT and RP are strategyproof and meet Strict Participation and Unanimous Fair Share.
- ii) Total utility at the CUT mixture is never below that at the RP mixture, and the former may Pareto dominate the latter. If RP picks an efficient mixture at some problem M , so does CUT.
- iii) The CUT rule is $(1 - \epsilon^{cut}(n))$ -inefficient with $\epsilon^{cut}(n) = O(n^{-\frac{1}{3}})$ and for all $n \geq 5$ we have

$$\epsilon^{cut}(n) \geq \frac{1}{n} + (1 - \frac{1}{n^{\frac{1}{3}}}) \frac{3}{n^{\frac{1}{3}}}. \tag{5}$$

The RP rule is $(1 - \epsilon^{RP}(n))$ -inefficient with $\epsilon^{RP}(n) \leq O(\frac{\ln(n)}{n})$.

Recall from the lemma in Section 3 that both CUT and RP are efficient if $n \leq 4$. For small values of n , the lower bound (5) implies a high guaranteed efficiency of CUT, a lower bound on $\epsilon^{cut}(n)$, and the computations in Step 2 of the proof below yield a much smaller worst case efficiency of RP, an upper bound on $\epsilon^{RP}(n)$:

n	6	8	12	32	64	1024	16384
$\epsilon^{cut} \geq$	91%	87%	82%	68%	58%	27%	11%
$\epsilon^{RP} \leq$	83%	72%	64%	40%	24%	3%	0.12%

⁷A consequence of this remark is that CUT and RP pick the same utility profile at problem M if and only if all undominated outcomes of M are liked by the same number of agents.

⁸If some mixture z is efficient ex-post but not ex-ante, then it will be present in each RP mixture with a positive weight, and prevent RP from being efficient.

Statements *ii*) and *iii*) make a very strong case that in our model the CUT rule is a much more efficient interpretation of the random dictator idea than RP. Statement *iii*) is proved in the appendix.

7 EFFICIENCY AND FAIRNESS; THE MAX NASH PRODUCT RULE

Our last rule of interest is a familiar compromise between the Utilitarian and Egalitarian rules:

$$\text{Max Nash Product rule (MNP): } F^{mnp}(M) = \arg \max_{U \in \Phi(M)} \sum_{i \in N} \ln U_i.$$

This rule is well-defined because it solves a strictly convex program, and obviously efficient. We make a few remarks about the computational aspects of MNP. In contrast to the other rules discussed above, the MNP outcome can be irrational (see, e.g., [2]). The problem is to maximize a convex objective $\sum_{i \in N} \log(u_i \cdot z)$ where z is a feasible mixture. Using convex optimization techniques (such as the ellipsoid method; see the discussion by Vazirani [46]), a lottery that approximates the objective value of this convex program within an additive term of $\epsilon > 0$ can be computed in time that is polynomial in the size of the profile and $1/\epsilon$.

Recall that Unanimous Fair Share offers welfare guarantees only to coalitions of agents with identical preferences (clones). The first of our two new “Fair Share” axioms applies, much more generally, to any group who can find at least one outcome that everyone likes.

Average Fair Share (AFS)

$$\text{for all } S \subseteq N: \{\exists a \in A : u_{ia} = 1 \text{ for all } i \in S\} \implies \frac{1}{|S|} U_S \geq \frac{|S|}{n}.$$

In words, if there is an object (say, a), approved by all members of S , then the average utility of this group should be at least as large as its relative size $\frac{|S|}{n}$. Note that if all members of S would approve only of a , then UFS requirement would imply everyone in the group should get at least $\frac{|S|}{n}$. Therefore AFS is easily seen as a further strengthening of IFS whereby a principle applied to individuals is applied to groups.

The next property conveys the idea that, as each agent is endowed with $1/n$ -th of total decision power, any coalition of size s can cumulate these shares and impose that a mixture of their choice be chosen with probability at least $\frac{s}{n}$:

Core Fair Share (CFS)

$$\text{for all } S \subseteq N : \nexists z \in \Delta(A) \text{ s.t. } \forall i \in S, U_i \leq \frac{|S|}{n} (u_i \cdot z) \text{ and } \exists i, U_i < \frac{|S|}{n} (u_i \cdot z).$$

This is a familiar core stability property which is widely used in cooperative game theory [22]. Note that CFS is not logically related to AFS. That UFS follows from either AFS or CFS is clear because we only consider anonymous rules. Applying CFS to $S = N$ implies that the rule is efficient, therefore neither the CUT nor the RP rule meets CFS. In Example 3.1, the AFS property selects uniquely the mixture maximizing the Nash Product.⁹ therefore CUT and RP fail AFS as well.

Theorem 3

- i*) The MNP rule is efficient and meets Strict Participation, Average Fair Share, and Core Fair Share.
- ii*) The MNP rule is not excludable strategyproof.

Proof

⁹We leave the proof to the reader.

We first prove AFS and CFS. The separation inequality capturing the optimality of the Max Nash utility profile $U^* = F^{mnp}(M)$ at problem M writes as follows:¹⁰

$$\sum_{i \in N} \frac{U_i}{U_i^*} \leq \sum_{i \in N} \frac{U_i^*}{U_i^*} = n \text{ for all } U \in \Phi(M) \quad (6)$$

Fix $S \subseteq N$ and combine (6) with Cauchy's inequality as follows

$$\begin{aligned} nU_S^* &\geq \left(\sum_{i \in S} \frac{U_i}{U_i^*} \right) \cdot \left(\sum_{i \in S} U_i^* \right) \geq \left(\sum_{i \in S} \sqrt{U_i} \right)^2 \implies \\ U_S^* &\geq \frac{1}{n} \max_{U \in \Phi(M)} \left(\sum_{i \in S} \sqrt{U_i} \right)^2 \end{aligned} \quad (7)$$

The AFS property follows, because if there is some $a \in A$ such that $u_{ia} = 1$ for all $i \in S$, the maximum on the right hand side of 7 is $|S|^2$. To check CFS we assume there is a mixture z such that $U_i^* \leq \frac{|S|}{n} (u_i \cdot z)$ for all $i \in S$ and use again (6) to compute:

$$n \geq \sum_{i \in S} \frac{u_i \cdot z}{U_i^*} \geq \frac{n}{|S|} \sum_{i \in S} \frac{U_i^*}{U_i^*} = n$$

therefore none of the inequalities $U_i^* \leq \frac{|S|}{n} (u_i \cdot z)$ can be strict.

