

Combinatorial Assignment under Dichotomous Preferences

Guilherme Freitas *
Preliminary version.

March 30, 2010

Abstract

We consider the problem of assigning shares of an imperfectly divisible resource when preferences are dichotomous. One such problem is the problem of assigning bundles from a finite set of indivisible objects to a finite set of agents. When preferences are dichotomous, mechanisms that satisfy voluntary participation only require agents to report a set of acceptable bundles/shares. We characterize strategyproof mechanisms for such problems and provide a mechanism that is utilitarian-efficient, strategyproof and envy-free, thereby showing that impossibilities like the ones pointed out by Kojima (2009) can be circumvented if we assume dichotomous preferences. We also show that, unlike in the assignment problem with dichotomous preferences of Bogomolnaia and Moulin (2004), the existence of a Lorenz-dominant assignment is not guaranteed. We analyze real-world difficulties involved in using efficient mechanisms, both from a computational and a strategic point of view. In particular, we show that utilitarian-efficient mechanisms require computations that can have running times that are exponentially long in the number of agents, but we point out that some classes of problems can be solved faster. We also show that agents with general preferences facing a mechanism that is strategyproof and efficient in the dichotomous domain might have an incentive to misreport their acceptable shares/bundles, and in that case, the only profitable deviation is to report a smaller set of acceptable shares/bundles.

1 Introduction

It is the same story every end of term: students' schedules for extra activities change, and cello teacher Vanessa Sullivan has to change her teaching schedule accordingly. Finding a schedule

*Thanks to John Ledyard and Federico Echenique for helpful discussions. Thanks to J. Morgan Kousser for many suggestions that improved the readability of the text. Any errors are my own. Financial support from the California Institute of Technology is gratefully acknowledged.

that somehow satisfies all students is a frustrating process that involves a large number of emails and stress, and in the end, nobody really knows how good the chosen schedule is.

Companies, government agencies and different organizations face the same kind of problem every day: how to schedule the use of a scarce shared resource. High tech equipment, a modern conference room or a competent specialist are often expensive, or hard to find. In many cases, it is only sensible for an organization to acquire a limited amount of such a resource and divide it among its workers via time-sharing. Deciding how to assign time with the resource is the problem. The usual solution is analogous to the cello teacher's solution, with similar shortcomings: the parts involved communicate somehow and try to reach an agreement. Such negotiations can be very time consuming, and there is no guarantee that in the end a optimal solution will be found.

In principle, time shares of use of the resource can be assigned in any way. However, it is often the case that only a specific finite set of shares are assignable, be it for technological or institutional constraints. In those cases, time is an imperfectly divisible resource. If the set of shares is finite, then this scheduling is not fundamentally different from the combinatorial allocation problem, where bundles composed of a finite number of objects must be assigned to a finite number of agents. In this paper, we study all such problems under the umbrella of "imperfectly divisible resources."

When a resource is perfectly divisible, it is often called *cake*, and the problem of assigning shares of a cake to agents that have additive and continuous preferences (a non-atomic measure on the over some sigma algebra of subsets of the cake) has been extensively studied. These additivity and continuity assumptions have led to the proof of existence of mechanisms with very strong efficiency, incentive and fairness properties. In the cake cutting literature, actual procedures for such mechanisms can be found.

In contrast, by not assuming that the resource is perfectly divisible, but keeping the assumption of additive preferences, Kojima (2009) has shown that for the combinatorial allocation problem, there is no mechanism that is ordinally efficient, envy-free and weakly strategyproof. In this paper, we avoid that impossibility by assuming *dichotomous preferences*, an idea inspired by Bogomolnaia and Moulin (2004), which obtain very strong results for the problem of assigning at most 1 object among a finite set of objects to a finite set of agents. The main difference between this paper and Bogomolnaia and Moulin (2004) is that agents may be assigned multiple objects, or, more abstractly, that agents may have preferences over shares/bundles that are not jointly feasible.

In the following, we review a subset of the related literature, and present our main results: we characterize efficient assignments, strategyproof assignments, and provide a mechanism that is

efficient, strategyproof and envy-free. We provide an example that shows that unlike in the case of Bogomolnaia and Moulin (2004), the existence of a Lorenz-dominant random assignment is not guaranteed. We analyze the running times of computing efficient assignments, and argue that even though worst-case scenarios are not efficiently computable, there are a number of reasons to believe that in real-world applications my yield efficiently computable typical cases. We end by analyzing what would happen if a mechanism designer used a mechanism that is tailored for dichotomous preferences with agents that have more general preferences. In these mechanisms, agents are only asked to report which shares are better than nothing, and we show that in that case there are preferences for which agents have an incentive to misreport.

