Guarantees in Fair Division: beyond Divide and Choose and Moving Knives

Hervé Moulin

University of Glasgow and HSE St Petersburg

HSE and European University seminar St Petersburg October 17, 2019

.

joint work with Anna Bogomolnaia and Richard Stong

Divide & Choose (**D&C**$_2$): the ancestor of mechanism design

recent (1948) generalisation to any number n of agents:

- the Diminishing Share (**DS**) rule (one way to generalise D&C: Steinhaus 1948)

- the Moving Knife (**MK**) rule (Dubins and Spanier 1961)

attractive features

- decentralized implementation of the *Proportional Guarantee*: the utility of my share is at least $\frac{1}{n}$-th of that of the whole manna

- informational parsimony as *privacy preservation*: I reveal very little of my preferences, at most one cut and $n-1$ queries for DS, $n-1$ "queries" for MK

- informational parsimony as *small cognitive effort*: I do not need to form full preferences

unappealing features

- work for goods, or bads, but not for a mixed manna: both rules requires *co-monotone* utilities (else trimming or padding is ambiguous)

- both require *additive utilities/preferences*: otherwise the Proportional Guarantee is neither feasible nor ordinally meaningful

- (both pick inefficient allocations: a consequence of informational parsimony)

- we generalize D&C$_2$ to the n-person D&C$_n$ implementing the Proportional Guarantee when utilities are additive, but the sign of marginal utilities varies across the manna and across agents; it requires neither trimming nor padding, and is informationally parsimonious

- D&C$_n$ implements, *for the much, much larger class of continuous, not necessarily monotonic preferences*, the **minMax Guarantee**: the utility of my best share in the worst possible partition

- when preferences are co-monotone (all parts of the manna are desirable *goods,* or all are undesirable *bads*) we generalize MK$_n$ to the rich family of Bid & Choose (**B&C**) rules: they implement better guarantees than the minMax, though still below the unfeasible **Maxmin** utility

- parsimonious computation of an efficient allocation of resources: Reiter (1972)), e. g. the competitive equilibrium: Mount and Reiter (1972), Reischelstein and Reiter (1988)

- protective/prudent implementation: Moulin (1981), Barbera and Dutta (1982), and in the entire cake-cutting literature: Brams and Taylor (1996), (2000)

- identifying the best welfare bounds (upper or lower guarantees) in cooperative production: Moulin (1990, 1991), Fleurbaey and Maniquet (1996); in fair division: Moulin (1991)

- guarantees when we distribute indivisible objects Buddish (2011), Procaccia and Wang (2014), $\cdots$

additive utilities, continuous model

the manna Ω is a measurable set in an euclidian space

utility u_i of agent $i \in N$ is a non atomic *real valued* measure on Ω (u_i is absolutely continuous w. r. t. Lebesgue): $u_i(S) = \int_S du_i(x) = \int_S \dot{u}_i(x)dx$

compare: in most of the cake-cutting literature all u_i are positive, or all negative

Fair Guarantee

$$\text{Proportional Guarantee } (\textbf{Pro})\text{: } u_i(S_i) \geq \frac{1}{n} u_i(\Omega) \text{ for all } i$$

additive utilities, discrete model

Ω is a finite set of objects

utility of agent $i \in N$ is a vector in $u_i \in \mathbb{R}^{\Omega}$, $u_i(S) = \sum_{x \in S} u_{ix}$

Fair Guarantee: *proportional up to one object*

$$\textbf{Pro1}\text{: } \exists a \in \Omega \diagdown S_i : u_i(S_i + a) \geq \frac{1}{n} u_i(\Omega)$$

$$\text{and/or } \exists b \in \Omega : u_i(S_i - b) \geq \frac{1}{n} u_i(\Omega) \text{ for all } i$$

a combinatorial Lemma

Fix the sets A of p items, M of $p-1$ agents, and an arbitrary bipartite graph of "likes" in $M \times A$ described for all $B \subseteq A$ by

$$\ell(B) = \{i \in M | i \text{ likes at least one item in } B\}$$

There exists a non empty $B \subsetneq A$ and a possibly empty $L \subseteq M$ such that

$$\ell(B) = L \text{ ; } |L| = |B| - 1 \text{ and}$$

$$\text{we can assign all but one item in } B \text{ to an agent in } L \text{ who likes it}$$