The proof that MNP satisfies PART* is more technical and is relegated to the appendix. The proof that the MNP rule fails EXSP is also relegated to the appendix. \square

Remark 3. Let U_N^{mnp} be the total utility for the Max Nash allocations and U_N^* be the maximal achievable utility (utilitarian welfare). Then it can be proved that $\frac{U_N^{mnp}}{U_N^*} \geq \frac{1}{n} U_N^*$. So if the average utilitarian utility is $\lambda \in [0, 1]$, then the average utility of the outcome under MNP is at least λ^2 .

Remark 4. Another version of the group fair share requirement is proposed by Bogomolnaia et al. [15]. The same concept was independently proposed by Duddy [23] who referred to it simply as *proportional share* [23]. For the sake of consistency with our other notions, we will refer to it as *Group Fair Share* (GFS). Writing u^{*S} for the maximum of all utility functions in S ($u_a^{*S} = \max_{i \in S} u_{ia}$), this condition requires $U^{*S} \geq \frac{|S|}{n}$ for all S . It is clearly stronger than UFS, but strictly weaker than CFS or AFS. Both CUT and RP satisfy GFS.

Remark 5. It has been mentioned as an open problem, in the more general voting model with weak preference orders, whether there exists some rule that satisfies 'Very Strong Stochastic Dominance Participation' and 'Stochastic Dominance Efficiency' for weak orders [17, 19]. Because MNP satisfies both Strict Participation and Efficiency, we see that this question is resolved at least for the case of dichotomous preferences.

8 EXPERIMENTS

We ran experiments for small numbers of agents and outcomes. We focussed on the ratio of utilitarian welfare of the result to the maximum utilitarian welfare. The ratio gives a lower bound on ε as used to define the inefficiency of a mixture. For each combination of n and $|A|$ in $\{3, 5, 7, 10, 15, 20\}$ and for each rule, we examined under the impartial culture (1) the minimum of this ratio and (2) its average. For RP, we did not run the experiments for $n = 15$ and 20 because the computation becomes very slow. This illustrates the relative computational infeasibility of RP when we want the exact mixture, even for a relatively modest number of agents. The results are in the appendix.

¹⁰This is simply the first order optimality condition.

As the number of agents increase, the ratios start to get worse. But for a fixed number of agents, the ratios do not necessarily get worse as we increase the number of outcomes. We note that CUT seems to fare marginally but consistently better than MNP, RP, and EGAL in the utilitarian metric. This is especially so when we consider the average rather than the worst ratios. We note that MNP rule's fairness constraints also lead to loss of utilitarian welfare. On the other hand, it has been shown that on certain real-world participatory budgeting datasets, core fair outcomes often coincide with welfare maximizing ones [24]. Since the objective of EGAL is diametrically opposed to utilitarian objectives, it is not surprising that EGAL fares the worst in the utilitarian metric among the rules we consider. In particular its worst case ratios drop rapidly as we increase the number of agents and outcomes.

Since CUT performs better than the other competing rules in terms of welfare, we zoomed into the case of CUT and computed its average and worst case inefficiency n in $\{5, 10, 20, 30, 40, 50\}$ and $|A|$ in $\{5, 610, 1530, 40, 50, 60\}$ where agents approve a given percentage of outcomes where the percentages are 10, 25, 33, 50, 75. For each of the combinations of parameters, 200000 draws are taken. The level of inefficiency is negligible for all the combinations. By sampling using several probabilistic models, [34] confirms that CUT achieves very high utilitarian welfare on average.

9 DECENTRALIZATION

We also introduce a *Decentralization* (DEC) property for polarized societies. Say the agents and the deterministic outcomes are color-coded with the same set of colors: we call a profile of preferences *polarized* if each agent only likes outcomes of her own color. The requirement is that if an agent is red, the number of green agents will matter to her but not their preferences inside green outcomes. This natural independence property adds to the appeal of the MNP rule, but also of the CUT and RP rules.

Consider a problem $M = (N, A, u)$ and two partitions $\Gamma = (N^k)_{k=1}^K$ and $\Lambda = (A^k)_{k=1}^K$ of N and A respectively. We call this problem *polarized along the partitions* Γ, Λ if $u_{ia} = 0$ whenever $i \in N^k, a \in A^{k'},$ and $k \neq k'$. Then if u^k is the restriction of u to $N^k \times A^k$, problem M is captured by its K subproblems $M^k = (N^k, A^k, u^k)$. We write $\Pi(\Gamma, \Lambda)$ the set of polarized problems.

Decentralization (DEC): For all M, M' , for all u, u' , and for any Γ, Λ and k

$$\{M, M' \in \Pi(\Gamma, \Lambda) \text{ and } u_{ia} = u'_{ia} \text{ if } i \in N^k, a \in A^k\} \implies F_i(M) = F_i(M') \text{ for } i \in N^k.$$

Decentralization can be viewed as satisfying an extension of party-list proportional representation from multi-winner voting [25]. Combined with the UFS property, it implies that in a polarized problem, each colored subset N^k chooses the distribution in $\Delta(A^k)$ as if other colors were not present, then the selected outcome in $f(M^k)$ is weighted down in proportion of the size of N^k .

Theorem 4 *The Max Nash, Conditional Utilitarian, and Random Priority rules meet Decentralization. Moreover, for any polarized problem $M \in \Pi(\Gamma, \Lambda)$ they satisfy*

$$F(M) = \sum_{k=1}^K \frac{|N^k|}{n} F(M^k) \quad (8)$$

where the profile $F(M^k)$ is filled with zeros outside M^k .

On the other hand, the Utilitarian and Egalitarian rules violate DEC. Consider the two polarized problems along the partition $\{1\} \cup \{2, 3\}$:

$$M : \begin{array}{cccc} 1 & 1 & 0 & 0 \\ 2 & 0 & 1 & 0 \\ 3 & 0 & 0 & 1 \end{array} \quad M' : \begin{array}{cccc} 1 & 1 & 0 & 0 \\ 2 & 0 & 1 & 0 \\ 3 & 0 & 1 & 1 \end{array} .$$

Both UTIL and EGAL choose $z = (\frac{1}{3}, \frac{1}{3}, \frac{1}{3})$ at M , but at M' they pick respectively $z' = (0, 1, 0)$ and $z'' = (\frac{1}{2}, \frac{1}{2}, 0)$, in contradiction of DEC.