2 Related Literature

Economists, mathematicians, computer scientists and social scientists in general have long been interested in a family of problems that can be classified as assignment problems: problems where agents have to be given “one share” of a resource that is not privately owned by them. In the following section we present earlier attempts and results on different variants of the assignment problem.

2.1 Concepts to Classify the Literature

The problem that we presented in the introduction was related to the literature of three different problems, usually referred to as the assignment problem, which I will refer to as the *house assignment problem*¹ and the *cake-cutting problem*. The first refers to problems such as assignment of houses or jobs, and the second to problems such as assignment of time or land, when those are perfectly divisible.

There are four types of concepts that we will use to organize the literature around the assignment problem: first, the performance criteria to distinguish between “good” and “bad” assignments, namely Pareto efficiency, “fairness” and incentives to truth-telling; second, the divisibility of the resources; third, the continuity of preferences; and fourth, whether or not compensatory transfers are allowed.

2.2 Linear Programming and the Optimal Assignment Problem

A typical problem in the early literature was the *personnel assignment* problem: if there are n jobs and n workers *whose productivity at each of the n jobs is known*, what is the match of jobs and

¹Also known as the one-sided one-to-one matching problem

workers that maximizes the firm's profits, and how can we compute it? This problem illustrates the focus on some measure of efficiency, such as profits, and thus this family of problems was called the *optimal assignment* problem. Computational aspects of the problem were emphasized, with no concerns about fairness or incentives. Nonetheless, the problem was still of considerable economic interest. Problems like the optimal assignment of production facilities to different locations, and whether or not that can be achieved by a price system Koopmans and Beckman (1957) do not depend on any incentive-compatibility constraints, and do not call for fairness concerns.

The literature on the optimal assignment problem started in the early fifties as an application of nascent linear programming methods. The techniques used relied heavily on graph-theoretic arguments, initially developed by Hungarian mathematicians in the thirties. König-Hall's marriage theorem, the Birkhoff-von Neumann decomposition theorem, and Dantzig's simplex method were at the foundation of the results obtained at that time. See Kuhn (1955) for a list of early references, Berge (2001) and Papadimitriou et al. (1998) for textbook treatments.

As with the literature on the optimal assignment problem, we will be concerned with maximizing some measure of efficiency. However, we differ in a few key points: unlike houses, the resources we are assigning are divisible, and; we will be concerned with issues of fairness, efficiency, and incentive compatibility. Finally, while we try to present constructive solutions as much as possible, we will not attempt to describe and prove theorems about the performance of algorithms to implement our solutions; our focus will be on the properties of the solutions. We wait until the conclusion section to point out references that are relevant for computing some of our proposed solutions.

2.3 Fair ways to Cut a Cake

The problem is how to fairly allocate shares of a divisible and heterogeneous good, usually incarnated in the metaphor of a cake; one can cut a cake however one pleases, but different people might like different parts of the cake.

It is hopeless to try to summarize the literature on fair division in a few pages, so we will highlight some contributions that point to different directions in the literature and that are in some way related to our approach. The problem of defining fairness goes back to Plato; see the first chapters in Moulin (2004) for a concise and modern overview. Whenever it is not crucial, we will refer to "fairness" without specifying precisely which notion of fairness we are talking about.

Following the work of Steinhaus (1948), mathematicians have been drawn to the problem of fair division. That interest developed into what is often called the *cake-cutting* litera-

ture. See Barbanel (2005), Brams and Taylor (1996) and Robertson and Webb (1998) for textbook treatments. A classic performance requirement for a solution is that each agent is assigned at least his/her “fair share”. Another trademark of this literature is the assumption that utilities are measures, that is, additive set functions. The focus is on fairness and efficiency, not incentives. See Berliant et al. (1992) for an axiomatic treatment that includes incentives and uses Bewley’s Bewley (1972) result on the existence of general equilibrium in infinite dimensional economies to prove the existence of a group-envy-free and efficient allocation.

As with the cake-cutting problem, we are interested in the division of a heterogeneous and divisible good and we also care about fairness, in particular envy-freeness. However, we differ because we assume preferences to be discontinuous, which connects us to the next branch of the literature.

