

Fair Division in the Age of Internet

Hervé Moulin

University of Glasgow and Higher School of Economics, St Petersburg

November 5, 2018

Abstract

Fair Division, a key concern in the design of many social institutions, is for 70 years the subject of interdisciplinary research at the interface of mathematics, economics and game theory.

Motivated by the proliferation of moneyless transactions on the internet, the Computer Science community has recently taken a deep interest in fairness principles and practical division rules. The resulting literature brings a fresh concern for computational simplicity (scalable rules), and realistic implementation.

In this survey of the most salient Fair Division results of the past thirty years, we concentrate on division rules with the best potential for practical implementation. The critical design parameter is the message space that the agents must use to report their individual preferences. A simple preference domain is key both to realistic implementation, and to the existence of division rules with strong normative and incentive properties. We discuss successively the one-dimensional single-peaked domain, Leontief utilities, ordinal ranking, dichotomous preferences, and additive utilities. Some of the theoretical results in the latter domain are already implemented in the user-friendly SPLIDDIT platform (spliddit.org).

Keywords: manna, goods, bads, envy free, fair share, competitive division, egalitarian division, preference domains

Acknowledgments: Many constructive criticisms by the Editor, Matthew Jackson, and by Fedor Sandomirskyi, have greatly improved the readability of this survey.

1 The problem and the punchline

By *Fair Division* (FD), we mean the problem of allocating among a given set of participants a bundle of items called the *manna*, that may contain only desirable disposable *goods*, but sometimes non disposable undesirable *bads*, or even a combination of goods and bads. To fix ideas think of family heirlooms shared by the siblings, of the assets of a dissolving partnership (a marriage, a bankrupt firm), or of jobs (family chores, teaching loads, contracts) that a given set of substitutable workers is responsible for.

The participants (beneficiaries if we divide goods, liable agents if we allocate bads) have equal rights to the good manna (resp. responsibilities toward the bad one). The division rules we discuss can typically be adjusted to account for unequal rights, but we stick for simplicity to the paramount case of equal rights/liabilities for all participants.

Fair Division is a special case of the "commons" problem where the manna can be any kind of resources in common property, such as a technology, natural resources, human capital, etc.. The agents apply some effort to derive benefits from the resources, and individual allocations of the desirable outputs must depend "fairly" upon their profile of individual inputs: the literature on the exploitation of such a general commons (e. g., Moulin (1995)) is much less developed than the FD literature, where the manna shows up without any production effort.¹

Another variant of FD assume that, in addition to the items in the manna, cash is available in any amount and agents have quasi-linear utilities. The efficient division of the manna is then essentially unique, and the focus is on the fair compensations of agents who do not get much manna. A new rule applying the Shapley value to this model challenges the Competitive and Egalitarian rules (Moulin (1992)).

About individual preferences, we make three critical assumptions. First, agents are *selfish*: they only care about their own share of the manna, not at all about how the rest is divided among others; their preferences exhibit no altruism, or spite. Second, everyone bears full individual responsibility for his or her tastes; no one "needs" a bigger share of a particular item (or of the entire manna) because without it his welfare would be unjustly low; no one has a claim on a particular share of the resources.² Third, related to the second point, our division rules only take into account the profile of ordinal preferences, i. e. the way each agent orders all his/her potential allocations. We routinely use utility functions to represent conveniently such preferences, but their intensity is immaterial, it cannot be measured objectively hence must be ignored.³

The punchline What makes Fair Division an interesting and difficult problem is that the participants' preferences over the resources vary, so that a simple equal split of the resources is in general inefficient. In an efficient division, the shares must be unequal, so to call them "fair" requires a convincing - and not necessarily simple - definition of fairness. There is no single compelling definition but the extensive academic literature a) converged in the early 1980s on a handful of key tests, each capturing a precise normative statement that we can agree is relevant to fairness; and b) revealed the logical incompatibilities

¹Note that the website SPLIDDIT discussed in Section 6 offers solutions for two such commons problems: Split Fare, and Assign Credit.

²So we do not cover the celebrated and popular models of rationing and bankruptcy, (O'Neill (1982), Aumann Maschler (1985), Thomson (1985)) where agents are distinguished by their objective individual "claims" over the resources.

³If reporting my utility as u_1 or $2u_1$ did matter to my final share, I would have a transparent strategic manipulation: see e. g., Exercise 2.8 in Moulin (1988).

between these tests, efficiency and incentive compatibility.

Our subject here is the broad renaissance of interest in Fair Division rules that started about three decades ago, and expanded sharply twenty years ago when computer scientists joined microeconomists. Leaving aside the aggressively general models of the early mechanism design literature replete with impossibility results (Hurwicz (1978), Maskin (1977-published 1999)), the focus turns to allocation problems with much more structure, where participants can easily report their preferences because they vary in a domain of very low dimension, and where there is a chance to identify simple division rules with rich normative properties, including sometimes aligning incentives with both efficiency and fairness, a feat that is normally impossible in rich preference domains (Section 4). The theory of Fair Division is now much closer to practical applications.

Contents Steinhaus' (1948) paper on dividing a cake fairly is to my knowledge the oldest mathematical exercise of the mechanism design literature: it is reviewed in Section 2, in particular because computer scientists recently solved one of its long standing open questions.

From the late 1950s to the early 1990s economists discussed Fair Division in the context of Arrow Debreu's economies, and their conceptual insights followed closely the development of theories of Distributive Justice in political philosophy (Sen (1970), Rawls (1971), Dworkin (1981a,b), Roemer (1996), Fleurbaey (1996)). They ended up promoting two families of rules, one welfarist (the Egalitarian equivalent rules), the other resourcist (the Competitive rule): see Section 3. This work is very general because it allows arbitrary preferences in the vast Arrow Debreu domain; it is very short on applications.

Shortly after its birth in the early 1990s, the Internet fostered many peer to peer interactions to divvy computing resources, files, data, reputation scores, and many other commodities; to regulate such interactions, which mostly exclude the exchange of money, we need division rules that are both transparent and agreeable, in other words, fair. Thus concepts of fairness and their interaction with efficiency became relevant to the Computer Science (CS) community, that brought to the debate its own methodology and normative concerns. An example of the latter is the evaluation of the computational complexity of division rules; an example of the latter is the systematic quest for numerical evaluations of the tradeoffs between normative requirements proven to be incompatible.

General Arrow Debreu (AD) preferences were dismissed at once as impractical, because no real person can perceive, let alone formulate, the complex pattern of indifference surfaces that the AD domain allows. Instead the focus is on FD problems with enough structure that individual characteristics are described by a simple message (a single number, a simple ranking). The rule should also be scalable, i. e., easily computed even for a large number of participants or items to divide.

In Section 4 we describe three problems where a "perfect" rule can be found: it combines fairness, efficiency and incentive compatibility in the strong sense of

Strategyproofness. They are: the division of a single non disposable commodity (subsection 4.1); and the division of multiple goods when preferences are linear and dichotomous (subsection 4.2), or exhibit perfect complementarity (subsection 4.3). Interestingly in each case the ideal rule follows a simple egalitarian principle.

The next four Sections collect the most important new results and open questions of the recent Fair Division literature. Each Section rests on a "practical" restriction of the domain of individual preferences.

In Section 5 agents must be *assigned* to objects (one object at most per person) and individual reports are a simple ordering of the objects from best to worst. Fairness is achieved by randomization (or time-sharing). This model informs many applications: the allocation of rooms on campus, of students to high schools, and of organs for transplant. The theoretical results develop around two practical solutions: the Random Priority, and Probabilistic Serial rules.

In Section 6, preferences are linear, represented by additive utilities, thus ruling out complementarities between the commodities in the manna. The Competitive rule is compelling when the manna contains only goods, in particular because it maximizes the Nash product of utilities. However if the manna is made of bads (or a mixture of goods and bads) multiple competitive divisions is a frequent occurrence, and the Egalitarian approach offer plausible alternative rules.

The contribution of the CS community is especially prominent in Section 7, where the task is to allocate indivisible objects, without randomization or cash compensations, and utilities are simply additive over the objects (the cake division model of Section 2 is a limit case). We look for a good approximation of the Egalitarian and Competitive rules. This is easy for the former but not for the latter rule. Even approximating the Envy Free property proves challenging.

The final Section 8 collects open questions, in addition to the ones we encountered along the way, thus suggesting future directions for research.

2 Cake cutting: an algorithmic story

Divide and Choose (D&C) appears in the Bible (Chapter 13 of the Book of Genesis), and is likely the oldest Fair Division rule in recorded history. If the Divider splits the pile of manna in two parts she consider equally valuable, the outcome is fair to her because she gets as good a share as the Chooser's; it is fair as well to the Chooser because no matter what the Divider does, he can pick a share that is at least as good to him as the Divider's share. Each agent can ensure that he/she prefers his/her share to the other agent's share, and if this does not happen, he/she brought it upon him/her-self.

The property above is the well known *Envy Free* test which plays a key role throughout this review, and is closely related to the Competitive approach. On the practical side, the utter simplicity of the messages by the two participants is very appealing. Chooser needs only compare two shares, a *query*, while Divider

needs only to report a *cut*: neither of them must report anything more about own preferences.

The mathematician Hugo Steinhaus (1948) was the first to propose a cake-cutting rule among any number of agents preserving this simplicity of messages and implementing a weaker test of fairness than Envy Free. Let n agents divide the cake C , a measurable set in an Euclidian space, and let u_i be an atomless non negative measure on C representing the preferences of agent $i, i \in N$: so $u_i(z_i)$ is his utility for consuming the share z_i , where z_i any (measurable) subset of C . Importantly these utilities are *additive*: $u_i(z_i \cup z'_i) = u_i(z_i) + u_i(z'_i)$, just like in Sections 6 and 7 below.

We call *Fair Share* (FS) the test that Steinhaus' rule is designed to ensure. The measurable partition $(z_i)_{i \in N}$ of C guarantees FS iff

$$u_i(z_i) \geq \frac{1}{n}u_i(C) \text{ for all } i \in N$$

Envy Free (EF) requires instead

$$u_i(z_i) \geq u_i(z_j) \text{ for all } i, j \in N \tag{1}$$

and by additivity of $z_i \rightarrow u_i(z_i)$ we see that EF implies FS (but not vice versa).

Steinhaus' rule uses only cuts and queries by the agents themselves. *Round 1*. Agent 1 cuts a share z_1 which he claims is worth exactly $u_1(z_1) = \frac{1}{n}u_1(C)$; this share is offered to agent 2: if 2 says $u_2(z_1) > \frac{1}{n}u_2(C)$ she must claim a smaller share $z_2 \subset z_1$ that she says is worth $u_2(z_2) = \frac{1}{n}u_2(C)$, then this piece is offered to agent 3; if 2 says $u_2(z_1) \leq \frac{1}{n}u_2(C)$ then z_1 , still claimed by 1, is offered to agent 3; and so on until every agent has been offered (some reduction of) the initial share, at which point the remaining share goes to the last agent who claimed it (e. g. agent 1 gets z_1 if nobody touched it). *Round 2* repeats Round 1 but with a smaller cake and one less agent. And so on.

If in this algorithm agent i always answers queries and select cuts truthfully (following her actual preferences), she is guaranteed to end up with a share worth $\frac{1}{n}u_i(C)$ or more.⁴ This is a substantial incentive property (shared by D&C), which is however much weaker than Strategyproofness discussed in Sections 4 and 5: a well informed Divider can anticipate the Chooser's reactions and use this knowledge to achieve much more than his FS utility.⁵

Here is a long standing question in the mathematical theory of Fair Division: can we design in the cake-cutting model an algorithm working by cuts and queries and guaranteeing an Envy Free share to any truthful agent?

With three agents, the relatively simple algorithm invented by Selfridge and Conway (quoted by Robertson & Webb (1998)) achieves exactly this. It works as follows.

⁴Simply because if she does not get a share in the first round, the cake remaining in round 2 is worth at least $\frac{n-1}{n}u_i(C)$ to her.

⁵The *moving knife* procedure, due to Banach and Knaster and reported in the follow up paper Steinhaus (1949), implements similarly Fair Share but it is much less efficient than the above rule because the path of the knife is arbitrary.

Agent 1 cuts C in three parts he views as of equal value. Agents 2 and 3 tell which part they prefer from the three (reporting indifferences is OK). If they can both get (one of) their best share, they get it and agent 1 gets the remaining piece. Then nobody is envious (agent 1 included, if he did not lie). Now assume shares are $\{z_1, z_2, z_3\}$ and agents 2,3 both pick z_1 and nothing else. Also label z_2, z_3 so that $u_2(z_2) \geq u_2(z_3)$. Then we ask agent 2 to trim z_1 to $z'_1 \subset z_1$ such that $u_2(z'_1) = u_2(z_2)$. Next agent 3 picks a share in $\{z'_1, z_2, z_3\}$: if he picks z'_1 then agent 2 picks z_2 ; if 3 picks z_2 or z_3 then 2 must choose z'_1 . We have now allocated z'_1, z_2 and z_3 and nobody is envious in the division of this smaller cake.

We must still divide carefully the trimmed part $z_0 = z_1 - z'_1$. Call i the agent in 2,3 who did get z'_1 , and j the other in 2,3. We ask agent j to divide z_0 in three parts of equal value, then agents $i,1$, and j pick one of those parts, in that order. Note that 1 is not envious of i because i got a subset of z_1 while 1 got more than either z_2 or z_3 , that he values more than z_1 . That neither i nor j is envious is clear.

This algorithm delivers a non envious division of the cake in at most five cuts and six queries. It stands to reason that the number of cuts and queries will grow fast when we attempt to generalize it to an arbitrary number of agents: in fact this growth is at least exponential as shown by Branzei (2015).

Brams & Taylor (1995,1996), a political scientist and a mathematician, constructed the first such generalization, in which the length of the algorithm can be unbounded. Recently Aziz & Mc Kenzie (2016), two computer scientists, found an algorithm with a finite (albeit enormous) number of cuts and queries.

One serious problem of these algorithms, already apparent in Selfridge and Conway, is that they generate "crumbs". The S&C algorithm cuts the cake in six pieces and everyone gets two pieces. The number of such pieces grows very fast with the number of agents and some individual pieces end up microscopic. Thus these algorithms are useless to divide a piece of land, which motivates a stream of research about cutting the cake in connected pieces, or pieces with specific geometric properties (Berliant et al. (1992), Aumann & Dombb (2010), Segal-Halevi et al. (2017)).