10 CONCLUSION AND OPEN QUESTIONS

We compared the relative merits of some well-known rules (EGAL, RP, MNP) and of an (essentially) new one (CUT), for the model of probabilistic/fractional voting under dichotomous preferences. The two rules that are especially desirable in the instances where protection of minorities and participation concerns matter most are CUT and MNP. The *Conditional Utilitarian* rule is strongly incentive compatible, but in extreme cases it may be severely inefficient (see statement iii) in Theorem 2). The *Max Nash Product* rule is efficient and gives much better guarantees to groups of agents than CUT, but it even fails the weak form of strategyproofness where outcomes are excludable.

Our study helps identify two especially interesting open questions. We know from [16] that efficiency, Individual Fair Share and strategyproofness are incompatible. If we are content to achieve only the excludable version of strategyproofness, this incompatibility disappears, and the Egalitarian rule is an example. The unpalatable feature of this rule is that it pays no attention to clones (subgroups of agents with identical preferences) hence offers no protection to sizable minorities. But can a rule combine efficiency, excludable strategyproofness and Strict Participation; or efficiency, excludable strategyproofness and Unanimous Fair Share? Such a rule would be a serious new contender in our fair mixing model.

Bogomolnaia et al. [15, 16] defined and studied a family of welfarist rules directly borrowed from classical social choice theory. Fix an increasing, strictly concave, and continuous function h on $[0, 1]$. A rule in the sense of Definition 1 is obtained by maximizing the sum of individual utilities weighted by h : **h -rule**: $f(M) = \arg \max_{U \in \Phi(M)} \sum_{i \in N} h(U_i)$. This maximization has a unique solution in $\Phi(M)$. The MNP rule is of course a paramount example. All h -rules are efficient, and by mimicking Step 2 in the proof of Theorem 3, we see that they satisfy PART* provided $h'(0) = \infty$. They satisfy (resp. fail) IFS if h is at least as concave as (resp. less concave than) the *log* function; but MNP is the only h -rule meeting UFS (these two facts are already proven in [15]). Finally, all h -rules fail EXSP and only MNP meets DEC. Thus they do not add much to our axiomatic discussion. However, once we observe that the EGAL and UTIL rules are the two end points of the family of h -rules¹¹ the following intriguing fact emerges: most h -rules meet PART* but neither EGAL nor UTIL does; EGAL and UTIL meet EXSP, but none of the h -rules does.

ACKNOWLEDGEMENTS

Haris Aziz is supported by a UNSW Scientia Fellowship. The authors thank Edward Lee and Yekaterina Rzhewskaya for assistance with the experiments. The comments of seminar participants at Seoul National University, Université Paris 1, University Technology Sydney, and the Paris School of Economics are gratefully acknowledged. The authors thank the detailed comments by the anonymous reviewers of ACM EC and TEAC.

¹¹When h converges pointwise to a linear function, e.g. $h(x) = x^q$ with $q \uparrow 1$, the h -rule converges pointwise to UTIL; when h becomes infinitely concave, e.g. $h(x) = -x^q$ with $q \downarrow -\infty$, it converges pointwise to EGAL.

