Fair Mixing of Public Outcomes under Dichotomous Preferences

Hervé Moulin

University of Glasgow and HSE St Petersburg

joint work with

Haris Aziz, Data61, CSIRO and University of New South Wales

Anna Bogomolnaia, University of Glasgow and HSE St Petersburg

the problem

- voting to choose a *mixture* of (mutually incompatible) public outcomes

- mixing implemented by: lottery, time sharing, division of resources

(we use the probabilistic terminology)

- mixing to achieve a *fair compromise* between conflicting preferences, hence *not* as a second best to a *pure* outcome a familiar test of fairness: **protection of minorities**

a pure (deterministic) outcome, e. g. the Condorcet winner, typically ignores the preferences of many voters

general principle: everyone is entitled to some benefit from the public resources

$\longrightarrow$ *cumulative voting* protects ethnic minorities in political elections, minority stockholders in corporate governance

$\longrightarrow$ proportional representation

$\longrightarrow$ proportional veto power small or large scale examples:

scheduling a periodic club meeting

balancing the program of a public TV station

budgetary participation: spreading a fixed budget over different projects

our model

key simplifying assumption: *dichotomous preferences*

simple Facebook-style like/dislike reports $\Longrightarrow$ easy elicitation

we extend the discussion started in Bogomolnaia, Moulin, Stong (2002), (2005)

using more nuanced concepts of *Fairness* and *Incentive Compatibility,* as in Aziz, Brandl, Brandt, Brill (2016) and Brandt, Brandt, Hofbauer (2015) *we explore the tradeoffs between*

- Fairness: Individual and Group guarantees; two concerns: *protection of minorities* and *numbers matter*

- Incentives to report preferences: *two versions* of Strategyproofness (outcomes excludable or not)

- Incentives to participate: no No Show Paradox

- Efficiency

we find

$\longrightarrow$ a simple Efficient rule, weakly Strategyproof and protecting individual welfare: the *Egalitarian* rule

$\longrightarrow$ two good rules Fair and Incentive Compatible: *the Conditional Utilitarian*, and the familiar *Random Priority* rules

$\longrightarrow$ but CUT is less inefficient and much easier to compute than RP

$\longrightarrow$ one good rule Fair, Efficient, but not IC: the *Nash Max Product* rule

we do not find any characterization or impossibility results

several challenging open questions emerge around the tradeoffs between these normative properties

the model

a *problem* $M =$

$N \downarrow A \rightarrow$	a	b	c	d	e
1	0	0	0	1	0
2	0	0	1	1	1
3	1	1	0	0	0
4	1	0	1	0	0
5	0	1	0	1	1

a *mixture/lottery*: $z = (z_a, z_b, z_c, z_d, z_e) \geq 0$, $z_a + z_b + z_c + z_d + z_e = 1$

the corresponding *utilities*: $U_1 = z_d$, $U_2 = z_c + z_d + z_e$, etc..

Problem: $M = (N, A, u)$; $\Phi(M) = \{$feasible utility profiles $U\}$

$\widetilde{\Phi}(M) = \{$"deterministic" feasible utility profiles $U\}$

a **rule** assigns to each problem M

a set of mixtures $f(M) = \{z \in \Delta(A)\}$ with a single valued utility profile $U = F(M) \in \Phi(M)$

hard wired in the definition of a rule:

$\longrightarrow$ it ignores null rows (indifferent agents), null columns (useless outcomes), and clone outcomes

$\longrightarrow$ it is *Anonymous* (ANO) and *Neutral* (NEUT)

Efficiency

$$F(M) \leq U \Longrightarrow F(M) = U, \text{ for all } M \text{ and } U \in \Phi(M)$$

in the example $z \in \Delta(A)$ is inefficient *iff* $z_e > 0$ and/or $z_b \cdot z_c > 0$

the rule F is ε-**inefficient** if there exists M and $U \in \Phi(M)$ such that

$$F(M) \leq \varepsilon U$$

the classic impossibility result with *full-fledged vNM preferences*

$$\text{Efficiency} + \text{Strategyproofness (SP)} + \text{Anonymity} = \varnothing$$

disappears in the dichotomous domain

Utilitarian rule (UTIL) (*Approval voting*)

$$F^{ut}(M) = avg\{U | U \in \arg\max_{U' \in \widetilde{\Phi}(M)} \sum_{i \in N} U'_i\}$$

average of all *deterministic utility profiles* with largest approval

Lemma: *UTIL is Efficient, Strategyproof, Anonymous and Neutral*

$\longrightarrow$ but UTIL ignores minority entirely

a familiar welfare guarantee, implementing a fair share of decision power:

$$\textbf{Individual Fair Share}\text{ (IFS): } U_i \geq \frac{1}{n} \text{ for all } i$$