2.4 House Assignment Problems and Extensions

The classic house exchange market was introduced by Shapley and Scarf Shapley and Scarf (1974) in 1974: each agent has one house, and can only have one house. Additionally, agents have preferences over houses. How can we find an efficient allocation? Given the ownership structure, is the core of this game nonempty? If so, how can we find it?

The canonical answer was given by Gale’s *top trading cycles (TTC)* algorithm.² Assuming preferences are strict, pick an arbitrary agent and let him point to his favorite house; the owner in that house in turn points to his favorite house, so on and so forth. Because the number of houses/agents is finite, this problem will eventually come to a cycle. Let the agents in this cycle trade their houses according to their preferences. Remove these agents from the economy. Restart the procedure with the remaining agents. Cycles will be formed, and trades will occur until there are no more agents left, which signals the end of the procedure. The assignment obtained in the end is efficient, and it is the unique element in the core Roth and Postlewaite (1977). Moreover, TTC provides the basis for the following *strategyproof* direct mechanism Roth (1982): assign each house to one agent; ask for the agents’ preferences; apply the TTC procedure to obtain the final assignment.³

A related mechanism is the *serial dictatorship*, or *priority mechanism*, introduced by Satterthwaite and Sonnenschein Satterthwaite and Sonnenschein (1981): fix a strict priority order among the agents; ask for the preferences; give the agent at the top of the priority ranking his top choice, the second agent in the priority ranking his favorite object among the remaining ones,

²First presented in Shapley and Scarf’s Shapley and Scarf (1974) paper, but attributed to Gale.

³Roth’s Roth (1982) result does not require preferences to be strict, and then the mechanism has to be modified by requiring that ties in the preferences be broken by a fixed rule.

etc. This direct mechanism is strategyproof, and if preferences are strict, it is also efficient.

It is clear that for every outcome of a TTC procedure, there exists a priority ranking of the agents such that the outcome of the corresponding priority mechanism is the same as the one given by the TTC. In this sense, TTC and priority mechanisms are “equivalent”.⁴

Priority mechanisms or TTC procedures have great efficiency and incentives properties, but they can yield severely unfair outcomes. This is not a fault of the mechanism, but an inherent property of deterministic mechanisms for assignment problems with non-transferable utility. Think of the setting where all preferences are the same; any allocation is efficient, and someone will face the worst possible outcome.

One way to restore fairness is to allow for *random assignments*. For example, a *random priority* mechanism associates with each preference profile a probability distribution over priority mechanisms and assigns houses according to a mechanism drawn from this distribution; a TTC with random endowments assigns house-endowments randomly and then applies a TTC procedure. A random assignment can also be viewed as a mapping from agents to lotteries over houses.

Abdulkadiroglu and Sönmez Abdulkadiroglu and Sonmez (1998) prove that, in the strict preference domain⁵ random priority assignments are “equivalent” to TTC with random endowments. If the probabilities are chosen uniformly, both mechanisms are strategyproof, *ex-post* efficient, and fair in the sense of equal treatment of equals.

However, random priority assignments are still unambiguously undesirable in the following ways. First, there is no guarantee that for every preference profile there will be no envy among agents with respect to the lotteries they are assigned. Second, if agents have von Neumann-Morgenstern expected utilities, the outcome of these two mechanisms may be *ex-ante* inefficient, as conjectured by David Gale and proved by Lin Zhou Zhou (1990). Third, Bogomolnaia and Moulin Bogomolnaia and Moulin (2001) provide an example that shows that the outcome of a random priority mechanism may be first-order stochastically dominated by another assignment for all agents in the economy.

To address this problem, Bogomolnaia and Moulin Bogomolnaia and Moulin (2001) propose the concept of *ordinal efficiency* and the *probabilistic serial* mechanism that implements ordinally efficient assignments. An assignment is *ordinally efficient* when it is not first-order stochastically dominated by another assignment for all agents in the economy.⁶ The *probabilistic serial* mecha-

⁴It is a trivial fact that, for a given outcome of a priority mechanism, there is always an initial assignment and TTC procedure that leads to the same allocation; just make the initial assignment equal to the outcome of the priority assignment.

⁵Their results allow for indifferences, but the gist of their results is captured in the strict preference case.