REFERENCES

- [1] A. Abdulkadiroğlu and T. Sönmez. 1998. Random Serial Dictatorship and the Core from Random Endowments in House Allocation Problems. *Econometrica* 66, 3 (1998), 689–701.
- [2] Airiau, S.; Aziz, H.; Caragiannis, I.; Kruger, J.; Lang, J.; and Peters, D. 2019. Portioning using ordinal preferences: Fairness and efficiency. In *Proceedings of the 28th International Joint Conference on Artificial Intelligence (IJCAI)*.
- [3] H. Aziz. 2013. Maximal Recursive Rule: A New Social Decision Scheme. In *Proc. of 23rd IJCAI*. AAAI Press, 34–40.
- [4] H. Aziz. 2019. A Probabilistic Approach to Voting, Allocation, Matching, and Coalition Formation. In *The Future of Economic Design*, J.-F. Laslier, H. Moulin, R. Sanver, and W. S. Zwicker (Eds.). Springer.
- [5] H. Aziz, F. Brandl, F. Brandt, and M. Brill. 2018. On the Tradeoff between Efficiency and Strategyproofness. *Games and Economic Behavior* 110 (2018), 1–18.
- [6] H. Aziz, F. Brandt, and M. Brill. 2013. The Computational Complexity of Random Serial Dictatorship. *Economics Letters* 121, 3 (2013), 341–345.
- [7] H. Aziz, M. Brill, V. Conitzer, E. Elkind, R. Freeman, and T. Walsh. 2017. Justified Representation in Approval-Based Committee Voting. *Social Choice and Welfare* 48, 2 (2017), 461–485.
- [8] H. Aziz, E. Elkind, S. Huang, M. Lackner, L. Sánchez-Fernández, and P. Skowron. On the complexity of Extended and Proportional Justified Representation. In *Proceedings of the 32nd AAAI Conference on Artificial Intelligence (AAAI)*, pages 902–909. AAAI Press, 2018.
- [9] H. Aziz and P. Stursberg. 2014. A Generalization of Probabilistic Serial to Randomized Social Choice. In *Proc. of 28th AAAI Conference*. AAAI Press, 559–565.
- [10] G. Benade, S. Nath, A. D. Procaccia, and N. Shah. 2017. Preference Elicitation For Participatory Budgeting. In *Proc. of 31st AAAI Conference*. AAAI Press, 376–382.
- [11] A. Bogomolnaia. 2018. The most ordinally egalitarian of random voting rules. *Journal of Public Economic Theory* 20, 2 (2018), 271–276.
- [12] A. Bogomolnaia and H. Moulin. 2001. A New Solution to the Random Assignment Problem. *Journal of Economic Theory* 100, 2 (2001), 295–328.
- [13] A. Bogomolnaia and H. Moulin. 2004. Random Matching under Dichotomous Preferences. *Econometrica* 72, 1 (2004), 257–279.
- [14] A. Bogomolnaia, H. Moulin, F. Sandomirskiy, and E. Yanovskaya. 2017. Competitive division of a mixed manna. *Econometrica* 85, 6 (2017), 1847–1871.
- [15] A. Bogomolnaia, H. Moulin, and R. Stong. 2002. Collective choice under dichotomous preferences. (2002).
- [16] A. Bogomolnaia, H. Moulin, and R. Stong. 2005. Collective choice under dichotomous preferences. *Journal of Economic Theory* 122, 2 (2005), 165–184.
- [17] F. Brandl, F. Brandt, and J. Hofbauer. 2015. Incentives for Participation and Abstention in Probabilistic Social Choice. In *Proc. of 14th AAMAS Conference*. IFAAMAS, 1411–1419.
- [18] F. Brandl, F. Brandt, and H. G. Seedig. 2016. Consistent Probabilistic Social Choice. *Econometrica* 84, 5 (2016), 1839–1880.
- [19] F. Brandt. 2017. Rolling the Dice: Recent Results in Probabilistic Social Choice. In *Trends in Computational Social Choice*, U. Endriss (Ed.). AI Access, Chapter 1, 3–26.
- [20] Y. Cabannes. 2004. Participatory budgeting: a significant contribution to participatory democracy. *Environment and Urbanization* 16, 1 (2004), 27–46.
- [21] I. Caragiannis, D. Kurokawa, H. Moulin, A. D. Procaccia, N. Shah, and J. Wang. 2016. The Unreasonable Fairness of Maximum Nash Welfare. In *Proc. of 17th ACM-EC Conference*. 305–322.
- [22] T. Driessen. 1988. *Cooperative Games, Solutions and Applications*. Kluwer.
- [23] C. Duddy. 2015. Fair sharing under dichotomous preferences. *Mathematical Social Sciences* 73 (2015), 1–5.
- [24] B. Fain, A. Goel, and K. Munagala. 2016. The Core of the Participatory Budgeting Problem. In *Web and Internet Economics - 11th International Conference, WINE 2016, Montreal, Canada, December 11-14, 2016, Proceedings*. 384–399.
- [25] P. Faliszewski, P. Skowron, A. Slinko, and N. Talmon. 2017. Multiwinner Voting: A New Challenge for Social Choice Theory. In *Trends in Computational Social Choice*, U. Endriss (Ed.). Chapter 2.
- [26] P. C. Fishburn. 1984. Probabilistic Social Choice Based on Simple Voting Comparisons. *Review of Economic Studies* 51, 4 (1984), 683–692.
- [27] P. C. Fishburn and S. J. Brams. 1983. Paradoxes of Preferential Voting. *Mathematics Magazine* 56, 4 (1983), 207–214.
- [28] T. Fluschnik P. Skowron, M. Triphaus and K. Wilker. 2019. Fair Knapsack. In *Proc. of 33rd AAAI Conference on Artificial Intelligence*. AAAI-19, 1941–1948.
- [29] A. Gibbard. 1977. Manipulation of schemes that mix voting with chance. *Econometrica* 45, 3 (1977), 665–681.
- [30] J. Gordon. 1994. Institutions as Relational Investors: A New Look at Cumulative Voting. *Columbia Law Review* 94, 1 (1994), 124–192.
- [31] U. Grandi. 2017. Agent-Mediated Social Choice. In *The Future of Economic Design*, J.F. Laslier, H. Moulin, R. Sanver, and W. Zwicker (Eds.). Springer.

- [32] J. Hughes and G. Sasse. 2003. Monitoring the monitors : EU enlargement conditionality and minority protection in the CEECs. *Journal on ethnopoltics and minority issues in Europe* 1 (2003), 35.
- [33] G. Laffond, J.-F. Laslier, and M. Le Breton. 1993. The Bipartisan Set of a Tournament Game. *Games and Economic Behavior* 5, 1 (1993), 182–201.
- [34] M. Michorzewski, D. Peters and P. Skowron. 2020. Price of Fairness in Budget Division and Probabilistic Social Choice. In *Proc. of 34rd AAAI Conference on Artificial Intelligence*. AAAI-20.
- [35] H. Moulin. 1981. The Proportional Veto Principle. *Review of Economic Studies* 48, 3 (1981).
- [36] H. Moulin. 1982. Voting with Proportional Veto Power. *Econometrica* 50, 1 (1982), 145–162.
- [37] H. Moulin. 1988. *Axioms of Cooperative Decision Making*. Cambridge University Press.
- [38] H. Moulin. 2003. *Fair Division and Collective Welfare*. The MIT Press.
- [39] H. Moulin. 2019. Fair Division in the Internet Age. *Annual Review of Economics* 11 (2019), 1–37.
- [40] J. F. Nash. 1950. The Bargaining Problem. *Econometrica* 18, 2 (1950), 155–162.
- [41] R. La Porta, F. Lopez de Silanes, A. Schleifer, and R. Vishny. 2000. Investor protection and corporate governance. *Journal of Financial Economics* 58, 1–2 (2000), 3–27.
- [42] J. Sawyer and D. MacRae. 1962. Game Theory and Cumulative Voting in Illinois: 1902-1954. *The American Political Science Review* 56, 4 (1962), 936–946.
- [43] H. Steinhaus. 1948. The problem of fair division. *Econometrica* 16 (1948), 101–104.
- [44] W. Thomson. 2016. Introduction to the Theory of Fair Allocation. In *Handbook of Computational Social Choice*, F. Brandt, V. Conitzer, U. Endriss, J. Lang, and A. D. Procaccia (Eds.). Cambridge University Press, Chapter 11.
- [45] H. R. Varian. 1974. Equity, Envy, and Efficiency. *Journal of Economic Theory* 9 (1974), 63–91.
- [46] Vazirani, V. V. 2012. Rational convex programs and efficient algorithms for 2-player Nash and nonsymmetric bargaining games. *SIAM Journal on Discrete Mathematics* 26(3):896–918.
- [47] G. Young. 1950. The case for cumulative voting. *Wisconsin Law Review* 49–56 (1950).