Another drawback of D&C, Selfridge Conway, and the two general algorithms just mentioned, is that they pay no attention to efficiency. The Divider needs no information about the Chooser's preferences to cut the cake, and after his cut the only two possible partitions are typically inefficient. Of course eliciting an efficient division of the cake would require interpersonal comparison of preferences, which the "cut and query" format does not allow. The existence of algorithms delivering an efficient, or nearly efficient cake cutting with "simple" messages is a widely open question.

The mathematical theory of Cake cutting offers many more deep and more abstract results, early instances of which are in Dubins & Spanier (1961). A consequence of the non atomicity of the utility measures and Lyapounov's theorem is the existence of a "perfect" division of the cake, where everyone finds *all* pieces of equal value. In a different vein, Su (1999) proves the existence of a non envious division of a linear cake with only a mild continuity assumption on preferences. A good survey is Procaccia (2013).

3 Two theories of Distributive Justice

We review the key findings of the microeconomic theory of Fair Division until the end of the 1980s.

3.1 Envy Free and the Competitive rule

In an Arrow Debreu (AD) economy, the manna is a bundle of desirable and divisible commodities ω in \mathbb{R}_+^K , where K is the set of commodities. Agent i 's share z_i is non negative and the division is feasible: $\sum_N z_i = \omega$. Agent i 's preferences are convex, continuous, and monotonic (MasColell et al. (1995)) and represented by a utility function u_i .

Fair Share means

$$u_i(z_i) \geq u_i\left(\frac{1}{n}\omega\right) \text{ for all } i \in N$$

and is always feasible.⁶ The original Envy Free test (1), due to the economist D. Foley (1967), was quickly adopted by mathematicians (Section 2), political philosophers and other social scientists for its normative simplicity. To economists, the key observation is that it is easy (under the assumptions above on preferences and divisibility of the goods) to combine Efficiency and Envy Free: simply select a competitive equilibrium in the exchange economy where each agent is endowed with the initial allocation $\frac{1}{n}\omega$ (Varian (1974)). For brevity we call such allocations *Competitive (C)*.⁷

The three concepts, Fair Share, Envy Free, and Competitive allocations, interpret fairness as the allocation of *virtual rights* to each participants. Fair Share protects my welfare by giving me the right to claim the default "equal split allocation". If, after the manna is distributed, I can claim anyone else's share in lieu of the share I received, Envy Free allocations are the only ones on which we can agree. In a Competitive division, I own the bundle $\frac{1}{n}\omega$ and can trade it freely at the competitive price. Finally a Competitive division is *Core Stable*: no coalition of agents could pool its endowment and, standing alone, distribute better shares to each member of the coalition.

We keep in mind that these rights are only virtual, because a division rule is a simple black box that delivers a final allocation and bypasses direct interactions between agents, strategic or otherwise. The role of properties like FS, EF, C or the Core, is to convince the agents to accept the outcome of the rule, thus avoiding the potentially large transaction costs of direct negotiations.

A Competitive allocation always exists in the general AD domain. However its computation is difficult and, more importantly, we may have multiple Competitive allocations with very different welfare consequences (some examples are

⁶If utilities are concave and $u_i(0) = 0$, this requirement is stronger than $u_i(z_i) \geq \frac{1}{n}u_i(\omega)$. If utilities are only quasi-concave, the latter inequalities may not be feasible.

⁷That is, an allocation z such that $\sum_i z_i = \omega$, and for some non zero price $p \in \mathbb{R}_+^K$, agent i can afford to exchange $\frac{1}{n}\omega$ for z_i , and weakly prefers z_i to any allocation y_i she can also afford.

in Section 6). It is impossible to identify a reasonable selection, for instance one that varies continuously when the endowment ω or individual preferences vary. Our goal is to define single valued division rules, at least welfarewise: the Competitive rule is a compelling answer in any subdomain of the AD domain where it is unique and continuous in the parameters of the problem; in other domains, it is still a useful approach to fairness, but may need an additional selection step to be guided by additional normative principles.

3.2 Egalitarian Equivalence

The Competitive approach is a theory of Justice based on equal opportunity *ex ante* (the common budget set of all individual allocations) and virtual free trade (Dworkin (1981b)). Influenced by Rawls' Theory of Justice, Pazner and Schmeidler (1978) propose an alternative theory in which we equalize ex post some virtual measure of individual welfare. The trick is to choose a neutral calibration of "welfare" that does not discriminate against particular agents.

A feasible division $z = (z_i)_{i \in N}$ of the manna is *Egalitarian Equivalent* (EE) if there exists an individual allocation z_0 such that each agent i is indifferent between z_i and z_0 .

Many EE and efficient allocations exist. To extract a single valued profile of welfares, we choose a calibration line in the consumption space, and restrict the benchmark allocations z_0 to that line. The most popular choice (as explained below) is the line borne by the manna ω . It defines a division rule that we write $EE(\omega)$, selecting a feasible efficient division z such that

$$\text{for some } \lambda > 0: u_i(z_i) = u_i(\lambda\omega) \text{ for all } i \in N \quad (2)$$

Existence of such allocation only requires preferences to be continuous and (strictly) monotonic. Uniqueness in welfare is then guaranteed.⁸

Alternative choices of the benchmark vector lead to the rules $EE(e^k)$ (where e^k is the coordinate vector of good k , interpreted as "numeraire") or $EE(\theta)$ for an arbitrary positive bundle θ . In equation (2) above, simply replace ω by e^k or θ .

So the Egalitarian rules, unlike the Competitive one, adjust well to non convex preferences, but rely critically on their monotonicity. Contrast this with the Competitive rule, well defined for non monotonic preferences, as we illustrate in Section 6, and even more general preferences.⁹

Beyond the differences in their existence properties, the two approaches capture radically different interpretation of FD: equal ownership of the manna, versus equal right to consume the manna. The following simple example illustrate this contrast.

⁸With monotonic but not strictly monotonic preferences, we can still define agent i 's virtual welfare at z_i as the smallest λ_i such that $u_i(z_i) = u_i(\lambda_i\omega)$, and maximize the leximin ordering over all feasible virtual welfare profiles. . Then $EE(\omega)$ is single-valued if the set of such profiles is convex.

⁹The most general existence results require only an improvement correspondence with an open graph: see Mas-Colell (1992) and McLennan (2018).

The manna contains \$210 and one indivisible object with cash value $90 > 75 > 15$ to the three agents 1,2,3 respectively. Utilities are linear in money. Efficiency requires that agent 1 gets the object, but how should the money be divided, so as to give a fair compensation to 2 and 3?

The Competitive rule picks a price p for the object between 90 and 75 (the price of cash is 1);¹⁰ 2 and 3 each receive, in cash, one third of $210 + p$, the competitive value of the manna, so between \$95 and \$100, and agent 1 keeps between \$20 and \$10 of the cash. Even though agent 3 derives little utility from the object, he is entitled to a "rent" of at least \$25. On the other hand the Egalitarian solution $EE(\omega)$ gives cash to 2 and 3 in proportion to \$285 and \$225, their total valuations of the manna, and agent 1 gets the same proportion of \$300: so 2 gets \$105.5, 3 gets \$83.3 and 1 keeps \$20.7.

3.3 Solidarity properties

A Fair Division rule makes a recommendation for every possible choice of the manna, of the set of beneficiaries, and of their preferences. From the next Section onward the set of preferences that a given rule can accommodate, its *preference domain*, is its main defining characteristic. The domains we consider are much smaller than the infinite dimensional Arrow Debreu domain.

The fairness tests discussed so far apply to a fixed FD problem (a single manna and profile of preferences), they do not compare the recommendation of the rule across different problems. But suppose that two problems are "clearly comparable", by which we mean that a single parameter of the problem changes, and this move has clear welfare consequences. Then the corresponding recommendations of the rule better be comparable as well, lest an objection arises against the division in problem 2 based on the precedent of the division in problem 1. Moreover, properties like the two we now introduce provide insights into the logic behind how a rule works, instead of placing restrictions on the outcome for a particular set of parameters like the Envy Free test.

One instance of clear comparability is when the beneficiaries of the manna do not change but there is more (good) manna to share in one problem than in the other. This is good news for the agents as a group, so if an agent ends up with a worse share of the better manna, she has a legitimate objection: we all have equal rights to the manna therefore the happy shock on the resources should be good news for everyone. Moulin & Thomson (1988) introduce the corresponding fairness test in the case of a good manna.

Resource Monotonicity (RM): $\omega \leq \omega' \implies u_i(z_i) \leq u_i(z'_i)$ for all $i \in N$

where the rule picks z or z' when the manna is ω or ω' . Of course if we divide a bad manna, the second inequality should simply be reversed.

The second monotonicity test compares the outcomes of two problems that only differ in the presence of one particular agent. As there are fewer of us to

¹⁰It is not unique because the object is indivisible.

share the same cake (good manna), this should be good news for each of us:

Population Monotonicity (PM): $N' = N \setminus \{i^*\} \implies u_i(z_i) \leq u_i(z'_i)$ for all $i \in N'$

where the rule picks z or z' when agent i^* is present or not. (we reverse the inequality if the manna is bad). The axiom is due to Thomson (1983).

These two properties capture quite nicely the idea of solidarity between our agents: the adjustment to exogenous shocks in the data of the problem should not create winners and losers.

Although RM and PM are quite strong requirements there are many efficient division rules meeting both, in particular the Egalitarian rule $EE(\theta)$ just defined, for any benchmark vector θ . And there are many more, for instance the *money metric* rules $MM(p)$ that we now define.

Fix the strictly positive price vector p , and find the largest *virtual budget* b such that for some feasible allocation $z = (z_i)_{i \in N}$, we have

$$u_i(z_i) = \max_{y: p \cdot y_i \leq b} u_i(y_i) \text{ for all } i \in N$$

Then $MM(p)$ outputs allocation z (single valued at least welfarewise).

To check that $EE(\theta)$ and meet RM, observe that if $z = (z_i)_{i \in N}$ is feasible at ω and $u_i(z_i) = u_i(\lambda\theta)$ (resp. $u_i(z_i) = \max_{p \cdot y_i \leq b} u_i(y_i)$) for all i , then at ω' the same benchmark level λ (resp. the same budget b) is feasible as well, so the largest feasible λ' (or the largest feasible b') increases. The proof of PM is just as easy.

But all rules $EE(\theta)$ and $MM(p)$ fail Fair Share!

Consider a two-agent, two-good problem with the manna $\omega = (1, 3)$, the benchmark bundle $\theta = (3, 1)$ and the utilities

$$u_1(x_1, y_1) = x_1 + 3y_1 ; u_2(x_2, y_2) = 3x_2 + y_2$$

Efficiency requires $x_1 \cdot y_2 = 0$ and $EE(\theta)$ chooses a division such that

$$\frac{3y_1}{u_1(3, 1)} = \frac{3x_2 + y_2}{u_2(3, 1)} \implies z_1 = (0, 1), z_2 = (3, 2)$$

so that $u_1(z_1) = u_1(\frac{3}{10}\omega) < u_1(\frac{1}{2}\omega)$.

Next choose $p = (1, 3)$ and check that the maximal virtual budget is $b = \frac{9}{5}$, and that the corresponding money metric allocation gives $z_1 = (0, \frac{3}{5})$ to agent 1, violating FS.

On the other hand the rule $EE(\omega)$ meets PM and FS but violates RM. The deeper reason is a systematic impossibility result: *if the domain of preferences contains Leontief preferences (goods are perfect complements: see below and Subsection 4.3), no efficient and Resource Monotonic rule can guarantee Fair Share. (Moulin and Thomson (1988)).*

The proof uses again a simple two-agent, two-good problem with $\omega = (3, 3)$ and Leontief utilities

$$u_1(x_1, y_1) = \min\{\frac{1}{2}x_1, y_1\} ; u_2(x_2, y_2) = \min\{x_2, \frac{1}{2}y_2\}$$

Suppose a rule meets FS and RM and selects the division z . Consider the problem with a larger manna $\omega' = (6, 3)$ where it picks z' . By FS and feasibility we have

$$u_1(z'_1) \geq u_1(3, \frac{3}{2}) \implies z'_1 \geq (3, \frac{3}{2}) \implies z'_2 \leq (3, \frac{3}{2}) \implies u_2(z'_2) \leq \frac{3}{4}$$

Now RM implies $u_2(z_2) \leq \frac{3}{4}$. A symmetrical argument yields $u_1(z_1) \leq \frac{3}{4}$ so the division z is inefficient because $\tilde{z}_1 = (2, 1), \tilde{z}_2 = (1, 2)$ is feasible and Pareto superior to z .

In the linear preference domain of Section 6 and in the domain of Cobb Douglas utilities, the Competitive rule (for goods) meets FS, RM and PM. Interestingly in the subdomain of Leontief preferences, although FS and RM are incompatible as we just showed, the incentive compatibilities of the Egalitarian rules $EE(\omega)$ and $EE(\theta)$ are very attractive: see subsection 4.3.

4 Practical division rules: three compelling examples

The two theories above aim at great generality, but we already mentioned that real participants in a Fair Division exercise cannot form preferences as complex as the AD domain allows, let alone report them.

We describe here three problems with a preference domain simple enough for their practical implementation: respectively a positive number, a subset of a common finite set, and a line in the positive orthant. In each case we find one or more division rule combining Efficiency, Fairness and Incentive Compatibility, a rare treat in the mechanism design literature.

In the first two models the canonical rule is Competitive and has an egalitarian interpretation, though not in the sense of the Egalitarian Equivalent family $EE(\theta), EE(\omega)$ discussed in Section 3. In the third model the rules of interest are precisely the Egalitarian Equivalent ones of Section 3, and they are not Competitive.

4.1 One non disposable item

The manna ω is a positive number, the amount of a *non disposable* item: 20 hours of baby sitting between family members, 200 shares of a stock between investors, 20 students between teachers, 200 identical cars between car dealers, etc..

Agents' preferences over their share are convex over $[0, \omega]$, with a single maximum, which gives the familiar single-peaked shape. Agent i 's most preferred share is $\pi_i, 0 \leq \pi_i \leq \omega$; her preferences go up from the share 0 to π_i , then down from π_i to ω . So $\pi_i = \omega$ means that i wants as much of the item as possible, and $\pi_i = 0$ as little as possible: the item can be a genuine good for some, and a genuine bad for others.