⁶Alternatively, an assignment is ordinally efficient when there exists a utility profile such that the assignment is Pareto efficient for that utility profile McLennan (2002). See Manea (2008) for a constructive proof.

nism lets each agent “eat”, with equal speed, shares of one unit of their favorite object; as soon as one object is fully eaten, the agents that were eating that object move to their second-best object, and the process is repeated. After all objects have been eaten, each agent will have eaten shares of different objects. The Birkhoff-von Neumann theorem then guarantees that there exists a random assignment that gives to each agent a probability of getting a certain object equal to the share he has eaten, and there are constructive procedures to accomplish this. The resulting random assignment is ordinally efficient and envy-free; however, it is only weakly strategyproof, that is, an agent cannot obtain an allocation that first-order stochastically dominates the probabilistic serial allocation by misreporting his preferences.

We want solutions to the cake assignment problem satisfying strong incentive, fairness, and efficiency properties, such as the ones obtained by random priority or by the probabilistic serial mechanism in the house assignment problem. However, the divisional structure of the resource we are interested in is different: a cake is divisible, houses are not. We will see in examples that this feature completely changes the problem. In particular, efficient allocations in the housing problem always have someone obtaining his top choice; this will be no longer true in the cake assignment problem.

2.5 There and Back Again

Unfortunately, none of the mechanisms mentioned above work well in the case of the combinatorial assignment problem. Random priority mechanisms are known to be inefficient in the preferences of preferences with large indifference sets, even in the house assignment problem. In the case of strict preferences but in a combinatorial allocation problem, because one agent’s allocation can block two agents’ allocations, random priority can lead to allocations that allocate something to a small number of agents, which might be undesirable. And not even weak-strategyproofness is possible Kojima (2009) if we also require the random assignments to be ordinally efficient and envy-free. We present some examples to illustrate the failure of priority mechanisms and top trading cycles when desired resources may overlap.

The negative results above have a common feature: they try to solve the time assignment problem in a relatively large preference domain.⁷ In face of that difficulty, we will restrict our focus to preferences where, for each agent, the set of assignable shares is partitioned in a set of acceptable time slots and a set of unacceptable time slots. That is, there are only two indifference sets: the acceptable shares, and the non acceptable ones. As we focus on mechanisms that satisfy voluntary participation, these preferences are equivalent to dichotomous preferences.

⁷Kojima’s result about the impossibility of envy-free weak strategyproofness assumes additive preferences.

In the dichotomous preference domain, Bogomolnaia and Moulin (2004) obtain very strong results for the house assignment problem: their *egalitarian solution* is efficient, group-strategy proof, envy-free, and Lorenz-dominant among all efficient assignments. Their solution is not directly applicable to our problem. The reason is the following: in the house assignment problem, a dichotomous preference profile can be represented by a bipartite graph connecting agents to their acceptable houses. Efficient assignments are those that correspond to maximal matchings of this graph. *All* such matchings correspond to feasible assignments; in our case, where desired resources may overlap, some maximal matchings are not feasible.

On a final note, we must mention some other directions in the literature that attack similar problems, but with techniques or assumptions very removed from ours. The seminal paper of Hylland and Zeckhauser (1979) introduced the house assignment problem in the economics literature from the point of view of mechanism design. However, they worked with cardinal preferences, while we only work with properties that hold for all cardinal representation. Also in the domain of cardinal preferences, Ledyard, Noussair and Porter (1996) present a framework for dealing with a time allocation problem with capacities (like a many-to-one matching problem) in NASA's Deep Space Network. They recommend the use of an ascending-bid auction, the Adaptive User Selection Mechanism (AUSM) with tokens. They conduct some experiments to evaluate the performance of AUSM with tokens, AUSM with money and a random mechanism (the sequential dictator algorithm) and conclude that the AUSM with tokens is a better mechanism for situations with high level of conflict. We focus on mechanisms that do not depend on any sort of transfer, and thus the AUSM mechanism cannot be applied.

3 The Model

Consider the problem of assigning shares of a divisible resource to agents that only care whether or not they get an "acceptable" share or not. As we will focus on mechanisms that satisfy voluntary participation, there is no need to distinguish between the case where agents find unacceptable shares strictly worse than the empty set (which means getting nothing from the mechanism) and the case where they are indifferent between the empty set and unacceptable shares. Such preferences are characterized by one indifference set, the set of acceptable shares, and are called *dichotomous preferences*.⁸

An instance of this problem is characterized by:

- a set X of *resources*.