11 APPENDIX

Proof of Theorem 2, part iii)

Step 1 Worst case inefficiency of CUT

Step 1.a: We construct a problem with large n where the CUT profile is $(1 - O(n^{-\frac{1}{3}}))$ -inefficient. We fix N of size n , a partition $N = N_1 \cup N_2$, and an integer p such that

$$p < n_1, n_2 \text{ and } n_1 \text{ divides } (p-1)n_2 \text{ where } n_i = |N_i|, i = 1, 2.$$

Problem M has $2n_2 + 1$ outcomes labeled as $A = \{a\} \cup B \cup C$, where $B = \{b_j, j \in N_2\}$ and $C = \{c_j, j \in N_2\}$. Setting $(p-1)n_2 = qn_1$, each agent $i \in N_1$ likes a , exactly q outcomes in B , and none in C ; and each $j \in N_2$ dislikes a , likes only outcome b_j in B , and exactly $p-1$ outcomes in C . Moreover, the problem is symmetric in N_1 and in N_2 , which can be achieved by arranging cyclically the approval sets of the N_1 agents in B and the approval sets of the N_2 agents in C . Here is an example with $n_1 = n_2 = 5, p = 4$ and $q = 3$, and the top five agents form N_1 :

a	b_1	b_2	b_3	b_4	b_5	c_1	c_2	c_3	c_4	c_5
1	1	0	0	1	1	0	0	0	0	0
1	1	1	0	0	1	0	0	0	0	0
1	1	1	1	0	0	0	0	0	0	0
1	0	1	1	1	0	0	0	0	0	0
1	0	0	1	1	1	0	0	0	0	0
0	1	0	0	0	0	1	0	0	1	1
0	0	1	0	0	0	1	1	0	0	1
0	0	0	1	0	0	1	1	1	0	0
0	0	0	0	1	0	0	1	1	1	0
0	0	0	0	0	1	0	0	1	1	1

Note that each outcome b_j is liked by exactly p agents, all but one of them in N_1 , and each c_j is liked by exactly $p-1$ agents, all in N_2 .

Under the CUT rule, each agent $i \in N_1$ loads only a because $n_1 > p$, so $z_a = \frac{n_1}{n}$, and each $j \in N_2$ loads only b_j so $z_{b_j} = \frac{1}{n}$; there is no weight on C . The total utility in each group is

$$U_{N_1} = \frac{(n_1)^2}{n} + \frac{n_2}{n}(p-1); U_{N_2} = \frac{n_2}{n}$$

and by the symmetries these are equally shared in N_1 and N_2 respectively.

Now consider the mixture z' : $z'_a = \frac{2}{3}$, $z'_{c_j} = \frac{1}{3n_2}$ for all $j \in N_2$, and zero weight on B , resulting in the total utilities

$$U'_{N_1} = \frac{2}{3}n_1; U'_{N_2} = \frac{1}{3}(p-1)$$

again equally shared in each N_i .

For n large enough we can pick n_1 and p such that $n_1 \approx n^{\frac{2}{3}}$ and $p-1 \approx n^{\frac{1}{3}}$ (if n is a cube these values are exact and $q = n^{\frac{2}{3}} - n^{\frac{1}{3}}$) so that $\frac{n_2}{n} \approx 1$. This yields the ratios

$$\frac{U'_{N_1}}{U_{N_1}} \approx \frac{\frac{2}{3}n^{\frac{2}{3}}}{2n^{\frac{1}{3}}} = \frac{1}{3}n^{\frac{1}{3}} = \frac{U'_{N_2}}{U_{N_2}}$$

and completes the proof of Step 1.a.

Step 1.b. For an arbitrary problem M we give a upper-bound of the inefficiency of the CUT mixture.

We fix a problem M and partition the agents according to their scores $\max_{U \in \Phi^p(M; i)} U_N$, that is, the utilitarian score of the outcomes on which they spread their weight under the CUT rule. Let $p_1 > p_2 > \dots > p_K > 0$ be the sequence of such scores and N_k the subset of agents who load

outcomes with score p_k . Note that $n_1 \geq p_1$. Set A_k to be the set of outcomes loaded by at least one agent in N_k : they all have the same score p_k so the A_k -s are pairwise disjoint. Note also that agents in N_k do not like any outcome in A_ℓ for $\ell < k$.

Consider finally the outcomes b in $B = A \setminus (\cup_1^K A_k)$, if any. Their utilitarian score u_{N_b} is at most $p_1 - 1$. We partition B by gathering in B_k all the outcomes with a score in $[p_{k+1}, p_k]$, with the convention $p_{K+1} = 0$. Therefore the agents in N_k do not like any outcome in B_ℓ for $\ell < k$.

We prove first that for any feasible profile $U \in \Phi(M)$, we can find convex weights π_1, \dots, π_K such that

$$U_{N_k} \leq \pi_k p_k \text{ for } k = 1, \dots, K \quad (9)$$

Pick $z \in \Delta(A)$ implementing U and write for simplicity $z_{A_k} = x_k$ and $z_{B_k} = y_k$. The total contribution¹² $U_{N_k A_k} = x_k p_k$ of A_k to U_N is shared between the agents of $\cup_1^k N_\ell$ only, so there are some convex weights $\gamma_\ell^k, 1 \leq \ell \leq k$, such that

$$U_{N_\ell A_k} = \gamma_\ell^k x_k p_k \text{ for all } 1 \leq \ell \leq k \leq K.$$

Similarly, the contribution $U_{N_k B_k}$ of B_k is shared in $\cup_1^k N_\ell$ and $U_{N_k B_k} \leq y_k p_k$. So we can find convex weights $\delta_\ell^k, 1 \leq \ell \leq k$, such that

$$U_{N_\ell B_k} \leq \delta_\ell^k y_k p_k \text{ for all } 1 \leq \ell \leq k \leq K.$$

Combining the above equality and inequality we have for all k

$$U_{N_k} = \sum_{\ell=k}^K (U_{N_k A_\ell} + U_{N_k B_\ell}) \leq \sum_{\ell=k}^K (\gamma_\ell^k x_\ell + \delta_\ell^k y_\ell) p_\ell \leq p_k \sum_{\ell=k}^K (\gamma_\ell^k x_\ell + \delta_\ell^k y_\ell)$$

so the weights $\pi_k = \sum_{\ell=k}^K (\gamma_\ell^k x_\ell + \delta_\ell^k y_\ell)$ are indeed convex and satisfy (9).