If in a feasible allocation $z = (z_i)_{i \in N}$ of ω between the agents in N , we can find i, j such that $z_i < \pi_i$ and $z_j > \pi_j$, then a (small) transfer from j to i is Pareto improving. Hence the efficient allocations are described very simply. We say that the manna is *overdemanded* if $\sum_i \pi_i > \omega$, in which case z is efficient if and only if $z_i \leq \pi_i$ for all i . If the manna is underdemanded, $\sum_i \pi_i < \omega$, the efficient allocations are characterized by $z_i \geq \pi_i$ for all i . And if $\sum_i \pi_i = \omega$ the unique efficient allocation gives his/her peak to everyone.

The canonical *uniform division rule* proposed by Sprumont (1991) starts from the equal split (fair share) allocation $\bar{z}_i = \frac{1}{n}\omega$ for all i ; it stays there if the peaks π_i are all on the same side of the fair share. If this is not the case, write N^{under} and N^{over} for the non empty sets of *underdemanding* agents, $\pi_i < \frac{1}{n}\omega$, and of *overdemanding* agents, $\pi_j > \frac{1}{n}\omega$, respectively. Finally N^{fs} is the set of agents k whose peak is “just right”: $\pi_k = \frac{1}{n}\omega$.

The uniform rule gives $z_k = \frac{1}{n}\omega$ to each $k \in N^{fs}$, as required by Fair Share. It is feasible and Pareto improving to decrease z_i for underdemanding agents, and increase z_j for the overdemanding ones. If the manna is overdemanded, efficiency and feasibility imply that each $i \in N^{under}$ gets her peak, $z_i = \pi_i$, while agents j in N^{over} get $z_j \in [\frac{1}{n}\omega, \pi_j]$: the latter agents cannot all get their peak. The uniform rule equalizes the gains $(z_j - \frac{1}{n}\omega)$ for those agents, as much as allowed by the constraint $z_j \leq \pi_j$: so $z_j = \min\{\lambda, \pi_j\}$, where λ is determined by feasibility.¹¹

Symmetrically, if the manna is underdemanded, agents in $N^{over} \cup N^{fs}$ get their peak while those in N^{under} get $z_i = \max\{\mu, \pi_j\}$, where μ is determined by feasibility.

The uniform rule is Competitive: we set a price of 1 and a budget of λ in the case of overdemand, and a price of -1 with a budget $-\mu$ (I can buy μ units or more) if there is underdemand. In particular, the allocation is Envy Free.

The uniform rule has two attractive features. First, the message of each agent is a single number, the peak of one’s preferences: he does not need to report, or even conceptualise, how allocations compare across the peak of his preferences. The second property is stated in the Theorem below.

A division rule elicits messages from the agents (related to their preferences) and outputs an allocation for each agent. The rule is *Strategyproof* (SP) if, once the messages of other agents are fixed, reporting my “truthful” message (here the peak of my preferences) never results in a worse outcome for me than any *misreport*. It is *GroupStrategyproof* (GSP) if no subgroup of agents can coordinate their misreport so that they all end up no worse than from telling the truth, and at least one of them is strictly better off.

The rule is Anonymous if it does not discriminate among agents on the basis of their names: formally, if two agents exchange their messages, ceteris paribus, their allocations are exchanged as well and other agents’ allocations do not change. All the rules showcased in this Section and later are Anonymous.

Theorem (Sprumont (1991), Ching (1994)): *The uniform rule is Anonymous and GroupStrategyproof. No other Efficient and Anonymous rule is Group-*

¹¹It solves the equation $\sum_{i \in N^{under} \cup N^{fs}} \pi_i + \sum_{j \in N^{over}} \min\{\lambda, \pi_j\} = \omega$.

Strategyproof.

Strategyproofness is clear from the description of the rule above. Assume underdemand: only an underdemanding agent may not get her peak allocation; pretending to be overdemanding she gets an allocation no smaller than $\frac{1}{n}\omega$, therefore worse than the truthful allocation z_i . If she still reports as underdemanding, the only way to change z_i is by reporting π'_i in $]z_i, \frac{1}{n}\omega[$, then z'_i is also in $]z_i, \frac{1}{n}\omega[$ and she is worse off.

There are many other Efficient and Anonymous rules. If M is the set of agents on the long side (N^{over} – resp N^{under} – if the manna is over – resp under – demanded) their shares could be proportional to $|\pi_i - \frac{1}{n}\omega|$, or we could equalize as much as possible $|\pi_i - z_i|$ and so on. But all such rules are vulnerable to strategic misreports.

Note finally that neither RM not PM apply to this model as preferences are not monotonic.

The literature provides alternative characterizations of the uniform rule and explores the large family of Efficient and GSP, but non Anonymous rules: Barbera et al. (1997), Moulin (1999)). See also a generalization of the model encompassing Voting, Fair Division, and more in Moulin (2017).

4.2 Dichotomous additive utilities

As mentioned before the Theorem above, the uniform rule only requires agents to identify their top share of the manna, not the potentially complex comparisons across the peak allowed in the single peaked domain. In our second model the manna is a set of semi-homogenous goods, in the sense that each agent likes only a subset of these, but does not distinguish between the goods that she likes. This subset fully describes her preference relation.

Examples include patients of different blood types sharing blood also of different types; workers sharing time on several machines delivering identical service but requiring specific handling skills, partners dividing clients who speak different languages, as do the partners, assignment of classrooms to charter schools (Kurokawa et al. (2015)), and organs for transplants (Roth et al. (2004)).

Let A be the set of goods. The profile of preferences is represented by a matrix $u = [u_{ia}]_{N \times A} \in \{0, 1\}^{N \times A}$, where $u_{ia} = 1$ (resp. 0) means that agent i “likes” (dislikes) good a . Recall the cardinal intensity of utility has no meaning in this or other models of the paper: but the common calibration at $u_{ia} = 1$ proves very useful for describing the canonical rule.

There is a positive amount of each good so the manna is $\omega \in \mathbb{R}_{++}^A$. A feasible division of ω is, as usual, $z = (z_i)_{i \in N}$ such that $\sum_N z_i = \omega$ and $z_i \geq 0$.

Notation: for any $S \subseteq N, X \subseteq A$, $\zeta_S = \sum_{i \in S} \zeta_i$, $\xi_X = \sum_{a \in X} \xi_a$ and $\chi_{SX} = \sum_{i \in S, a \in X} \chi_{ia}$.

It is convenient to represent agent i 's preferences by his “like set” $L_i = \{a \in A | u_{ia} = 1\}$. Without loss we assume that each L_i is non empty and each good is liked by at least one agent. Then agent i 's utility at z is $U_i = u_i \cdot z_i = z_{iL_i}$.

If the utility profile U is feasible at ω , it satisfies the inequalities

$$U_S \leq \omega_{L(S)} \text{ for all } S \quad (3)$$

where $L(S) = \cup_{i \in S} L_i$. By Hall's theorem these inequalities together characterize feasible profiles U .

Efficiency of the division z requires $z_{ia} = 0$ if $a \notin L_i$: goods are eaten only by agents who like them. This implies $U_N = \omega_A$. Conversely, U is efficient if and only if it is feasible ((3)) and $U_N = \omega_A$.

Therefore the efficient frontier of our problem is described as the core of the cooperative game (N, v) where $v(S) = \omega_{L(S)}$ for all S . This game is clearly concave¹² therefore its core contains a canonical *Lorenz dominant* profile U^* (Dutta & Ray (1989)) equalizing the utilities U_i as much as permitted by the feasibility constraints (3). So U^* maximizes the smallest individual utility, and for all $s, 1 \leq s \leq n-1$, U^* it also maximizes the smallest sum U_S when S is of size s :¹³

$$\min_{S:|S|=s} U_S^* = \max_{U \text{ feasible}} \min_{S:|S|=s} U_S$$

There is a simple algorithm to compute U^* and the corresponding division of the goods, that also reveals why U^* is the Competitive allocation. Find first the largest set of agents S^1 in N such that

$$S^1 \in \arg \min_{S \subseteq N} \frac{\omega_{L(S)}}{|S|} \quad (4)$$

If S and T both solves the program above, so does $S \cup T$, therefore the ‘‘largest’’ solution is well defined. If $S^1 = N$, property (4) means that the equal split profile $U_i^* = \frac{1}{n}\omega_A$ is feasible, and it is of course the most egalitarian division. If, on the contrary $S^1 \subsetneq N$, then $\frac{\omega_{L(S^1)}}{|S^1|} < \frac{1}{n}\omega_A$. Agents in S^1 share all they can eat (the goods in $L(S^1)$), and it is feasible (by Hall's Theorem) to divide those goods so that each end up with the same utility

$$U_i^* = U^{*1} = \frac{\omega_{L(S^1)}}{|S^1|} \text{ for each } i \in S^1$$

Next we find S^2 , the largest set such that

$$S^2 \subseteq N \setminus S^1 \text{ and } S^2 \in \arg \min_{S \subseteq N \setminus S^1} \frac{\omega_{L(S) \setminus L(S^1)}}{|S|}$$

and those agents can (by Hall's Theorem) share the goods in $L(S^2) \setminus L(S^1)$ to achieve the common utility

$$U_i^* = U^{*2} = \frac{\omega_{L(S^2) \setminus L(S^1)}}{|S^2|} > \frac{\omega_{L(S^1)}}{|S^1|} \text{ for all } i \in S^2$$

¹²For any $S \subset T \subset N$ and any i we have $v(T \cup i) - v(T) \leq v(S \cup i) - v(S)$.

¹³It also maximizes the sum $\sum_N f(U_i)$ for any concave function f .

The last inequality is critical. If the opposite inequality was true, it would imply $\frac{\omega_{L(S^2) \setminus L(S^1)} + \omega_{L(S^1)}}{|S^2| + |S^1|} \leq \frac{\omega_{L(S^1)}}{|S^1|}$ and contradict the definition of S^1 .

Repeating this argument we build a partition S^1, \dots, S^K of N and give to S^k all the goods in $X^k = L(S^k) \setminus L(\cup_{\ell=1}^{k-1} S^\ell)$ (and nothing else), which they (can) divide to achieve equal utility $U^{*k} = \frac{\omega_{X^k}}{|S^k|}$. The sequence U^{*k} increases strictly.

Give now a budget of 1 to each agent and set the price $p^k = \frac{1}{U^{*k}}$ for the goods in X^k . An agent in S^k does not like any of the goods in $\cup_{\ell=k+1}^K X^\ell$, and those in X^k are the cheapest in $\cup_{\ell=1}^k X^\ell$: therefore this agent buys exactly U^{*k} units in X^k , so the canonical division is indeed the Competitive one.

In the following five-agent, four-good example

	a	b	c	d	
u_1	1	0	1	1	
u_2	1	0	0	0	
u_3	0	1	1	0	(5)
u_4	1	0	1	0	
u_5	0	0	1	1	
ω	1	2	3	6	

the partition of N is $S^1 = \{2\}, S^2 = \{3,4\}, S^3 = \{1,5\}$ and that of A is $X^1 = \{a\}, X^2 = \{b, c\}, X^3 = \{d\}$, with the utility profile $U^* = (3, 1, 2.5, 2.5, 3)$.

The Shapley value of the game (N, v) above is another natural selection in the set of efficient utility profiles, with quite different welfare consequences. In the example it delivers the profile $U^\sim = (4.08, 0.33, 2.75, 1.09, 3.75)$. But the Shapley division rule is not Strategyproof and can generate Envy.

Theorem: *The Competitive rule just described is GroupStrategyproof, picks an Envy Free allocation, and guarantees Fair Share. It is also Resource and Population Monotonic.*

Envy Free is true as always for the Competitive division. RM and PM follow from the same property of the Competitive division in the more general model of Section 6. GSP is proven much like in the assignment variant of the model described below.

The assignment problem Bogomolnaia & Moulin (2004) discuss a familiar variant of the model above with $|A| = |N|$, $\omega = e^A$ and the feasibility constraint $z_{iA} = 1$ for all i . Think of the objects as one day jobs and of agents as substitutable workers who can complete one job per day (whether they like the job or not). An allocation $z = [z_{ia}]_{N \times A}$, written as a $N \times A$ matrix, is feasible if and only if it is bi-stochastic. The entry z_{ia} represent the amount of time worker i is assigned to job a , or the probability that i is assigned to a .

Preferences are as before, and the analysis is essentially the same. The utility profile U is feasible and efficient if and only if

$$U_S \leq \min\{|S|, |L(S)|\} \text{ for all } S \subset N \text{ and } U_N = |A|$$

The new cooperative game (N, w) , $w(S) = \min\{|S|, |L(S)|\}$ is still concave, and the Lorenz dominant profile U^* in its core defines a GSP division rule with a similar competitive interpretation (hence ruling out Envy).

In the next Section we come back to the assignment model with much more complex preferences. The miraculous compatibility of Efficiency, GSP and fairness is then lost.

4.3 Complementary goods: Leontief preferences

In the previous problem two goods that an agent “likes” are perfect substitute. We turn to the polar opposite case where the goods are perfect complements: each agent “needs” all goods to generate utility from the manna. For instance entrepreneurs share a manna made of capital goods, raw materials and labor, each needs all three to open shop but not necessarily in the same proportions. In cloud computing, each user needs a personal combination of memory, computing resources and bandwidth to perform his task (Ghodsi et al. (2011)). And so on.

Formally, A is still the set of goods, and agent i 's preferences are represented by the following utility function, where $(\kappa_{ia})_{a \in A}$ are non negative parameters:

$$u_i(z_i) = \min_{a \in A} \left\{ \frac{z_{ia}}{\kappa_{ia}} \right\}$$

Note that $\kappa_{ia} = 0$ simply means that i does not need good a .

Thus agent i reports the vector $(\kappa_{ia})_{a \in A}$, up to a multiplicative constant. This is more complex than reporting a single number in Subsection 4.1, but still manageable if the number of goods is not large.

The rules of interest in this problem are the Egalitarian ones, $EE(\theta)$ and $EE(\omega)$. Computing them is relatively simple.