⁸As defined, for example, in Bogomolnaia and Moulin (2004).

- a finite collection \mathcal{F} of subsets of X , including the empty set \emptyset , that denote the assignable *shares* of the resource;
- a set $N = \{1, 2, \dots, n\}$ of *agents*;
- for each agent $i \in N$, an *acceptable set* $A(\succsim_i) \subset \mathcal{F}$.

We denote the set of all dichotomous preferences over \mathcal{F} by 2Pref , and we write $\succsim_i \in 2\text{Pref}$ to indicate the preference relation of agent i that has A_i as the acceptable set. A preference profile $(\succsim_1, \dots, \succsim_n)$ is denoted \succsim .

A *deterministic assignment* is a mapping $\mu : N \rightarrow \mathcal{F}$ subject to a feasibility restriction. This feasibility condition depends on the problem. In the case of the cello teacher, it would be

$$i \neq j \implies \mu(i) \cap \mu(j) = \emptyset. \quad (1)$$

In the case of a set of resources that can be used by 5 people at the same time, like a room with 5 machines, the feasibility condition would be:

$$\sum_{i \in N} \mathbb{1}_{\mu(i)}(x) \leq 5 \quad \forall x \in X, \quad (2)$$

where $\mathbb{1}_{\mu(i)} : X \rightarrow \{0, 1\}$ is an indicator function, assuming the value 1 when $x \in \mu(i)$ and the value 0 otherwise.

Note that if for every B, C in \mathcal{F} we have $B \cap C = \emptyset$, then we are back to the classic *housing assignment problem* Shapley and Scarf (1974), Hylland and Zeckhauser (1979). We reserve the symbols μ and η for deterministic assignments, whose set we denote \mathcal{M} .

We say a deterministic assignment *respects preferences* \succsim_i if $\mu(i) \neq \emptyset \implies \mu(i) \in A_i$. We denote $r(\succsim_i) \subseteq \mathcal{M}$ the set of all deterministic assignments that respect \succsim_i , and we define $r(\succsim) = \bigcap_{i \in N} r(\succsim_i)$ the set of all deterministic assignments that respect all the preferences in \succsim .

We need to define some more notation. For any function f , and binary relation R in the codomain of f , we write $[fRz]$ to indicate those elements x in the domain of f such that $f(x)Rz$, where z is in the codomain of f . For example, $[P > 0]$ is the support of the probability distribution P and $[\mu \neq \emptyset]$ is the set of all *assigned agents*, while $[\mu = \emptyset]$ is the set of *unassigned agents*.

A solution to an *instance* of our assignment problem is a *random assignment*: a probability distribution over deterministic assignments, that is, an element of $\Delta\mathcal{M}$. A solution to our assignment problem is a *direct mechanism* $g : 2\text{Pref}^n \rightarrow \Delta\mathcal{M}$, which maps every preference profile \succsim to a random assignment $g(\succsim) \in \Delta\mathcal{M}$. We want a solution to satisfy some minimum performance requirements, namely *efficiency*, *voluntary participation*, *strategyproofness* and *fairness*. Voluntary participation is obtained by focusing on mechanisms that only put positive probability on assignments μ that respect preferences, that is, $g(\succsim)(\mu) > 0 \implies \mu \in r(\succsim)$ for all $\succsim \in 2\text{Pref}^n$. We define and characterize the other concepts in the next sections.

4 Efficiency

We say that a deterministic assignment η *Pareto dominates* another assignment μ when $\eta(i) \succ_i \mu(i)$ for every $i \in N$ and $\eta(i) \succ_i \mu(i)$ for at least one $i \in N$. We say that μ is Pareto efficient, or simply an *efficient assignment*, if there is no other assignment that Pareto dominates μ .

With dichotomous preferences, efficient assignments are in some sense maximal. Define the preorder \supseteq on the set of assignments \mathcal{M} as follows:

$$\mu \supseteq \eta \Leftrightarrow [\mu \neq \emptyset] \supseteq [\eta \neq \emptyset],$$

in which case we say μ *contains* or *includes* η . It should be clear that the \supseteq symbol on the left side is the one we are defining, and the one on the right side refers to the usual set inclusion. The relations \supsetneq , \subseteq and \subsetneq on \mathcal{M} are defined as one would expect. The following lemma is an immediate conclusion of these definitions.