Next we evaluate the blocks of the profile U^{cut} in the same fashion. Agents in N_k load exclusively A_k therefore if z implement U^{cut} we have $z_{A_k} = \frac{n_k}{n}$ and $U_{N_k A_k}^{cut} = \frac{n_k}{n} p_k$. We can find as above convex weights $\theta_\ell^k, 1 \leq \ell \leq k$, such that

$$U_{N_\ell A_k}^{cut} = \theta_\ell^k \frac{n_k}{n} p_k \text{ for all } 1 \leq \ell \leq k \leq K$$

and then as above we get

$$U_{N_k}^{cut} = \sum_{\ell=k}^K \theta_\ell^k \frac{n_\ell}{n} p_\ell.$$

Assume now the profile U^{cut} is $(1 - \varepsilon)$ -inefficient: $U^{cut} \leq \varepsilon U$ for some feasible U . From (9) we find convex weights π such that $U_{N_k}^{cut} \leq \varepsilon \pi_k p_k$ for all k , which implies

$$\varepsilon \geq \sum_{k=1}^K \frac{1}{p_k} U_{N_k}^{cut} = \sum_{\ell=1}^K \frac{n_\ell}{n} \sum_{k=1}^{\ell} \theta_k^\ell \frac{p_\ell}{p_k}.$$

The key inequality is $U_{N_k A_k}^{cut} \geq \frac{n_k}{n}$ because agent $i \in N_k$ loads only A_k containing her approval set: this implies $\theta_k^k \geq \frac{1}{p_k}$. Moreover, in the sum $\sum_{k=1}^{\ell} \theta_k^\ell \frac{p_\ell}{p_k}$ the terms $\frac{p_\ell}{p_k}$ increase in k . Combining these two observations we have for any $\ell \geq 2$:

$$\sum_{k=1}^{\ell} \theta_k^\ell \frac{p_\ell}{p_k} \geq \left(\sum_{k=1}^{\ell-1} \theta_k^\ell \right) \frac{p_\ell}{p_1} + \theta_\ell^\ell \geq \left(1 - \frac{1}{p_\ell}\right) \frac{p_\ell}{p_1} + \frac{1}{p_\ell} = \frac{p_\ell - 1}{p_1} + \frac{1}{p_\ell}.$$

¹²Recall our notation $u_{SB} = \sum_{i \in S} \sum_{a \in B} u_{ia}$.

We invoke now the inequality $\frac{\alpha-1}{p_1} + \frac{1}{\alpha} \geq \frac{2}{\sqrt{p_1}} - \frac{1}{p_1}$, for any $\alpha > 0$, that we apply to each $\alpha = p_\ell$, $\ell \geq 2$, and combine with the two inequalities above as well as $\theta_1^1 = 1$:

$$\varepsilon \geq \frac{n_1}{n} + (1 - \frac{n_1}{n})(\frac{2}{\sqrt{p_1}} - \frac{1}{p_1}).$$

Finally, the term $\frac{2}{\sqrt{p_1}} - \frac{1}{p_1}$ decreases in p_1 and we know $p_1 \leq n_1$, so we get

$$\varepsilon \geq \frac{1}{n}(n_1 + (n - n_1)(\frac{2}{\sqrt{n_1}} - \frac{1}{n_1})).$$

It remains to compute the minimum of the above expression for fixed n and variable $n_1 \in [1, n]$. With the real variable x instead of n_1 the right hand term and its derivative are

$$\varphi(x) = \frac{1}{n}(1 + x - 2\sqrt{x}) + (\frac{2}{\sqrt{x}} - \frac{1}{x}) \implies \varphi'(x) = (1 - \frac{1}{\sqrt{x}})(\frac{1}{n} - \frac{1}{x^{\frac{3}{2}}}).$$

therefore $x = n^{\frac{2}{3}}$ achieves the minimum and we compute

$$\varepsilon \geq \varphi(n^{\frac{2}{3}}) = \frac{1}{n} + (1 - \frac{1}{n^{\frac{1}{3}}})\frac{3}{n^{\frac{1}{3}}}.$$

which is inequality (5).

Step 2: Lower bounding the worst case inefficiency of RP

Fix N and integers k, d, ℓ such that $n = kd$ and $2 \leq \ell < k$. Fix a partition $N^1 \cup \dots \cup N^d$ of N where each subset contains k agents. This construction requires $n \geq 6$ and is not feasible for all n .

We consider the problem with $A = D \cup C$ where $D = \{1, \dots, d\}$ and each $\delta \in D$ is liked exactly by the k agents in N^δ ; also $|C| = \binom{n}{\ell}$ and each outcome in C is liked exactly by a different subset of ℓ agents.

The symmetric (egalitarian) and efficient outcome is the uniform distribution in D and yields the utility profile $U_i^* = \frac{1}{d}$ for all i . We compute now the symmetric profile U implemented by RP.

Fix an ordering $\sigma \in \Theta(N)$ and let L be the set of its ℓ highest priority agents. In the resulting profile U^σ , the first ℓ agents have full utility (because there is $a \in C$ where they all do). Two cases arise. In the favourable case L is contained in some set N^δ : then δ is the only efficient pure outcome liked by all agents in L , thus it must be chosen by the σ -priority rule and $U_N^\sigma = k$. In the unfavourable case L straddles two or more sets N^δ and there is only one outcome (in C) that everyone in L like, so that $U_N^\sigma = \ell$. Therefore

$$\begin{aligned} U_N &= \frac{d \binom{k}{\ell}}{\binom{n}{\ell}} \cdot k + (1 - \frac{d \binom{k}{\ell}}{\binom{n}{\ell}})k \cdot \ell = (k - \ell)\frac{n}{k} \frac{\binom{k}{\ell}}{\binom{n}{\ell}} + \ell. \\ \implies \varepsilon(n) &\leq \frac{U_N}{U_N^*} = (1 - \frac{\ell}{k})\frac{(k-1) \cdots (k-\ell+1)}{(n-1) \cdots (n-\ell+1)} + \frac{\ell}{k} \end{aligned} \quad (10)$$

For the asymptotic statement we use the inequality $\frac{\binom{k}{\ell}}{\binom{n}{\ell}} \leq (\frac{k}{n})^\ell$ and compute

$$\implies \frac{U_i}{U_i'} = \frac{U_N}{U_N'} \leq (\frac{k}{n})^{\ell-1} + \frac{\ell}{k}.$$

Then we choose $k \simeq \frac{n}{e}$ and $\ell \simeq \ln(n)$ so that $(\frac{k}{n})^{\ell-1} + \frac{\ell}{k} \simeq e^{\frac{\ln(n)}{n}}$. The systematic inequality $\varepsilon^{rp}(n) \leq 6\frac{\ln(n)}{n}$ is obtained by numerical estimations of (10), omitted for brevity. ■

Remark 2 The proof of Step 2 improves upon, with a similar proof technique, Example 1 in [15] establishing that RP is $(1 - \frac{2}{\sqrt{n}})$ -inefficient.