A problem is described by N, A , the matrix of parameters $[\kappa_{ia}]_{N \times A} \in \mathbb{R}_+^{N \times A}$, and the manna $\omega \in \mathbb{R}_+^A$. We fix $\theta \in \mathbb{R}_+^A$ and compute $EE(\theta)$. It is convenient to normalize the preference parameters so that for each agent

$$1 = u_i(\kappa_i) = u_i(\theta) \iff \max_a \frac{\kappa_{ia}}{\theta_a} = 1 \quad (6)$$

We are looking for the largest parameter λ such that for some feasible allocation z and all i :

$$u_i(z_i) \geq u_i(\lambda\theta) = \lambda \iff z_i \geq \lambda\kappa_i$$

Feasibility requires $\lambda \sum_j \kappa_{jb} \leq \omega_b$ for all b , therefore the largest λ correspond to a bottleneck good a^* (typically unique):

$$\begin{aligned} \lambda^* &= \frac{\omega_{a^*}}{\kappa_{Na^*}} \text{ where } \frac{\omega_{a^*}}{\kappa_{Na^*}} = \min_b \frac{\omega_b}{\kappa_{Nb}} \\ \implies z_i^* &= \frac{\omega_{a^*}}{\kappa_{Na^*}} \kappa_i ; U_i = u_i(z_i^*) = \frac{\omega_{a^*}}{\kappa_{Na^*}} \text{ for all } i \end{aligned} \quad (7)$$

Note that the allocation z^* is feasible but does not exhaust the non-bottleneck goods: the excess supply of these goods cannot help improve the utility of

anyone, so we may simply dispose of them. This last feature is important for the following incentive statement.

Theorem (Ghods et al. (2011)): *Each division rule $EE(\theta)$ and $EE(\omega)$ with disposal of redundant goods is GroupStrategyproof, and picks an Envy Free allocation.*

Recall from Subsection 3.3 that each $EE(\theta)$ meets RM and PM but fails FS, while $EE(\omega)$ meets FS and PM but fails RM.

Checking Envy Free is easy. By equations (7), agent i envying j requires $z_i \ll z_j \iff \kappa_{ib} \ll \kappa_{jb}$ for all b , contradicting the normalisation (6).

Strategyproofness requires more work. Assume without loss $\omega = e^A$ (the unit vector), and consider agent 1 reporting κ'_1 in lieu of κ_1 . Set $\kappa'_i = \kappa_i$ for other agents, and let z, z' be the corresponding allocations. Recall a^* is the "truthful" bottleneck good, and let b^* be a bottleneck good at κ' . The misreport κ'_1 by the single agent 1 is successful iff $z'_{1b} > z_{1b}$ for all b , which we assume. By (7) and $\omega = e$ this means

$$\frac{\kappa'_{1b}}{\kappa'_{Nb^*}} > \frac{\kappa_{1b}}{\kappa_{Na^*}} \text{ for all } b \quad (8)$$

The normalisation of κ_1 and κ'_1 imply $\kappa'_{1b} \leq \kappa_{1b}$ for some b (possibly a^* itself). This and (8) gives

$$\frac{\kappa_{1b}}{\kappa'_{Nb^*}} \geq \frac{\kappa'_{1b}}{\kappa'_{Nb^*}} > \frac{\kappa_{1b}}{\kappa_{Na^*}} \implies \kappa'_{Nb^*} < \kappa_{Na^*}$$

Now $\kappa'_{Na^*} \leq \kappa'_{Nb^*}$ by definition of b , so we get $\kappa'_{1a^*} < \kappa_{1a^*}$. This, and (8) applied to a^* , gives

$$\frac{\kappa'_{1a^*}}{\kappa'_{Nb^*}} > \frac{\kappa_{1a^*}}{\kappa_{Na^*}} > \frac{\kappa'_{1a^*}}{\kappa'_{Na^*}} \implies \kappa'_{Nb^*} < \kappa'_{Na^*}$$

contradicting the definition of b^* .

The proof of the GSP property is similar.

Note that if the rule must allocate redundant supply of non bottleneck goods to some agents, Strategyproofness is no longer compatible with Efficiency and Envy Freeness (Nicolo (2004)).

Li & Xue (2013) discuss many other Efficient and GSP rules, in particular we can replace the benchmark line along θ by an arbitrary increasing path in \mathbb{R}_{++}^A .

Finally we compute the Competitive division: it is easy in this model because a general result by Eisenberg and Gale (discussed in Section 6) implies that this allocation maximizes the Nash product of utilities. Agent i 's allocation takes the form $z_i = \lambda_i \kappa_i$ (plus some redundant goods) with corresponding utility λ_i . We must solve the program

$$\max \sum_i \ln \lambda_i \text{ such that } \forall a : \sum_i \kappa_{ia} \lambda_i \leq \omega_a$$

The tight constraint(s) $\sum_i \kappa_{i\tilde{a}} \lambda_i = \omega_{\tilde{a}}$ identifies the bottleneck good(s) \tilde{a} for this allocation, then the First Order Conditions imply that $\kappa_{i\tilde{a}} \lambda_i$ is independent of i . Thus the CEEI splits the bottleneck good(s) equally:

$$z_{i\tilde{a}} = \frac{1}{n} \omega_{\tilde{a}} ; z_{ib} = \frac{1}{n} \frac{\kappa_{ib}}{\kappa_{i\tilde{a}}} \omega_{\tilde{a}} ; U_i = \frac{1}{n} \frac{1}{\kappa_{i\tilde{a}}} \omega_{\tilde{a}} = \lambda_i$$

It is easy to check, in a two-agent two-good example, that this division rule is not Strategyproof.

5 Random assignment

Lotteries are the time honored device achieving ex ante fairness when ex post fairness is impossible, in particular when the manna is made of indivisible items. They are routinely used to allocate seats in school, in universities, administrative jobs in ancient China, rooms on campus, overdemanded tickets for shows and games, etc.. In fact, dividing the manna in lots that are approximately equal (in the eyes of the division manager) and assigning them with equal probability to each beneficiary, is a simple and well known Fair Division rule. But it ignores the agents' preferences, hence can be grossly inefficient and does not guarantee Fair Share.

In this Section we discuss the *assignment* problem, in which each agent is to receive, ex post, exactly one "object": students are assigned to one school, get one campus room, workers get one job and so on. If agents can receive any number of objects, we are in the model of the next Section.

An important alternative interpretation of the assignment model is *time sharing*. We have n machines (projects) and n workers; agent i is assigned for a share z_{ia} of the month to machine (project) a : this problem involves no lottery but is formally equivalent to random assignment. To fix ideas we retain the probabilistic terminology.

The meaning of ex ante efficiency and fairness of an allocation, or a rule, is dictated by the domain of preferences over random allocations. One approach, briefly discussed in the last Section 8, assumes von Neuman Morgenstern utilities: each agent reports a cardinal utility function, and compares lotteries by their expected utilities. In many contexts this is too demanding: in the school choice example, parents can not be expected to compare a handful of potential schools in such sophisticated fashion, especially if they have only a vague idea of the probabilities that their student will be accepted in the different schools they apply to. But it is a relatively simple exercise to rank those schools from best to worst.

Here we discuss this latter version of the fair assignment problem. We have n agents in N and n objects in A (it does not matter if they are goods or bads) and everyone is to receive, ex post, one object.¹⁴ An *ex post* assignment ς is a one-to-one mapping of A into N . An *ex ante* assignment z is a bi-stochastic

¹⁴The model and results are easily extended to arbitrary numbers of objects and agents, as long as we maintain the assumption that each agent receives at most one object.

matrix $[z_{ia}] \in [0, 1]^{N \times A}$ in which z_{ia} is the probability that agent i gets object a . By the classic Birkhoff-von Neuman theorem (BvN), every ex ante assignment z obtains as some probability distribution over ex post assignments ς .

Agent i only reports an *ordinal ranking* \succ_i of the objects, and, with the exception of example (9) below, we assume that this ordering is strict (in sharp contrast to Subsection 4.2). Note that the analysis can be extended at some technical cost to allow for indifferences in individual preferences. The “one object per person” constraint is what makes the report of such preferences simple: if agents could get any number of objects they would need to compare nearly 2^n subsets of A .

Of course an ordinal ranking \succ_i of A is not enough to decide how agent i compares all probability distributions z_i on A . It does however induce a partial ordering of $\Delta(A)$, the set of all such distributions.¹⁵ Write $U(\succ_i, a)$ for the upper contour of \succ_i at a . We define the statement “allocation z_i *stochastically dominates* (SD) allocation z'_i ”:

$$z_i \succ_i^{sd} z'_i : z_{iU(\succ_i, a)} \geq z'_{iU(\succ_i, a)} \text{ for all } a \in A$$

$$z_i \succ_i^{sd} z'_i : z_i \succ_i^{sd} z'_i \text{ and } z_{iU(\succ_i, a)} > z'_{iU(\succ_i, a)} \text{ for some } a$$

With this partial ordering in hand, we adapt in the obvious way our concepts of efficiency, fairness and incentives. The assignment z meets Fair Share if $z_i \succ_i^{sd} \frac{1}{n} e^A$ (the unit vector) for all i ; Envy Free if $z_i \succ_i^{sd} z_j$ for all i, j , and Efficiency¹⁶ if $\{z_i \succ_i^{sd} z'_i \text{ for all } i\} \implies z = z'$. The rule is Strategyproof if $z_i \succ_i^{sd} z'_i$ where z_i obtains from the truthful report and z'_i , ceteris paribus, from any other report.

When agents have different status (due to their performance at some competitive exam, seniority, or any other exogenous parameter), they can be ordered accordingly, and served by the corresponding *Priority rule*: the first in line picks her top ranked object, then the next in line takes the best remaining object, and so on. This deterministic rule is clearly Strategyproof, even GroupStrategyproof (because we rule out indifferences between objects). It is also Efficient but not Fair.

To restore fairness while preserving some degree of efficiency, the *Random Priority* (RP)¹⁷ rule draws randomly an ordering σ of N with uniform probability $\frac{1}{n!}$ on each ordering, and averages the corresponding deterministic assignments ς^σ .

The RP rule inherits the Strategyproofness property of the Priority rules. However it is challenged on Efficiency and Fairness grounds by the *Probabilistic Serial* rule, defined by the familiar *eating algorithm* that we illustrate first in

¹⁵If preferences are dichotomous as in subsection 4.2, SD is a complete ordering of $\Delta(A)$.

¹⁶The usual name is Ordinal Efficiency, which is stronger than Ex Post Efficiency: z is a convex combination of efficient deterministic assignments.

The latter property has no meaning in the time-sharing context.

¹⁷Aka *Random Serial Dictatorship*, an unfortunate terminology because no one is a dictator (let alone a serial killer) in the RP rule.

the following four-agent example:

$$\begin{aligned}
 \succ_1 &: a \succ b \succ c \succ d \\
 \succ_2 &: a \succ b \succ d \succ c \\
 \succ_3 &: b \succ a \succ c \succ d \\
 \succ_4 &: c \succ d \succ a \succ b
 \end{aligned}$$

Agents eat their favorite object at constant speed of 1. At time $\frac{1}{2}$ agents 1 and 2 have eaten $\frac{1}{2}$ of a each, while agent 3 has $\frac{1}{2}$ of b and 4 has $\frac{1}{2}$ of c ; object a being exhausted, agents 1 and 2 start eating object b (of which $\frac{1}{2}$ is left), together with 3, while 4 keeps eating c . At time $\frac{1}{2} + \frac{1}{6}$ object b is exhausted and both 1 and 2 get $\frac{1}{6}$ each, the rest going to 3; we are left with $\frac{1}{3}$ of c not yet eaten, and a full object d : agents 1,3 and 4 now eat the rest of c while agent 2 starts eating d . At time $\frac{1}{2} + \frac{1}{6} + \frac{1}{9}$ object c is finished, agents 1 and 3 get $\frac{1}{9}$ each, and 4 gets the rest. In the remaining $\frac{2}{9}$ units of time, they all eat d and agent 2 gets the biggest share $\frac{1}{9} + \frac{2}{9}$. The resulting assignment is

	a	b	c	d
z_1	1/2	1/6	1/9	2/9
PS: z_2	1/2	1/6	0	1/3
z_3	0	2/3	1/9	2/9
z_4	0	0	7/9	2/9

The general definition of the PS assignment matrix (Bogomolnaia & Moulin (2001)) is similar: agents eat for one unit of time, at the same speed, from their best object among those not yet exhausted.¹⁸

In order to implement this assignment (whether by lotteries or time-sharing) we need to find a BvN decomposition, which is a tedious process but only of polynomial complexity.

By contrast, to find the RP matrix in our example, we need to average the 24 deterministic assignments corresponding to each ordering of the agents: e. g., ordering 2314 yields $1 \leftarrow c, 2 \leftarrow a, 3 \leftarrow b, 4 \leftarrow d$. This is, again, tedious, but unfortunately of exponential complexity in the size of our problem. In the example we find

	a	b	c	d
z_1	1/2	1/6	1/12	1/4
RP: z_2	1/2	1/6	0	1/3
z_3	0	2/3	1/12	1/4
z_4	0	0	5/6	1/6

Note that it is enough to roll a dice to implement RP in an expected sense. This works for the random interpretation of this rule, but for the time-sharing interpretation we need to compute the matrix above.

¹⁸When individual preferences allow for indifferences, the definition of PS is a little more involved: see Katta Sethuraman (2006) and Bogomolnaia (2015). The Priority rules and RP are easily adjusted for indifferences with the help of the leximin ordering, as illustrated in the example following the next Theorem.

In the example, both agents 1 and 3 prefer PS over RP (in the Stochastic Dominance sense), agent 4 prefer RP over PS, and agent 2 is indifferent. There are cases (see below) where the PS assignment is Pareto superior to the RP one, but the reverse cannot happen: this follows from statement *ii*) in our next result, comparing our two mechanisms and uncovering a severe impossibility result.

Theorem (Bogomolnaia & Moulin (2001)) *Assume $n \geq 4$.*

- i) RP is Strategyproof and meets Fair Share; but it is not Efficient and fails Envy Free;*
- ii) PS is Efficient and Envy Free (hence meets Fair Share); but it is not Strategyproof;*
- iii) No Efficient and Anonymous rule can be Strategyproof as well.*

Strategyproofness is a strong point in favor of RP against PS;. The only incentives advantage of PS is the fact that the PS assignment is *core stable* in the cooperative game where agents are endowed with a fair share of every object. That RP is not core stable is implied by the fact that it is inefficient.

When we only have three agents and three objects, the RP matrix is in fact Efficient.¹⁹ When $n \geq 4$, the RP rule is not GroupStrategyproof, a consequence of the inefficiency of some RP assignments.

The simplest example of an inefficient RP assignment is a 4x4 scheduling problem where all four agents weakly prefer an early slot to a later one; agent 1 is indifferent between the last three slots, and agent 2 between the last two; agents 3 and 4 strictly prefer an earlier slot:

$$\begin{aligned}
 \succeq_1: a &> \{b, c, d\} \\
 \succeq_2: a &> b > \{c, d\} \\
 \succeq_3: a &> b > c > d \\
 \succeq_4: a &> b > c > d
 \end{aligned} \tag{9}$$

To compute RP we make sure it only chooses efficient deterministic assignments: for instance if the priority is 3124, agent 3 grabs a first, then by assigning d to agent 1 we can give b to 2 and c to 4, without prejudice to 1. This and similar computations deliver the RP matrix:

	a	b	c	d
z_1	1/4	0	1/8	11/24
RP: z_2	1/4	1/3	1/8	11/24
z_3	1/4	1/3	3/8	1/24
z_4	1/4	1/3	3/8	1/24

where the shares of 1 and 2 in c and d can be changed as long as z_{1cd} and z_{2cd} remain constant.