$\longrightarrow$ the main result in BMS 05: if $n \geq 5$ and $|A| \geq 17$

Efficiency + Strategyproofness + Individual Fair Share + ANO+ NEUT $= \varnothing$

first result: *this impossibility disappears if we weaken SP*

if the public outcomes are non rival but excludable: club meeting, cable TV broadcast, .. then we can weaken the truth-telling requirement from

$$\textbf{Strategyproofness}(\text{SP}) : u_i \cdot f(u) \geq \max_{z' \in f(u|^i u_i')} u_i \cdot z' \text{ for all} i \text{ and } u_i'$$

to

Excludable Strategyproofness (EXSP):

$$u_i \cdot f(u) \geq \max_{z' \in f(u|^i u_i')} (u_i \wedge u_i') \cdot z' \text{ for all } i \text{ and } u_i'$$

Egalitarian rule (EGAL)

$$F^{eg}(M) = \text{arg} \max_{U \in \Phi(M)} \succ_{leximin}$$

it equalizes utilities in the leximin sense $\Longrightarrow$ IFS is clear

in the example EGAL picks $\frac{1}{2}a + \frac{1}{2}d$, while UTIL picks d

Proposition: *The Egalitarian rule is Efficient, Excludable Strategyproof, and guarantees Individual Fair Shares* but EGAL is “extremist” in that *numbers do not matter* at all, the rule is

CLONE INVARIANT

ignoring $i \in N$ has no effect if $u_i = u_j$ for some other $j \in N$

$\longrightarrow$ the familiar objection to single-minded egalitarianism is especially compelling

to make numbers matter we impose two natural requirements (that both UTIL and EGAL fail)

$\longrightarrow$ an incentive property

Strict Participation (PART*)

$$U_i(N) \geq \max_{z \in f(N \setminus i)} u_i \cdot z \; ; \; U_i(N) > \min_{z \in f(N \setminus i)} u_i \cdot z \quad \text{if} \quad \min_{z \in f(N \setminus i)} u_i \cdot z < 1$$

(No Show always hurts)

$\longrightarrow$ a fairness property

Unanimous Fair Share (UFS)

for all coalition S: $u_i = u_j$ for $i, j \in S \Longrightarrow U_i \geq \dfrac{|S|}{n}$ for all $i \in S$

individual welfare guarantees add up *among clones*

Open questions

- is there a rule meeting Efficiency, Excludable Strategyproofness and Strict Participation ?

- is there a rule meeting Efficiency, Excludable Strategyproofness, and Unanimous Fair Share ?

we define two variants of the familiar "Random Dictator" idea

$\sigma \in per(N)$ is an ordering of the agents

the σ-*Priority rule* F^{σ}: ensures $u_{\sigma(1)} = 1$; $u_{\sigma(2)} = 1$ as well if 1 and 2 like a common outcome; $u_{\sigma(3)} = 1$ if $\sigma(3)$ likes an outcome common with all happy agents before her; and so on ..

Random Priority rule (RP)

$$F^{rp}(M) = \frac{1}{n!} \sum_{\sigma \in per(N)} U^{\sigma}$$

Conditional Utilitarian rule (CUT)

$$\widetilde{\Phi}(M;i) = \{U \in \widetilde{\Phi}(M) | U_i = 1\}$$

$$F^{cut}(M) = \frac{1}{n} \sum_{i \in N} avg\{U | U \in \arg \max_{U' \in \widetilde{\Phi}(M;i)} \sum_{i \in N} U'_i\}$$

(recall we drop indifferent agents)

each agent spreads her share $\frac{1}{n}$ equally between the outcomes with maximal support among those she likes (Duddy (2015)): "charitée bien ordonnée commence par soi même" "charity begins at home"

RP and CUT are inefficient

$N \downarrow A \rightarrow$	a	b	c	d	e
1	0	0	0	1	0
2	0	0	1	1	1
3	1	1	0	0	0
4	1	0	1	0	0
5	0	1	0	1	1