Proof of Theorem 3 i): MNP satisfies PART*

In a preliminary result we fix $S \subset \mathbb{R}_+^N$ convex and compact, and write $S(-1)$ for its projection on $\mathbb{R}_+^{N \setminus 1}$. Define

$$U^* = \arg \max_{U \in S} \sum_{i \in N} \ln(U_i)$$

$$\bar{U}_{-1} = \arg \max_{U_{-1} \in S(-1)} \sum_{i \in N \setminus 1} \ln(U_i) \text{ and } \bar{U}_1 = \max_{(U, \bar{U}_{-1}) \in S} U_1.$$

Inequality $U_1^* < \bar{U}_1$ brings a contradiction as follows

$$\sum_{i \in N} \ln(\bar{U}_i) \geq \ln(\bar{U}_1) + \sum_{i \in N \setminus \{1\}} \ln(U_i^*) > \sum_{i \in N} \ln(U_i^*).$$

Assume next $U_1^* = \bar{U}_1$. The right hand inequality above becomes an equality, so we get $\sum_{i \in N} \ln(\bar{U}_i) = \sum_{i \in N} \ln(U_i^*)$ and finally $\bar{U} = U^*$. Summing up, we have just proven:

$$U_1^* \geq \bar{U}_1; \text{ and if } U_1^* = \bar{U}_1 \text{ then } U_{-1}^* = \bar{U}_{-1} \quad (11)$$

Applying (11) to $S = \Phi(M)$, $U^* = F^{mnp}(M)$, $\bar{U}_{-1} = F^{mnp}(M(-1))$ gives $\bar{U}_1 = U_1(-1)$ and $U_1^* \geq \bar{U}_1$, the first inequality in PART* (i.e., PART). To check the second we can assume that any two columns of u are different, for if two columns are identical one of them can be eliminated as a redundant outcome. Also recall that no row of u is null.

Because $U_i^* > 0$ for all i , the statement is true if $\bar{U}_1 = 0$. We assume now $0 < U_1^* = \bar{U}_1 < 1$ and derive a contradiction. Property (11) implies $U^* = \bar{U}$, therefore there is some $z \in \Delta(A)$ solving both problems: $z \in f^{mnp}(M) \cap f^{mnp}(M(-1))$.

As $0 < U_1^* < 1$ the mixture z cannot be deterministic, moreover there exists two outcomes a, b in the support $[z]$ of z such that $u_{1a} = 1, u_{1b} = 0$. Writing $N(x; y)$ for the set of agents in N who like x and dislike y , this means $1 \in N(a; b)$.

Note that $N(b; a)$ must contain at least one $i \in N \setminus 1$: otherwise the column U_a dominates column U_b (outcome b is Pareto inferior to a) which contradicts the efficiency of z in M . We claim that $N(a; b)$ as well contains some $j \in N \setminus 1$: suppose not, then the restriction of column U_b to $N \setminus 1$ either dominates the corresponding restriction of U_a , or these two restricted columns are equal; the former case contradicts efficiency of z in $M(-1)$, the latter contradicts its efficiency in M .

We have shown that $N(a; b)$ and $N(b; a)$ both contains at least one outcome in $N \setminus 1$. Recalling that z_a, z_b are both positive, we define $z(\varepsilon)$ by shifting the weight ε from a to b : this outcome is well defined for ε small enough and of arbitrary sign; such a shift does not affect agents outside $N(a; b) \cup N(b; a)$. From $z \in f^{mnp}(M(-1))$ we see that the strictly concave function $\varphi(\varepsilon) = \sum_{i \in (N(a; b) \cup N(b; a)) \setminus 1} \ln(u_i \cdot z(\varepsilon))$ reaches its maximum at $\varepsilon = 0$. And $z \in f^{mnp}(M)$ implies that the function $\varphi(\varepsilon) + \ln(u_1 \cdot z(\varepsilon))$ is also maximal at $\varepsilon = 0$: this is a contradiction because $\ln(u_1 \cdot z(\varepsilon))$ decreases strictly in ε . ■

Proof of Theorem 3 ii): The MNP rule fails EXSP

A numerical example. The following small example with 7 voters and 4 outcomes is due to Dominik Peters.

	<i>a</i>	<i>b</i>	<i>c</i>	<i>d</i>
No. of agents types				
1	1	0	0	0
2	1	1	0	0
3	1	1	1	0
4	1	1	1	0
5	1	0	1	0
6	0	1	0	1
7	0	0	1	1

MNP returns the following mixture: [$a : 0.47619047619951582$, $b : 0.23809523808330549$, $c : 0.23809523808330565$, $d : 0.047619047633873257$]

Voter 1 manipulates by additionally liking b .

	<i>a</i>	<i>b</i>	<i>c</i>	<i>d</i>
No. of agents types				
1	1	1	0	0
2	1	1	0	0
3	1	1	1	0
4	1	1	1	0
5	1	0	1	0
6	0	1	0	1
7	0	0	1	1

MNP returns the following mixture: [$a : 0.53169166954599634$, $b : 0.11707708261768104$, $c : 0.11707708261768078$, $d : 0.23415416521864219$]. The misreporting agent 1 gets more utility (equivalently more share for outcome a) by additionally liking b .

Experiments: welfare achieved by the rules

	A	3	5	7	10	15	20
N							
3		0.8314	0.8155	0.8069	0.8005	0.781	0.7149
5		0.7777	0.7778	0.7322	0.7531	0.7072	0.7172
7		0.7678	0.80790	0.7373	0.695	0.7581	0.7109
10		0.7524	0.7334	0.808	0.7843	0.7857	0.7204
15		0.7862	0.8029	0.7561	0.7801	0.7747	0.7737
20		0.792	0.8234	0.7764	0.8155	0.7505	0.7896

Table 2. Minimum ratio of utilitarian welfare under the MNP rule to maximum utilitarian welfare for 100 profiles draws under impartial culture assumption for each combination of # agents and # outcomes.