¹⁹And the RP rule is characterized by Efficiency, Anonymity, and SP: Bogomolnaia & Moulin (2001).

In the PS eating algorithm a is eaten first by all agents, then b by agents 2,3,4, then agents 3,4 need only $1 - \frac{1}{4} - \frac{1}{3} = \frac{5}{12}$ units of c to fill their share:

$$\begin{array}{rcccc}
 & a & b & c & d \\
 \text{PS: } z_1 & 1/4 & 0 & 1/12 & 1/2 \\
 z_2 & 1/4 & 1/3 & 1/12 & 1/2 \\
 z_3 & 1/4 & 1/3 & 5/12 & 0 \\
 z_4 & 1/4 & 1/3 & 5/12 & 0
 \end{array} \tag{10}$$

We see that agents 1,2 are indifferent between the two assignments RP and PS, but 3 and 4 strictly prefer (in the SD sense) their PS to their RP allocation. The inefficiency in RP comes from the two orderings 1234 and 1243, giving to 4 and 3 a fraction of slot d , which could be passed on to 1 and 2 in exchange for a bigger share of c to 4 or 3.

Next we check that RP is not Envy Free in the following 3x3 example

$$\begin{array}{l}
 \succ_1 : a \succ b \succ c \\
 \succ_2 : a \succ c \succ b \\
 \succ_3 : b \succ a \succ c
 \end{array}$$

$$\begin{array}{rcccc}
 & a & b & c & \\
 \text{RP: } z_1 & 1/2 & 1/6 & 1/3 & \\
 z_2 & 1/2 & 0 & 1/2 & \\
 z_3 & 0 & 5/6 & 1/6 & \\
 \text{PS: } z_1 & 1/2 & 1/4 & 1/4 & \\
 z_2 & 1/2 & 0 & 1/2 & \\
 z_3 & 0 & 3/4 & 1/4 &
 \end{array}$$

Under RP agents 2 and 3 are not envious: $z_2 \succ_2^{sd} z_1, z_3$, and $z_3 \succ_3^{sd} z_1, z_2$. Agent 1 does not envy 2, $z_1 \succ_1^{sd} z_2$, however $z_1 \not\succeq_1^{sd} z_3$ because the total probability of $\{a, b\}$ is larger in z_3 than in z_1 . So for some refined preferences over lotteries, e. g., a vNM utility $u_1(a) = 1, u_1(b) = 0.9, u_1(c) = 0$, she envies agent 3.

No one is envious in the PS assignment above, however to agent 3 the misreport $\succeq_3^* : a \succ b \succ c$ is potentially profitable (again, in some refinement of his stochastic dominance relation): it results in the new allocation $z_3^* = \frac{1}{3}a + \frac{1}{2}b + \frac{1}{6}c$, where the share of the top two objects $\{b, a\}$ is higher. Thus PS is not Strategyproof.²⁰

Characterization results Our two rules have received a lot of theoretical attention, and are not seriously challenged by other mechanisms such as the *Boston mechanism* poised to maximize the number of agents getting, ex post, their first choice of school. Several results confirm their prominence in the random assignment problem.

The earliest characterization result is a striking alternative definition of the RP rule. Start with an arbitrary deterministic assignment of the objects to the agents: in this exchange economy, there is a single competitive and core

²⁰In a scheduling model generalizing example (10) the PS rule is actually Strategyproof, and easy to characterize (Bogomolnaia & Moulin (2002)).

stable allocation obtained by the celebrated *Top Trading Cycle* algorithm.²¹ If the initial endowment is random with uniform probability, and we average the corresponding TTC assignments, the result is precisely the Random Priority assignment (Abdulkadiroglu & Sonmez (1998)).

The only other characterization of RP bears on the dynamic version of this rule: an agent is drawn at random, picks his object, then a second agent is drawn, and so on. In this dynamic game the sincere report of preferences is “obviously” strategyproof (in the sense given this term by Li (2017)), which together with Ex Post Efficiency and Anonymity singles out the RP rule (Pycia & Troyan (2018)).

The PS assignment has an alternative Egalitarian definition, valid even when indifferences are allowed (Bogomolnaia (2015)). We describe the latter when preferences are strict. For each assignment z , agent i , and index k , $1 \leq k \leq n-1$, compute the total weight τ_i^k in z_i of the k best alternatives for i ; ordering these numbers increasingly gives a vector $\tau(z)$ of dimension $n(n-1)$. The vector $\tau(z^{PS})$ maximizes the leximin ordering over all feasible vectors $\tau(z)$.

Finally the PS rule is characterized by Efficiency, Envy Free, and a powerful invariance property: when I scramble my preferences over my last k objects, this does not affect my share of my top $n-k$ objects (Bogomolnaia & Heo (2015); Hashimoto et al (2014)).

Finally we note that in a large society where each type of preferences is represented by a positive fraction of the agents, the two assignments RP and PS converge asymptotically (Che & Kojima (2010)), and this common limit unsurprisingly combines all the efficiency, fairness and incentives properties discussed above.

6 Additive utilities

In Subsection 4.3 the goods in the manna are perfect complements for the agents, and we found that the Egalitarian rule(s) is more appealing than the Competitive one. Here we make the polar opposite assumption that the goods are perfect substitutes (equivalently, preferences are represented by additive utilities), and we find that the Competitive rule normatively dominates the Egalitarian ones.

Such preferences are realistic when we divide the manna into truly “unrelated” goods such as family heirlooms including a computer, a bicycle and a family portrait; of course the pair of matching chandeliers must be counted as one item. Such a partition of the manna in unrelated assets is plausible in a divorce, or the dissolution of a professional partnership.

Goods (or bads) are divisible in this Section, either physically (a pile of cash, a large number of bottles), or by means of lotteries or time-sharing. The later two types of division are not always easy to implement, therefore it will be desirable to use them parsimoniously: a division involving fewer split goods is simpler to describe and to understand.

²¹Exchange objects along all the cycles of the graph where each agent points to the owner of his favorite object, possibly herself; repeat among the agents who were not part of a cycle.

To elicit additive utilities over the items in the manna requires more sophistication from the participants than in the assignment problem of Section 5. The practicality of such reports is illustrated by Pratt & Zeckhauser (1990) for the division of family silver, and more recently by the success of two user-friendly platforms where users report additive utilities by dividing 100 points over the items in the manna. SPLIDDIT was designed by Goldman & Procaccia (2014) and ADJUSTED WINNER by Brams and Taylor.²² Thousands of visits and usage of these sites confirm that the practicality of this approach.

We start by the "standard" model where the manna contains only goods. It generalizes the model of subsection 4.2 by allowing arbitrary non negative marginal utilities. A problem specifies as usual the set N of agents, the set A of objects, the manna $\omega \in \mathbb{R}_+^A$ and the marginal utility matrix $u = [u_{ia}] \in \mathbb{R}_+^{N \times A}$. A feasible allocation $z = [z_{ia}]$ is as usual a matrix in $\mathbb{R}_+^{N \times A}$, such that $z_{Na} = \omega_a$ for all a ; their set is written Φ . Allocation z results in the utilities $U_i = u_i \cdot z_i$. Multiplying the row u_i by a positive scalar does not change i 's preferences, so our axioms and division rules will be accordingly invariant.

The Competitive rule has a striking connection to the Nash bargaining theory (Nash (1950)).

Theorem (Eisenberg & Gale (1959)): *The two following statements are equivalent:*

- i) The feasible allocation z is competitive;*
- ii) The utility profile $U = (u_i \cdot z_i)_{i \in N}$ maximizes the Nash product $\prod_{i \in N} U_i$ over all feasible profiles.*

The Competitive allocations all have the same welfare, and the same price.

The Egalitarian allocations $EE(\theta)$ and $EE(\omega)$ are by design welfarist, they equalize utilities calibrated along a benchmark bundle of goods. The welfarist interpretation of the Competitive allocation is much less obvious, and in fact applies to a much broader domain than linear preferences.

The Eisenberg Gale theorem is about the agregation of competitive demands in the so called *Fischer economies* (Cole et al. (2016)), where agents are endowed with shares of fiat money that can only be used to buy the goods in the manna (no one initially owns any of the goods). If $\lambda = (\lambda_i)_{i \in N}$ is the distribution of shares summing up to one and p the price, then in the corresponding Competitive allocation(s) agent i 's budget is $\lambda_i(p \cdot \omega)$. The Theorem says that such allocation(s) maximizes $\prod_i U_i^{\lambda_i}$ over all feasible utility profiles.²³ Proven first for linear preferences, this result was generalized (Chipman (1974)) to all economies with homothetic preferences, represented by a 1-homogenous non negative utility (canonical up to a constant factor). This includes all the standard families of utility functions: additive, Constant Elasticity of Substitution, CobbDouglas, Leontief, and arbitrary convex combinations of these. On such domain the Competitive rule is blissfully singlevalued, unlike in the general

²²Respectively www.spliddit.org/ and www.nyu.edu/projects/adjustedwinner/.

²³Note that the Competitive allocation in the economy with endowments $\omega = (\omega_i)_{i \in N}$ does not in general maximize the Nash product of utilities, but in our Fair Division model equal split of the manna or of fiat money is the same thing.

Arrow Debreu domain.

Moreover, in the additive domain the Competitive rule, in addition to Envy Free (and FS), meets our two monotonicity properties.²⁴

Theorem: *In the additive domain, the Competitive utility profile is continuous in the utility matrix u and the manna ω . It is also Resource Monotonic and Population Monotonic.*

A proof of RM and PM is in (Bogomolnaia et al. (2017)) and in Segal Halevy & Sziklai (2018) for the more general Cake cutting model.

By contrast the Egalitarian rules $EE(\theta)$ and $EE(\omega)$ behave in the additive domain exactly like in the general AD domain. They are unique welfare-wise, continuous in the parameters and meet PM; but $EE(\theta)$ meets RM and violates FS, while $EE(\omega)$ meets FS but fails RM.

In the familiar domain of Cobb Douglas (CD) utilities it is easy to check that the Competitive rule also meets RM and PM. Moreover, unlike in the additive domain, the unique Competitive allocation is given by closed form expressions.²⁵ But the Cobb Douglas utilities force each agent to consume a positive amount of every good, lest they end up with no benefit at all: in other words, each good must be split n ways. Like the Leontief utilities of subsection 4.3 the CD utilities capture a strong complementarity between all objects that make little sense in many practical FD problems involving heterogenous items.

A key advantage of the additive domain is that *every efficient utility profile is achieved by an allocation where at most $n - 1$ goods are split*. In fact the Competitive division often involves fewer split goods, as in the following three examples, where we compare the C and $EE(\omega)$ allocations.

First we have two agents, three goods, and $\omega = (1, 1, 1)$:

$$\begin{array}{cccc} & a & b & c \\ u_1 & 6 & 3 & 1 \\ u_2 & 1 & 2 & 1 \end{array} \implies z^C = \begin{array}{cccc} & a & b & c \\ z_1 & 1 & 0 & 0 \\ z_2 & 0 & 1 & 1 \end{array} ; z^{EE(\omega)} = \begin{array}{cccc} & a & b & c \\ z_1 & 1 & 0.19 & 0 \\ z_2 & 0 & 0.81 & 1 \end{array} \quad (11)$$

so the Competitive allocation does not split any good. With only two agents this is a “frequent” occurrence, while the $EE(\omega)$ and $EE(\theta)$ allocations typically require to divide one good (but no more, as just mentioned).

The next example underlines the sharp normative differences between our two rules. We have four agents and three goods, and $\omega = (1, 1, 1)$. The first three agents are *single-minded*, agents i likes only good a_i , while agent 4 on the contrary is *flexible*, he likes all goods equally:

$$\begin{array}{cccc} & a_1 & a_2 & a_3 \\ u_1 & 1 & 0 & 0 \\ u_2 & 0 & 1 & 0 \\ u_3 & 0 & 0 & 1 \\ u_4 & 1 & 1 & 1 \end{array} \quad (12)$$

²⁴Recall from subsection 3.3 that RM and FS are incompatible in any domain containing Leontief utilities.

²⁵Set $U_i(z_i) = \sum_{a \in A} \delta_{ia} \ln(z_{ia})$ where $\delta_i \in \Delta(A)$, then the Competitive allocation is $z_{ia} = \frac{\delta_{ia}}{\sum_j \delta_{ja}} \omega_a$, from which the statement is immediate, as well as continuity in the parameters.

The competitive price is $\frac{4}{3}$ for each good; each agent 1, 2, or 3 buys $\frac{3}{4}$ units of “his” good while agent 4 gets $\frac{1}{4}$ of each good. Contrast this with the Egalitarian division that splits each good a_i equally between agent i and agent 4, so that everyone ends up with a share worth one half of the entire manna:

$$z^C = \begin{array}{c} \\ z_1^c \\ z_2^c \\ z_3^c \\ z_4^c \end{array} \begin{array}{ccc} 1 & 2 & 3 \\ 3/4 & 0 & 0 \\ 0 & 3/4 & 0 \\ 0 & 0 & 3/4 \\ 1/4 & 1/4 & 1/4 \end{array} \quad z^{EE(\omega)} = \begin{array}{c} \\ z_1^e \\ z_2^e \\ z_3^e \\ z_4^e \end{array} \begin{array}{ccc} 1 & 2 & 3 \\ 1/2 & 0 & 0 \\ 0 & 1/2 & 0 \\ 0 & 0 & 1/2 \\ 1/2 & 1/2 & 1/2 \end{array}$$

The Egalitarian rule focuses on the (relative) benefits of consuming each good a_i , and in this example shares them equally between the two relevant agents, i and 4. The Competitive rule is much more generous to the single-minded agents: agent 1 for instance “owns” a quarter of each good, so she is entitled to $\frac{1}{4}$ of the surplus generated by goods a_2 and a_3 .