Proof: simple application of the Marriage Lemma the D&C$_n$ rule (continuous or discrete model)

order the agents $1, \cdots, n$; agent 1 partitions Ω in n shares S_k she is is reputed to *like* (Fair Guarantee); other agents report **all** the shares they *like*, and **at least one**

find a subset B of shares and L of agents in $\{2, \cdots, n\}$ s. t. we can assign to everyone in $L+1$ a share he likes, and nobody outside $L+1$ likes any share in B

repeat with the remaining manna and agents: the lowest agent in $[n]-(L+1)$ partitions $\Omega - \cup_B S_k$ in $n-|L|-1$ shares she is reputed to like, etc..

$\rightarrow$ important privacy feature: I do not report which *individual* objects I like or dislike

Theorem: *additive utilities*

continuous model: an agent who cuts shares of equal value when called to cut, and otherwise reports at each step all shares worth at least $\frac{1}{n}u_i(\Omega)$ (even if we divide less than Ω among less than n agents), ensures that her share meets Pro.

discrete model: an agent who cuts shares meeting Pro1 when called to cut, and otherwise reports at each step all shares meeting Pro (*not Pro1 !*) in the entire Ω, ensures that his share meets Pro1.

proof in the continuous model

1. every share in an *equi-partition* gives the utility $\frac{1}{n}u_i(\Omega)$

2. at each step where i is not served while q other agents are, the per capita value to i of the remaining cake increases weakly

proof in the discrete model: in any partition with per capita value at least $\frac{1}{n}u_i(\Omega)$, at least one share meets Pro1; in a *Maxmin* partition, maximizing the utility of her worst share, all shares meet Pro1; so she can at any step partition the remaining manna in shares meeting Pro1

the general continuous model

Ω is a compact set in an euclidian space s. t. $\Omega = \overline{\overset{\circ}{\Omega}}$; shares are the closed subsets $\varnothing \subseteq S \subseteq \Omega$ or some subfamilies of these (e. g., connected subsets); "partitions" allow for overlaps of lower dimension

individual utilities are real valued, continuous for the Hausdorff distance, and

$$u(\varnothing) = 0 \text{ and } u(S) = u(\overline{\overset{\circ}{S}})$$

example: the fair division of Arrow Debreu commodities

the general discrete model Ω is a finite set, u is real valued on 2^{Ω} and $u(\varnothing) = 0$

the hard question: under general utilities/preferences, what Guarantees are feasible, and parsomoniously implementable?

the Maxmin share (Buddish (2011)): a natural (ordinal) proposal in the spirit of D&C:

$$Maxmin(u;n) = \max_{P} \min_{1 \leq k \leq n} u(S_k)$$

where $P = (S_k)_{k=1}^{n}$ is a n-partition of Ω

in the continuous model with additive utilities $Maxmin(u; n) = \frac{1}{n}u_i(\Omega)$, but with general utilities the profile $(Maxmin(u_i; n))_{i \in N}$ is easily not feasible, already with two agents

Example agents 1 and 2 with utilities u and v divide $\omega = (1, 1)$

$$u(x, y) = \min\{x, y\} \; ; \; v(x, y) = \max\{x, y\}$$

$$Maxmin(u; 2) = u(\frac{1}{2}\omega) \; ; \; Maxmin(v; 2) = v(\omega)$$

consider instead the (ordinal) *minMax share* (Shams et al. (2019))

$$minMax(u;n) = \min_{P} \max_{1 \le k \le n} u(S_k)$$

in the continuous model, if P is an equi-partition of Ω for u we have

$$minMax(u;n) \le u(P) \le Maxmin(u;n)$$

Lemma *the continuity assumptions ensure that such an equi-partition exists*

Proof: if u is non negative (all shares desirable) this follows from Su (1999) or a simple application of the KKM lemma. If u is real valued, the proof is harder.