$\rightarrow$

	a	b	c	d
1	0	0	0	1
2	0	0	1	1
3	1	1	0	0
4	1	0	1	0
5	0	1	0	1

$z_e = 0$ because e is dominated

$z^{cut} = (\frac{1}{5}, \frac{1}{10}, \frac{1}{10}, \frac{3}{5})$ is dominated by $(\frac{1}{5} + \frac{1}{10}, 0, 0, \frac{1}{10} + \frac{3}{5})$

$z^{rp} = (\frac{1}{5}, \frac{1}{6}, \frac{1}{6}, \frac{7}{15})$ is dominated by $(\frac{1}{5} + \frac{1}{6}, 0, 0, \frac{1}{6} + \frac{7}{15})$

Proposition:

i) *Both CUT and RP are Strategyproof; they meet Strict Participation and Unanimity Fair Share*

ii) *CUT is very easy (polynomial) to compute, RP is #P-complete to compute (Aziz, Brandt, Brill (2013))*

iii) *Both CUT and RP mix undominated pure outcomes; they are efficient rules if and only if* $|A| \leq 3$ *and/or if* $n \leq 4$

iv) $F^{cut}(M)$ *is efficient if* $F^{rp}(M)$ *is efficient, for all* M *; the converse is not true*

v) *CUT is at most (and can be)* $O(n^{-\frac{1}{3}})$*-inefficient, while RP is at least* $O(\frac{\ln(n)}{n})$*-inefficient*

Open questions:

- what is the worst case inefficiency of RP?

- in the impartial culture, what is the probability that RP or CUT is efficient? what about some expected measure of their inefficiency?

a third appealing rule: efficient, more fair, but not EXSP

Nash Max Product rule (NMP)

$$f^{nmp}(M) = \arg \max_{z \in \Delta(A)} \sum_{i \in N} \ln(u_i \cdot z)$$

it solves a strictly convex program $\Longrightarrow$ well defined in the example $z^{nmp} = (\frac{2}{5}, 0, 0, \frac{3}{5})$ is efficient

the resulting utilities for our three rules are not comparable

$N \longrightarrow$	1	2	3	4	5
CUT	0.6	0.7	0.3	0.3	0.7
RP	0.47	0.63	0.37	0.37	0.63
NMP	0.6	0.6	0.4	0.4	0.6

$\longrightarrow$ *two group guarantee properties much more demanding than UFS*

Average Fair Share (AFS)

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

in the canonical example the inequalities

$$\frac{1}{3}(U_1 + U_2 + U_5) \geq \frac{3}{5} \text{ and } \frac{1}{2}(U_3 + U_4) \geq \frac{2}{5}$$

force the Nash utility profile!

AFS, unlike UFS, constrains the acceptable utility profiles in all problems

Core Fair Share (CFS)

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

a cooperative property of incentive compatibility: each coalition can cumulate its individual shares of decision power and form core objections

(AFS and CFS are not logically related)

(RP and CUT violate both AFS and CFS)

Proposition:

i) *The Nash rule is Efficient; it meets Strict Participation, Average Fair Share and Core Fair Share*

ii) *it is not Excludable Strategyproof*

iii) *its exact computational complexity is unknown, but it is easily approximated by C-plex methods*

$\longrightarrow$ checking that NMP violates SP is very easy in a 4×3 example; proving that it violates EXSP is much harder: we give a computer generated example with 36 agents and 4 outcomes, and a formal proof with 860 agents

Open question: for what sizes of N and A is the Nash rule EXSP ?

one more property: *a decomposition axiom with a fairness content*

a polarized electorate

$$(N, A, u) = \begin{array}{cccc} & A^1 & A^2 & A^3 \\ N^1 & u^1 & 0 & 0 \\ N^2 & 0 & u^2 & 0 \\ N^3 & 0 & 0 & u^3 \end{array}$$

Decentralization (DEC): $F(N, A, u) = \sum_k \frac{|N_k|}{|N|} F(N^k, A^k, u^k)$

Lemma: *the Conditional Utilitarian, Random Priority and Nash rules are Decentralizable*

Conclusion: the three contenders

- the single-minded Egalitarian rule offers individual guarantees and excludable strategyproofness, but ignores numbers

- the Conditional Utilitarian rule has the best incentives properties (PART* and SP), good group guarantees (UFS) and is decentralizable; it is less inefficient, and much easier to compute than RP

- the Nash Max Product rule guarantees high average welfare to groups who have similar preferences(AFS), and allows them to form objections with a proportional share of decision power (CFS); it ensures SP, and is decentralizable; but it is manipulable even if outcomes are excludable.

THANK YOU