Notice an unpalatable feature of the Competitive allocation z^C : the flexible agent 4 gets exactly his Fair Share utility $u_4 \cdot \frac{1}{4}\omega = \frac{3}{4}$, and is on the verge of envying the other three ($u_4 \cdot z_4 = u_4 \cdot z_i$ for $i = 1, 2, 3$), while the others strictly improve upon their fair share $\frac{1}{4}\omega$, and are strictly non envious. The C rule does not reward uses the flexibility of agent 4’s preferences to his disadvantage. This situation always happens in the C rule for any agent who eats a positive share of each good. By contrast *everybody always* prefers her $EE(\omega)$ allocation to her Fair Share, unless *nobody* can get more than her FS utility.

On the other hand we can argue that the Egalitarian allocation gives *too much* to agent 4, because he gets (much) more than his fair share of *every* good. By contrast in *any* Competitive allocation, everyone gets *at most* a $\frac{1}{n}$ -th share of *at least* one good:²⁶

$$\min_{a \in A} z_{ia} \leq \frac{1}{n} \text{ for all } i \tag{13}$$

Our last example shows why the rule $EE(\omega)$ violates RM, and also illustrates its normative difference with the C rule. The example requires three or more agents.²⁷ We compare two problems with $B = \{a_1, a_2, a_3\}$ and $A = \{a_1, a_2, a_3, d\}$ respectively, and one unit of each good:

$$u^B = \begin{array}{c} \\ u_1 \\ u_2 \\ u_3 \end{array} \begin{array}{ccc} a_1 & a_2 & a_3 \\ 3 & 1 & 1 \\ 1 & 3 & 1 \\ 1 & 1 & 3 \end{array} \quad ; \quad u^A = \begin{array}{c} \\ u_1 \\ u_2 \\ u_3 \end{array} \begin{array}{cccc} a_1 & a_2 & a_3 & d \\ 3 & 1 & 1 & 0 \\ 1 & 3 & 1 & 4 \\ 1 & 1 & 3 & 4 \end{array}$$

The B -problem is symmetric. Any efficient and symmetric rule, in particular C and $EE(\omega)$, allocates goods diagonally: agent i gets all of a_i and so on; normalized utilities are $\frac{3}{5}$.

²⁶If $z_{ia} > \frac{1}{n}$ for all a the competitive price must be parallel to u_i but the equal budget equation $p \cdot z_i = p \cdot (\frac{1}{n}e^A)$ gives $u_i \cdot z_i = u_i \cdot (\frac{1}{n}e^A)$, contradiction.

²⁷The $EE(\omega)$ rule meets RM for $n = 2$.

In the A -problem a natural idea is to keep the same allocation of a_1, a_2, a_3 and divide d equally between agents 2 and 3, because agent 1 does not care for d . This is exactly what the Competitive rule recommends (prices are $(1, \frac{3}{5}, \frac{3}{5}, \frac{4}{5})$). But the normalized utilities at this allocation are $(\frac{3}{5}, \frac{5}{9}, \frac{5}{9})$, so the Egalitarian rule must compensate agents 2, 3 for the *loss* in normalized utilities caused by the *gain* of some new good! Equality is restored at the allocation

$$z^{EE(\omega)} = \begin{array}{cccc} & a & b & c & d \\ & 55/59 & 0 & 0 & 0 \\ & 2/59 & 1 & 0 & 1/2 \\ & 2/59 & 0 & 1 & 1/2 \end{array}$$

where agent 1's welfare has decreased.

Finally we note that computing the C allocation, or equivalently maximizing the Nash product of utilities, is solvable in polynomial time. A weakly polynomial algorithm was proposed by Devanur et al. (2008), and recently improved to a strongly polynomial one by Orlin (2018). Thus, from a complexity point of view, it is not harder than computing the Egalitarian allocation (Kurokawa et al. (2015)) though of course the algorithms are less simple.

Incentives We know that in the additive domain, no Efficient and Anonymous rule can be Strategyproof. Zhou (1991) shows that with only two agents, Effy and SP together imply that one of them is a dictator; see also Schummer (1997). This impossibility disappears in the dichotomous subdomain (subsection 4.2).

Interestingly Cole & Gkatzkelis (2015) propose a simple but *wasteful* approximation of the Competitive rule achieving Strategyproofness at the cost of Efficiency. The rule computes the C allocation $z^* = (z_j^*)$ but gives to agent i only a *fraction of z^** : this fraction is the ratio of the Nash product of utilities of other agents in z^* to their optimal Nash product when i is absent (and consumes no manna). So agent i 's allocation is

$$z_i = \frac{\prod_{j \neq i} u_j \cdot z_j^*}{\max_{z_{N \setminus i} = \omega} \prod_{j \neq i} u_j \cdot z_j} z_i^* \quad (14)$$

The discount rate on the C share is the relative externality of agent i 's presence on the collective welfare of the other agents, measured as the Nash product of utilities.

If agent i reports another utility \tilde{u}_i and $\tilde{z} = (\tilde{z}_j)$ maximizes the corresponding Nash product, his actual utility \tilde{U}_i (after throwing away part of \tilde{z}_i) is

$$\tilde{U}_i = \frac{\prod_{j \neq i} u_j \cdot \tilde{z}_j}{\max_{z_{N \setminus i} = \omega} \prod_{j \neq i} u_j \cdot z_j} u_i \cdot \tilde{z}_i \leq \frac{\prod_{j \neq i} u_j \cdot z_j^*}{\max_{z_{N \setminus i} = \omega} \prod_{j \neq i} u_j \cdot z_j} u_i \cdot z_i^* = U_i$$

so the misreport is not profitable. This proves SP. Just like the Eisenberg Gale Theorem, this argument holds for any profile of 1-homogenous utilities.

This division rule imitates, in multiplicative form, what the celebrated Vickrey-Clarke-Groves (VCG) mechanisms do in additive form. VCG mechanisms are

also Strategyproof and wasteful, in the sense that the monetary transfers ensuring SP are not balanced, they either burn some money, or require an exogenous subsidy. But if VCG mechanisms can often be fine tuned to achieve near budget balance (as in Guo Conitzer (2007) and Moulin (2009)), the Cole & Gkatzkelis rule must drop a very substantial chunk of manna. For instance in Example (11), equation (14) says that agent 1 keeps only 75% of her Competitive share, and agent 2 only 60%. In Example (12) every agent keeps only 42% of his share. It turns out that the smallest possible share an agent can keep is $\frac{1}{e} \simeq 37\%$ (Cole & Gkatzkelis (2015)).

Bad manna: limits of the Competitive approach The fair division of *bads* such as chores among family members, jobs in the work place, tasks between co-authors, and so on, is a different context than that of goods, but not a conceptually different problem. Agents are trying to spread the manna, now a burden, fairly and efficiently, and the normative properties discussed before are easily adapted.

A problem is as usual a list $N, A, \omega \in \mathbb{R}_+^A, u = [u_{ia}] \in \mathbb{R}_+^{N \times A}$, where u_{ia} is now i 's marginal **dis**utility for bad a . Efficiency means that we cannot lower everyone's disutility weakly and someone's strictly.

The rules $EE(\omega)$ and $EE(\theta)$ select an efficient allocation z satisfying property (2) (respectively for ω and for θ), which is always possible if the u_{ia} -s are all positive (see Footnote 8 for the case when some u_{ia} are zero). These rules are single-valued (welfarewise), continuous in the parameters u and ω , and, once the definitions of FS, RM and PM are reversed in the obvious way, behave as if the manna is made of goods.

A non zero price $p \in \mathbb{R}_+^A$ defines a Competitive allocation z if z is feasible and for each agent i , z_i minimizes $u_i \cdot y_i$ in the budget set $B(p) = \{y \in \mathbb{R}_+^A | p \cdot y \geq 1\}$: each agent must buy a bundle of bads costing *at least* 1. Efficiency and Envy Free follow. However Bogomolnaia et al. (2018) show that the Competitive rule is no longer single-valued welfarewise, and no single-valued selection of the Competitive correspondence is continuous in u or ω . Except in the simple case of two-agent or two-good problems, it is not known how many different C allocations may typically coexist.

The simplest example of this surprisingly bad news involves dichotomous utilities (as in Subsection 4.2), i. e., each entry u_{ia} of u is 0 or 1. Write $B(u)$ for the set of bads a s. t. some agent i does not mind the chore a : $u_{ia} = 0$. The allocation z is efficient if and only if every bad in $B(u)$ is eaten by such agent. Bads in $A \setminus B(u)$ are disliked by everybody and it does not matter to efficiency how they are divided.

One C allocation gives a zero price to every $a \in B(u)$ and $p_a = \frac{n}{|A \setminus B(u)|}$ for every other bad. Thus agents who take care painlessly of the bads in $B(u)$ are not rewarded for this. For another C allocation, take any $a^* \in B(u)$ and suppose the set M of agents s. t. $u_{ia^*} = 0$ is not N ; set a price of $|M|$ for a^* , so that each i in M eats $\frac{1}{|M|}$ of a^* and nothing else; other bads in $B(u)$ have price zero and the rest cost $p_a = \frac{n - |M|}{|A \setminus B(u)|}$: this gives a free lunch to agents in

M , and only those; there are also C allocations where a subset of bads in $B(u)$ have a positive price; in any of those, some agents get a free lunch, and the rest share equally $B(u)$.

In Example (11), where the matrix u represents now disutilities, we find three C allocations

$$\begin{array}{rcc}
 & a & b & c \\
 z^{C_1} = & z_1 & 1/6 & 1 & 1 \\
 & z_2 & 5/6 & 0 & 0 \\
 z^{C_2} = & z_1 & 0 & 1 & 1 \\
 & z_2 & 1 & 0 & 0 \\
 z^{C_3} = & z_1 & 0 & 7/12 & 1 \\
 & z_2 & 1 & 5/12 & 0
 \end{array}$$

with as usual budget of 1 and respective prices

$$p^{C_1} = \left(\frac{6}{5}, \frac{3}{5}, \frac{1}{5}\right); p^{C_2} = \left(1, \frac{3}{4}, \frac{1}{4}\right); p^{C_3} = \left(\frac{6}{11}, \frac{12}{11}, \frac{4}{11}\right)$$

In the following "dual" of Example (12), we have no less than seven C allocations:²⁸

$$\begin{array}{rccc}
 & a_1 & a_2 & a_3 \\
 u_1 & 1 & 3 & 3 \\
 u_2 & 3 & 1 & 3 \\
 u_3 & 3 & 3 & 1 \\
 u_4 & 1 & 1 & 1
 \end{array}$$

There is still a close connection between the Competitive allocations and the corresponding Nash product of utilities, which rests on the notion of critical point. A point w in a convex set Γ is a critical point of the (smooth) function f if the tangent hyperplane to the level curve of f supports Γ at w . Bogomolnaia et al. (2017) show that *the feasible allocation z is competitive if and only if the disutility profile $U = (u_i \cdot z_i)_{i \in N}$ is a critical point of the Nash product $\prod_{i \in N} U_i$ over all **efficient** feasible profiles.*

One technical explanation of the sharp difference in the behavior of the C rule when the manna is either good or bad compares the budget sets $B(p; -\omega)$ in both cases

$$B^+(p; \omega) = \{y_i \in \mathbb{R}_+^A | p \cdot y_i \leq \frac{1}{n} p \cdot \omega\}; B^-(p; \omega) = \{y_i \in \mathbb{R}_+^A | p \cdot y_i \geq \frac{1}{n} p \cdot \omega\}$$

Note that $B^-(p; \omega)$ is bounded while $B^+(p; \omega)$ is not, so a C division of bads is effectively a *Constrained Competitive* allocation in the sense discussed in Section 8: Eisenberg Gale theorem does not apply and multiple CC allocations are routinely expected.

The fact that the Competitive approach forces us to use a discontinuous division rule is a serious normative concern, that does not go away when we replicate the problem. Even if we deal with a fixed disutility matrix, the selection of a particular C allocation raises additional normative concerns (as in the dichotomous example above).

²⁸The competitive prices are $p = (\frac{4}{3}, \frac{4}{3}, \frac{4}{3}), (\frac{3}{2}, \frac{3}{2}, 1), (2, 1, 1)$, and their permutations.

Mixed manna: failure of the Egalitarian approach The manna may consist of some goods and some bads, it may contain assets as well as liabilities. We speak then of a *mixed manna* and a problem is a list $N, A, \omega \in \mathbb{R}_+^N, u = [u_{ia}] \in \mathbb{R}^{N \times A}$, where u_{ia} is i 's marginal utility for item a if $u_{ia} > 0$, or disutility if $u_{ia} < 0$.

Consider for instance the two-agent three-item problem

	a	b	c
u_1	3	1	-3
u_2	1	2	-5

with one unit of each good a, b and of the bad c . Here the canonical rule $EE(\omega)$ makes no sense because $u_1 \cdot \omega = 1, u_2 \cdot \omega = -2$, and the unique efficient allocation meeting $\frac{u_1 \cdot z}{u_1 \cdot \omega} = \frac{u_2 \cdot z}{u_2 \cdot \omega}$ yields the utilities $U = (-1, 1.6)$ violating Fair Share (and with opposite signs than $U_i \cdot \omega$).

In order to restore an Egalitarian approach we can treat our Fair Division problem like a bargaining model with the profile of Fair Share utilities as the disagreement outcome, and then apply the Kalai-Smorodinsky solution (Kalai & Smorodinsky (1975)). In doing so we lose the connection of the final allocation with the physical characteristics of the manna, and even Population Monotonicity is lost.

By contrast, the Competitive approach is well defined in the usual fashion: the non zero price vector $p \in \mathbb{R}^A$ may have positive, negative, or null components; ditto for the common budget β , and agent i gets her best share z_i such that $p \cdot z_i \leq \beta$.

A three elements partition of all problems $P = (N, A, \omega, u)$ determines the behavior of the Competitive rule. We call P *positive* if there is at least one strictly positive feasible utility profile U ; *null* if $U = 0$ is an efficient profile; and *negative* if there is no non negative feasible profile (none in \mathbb{R}_+^N). For instance the problem above is positive as $z_1 = (1, 0, \frac{2}{3}), z_2 = (0, 1, \frac{1}{3})$ gives $U_1, U_2 > 0$. It becomes null if the manna is $\omega' = (1, 1, \frac{1}{5})$, and is negative if we have even more of the bad c to share.

Theorem (Bogomolnaia et al. (2017))

In a positive (resp. null) problem, the Competitive profile $U = (u_i \cdot z_i)$ is unique and maximizes the Nash product $\prod_{i \in N} U_i$ over all positive feasible profiles (resp. is $U = 0$).

In a negative problem the Competitive profiles $U = (u_i \cdot z_i)$ are the critical points of the Nash product $\prod_{i \in N} U_i$ over all efficient (strictly) negative feasible profiles.