Lemma 1. *Given two deterministic assignments $\mu, \eta \in r(\succ)$, we say μ Pareto dominates η if and only if $\mu \supsetneq \eta$. An assignment $\mu \in r(\succ)$ is efficient if and only if it is \supseteq -maximal in $r(\succ)$: for every assignment $\eta \in r(\succ)$, $\mu \supseteq \eta$ whenever $\eta \supseteq \mu$.*

Define $f : \mathcal{M} \rightarrow \mathbb{R}$ to be \supseteq -*monotonic* if $f(\mu) \geq f(\eta)$ whenever $\mu \supseteq \eta$.

Corollary 1. *If an assignment $\mu \in r(\succ)$ maximizes a function f which is \supseteq -monotonic, then μ must be efficient.*

We say that a random assignment ψ *Pareto dominates* a random assignment ϕ when the probability of getting something acceptable in ψ is at least as great as in ϕ for every agent $i \in N$, and strictly greater for some agent $i \in N$. Our definition is motivated by the fact that for any expected utility representation of a binary preference, the utility of a lottery ψ is greater than the utility of a lottery ϕ if and only if the probability of getting something acceptable in ψ is greater or equal than the probability of getting something acceptable in ϕ . Therefore, we use the probability of obtaining an acceptable share as the canonical definition of the utility of a random assignment.⁹

We say that ψ is an *efficient random assignment* if there is no other random assignment that Pareto dominates ψ . We say that g is a *efficient mechanism* if $g(\succ)$ is an efficient random assignment for every preference profile $\succ \in 2\text{Pref}^n$.

It is easy to see that a random assignment $\phi \in \Delta\mathcal{M}$ is *utilitarian-efficient* — that is, it maximizes the sum of the utilities — if and only if its support is composed of deterministic assignments that respect preferences and that assign the largest possible number of agents.

⁹Our definition of efficiency is then the standard notion of *ex ante* efficiency for random assignments setting the utility of an acceptable share as 1 and the utility of an unacceptable share as 0.

5 Incentives to Truth-Telling

To keep the notation simple, we will extend an agent's preference relation over the set of shares \mathcal{F} to the set of all random assignments $\Delta\mathcal{M}$ in the only natural way: for $\psi, \phi \in \Delta\mathcal{M}$, we say $\psi \succ_i \phi$ if and only if the probability that i gets an acceptable share under ψ is at least as large as the probability that i gets an acceptable share under ϕ . Similarly, we say $g(\succ'_i, \succ_{-i}) \succ_i g(\succ''_i, \succ_{-i})$ when $g(\succ'_i, \succ_{-i})(r(\succ_i)) \geq g(\succ''_i, \succ_{-i})(r(\succ_i))$. Remember that $g(\succ) \in \Delta\mathcal{M}$ is a random assignment, that is, $g(\succ)(\mu)$ is the probability under $g(\succ)$ that the assignment μ will be chosen.

We say that a mechanism is *strategyproof* if for every agent there is never incentive for unilateral manipulation:

$$g(\succ) \succ_i g(\succ'_i, \succ_{-i}) \quad \forall \succ \in 2\text{Pref}^n, \forall \succ'_i \in 2\text{Pref}.$$

In the following, we will need the following notion: given preferences \succ_i and \succ'_i , we say that \succ_i is *less flexible than* \succ'_i when $A(\succ_i) \subseteq A(\succ'_i)$. Alternatively, we say that \succ'_i is *more flexible than* \succ_i .

Given \succ_i and \succ'_i , we also define the *join* $\succ_i \vee \succ'_i$ as the preference for which $A(\succ_i \vee \succ'_i) = A(\succ_i) \cup A(\succ'_i)$. We also define the *meet* $\succ_i \wedge \succ'_i$ of these two preferences as the preference for which $A(\succ_i \wedge \succ'_i) = A(\succ_i) \cap A(\succ'_i)$. It follows that $\succ_i \vee \succ'_i$ is more flexible than both \succ_i and \succ'_i , which are both more flexible than $\succ_i \wedge \succ'_i$.