example 1

Ω is a square in $\mathbb{R}^2$ and $u(S)$ is the diameter of $\overset{\circ}{S}$: $\frac{Maxmin(u;2)}{minMax(u;2)} = \frac{\sqrt{2}}{2/\sqrt{5}} = 1.27$; $\frac{Maxmin(u;4)}{minMax(u;4)} = 2$

example 2: in the Arrow Debreu model: $minMax(u;n) \leq u(\frac{1}{n}\omega) \leq Maxmin(u;n)$

for $n = 2$ and $\omega = (1,1)$

$$u(x,y) = \min\{x,y\}\text{: } minMax(u;2) = 0 < u(\frac{1}{2}\omega) = Maxmin(u;2)$$

$$v(x,y) = \max\{x,y\}\text{: } minMax(v;2) = v(\frac{1}{2}\omega) < v(\omega) = Maxmin(v;2)$$

Theorem: *continuous model*

in the D&C$_n$ rule, an agent who cuts shares of equal value when called to cut, and otherwise reports at each step all shares worth at least $minMax(u_i; n)$ (over the entire Ω and with n agents), guarantees that utility level

note that agent 1 who cuts first is guaranteed $Maxmin(u_1; n)$, but other agent only $minMax(u_i; n)$

proof

1. at each step where i is not served, the shares served to the leaving agents are worth strictly less to i than $minMax(u_i; n)$

2. so if i is not cutting in the next step, at least one of the shares on offer is worth at least $minMax(u_i; n)$

3. and if i is cutting in the next step, any equi-partition of the remaining manna guarantees $minMax(u_i; n)$ as well

in the discrete model with general preferences, things are not so simple

- equi-partitions typically do not exist

- the $minMax$ and $Maxmin$ utilities are no longer comparable, e. g., if u is additive $Maxmin(u; n) \leq minMax(u; n)$

- neither guarantee is feasible, even under additive utilities (Procaccia and Wang (2014))

so the D&C$_n$ rule is not interesting any more

note: the $Maxmin$ Guarantee is at least $\frac{3}{4}$-feasible if utilities are additive (Ghodsi et al. (2017)), but the corresponding algorithm is anything but simple or informationally parsimonious

monotone preferences: *increasing* $\mathcal{M}^+(\Omega)$*, or decreasing* $\mathcal{M}^-(\Omega)$

$$\forall S \subset \mathbf{\Omega}, T \subseteq \mathbf{\Omega} \diagdown S : u(S) \leq u(S \cup T) \text{ (or } u(S) \geq u(S \cup T))$$

increasing: the manna is (weakly) desirable, *freely disposable*

decreasing: we divide non disposable “bads”, “chores”

result: *in each domain we can improve substantially the minMax Fair Guarantee* the Moving Knife rule inefficiently restricts the available shares

the Bid & Choose rules (**B&C**) generalize MK by running a "more inclusive" knife

the B&C$_2^\theta$ rule: definition for **two** agents(continuous or discrete model)

θ is an increasing and continuous *calibration* (benchmark utility) of the shares s. t. $\theta(\varnothing) = 0$, $\theta(\Omega) = 1$, and $\theta(S) = 0$ if S is not full dimensional

agent i bids $x_i \in [0, 1]$; (one of) the lowest bidder i can choose any share S_i such that $\theta(S_i) \leq x_i$; the loser j gets $\Omega \diagdown S_i$

Theorem *preferences in $\mathcal{M}^+(\Omega)$, continuous or discrete model*

i) in the B&C$_2^\theta$ rule each agent can guarantee the utility level Γ_2^θ

$$\Gamma_2^\theta(u_i) = \max_{1 \leq x \leq 0} \min\{u_i^+(x), u_i^-(x)\} \tag{1}$$

$$\text{where } u_i^+(x) = \max_{\theta(S) \leq x} u_i(S) \ ; \ u_i^-(y) = \min_{\theta(S) \leq x} u_i(\Omega \diagdown S)$$

ii) the bid(s) x_i^* securing this Guarantee solves the program (1) above

iii) the Guarantee Γ_2^θ is maximal (unimprovable) and we have

$$minMax(u;2) \leq \Gamma_2^\theta(u) \leq Maxmin(u;2) \text{ for all } u$$

the rule B&C$_2^\theta$ is anonymous like MK; it is MK if $\theta(S) = \max_{S \subseteq S(t)} t$, where $t \longrightarrow S(t)$ is the cut at time t

whether in prudent or in Nash equilibrium strategies, numerical simulations show that it collects a larger share of the surplus than MK