Like the original Eisenberg Gale Theorem for a good manna, this result generalizes to all homothetic preferences.

The multiplicity and continuity issues notwithstanding, we conclude that the Competitive approach, unlike the Egalitarian one, is readily adapted to the mixed manna case, and performs very well for positive problems.

7 Indivisible goods, and bads

We maintain in this Section the assumption that agents can form additive utilities over the objects in the manna. When these objects are not physically divisible (painting, bicycle, student, client, license, etc..) randomization achieves a fair compromise only in the *ex ante* sense, but *ex post* fairness is not guaranteed. In court, and in many other contexts, rolling a dice is not an option: objections and counter-objections bear on the actual allocation of the manna. Thus we need a normative framework to propose *deterministic* fair divisions of *indivisible* objects.

The CS community takes to heart the general project of extending the formal analysis of FD to the technically difficult case of indivisible objects. Most but not all of this fast growing literature posits additive utilities, as in the previous Section, and for brevity we will maintain this assumption here.

The standard fairness properties discussed so far may well be unfeasible. Think of an efficient (or simply non wasteful) division of several rocks and one diamond: Fair Share is clearly not feasible, and neither is Envy Free. Bouveret and Lang (2008) show that it is a hard problem to determine whether or not, in a given problem, an efficient division exists that meets one of these tests. The conceptual challenge is to define convincing *approximations* of these tests, and ensure that they are always feasible and compatible with Efficiency, i. e., prove an existence result; then we must design tractable (computationally simple) algorithms to implement an actual division meeting those tests. This research is very young, and replete with exciting open questions.

Formally a problem is a list $(N, A, u \in \mathbb{R}_+^{N \times A})$ where A is the set of objects: here $\omega = A$ and a feasible allocations $z = (z_i)_{i \in N}$ is simply a partition of A . *Notation:* for $y \subseteq A$ we write $u_{iy} = \sum_{a \in y} u_{ia}$ as in subsection 4.2, and $y+a, y-a$ instead of $y \cup \{a\}, y \setminus \{a\}$, etc..

Most papers discuss the case of goods, as we do first; the little we know in the case of bads is reviewed next.

Start with Fair Share.²⁹ A natural relaxation is *Fair Share up to the Least Valued Good*:

$$\text{FSX: } \forall i \in N, \forall a \notin z_i : u_{i(z_i+a)} \geq \frac{1}{n} u_{iA}$$

But this property is not always feasible, even when agents have identical utilities. Here is an example with three agents, three “large” goods a, a, b and ten “small” goods c . Each a brings utility $\frac{3}{2}$, b brings utility 4, and each c brings $\frac{3}{10}$: so $u_{iA} = 10$ and the FS level is $3\frac{1}{3}$. Check that if agent 1 gets only c -s, and agent 2 gets b , then FSX fails if agent 3 gets no c , because $3 + \frac{3}{10} < \frac{1}{3}u_{iA}$; but then FSX fails for 1 who gets at most nine c -s. If 1 gets b while 2 and 3 get one a each, then one of 2,3 gets at most five c -s and a total utility of 3 so FSX fails again.

A weaker approximation of FS, *Fair Share up to one Good*, is always feasible,

²⁹The CS literature mostly speaks of *Proportionality* instead of Fair Share; so does the cake-cutting literature.

and compatible with Efficiency:

$$\text{FS1: } \forall i \in N, \exists a \notin z_i : u_{i(z_i+a)} \geq \frac{1}{n} u_{iA}$$

The claim follows from the Theorem below.

Budish (2011) proposes an appealing alternative interpretation of Fair Share, inspired by the Divide & Choose procedure. Let \mathcal{P} be the set of n -partitions $P = \cup_{k=1}^n B^k$ of A . We compute agent i 's *MaxMinShare* (MMS) by selecting a partition where the utility of his least desirable share is as high as possible. Formally:

$$\text{MMS: } \forall i \in N : u_{iz_i} \geq \max_{P \in \mathcal{P}} \min_k u_{iB^k}$$

If goods are divisible, this is exactly FS. With indivisible goods and only two agents, an efficient partition where each agent gets at least his MMS utility always exists.³⁰

The existence of a partition of the goods meeting MMS for any n was intensely discussed for several years, until Procaccia & Wang (2014) came up with a sophisticated counterexample. They showed, however, that it is always feasible to guarantee $\frac{2}{3}$ of his MMS utility to each agent. Then Kurokawa et al. (2016) argue that, in practice, utilities reported by non adversarial agents will be compatible with MMS.

We turn to the search for tractable division rules adapting the Egalitarian and Competitive approaches to the indivisible world. Indivisibilities imply that a single-valued rule cannot be Anonymous, so it must be supplemented by a tie breaking convention, that we do not need to specify.

The compelling Egalitarian answer simply maximizes the leximin ordering over the feasible profiles of *normalized* utilities $\frac{1}{u_{iA}} u_i$. The corresponding LEX allocations are efficient and equalize (normalized) utilities across agents “up to one object”:

$$\text{EGAL1: } \forall i, j \in N, \exists a \in z_j : u_{iz_i} \geq u_{j(z_j-a)}$$

They are also easy to compute (Plaut & Roughgarden (2018)).

The Competitive approach starts with an approximation of the Envy Free test, that we can define in a strong or a weak form, like Fair Share above. The first condition is known as *Envy free up to the Least Valued Good* (Caragiannis et al. (2016)):

$$\text{EFX: } \forall i, j \in N, \forall a \in z_j : \text{if } u_{ia} > 0 \text{ then } u_{iz_i} \geq u_{i(z_j-a)} \iff u_{i(z_i+a)} \geq u_{iz_j}$$

With only two agents, $n = 2$, a LEX allocation meets EFX (Plaut & Roughgarden (2018)). *But for $n \geq 3$, despite much brainstorming and numerical experiments, we still do not know whether or not a division of the goods meeting EFX exists for all configurations of additive utilities.*

The second, weaker approximation of EF is *Envy free up to one Good*:

$$\text{EF1: } \forall i, j \in N, \exists a \in z_j : u_{iz_i} \geq u_{i(z_j-a)} \iff u_{i(z_i+a)} \geq u_{iz_j}$$

³⁰Pick an optimal partition for agent 1 and let agent 2 pick her best share.

Note that EF1 implies FS1.

The following mechanism, known as the “NFL draft mechanism”, picks an EF1 division of the goods. Fix a priority ordering of the agents and let them take turns, in that order, to pick their best object in the remaining pile (with n agents and m objects we need $\lceil \frac{m}{n} \rceil$ rounds).³¹ Say agent i has lower priority than j : after we take away from z_j the object j picked in the first round, agent i does not envy j ’s reduced share.

But the draft algorithm may not pick an efficient division of the goods, as in the following example where agent 1 chooses first:

	a	b	c	d
u_1	3	2	2	2
u_2	3	1	1	1

There is a surprising relation between the Nash bargaining solution and Envy Free up to one Good, reminiscent of the stronger connection uncovered by the Eisenberg Gale theorem when the goods are divisible.

Theorem (Caragiannis et al. (2016)): *Any division of the objects maximizing the Nash product of utilities is Efficient and meets EF1 (hence FS1 as well).*

Here is the main argument. Let z be such a division. Efficiency is clear. Now suppose that agent i envies agent j even after removing any object from z_j . Let a minimize the ratio $\frac{u_{jb}}{u_{ib}}$ in z_j : this definition and the envy assumption imply

$$\left\{ \frac{u_{ja}}{u_{ia}} \leq \frac{u_{jz_j}}{u_{iz_j}} \text{ and } u_{i(z_i+a)} < u_{iz_j} \right\} \implies \frac{u_{ja}}{u_{ia}} u_{i(z_i+a)} < u_{jz_j} \quad (15)$$

Let z' be the partition obtained by transferring a from j to i ; we compute

$$r = \frac{\prod_N u_{kz'_k}}{\prod_N u_{kz_k}} = \frac{u_{jz'_j} u_{iz'_i}}{u_{jz_j} u_{iz_i}} = \left(1 - \frac{u_{ja}}{u_{jz_j}}\right) \left(1 + \frac{u_{ia}}{u_{iz_i}}\right)$$

$$r > 1 \Leftrightarrow \frac{u_{ja}}{u_{jz_j}} \frac{u_{i(z_i+a)}}{u_{iz_i}} < \frac{u_{ia}}{u_{iz_i}}$$

which is the last inequality in (15), and we reach the desired contradiction.

One problem with the *Nash Max Product* (NMP) rule is its severe computational complexity (Lee (2017)). If we must deal with large problems, whether in terms of the number of agents or of objects, we may not be able to find a division exactly maximizing the Nash product: see Cole et al. (2016), Lee (2017). Fortunately polynomial-time approximations of NMP, in the number of agents and of objects, are available: Cole and Gkatzelis (2015) and Barman et al. (2018).

There are alternative combinatorial ways to construct an efficient partition meeting EF1. Barman et al (2018) provide such an approach, and in addition construct a partition that would remain efficient even if the goods were divisible.

³¹The Priority rule at the beginning of Section 5 is the special case where $n = m$.

The last frontier is the definition of a convincing approximation of the Competitive rule. Babaioff et al. (2017) propose an allocation competitive with respect to a common price, and allow small variations in the individual budgets, which does not guarantee Envy Free. Barman et al (2018) just mentioned reach a similarly competitive allocation that guarantees EF1 and simultaneously approximates the Maximum Nash Product.

Indivisible bads As when the items are divisible, we expect differences with the case of goods. Indeed this is already true for Fair Share. The analog of FSX is *Fair Share up to the Least Harmful Bad*:

$$\text{FSX: } \forall i \in N, \forall a \in z_i : u_{i(z_i-a)} \leq \frac{1}{n} u_{iA}$$

It is always possible to find a division of the bads meeting FSX. Here is a simple algorithm ensuring this.

Start with a problem (N, A, u) with normalized utilities ($u_{iA} = 1$). Define $\theta(a) = \min_{i \in N} u_{ia}$ and order the bads in decreasing order of $\theta(a)$ from $\theta(a^1)$ till $\theta(a^m)$. For each a^k we also pick an agent i^k such that $\theta(a^k) = u_{i^k a^k}$. Now we distribute a^1 to i^1 , a^2 to i^2, \dots , until the first k such that the total share of some agent i^* is strictly above $\frac{1}{n}$. We freeze then the share z_{i^*} : by construction FSX holds for i^* ; moreover in the restriction of our problem to $N \setminus i^*, A \setminus z_{i^*}$ we have $u_{i(A \setminus z_{i^*})} < \frac{n-1}{n}$, therefore an FSX partition there will meet FSX for all $i \in N \setminus i^*$ in the initial problem.

Another open question: can we always divide the bads *efficiently* and meet FSX?

If we weaken FSX to *Fair Share up to one Bad*:

$$\text{FS1: } \forall i \in N, \exists a \in z_i : u_{i(z_i-a)} \leq \frac{1}{n} u_{iA}$$

then it is always possible to divide the bads efficiently and meet FS1. A simple rounding procedure starting from the Competitive or the Egalitarian allocation of the associated problem with divisible bads will do this.

The strong approximation of EF, *Envy free up to the Least Harmful Bad*:

$$\text{EFX: } \forall i \in N, \forall a \in z_i, \forall j \in N : u_{i(z_i-a)} \leq u_{iz_j} \iff u_{iz_i} \leq u_{i(z_j+a)}$$

raises the same challenging open question as with goods: *can we always find a division of the bads meeting EFX?*

The weaker form of EF, *Envy free up to one Bad*:

$$\text{EF1: } \forall i \in N, \exists a \in z_i, \forall j \in N : u_{i(z_i-a)} \leq u_{iz_j} \iff u_{iz_i} \leq u_{i(z_j+a)}$$

is still implemented by the draft algorithm. But the Nash product of disutilities is useless to find an efficient division of the bads meeting EF1. *It is still an open question whether a division meeting EF1 exists at all.*

Finally if the manna is mixed, contains goods and bads, the very definitions of Fair Share, Envy Free, and even of the Egalitarian rule, are still work in progress.

8 More open problems and future directions

Constrained Competitive division of goods In the definition of a competitive allocation of divisible goods with price p and manna ω , agent i 's competitive demand z_i maximizes i 's utility over all allocations y_i affordable at price p with budget $\frac{1}{n}p \cdot \omega$, *whether or not y_i is feasible at ω* : for some good a we may have $y_{ia} > \omega_a$. Of course z_i itself is feasible. The Eisenberg Gale Theorem and the uniqueness of the profile of Competitive welfares do not hold anymore if the budget set excludes allocations unfeasible at ω .

An allocation (z, p) is *Constrained Competitive* (CC) if it is feasible, and for all i , z_i it maximizes i 's utility over all *feasible* and *affordable* allocations y_i :

$$\forall y_i \in \mathbb{R}_+^A \{y_i \leq \omega \text{ and } p \cdot y_i \leq \frac{1}{n}p \cdot \omega\} \implies u_i(z_i) \geq u_i(y_i)$$

This budget set is smaller than the Competitive budget set, therefore we have (many) more CC than C allocations. For instance with two agents and additive utilities, all efficient allocations meeting Fair Share are CC, and vice versa. Typically CC allocations contain a full dimensional subset of the set of efficient and Envy Free allocations.

The Constrained Competitive concept makes particularly good sense when we distribute indivisible goods. Say the family heirlooms comprise one grandfather clock, a bicycle and a I-phone: my standard budget set may allow me to buy two grandfather clocks, which does not make much sense as a counterfactual argument. Indeed the literature discussed in Section 7 always uses the Constrained version of the Competitive rule, without being explicit about it.

If the plausible conjecture below is true, the CC correspondence provides a solution to the difficulty encountered in example (12), Section 6. There in the C allocation one agent is stuck at his Fair Share utility level, while everyone else is strictly better off.³²

Conjecture: (Arrow Debreu preferences, divisible goods) if the profile of Fair Share utilities is not efficient, there is always at least one CC allocation where everyone gets strictly more than her Fair Share.

This is only a clue toward the widely open problem of identifying a normatively appealing selection of the CC correspondence.

Fair division of bads with additive utilities With divisible bads, additive utilities and n agents, the number of Competitive allocations welfare-wise distinct can be as high as $2^n - 1$. We see this by adapting the example in footnote 27: see Bogomolnaia et al. (2018). It is not known if this is the actual upper bound, but this observation is enough to make clear that the definition of a meaningful sub-correspondence of C is another challenging and important open problem.