Proposition 1. *A mechanism $g : 2\text{Pref}^n \rightarrow \Delta\mathcal{M}$ is strategyproof if and only if g is **monotonic** and **sub-additive** in the following sense: for all \succ_{-i} and for every \succ_i^+ more flexible than \succ_i we have respectively,*

$$g(\succ_i^+, \succ_{-i})(r(\succ_i^+)) \geq g(\succ_i, \succ_{-i})(r(\succ_i)),$$

and

$$g(\succ_i^+, \succ_{-i})(r(\succ_i)) \leq g(\succ_i, \succ_{-i})(r(\succ_i))$$

Proof. Until the end of the proof, fix the preferences of all agents but i at \succ_{-i} . Suppose that a mechanism g is monotonic and sub-additive. Let i 's true preferences be \succ_i . Consider an alternative report \succ'_i for i . If $A(\succ_i) \cap A(\succ'_i) = \emptyset$, then i is better off reporting his true preferences, as g respects preferences. Now, there are two cases. First, if $A(\succ'_i) \subsetneq A(\succ_i)$, then by monotonicity i is better off reporting the truth. Second, if $A(\succ'_i) \not\subseteq A(\succ_i)$, then by sub-additivity i is better off reporting $\succ_i \wedge \succ'_i$. But then, we are back to the first case, and i is better off saying the truth. Therefore, proves that a monotonic and sub-additive mechanism is strategyproof.

To prove the converse, suppose that a mechanism g is not monotonic, that is, there are \succ_i and a more flexible preference \succ_i^+ such that

$$g(\succ_i^+, \succ_{-i})(r(\succ_i^+)) < g(\succ_i, \succ_{-i})(r(\succ_i)).$$

Then if i 's true preference were \succsim_i^+ , he would have an incentive to report \succsim_i instead, and thus g cannot be strategyproof. Now suppose that the mechanism g is not subadditive, that is, there are \succsim_i and a more flexible preference \succsim_i^+ such that

$$g(\succsim_i^+, \succsim_{-i})(r(\succsim_i)) > g(\succsim_i, \succsim_{-i})(r(\succsim_i))$$

Then if i 's true preference were \succsim_i , he would have an incentive to report \succsim_i^+ instead, and thus g cannot be strategyproof. It follows that a strategyproof mechanism must be monotonic and sub-additive, and this completes the proof. \square

In other words, to check strategyproofness, we just have to insure that an agent will never profit from two types of deviations: reporting a larger acceptable set (discouraged by subadditivity) or reporting a smaller acceptable set (discouraged by monotonicity).

Consider a mechanism such that, for every agent with some fixed preference it is always the case that, for every possible report he could send, he is better off both dropping unwanted shares from his report, and adding desirable shares to his report. It is clear that such a mechanism is strategyproof. However, it is not the case that all strategyproof mechanisms have that property. While it is true that, when facing a strategyproof mechanism, an agent should drop unwanted shares from his report (due to subadditivity), it is not the case that he should always add desirable shares, as shown in the following example.

Example 1. Let $X = \{a, b, c\}$ be a set of objects and $\mathcal{F} = 2^X$ the set of assignable shares. Let $N = \{1, 2\}$ be the set of agents, and let g be the following mechanism: if it is possible to give both agents something acceptable under the reported preferences, the mechanism does so. Otherwise, the mechanism gives player 1 an acceptable share, unless player 1 only accepts the whole set X . We divide this remaining case in two, and we only give the winning probabilities of player 2; player 1 gets X with the complementary probability.

If player 2 wants only one object, he gets it with probability .5. If he accepts two shares with one object each (he may accept other shares, but they must have more objects), then he gets a given acceptable object with probability .3. If he accepts three shares with one object each, then he gets b with probability .46 and a or c with probability .12 each. If player 2 only accepts shares with 2 or more elements, then gets nothing and player 1 gets X with probability 1.

It is easy to check that this mechanism is monotonic and subadditive, and therefore strategyproof. It is also efficient. Now suppose that player 1 reports X as his only acceptable share, and player 2 accepts $\{a\}$ or $\{c\}$ and reports that he accepts $\{a\}$ or $\{b\}$. With that report, he would get something acceptable with probability .3. Now, if he added $\{c\}$ to the report as an acceptable share, he would get something acceptable with probability .24. Thus, it is not in his best interest to add $\{c\}$ to his report, even though it is a desirable share.