	A	3	5	7	10	15	20
N							
3		0.9451	0.9652	0.9722	0.9678	0.9759	0.9634
5		0.9171	0.9309	0.9421	0.9377	0.9335	0.9004
7		0.8926	0.9324	0.9171	0.9277	0.9121	0.8856
10		0.8921	0.9014	0.91	0.9094	0.9056	0.8873
15		0.893	0.9013	0.8911	0.9049	0.8984	0.8774
20		0.8948	0.9001	0.8909	0.9047	0.9049	0.8941

Table 3. Average ratio of utilitarian welfare the MNP rule to maximum utilitarian welfare for 100 profiles draws under impartial culture assumption for each combination of # agents and # outcomes.

	A	3	5	7	10	15	20
N							
3		0.75	0.6397	0.5333	0.4815	0.4333	0.3743
5		0.625	0.3919	0.4244	0.4592	0.4956	0.403
7		0.5833	0.492	0.3632	0.5102	0.5599	0.5799
10		0.5834	0.375	0.4952	0.4253	0.5689	0.5696
15		0.5129	0.5525	0.57	0.4361	0.5198	0.5817
20		0.6001	0.625	0.5927	0.5525	0.6425	0.5656

Table 4. Minimum ratio of utilitarian welfare under EGAL to maximum utilitarian welfare for 100 profiles draws under impartial culture assumption for each combination of # agents and # outcomes.

	A	3	5	7	10	15	20
N							
3		0.9325	0.9256	0.8838	0.8075	0.844	0.8408
5		0.8482	0.8484	0.781	0.8019	0.82	0.8175
7		0.8221	0.8131	0.7817	0.7978	0.7992	0.8118
10		0.8176	0.8049	0.7902	0.7639	0.8152	0.7803
15		0.8267	0.807	0.7805	0.7476	0.8259	0.8009
20		0.8414	0.8278	0.8121	0.7748	0.8265	0.8084

Table 5. Average ratio of utilitarian welfare under EGAL to maximum utilitarian welfare for 100 profiles draws under impartial culture assumption for each combination of # agents and # outcomes.

	A	3	5	7	10	15	20
N							
3		0.8333	0.8333	0.8333	0.8333	0.8333	0.8333
5		0.8	0.7333	0.8	0.8	0.8	0.8667
7		0.75	0.7619	0.8571	0.8214	0.8857	0.8571
10		0.8	0.8	0.8714	0.86	0.8667	0.8833
15		0.8	0.8444	0.8583	0.8417	0.8741	0.8815
20		0.8038	0.85	0.8773	0.9	0.8944	0.8727

Table 6. Minimum ratio of utilitarian welfare under CUT to maximum utilitarian welfare for 100 profiles draws under impartial culture assumption for each combination of # agents and # outcomes.

	A	3	5	7	10	15	20
N							
3		0.9333	0.9717	0.9717	0.9867	0.9867	0.995
5		0.9372	0.9452	0.959	0.9748	0.969	0.9757
7		0.9139	0.9468	0.9549	0.9624	0.969	0.9778
10		0.9194	0.9383	0.9502	0.9586	0.9576	0.965
15		0.9263	0.9276	0.9483	0.9483	0.9567	0.9634
20		0.9195	0.9332	0.9486	0.955	0.9588	0.9631

Table 7. Average ratio of utilitarian welfare under CUT to maximum utilitarian welfare for 100 profiles draws under impartial culture assumption for each combination of # agents and # outcomes.

	A	3	5	7	10	15	20
N							
3		0.8333	0.8333	0.8333	0.8333	0.8333	0.8333
5		0.7778	0.7	0.7778	0.7778	0.7	0.8
7		0.7679	0.75	0.8036	0.75	0.7943	0.7778
10		0.7778	0.7737	0.7596	0.8116	0.7684	0.8031

Table 8. Minimum ratio of utilitarian welfare under RP to maximum utilitarian welfare for 100 profiles draws under impartial culture assumption for each combination of # agents and # outcomes.

Experiments: inefficiency of CUT

	$ A $	3	5	7	10	15	20
$ N $	3	0.9483	0.9733	0.9883	0.99	0.9867	0.9933
5	0.8992	0.9302	0.9351	0.9471	0.9512	0.962	
7	0.8851	0.8952	0.9143	0.9182	0.929	0.9305	
10	0.8839	0.89	0.8911	0.8969	0.9	0.8997	

Table 9. Average ratio of utilitarian welfare under RP to maximum utilitarian welfare for 100 profiles draws under impartial culture assumption for each combination of # agents and # outcomes.

$ N \times A $	10%	25%	33%	50%	75%
5 × 6	0	0	0	0	0
10 × 5	0	0	0	0	0
5 × 15	0	0.045000	0.050000	0	0
10 × 10	0	0.037037	0.025000	0.012500	0
10 × 20	0.027778	0.041667	0.033333	0.008000	0
10 × 30	0.050000	0.039583	0.027619	0.007937	0
20 × 20	0.031250	0.009524	0.007000	0.001852	0
20 × 30	0.026316	0.023611	0.009387	0.001515	0
20 × 40	0.027018	0.012500	0.006206	0.001786	0
20 × 60	0.035227	0.019236	0.008758	0.001777	0
30 × 30	0.017857	0.003580	0	0	0
40 × 60	0.007232	0	0	0	0
50 × 50	0.001096	0	0	0	0

Table 10. The worst inefficiency of U^{cut} ; 200000 draws for each level.

$ N \times A $	10%	25%	33%	50%	75%
5 × 6	0	0	0	0	0
10 × 5	0	0	0	0	0
5 × 15	0	0	0.000257	0	0
10 × 10	0	0.000008	0.000006	0.000001	0
10 × 20	0.000005	0.000168	0.000069	0.000009	0
10 × 30	0.000163	0.000236	0.000086	0.000016	0
20 × 20	0.000010	0.000001	0	0	0
20 × 30	0.000051	0.000002	0.000002	0	0
20 × 40	0.000127	0.000004	0.000001	0	0
20 × 60	0.000344	0.000007	0.000001	0	0
30 × 30	0.000002	0	0	0	0
40 × 60	0	0	0	0	0
50 × 50	0	0	0	0	0

Table 11. The average inefficiency of U^{cut} ; 200000 draws for each level.