³²Recall with additive utilities this is always true for an agent consuming a positive amount of each good.

We can broaden this search to include selections from the CC correspondence, or even from the Efficient and Envy Free correspondence. But Bogomolnaia et al. (2017) show even the latter correspondence admits no single valued selection continuous in the parameters of the problem (marginal utilities and manna). So we must be looking for a selection that is single valued almost everywhere, and stands out on some normative grounds. One approach is proposed in Bogomolnaia et al. (2018) for the simple case of two agents and/or two bads.

Random assignment with vNM preferences In the random assignment problem of Section 5, participants may be able to compare any two lotteries in terms of their expected von Neuman Morgenstern utilities. This model was introduced by (Hylland & Zeckhauser (1979)), who showed that a *Constrained* Competitive division exists: here the budget set consists of all the probability distributions over the objects affordable in the usual way. Note that it does not matter if we think of the objects as goods, bads, or mixed, because vNM preferences are translation invariant, everyone must receive one random object.

Very recently several authors noticed that multiple Competitive allocations is a robust possibility in this problem as well (McLennan (2018), Yanovskaia (private communication, 2017)), therefore this approach also require a selection step to become operational, just like in the case of a bad manna above. Except in the case of three objects, we do not know how many different Competitive assignments may robustly coexist.

9 References

- Abdulkadiroglu A., Sönmez T.1998. Random serial dictatorship and the core from random endowments in house allocation problems, *Econometrica*, 66, 689–701
- Aumann Y, Dombb Y.2010. The Efficiency of Fair Division with Connected Pieces, *International Workshop on Internet and Network Economics WINE 2010: Internet and Network Economics*,26-37
- Aumann RJ, Maschler M.1985. Game theoretic analysis of a bankruptcy problem from the Talmud, *J. Econ. Theory*,36(2), 195-21
- Aziz A, McKenzie S.2016. A Discrete and Bounded Envy-free Cake Cutting Protocol for Any Number of Agents, *IEEE 57th Annual Symposium on Foundations of Computer Science (FOCS)*, New Brunswick, NJ, USA, 9-11 Oct. 2016
- Babaioff M, Nisan N, Talgam-Cohen I.2017. Competitive Equilibria with Indivisible Goods and Generic Budgets, [arXiv.org > cs > arXiv:1703.08150](https://arxiv.org/abs/1703.08150).
- Barbera S, Jackson M, Neme A.1997. Strategy-proof allotment rules. *Games Econ. Behavior*, 18, 1–21.
- Barman S, Krishnamurthy SK, Vaish R.2018. Finding Fair and Efficient

- Allocations, *Proceedings of the 2018 ACM Conference on Economics and Computation EC '18*, 557-574, Ithaca, Massachusetts, USA, June 18 - 22, 2018
- Berliant M, Dunz K, Thomson W.1992. On the fair division of a heterogeneous commodity, *J. Math. Econ.*, 21, 201–16
- Bogomolnaia A.2015. Random assignment: Redefining the serial rule, *J. Econ. Theory*, 158(A), 308-31
- Bogomolnaia A, Heo EJ.2012. Probabilistic assignment of objects: Characterizing the serial rule, *J. Econ. Theory*, 147(5), 2072-82
- Bogomolnaia A, Moulin H. 2001. A New Solution to the Random Assignment Problem. *J. Econ. Theory*, 100, 295–328
- Bogomolnaia A, Moulin H. 2002. A simple random assignment problem with a unique solution, *Econ. Theory*, 19, 298–317
- Bogomolnaia A, Moulin H. 2004. Random Matching under Dichotomous Preferences, *Econometrica* 72, 257–279
- Bogomolnaia A, Moulin H, Sandomirskiy F, Yanovskaia E.2018. Dividing Bads under Additive Utilities, forthcoming *Soc. Choice Welfare*, doi: 10.1007/s00355-018-1157-x.
- Bogomolnaia A, Moulin H, Sandomirskiy F, Yanovskaia E.2017. Competitive division of a mixed manna, *Econometrica*, 85(6), 1847-71
- Bouveret S, Lang J.2008. Efficiency and envy-freeness in fair division of indivisible goods: logical representation and complexity. *J. Artificial Intelligence Research* 32, 525–564.
- Brams SJ, Taylor AD.1995. An envy-free cake division protocol, *American Math. Monthly*, 102(1):9-18
- Brams SJ, Taylor AD.1996. *Fair Division: From Cake-Cutting to Dispute Resolution*, Cambridge University Press, Cambridge, Mass, USA
- Branzei S.2015. A note on envy-free cake cutting with polynomial valuations, *Information Processing Letters* 115, 2, 93–95
- Budish E. 2011. The combinatorial assignment problem: Approximate competitive equilibrium from equal incomes. *J. Polit. Econ.* 119, 6, 1061–103
- Caragiannis I, Kurokawa D, Moulin H, Procaccia AD, Shah N, Wang J.2016. The Unreasonable Fairness of Maximum Nash Welfare, *Proceedings of the 2016 ACM Conference on Economics and Computation, EC '16*, 305–22
- Che YK, Kojima F.2010. Asymptotic Equivalence of Probabilistic Serial and Random Priority Mechanisms, *Econometrica*, 78(5),1625-72
- Ching S .1994. An Alternative Characterization of the Uniform Rule, *Soc. Choice Welfare*, 40, 57-60
- Chipman JS.1974.Homothetic preferences and aggregation, *J. Econ. Theory*, 8, 26-38
- Cole R, Gkatzelis V.2015. Approximating the Nash Social Welfare with Indivisible Items, *Proceedings of the 47th Annual ACM Symposium on Theory of Computing (STOC)*, 371–80
- Cole R, Devanur NR, Gkatzelis V, Jain K, Mai T, Vazirani VV, Yazdanbod S. 2017. Convex Program Duality, Fisher Markets, and Nash Social Welfare, *Proceedings of the 2017 ACM Conference on Economics and Computation EC'17*, 459-460

- Crès H, Moulin H.2001. Scheduling with opting out: improving upon random priority, *Operations Research*, 49, 565–77
- Devanur NR , Papadimitriou CH, Saberi A, Vazirani VV.2008. Market equilibrium via a primal–dual algorithm for a convex program. *Journal of the ACM*, 55: 1-18
- Dubins LE, Spanier EH.1961. How to cut a cake fairly. *American Math. Monthly*, 68, 1–17
- Dutta B, Ray D.1989. A Concept of Egalitarianism under Participation Constraints, *Econometrica*, 59, 615-36
- Dworkin R.1981a. What is equality? Part I: equality of welfare. *Phil.Public Affairs*, 10, 185–246
- Dworkin R.1981b. What is equality? Part II: equality of resources. *Phil. Public Affairs*, 10, 283–345
- Eisenberg E, Gale D.1959. Consensus of subjective probabilities: The parimutuel method. *Annals Math. Stat.*, 30(1), 165–168
- Fleurbaey M.1996. *Théories économiques de la justice*, Economica, Paris, France
- Foley D.1967. Resource allocation and the public sector, *Yale Economics Essays* 7, 45–98
- Ghodsí A, Zaharia M, Hindman B, Konwinski A, Shenker S, Stoica I.2011. Dominant Resource Fairness: Fair Allocation of Multiple Resource Types, *Proceedings of the 8th USENIX Conference on Networked Systems Design and Implementation (NSDI)*, 24–37
- Goldman J, Procaccia AD.2014. Spliddit: Unleashing Fair Division Algorithms, *SIGecom Exchanges*, 13(2), 41–6
- Guo M, Conitzer V.2007. Worst case optimal redistribution of VCG payments, *Proceedings of the 8th ACM Conference on Electronic Commerce EC’07*, 30–9, San Diego, CA, June 2007
- Hashimoto T, Hirata D, Kesten O, Kurino M, Unver U.2014. Two axiomatic approaches to the probabilistic serial mechanism, *Theoretical Economics*, 9(1),253-77
- Hurwicz L.1978. On the interaction between information and incentives in organizations, *K. Krippendorff (Ed.), Communication and control in society*, 123–47, New York: Scientific Publishers
- Hylland A, Zeckhauser R.1979. The efficient allocations of individual to positions, *J. Polit. Econ.*, 87, 293–314
- Kalai E, Smorodinsky M.1975. Other Solutions to Nash’s Bargaining Problem, *Econometrica*, 43(3), 513-8
- Katta AK, Sethuraman J.2006. A solution to the random assignment problem on the full preference domain, *J. Econ. Theory*, 131(1), 231-50
- Kurokawa D, Procaccia AD, Shah N.2015. Leximin Allocations in the Real World, *Proceedings of the Sixteenth ACM Conference on Economics and Computation EC’15* , 345-362, Portland, Oregon, USA, June 15 - 19, 2015
- Kurokawa D, Procaccia AD, Wang J.2016. When can the maximin share guarantee be guaranteed?, *Proceedings of the 30th AAAI Conference on Artificial Intelligence (AAAI)*,523-9

- Lee E.2017. APX-hardness of Maximizing Nash SocialWelfare with Indivisible Items, *Information Processing Letters*, 122, 17–20
- Li S.2017. Obviously Strategy-Proof Mechanisms, *American Econ. Rev.*,107, 3257–87
- Li J, Xue J.2013. Egalitarian division under Leontief preferences, *Econ. Theory*, 54 (3), 597–622
- Mas-Colell A.1992. Equilibrium Theory with Possibly Sated Preferences, *Equilibrium and Dynamics: Essays in Honor of David Gale*, ed. by M. Majumdar, 201–213, St. Martin’s Press, New York
- Mas-Colell A, Whinston MD, Green JR.1995. *Microeconomic Theory*, Oxford University Press, New York, USA
- Maskin E.1999. Nash equilibrium and welfare optimality, *Rev. Econ. Studies*, 66, 83–114
- McLennan A.2018. Efficient Disposal Equilibria of Pseudomarkets, *mimeo*, The University of Queensland, Australia.
- Moulin H.1988. *Axioms of Cooperative Decision Making*, Monograph of the Econometric Society, Cambridge University Press, Cambridge, Mass., USA
- Moulin H.1992. An application of the Shapley value to fair division with money, *Econometrica*, 60, 1331–49
- Moulin H.1995. *Cooperative Microeconomics: a game-theoretic introduction*, Princeton University Press, Princeton, USA
- Moulin H.1999. Rationing a commodity along fixed paths, *J. Econ. Theory*, 84, 41–72
- Moulin H.2009. Almost budget-balanced VCG mechanisms to assign multiple objects, *J. Econ. Theory*, 144(1), 96–119
- Moulin H.2017. One-dimensional mechanism design, *Theoretical Economics*, 12(2), 587–619
- Moulin H, Thomson W.1988. Can Everyone Benefit from Growth? Two Difficulties, *J. Math. Econ.*, 17, 339–45
- O’Neill B.1982. A problem of rights arbitration in the Talmud, *Math. Soc. Sci.*, 2, 345–71
- Orlin JB.2018. Improved algorithms for computing Fisher’s market clearing prices. *Proceedings of the 42nd ACM symposium on Theory of computing, STOC ’10*, ACM Press, 2010. 291
- Nash J.1950. The bargaining problem, *Econometrica*, 18(2), 155–62
- Nicolo A.2004. Efficiency and truthfulness with Leontief preferences. A note on two-agent, two-good economies, *Rev. Econ. Design*, 8, 373–82
- Pazner E, Schmeidler D.1978. Egalitarian equivalent allocations: A new concept of economic equity, *Quar. J. Econ.*, 92(4), 671–87
- Plaut B, Roughgarden T.2018. Almost envy-freeness with general valuations, *Proceedings of the Twenty-Ninth Annual ACM-SIAM Symposium on Discrete Algorithms SODA ’18*, 2584–603, New Orleans, Louisiana, January 07 - 10, 2018
- Pratt JW, Zeckhauser RJ.1990. The Fair and Efficient Division of the Winner Family Silver, *Management Sci.*, 36, 1293–301

- Procaccia AD. 2013. Cake cutting: not just child's play, *Communications of the ACM (CACM Homepage archive)*, 56(7), 78-87
- Procaccia AD, Wang J.2014. Fair Enough: Guaranteeing Approximate Maximin Shares, *Proceedings of the 14th ACM Conference on Economics and Computation EC'14*, 675–692
- Pycia M, Troyan P.2018. Obvious Dominance and Random Priority, SSRN: <https://ssrn.com/abstract=2853563> or <http://dx.doi.org/10.2139/ssrn.2853563>
- Rawls J.1971. *A Theory of Justice*, Belknap Press of Harvard University Press, Cambridge, Mass., USA
- Robertson JM, Webb WA.1998. *Cake Cutting Algorithms: Be Fair If You Can*, A. K. Peters
- Roemer J.1996. *Theories of distributive justice*, Harvard University Press, Cambridge, Mass, USA
- Roth AE, Sönmez T, Ünver MU.2004. Kidney Exchange, *Quart. J. Econ.*, 119 (2), 457–88
- Schummer J.1997. Strategy-proofness versus efficiency on restricted domains of exchange economies, *Social Choice and Welfare*, 14, 47–56
- Sen A.1970. *Collective Choice and Social Welfare*, Holden Day, San Francisco, USA
- Segal-Halevi E, Nitzan S, Hassidim A, Aumann Y.2017. Fair and square: Cake-cutting in two dimensions, *J. Math. Econ.*, 70, 1-28
- Segal-Halevi E, Sziklai B.2018. Monotonicity and Competitive Equilibrium in Cake-cutting, [arXiv.org > cs > arXiv:1510.05229](https://arxiv.org/abs/1510.05229)
- Sprumont Y.1991. The Division Problem with Single-Peaked Preferences: A Characterization of the Uniform Allocation Rule, *Econometrica*, 59, 509-19
- Steinhaus H.1948. The Problem of Fair Division, *Econometrica*, 16, 101–104
- Steinhaus H.1949. Sur la division pragmatique *Econometrica* (supplement), 17, 315-319
- Su FE.1999. Rental harmony: Sperner's lemma in fair division. *American Math. Monthly*, 10, 930–42
- Thomson W.1983. The fair division of a fixed supply among a growing population. *Math. of Oper. Res.*, 8, 319–26
- Thomson W.1995. *The theory of fair allocation*. forthcoming, Princeton University Press, Princeton NJ, USA
- Varian H.1974. Equity, envy and efficiency. *J. Econ. Theory*, 9, 63–91
- L. Zhou.1991. Inefficiency of strategy-proof allocation mechanisms for pure exchange economies, *Soc. Choice Welfare*, 8,247-57