alternative definitions: the lowest bidder i chooses any S_i such that $\theta(S_i) \leq x_j$: achieves the same guarantees and is more balanced

version for bads $\mathcal{M}^-(\Omega)$: the highest bidder i can choose any share S_i such that $\theta(S_i) \geq x_i$

example $n = 2$; $\omega = (1, 1)$ with $\theta(x, y) = \frac{1}{2}(x + y)$

for $u(x, y) = \min\{x, y\}$:

$$minMax(u; 2) = 0 < \Gamma_2^\theta(u) = u(\frac{1}{3}\omega) \leq Maxmin(u; 2) = u(\frac{1}{2}\omega)$$

for $v(x, y) = \max\{x, y\}$:

$$minMax(v; 2) = u(\frac{1}{2}\omega) \leq \Gamma_2^\theta(v) = u(\frac{2}{3}\omega) \leq Maxmin(v; 2) = u(\omega)$$

so that $(\Gamma_2^\theta(u), \Gamma_2^\theta(v))$ is a fair Pareto optimal utility profile at the profile (u, v).

Agent 1 bids $\frac{1}{3}$: if she wins she picks $(\frac{1}{3}, \frac{1}{3})$, if she loses she is guaranteed at least $\frac{1}{3}$ of each good

Agent 2 also bids $\frac{1}{3}$ and chooses $(\frac{2}{3}, 0)$ or is guaranteed at least $\frac{2}{3}$ of some good.

a discrete example: $n = 2$, eight balls $a, b, \cdots, h$

agent 1's utility: largest number of lexicographically adjacent balls

agent 2's utility: largest number of adjacent balls for the order a, e, c, g, b, f, d, h

benchmark: $\theta(S) = |S|$ who needs the smallest number of balls

for both agents, the prudent bid is 3 or 2 and

$$minMax(u_i) = 1 < \Gamma_2^{\theta}(u_i) = 2 \leq Maxmin(u_i) = 4$$

here $(3, 3)$ is a Pareto optimal utility profile

the B&C$_n^\theta$ rule: definition for n agents

Step 1 each agent i bids $x_i^1 \in [0,1]$

(one of) the winners (lowest bidders), 1, chooses S_1 s. t. $\theta(S_1) \leq x_1^1$

Step 2 each agent $i \geq 2$ bids $x_i^2 \in [x_1^1, 1]$

(one of) the winners, 2, chooses $S_2 \subseteq \Omega \diagdown S_1$ s. t. $\theta(S_1 \cup S_2) \leq x_2^2$

Step $n-1$: the two remaining agents bid $x_i^{n-1} \in [\sum_1^{n-2} x_j^j, 1]$

the winner, $n-1$, chooses $S_{n-1} \subseteq \Omega \diagdown \cup_1^{n-2} S_j$ s. t. $\theta(\cup_1^{n-1} S_j) \leq x_{n-1}^{n-1}$

the last agent gets $S_n = \Omega \diagdown \cup_1^{n-1} S_j$

the B&C Welfare Guarantee for general n

$$\Gamma_n^\theta(u,\Omega) = \max_X \min_{1\le k\le n} u^*(x^{k-1}; x^k)$$

where X is any weakly increasing sequence $0 = x^0 \le x^1 \le x^2 \le \cdots \le x^n = 1$, and for any y, z s.t. $0 \le y \le z \le 1$, we define

$$u^*(y;z) = \min_{\theta(S)\le y} \max_{T\subseteq\Omega\setminus S, \theta(S\cup T)\le z} u(T)$$

for instance

$$\Gamma_3^\theta(u,\Omega) = \max_{0\le x^1\le x^2\le x^3\le 1} \min\{u^+(x^1), u^*(x^1;x^2), u^-(x^2)\}$$

Theorem repeating the same three points take home points

- we generalise both the Divide and Choose and the Moving Knife rules

- our D&C$_n$ rule requires only equi-partitions and "like" reports; it applies to the maximal domain of continuous utilities and respects the privacy of preferences just like D&C$_2$

- our versatile B&C$_n^\theta$ rules allow great flexibility in the choice of θ, and preserve the simplicity and anonymity of MK; they only apply to co-monotone preferences

Thank You