6 Fairness

First, let us define the concept of no envy and Lorenz dominance. The output of a mechanism is a random assignment $g(\succsim)$, and associated to this random assignment are individual lotteries $\psi_{g(\succsim)} : N \rightarrow \Delta\mathcal{F}$ for $i \in N$. We say that a mechanism is *envy-free* when for every preference profile \succsim every agent $i \in N$ prefers his lottery $\psi_{g(\succsim)}(i)$ to any lottery $\psi_{g(\succsim)}(j)$ of another agent $j \in N$. For every vector $x \in \mathbb{R}^n$, let $\bar{x} = x_{(1)}, x_{(2)}, \dots, x_{(n)}$ be a permutation of the coordinates of x where $x_{(i)} < x_{(i+1)}$. Given utility profiles $u, v \in \mathbb{R}^n$ we say that u *Lorenz-dominates* v when $\sum_{i=1}^k (\bar{u}_i - \bar{v}_i) \geq 0$ for all $k \in \{1, \dots, n\}$.

Kojima (2009) shows that there is no mechanism that is ordinally efficient, envy-free and weakly strategyproof for the problem of randomly assigning arbitrary bundles of a finite number of objects when preferences are additive. On the other hand, (Bogomolnaia and Moulin (2004)) show that the problem of randomly assigning a finite number of objects (not bundles) to agents that have dichotomous preferences admits a mechanism that is group-strategyproof, envy-free and always yield a random assignment that is Lorenz-dominant in the class of all efficient random assignments.

The results we will present fall between the aforementioned results in the literature. By requiring preferences to be dichotomous, we open the door for mechanisms that are not only efficient and envy-free, but also strategyproof. However, as the example below shows, it is not always possible to obtain Lorenz dominant random assignments.

Example 2. Let $N = \{1, 2, 3, 4, 5\}$ and suppose we obtain the following utilitarian-efficient support $uEff(\succsim)$:

$$a_1 = (1, 1, 1, 0, 0) \tag{3}$$

$$a_2 = (0, 0, 1, 1, 1) \tag{4}$$

$$a_3 = (0, 1, 0, 1, 1) \tag{5}$$

Note that for a random assignment to be Lorenz-dominant in the class of all efficient assignments, it has to assign the largest number of people, that is, it has to be utilitarian-efficient.

Let a random assignment with support in $uEff(\succsim)$ be represented by a triple $p = (p_1, p_2, p_3)$ where p_i is the probability that deterministic assignment a_i will be selected. If there was a Lorenz-dominant assignment p for these preferences, then this random assignment would have to maximize any social welfare function $u \mapsto \sum_{i=1}^n f(u_i)$ where f is strictly increasing and strictly concave (see Olkin and Marshall (1980)). However, if we select p so as to maximize the Nash social

welfare function and the Rawlsian welfare function (that maximizes the utility of the lowest-utility agent in society), we obtain different results. Therefore there cannot be a Lorenz-dominant assignment in this example.

Before we proceed, we need to define some more notation. We say that two assignments μ and η in \mathcal{M} are *equivalent* when $[\mu \neq \emptyset] = [\eta \neq \emptyset]$, that is, when they assign the same set of agents. For every set of assignments $M \subseteq \mathcal{M}$ let \widetilde{M} be the set of the corresponding equivalence classes of assignments in M . Note that the utility profiles of the agents are the same across assignments in the same equivalence class. Thus, when we say we will choose an element from \widetilde{M} at random, we mean that we choose an equivalence class at random, and then choose an arbitrary assignment in it for implementation purposes.

Proposition 2. *The mechanism where $g(\succ)$ is the uniform distribution over $\widetilde{uEff}(\succ)$ is utilitarian-efficient, strategyproof and envy-free.*

Proof. Fix arbitrary representatives for each equivalence class in $\widetilde{uEff}(\succ)$. We already argued that g is utilitarian efficient and strategyproof. It remains to show that g is envy-free. Let i and j be two different agents in N , and suppose that i weakly prefers j 's random lottery to his. Precisely, suppose

$$\psi_{g(\succ)(j)} \succsim_i \psi_{g(\succ)(i)}. \quad (6)$$

We will show that in fact, i must be indifferent between the two lotteries.

Condition (6) holds if and only if there is a set of shares S acceptable to i such that the probability that j obtains a share from S under $g(\succ)$ is greater or equal than the probability that i obtains an acceptable share under $g(\succ)$. Let M_j be the set of assignments in $\widetilde{uEff}(\succ)$ such that j gets a share from S and i is assigned the empty set. For every assignment $\mu \in M_j$, we can construct an assignment $\eta \in \mathcal{M}$ where $\eta(l) = \mu(l)$ for every $l \in N \setminus \{i, j\}$ and $\eta(i) = \mu(j)$. Denote the set of such assignments η by M_i . By construction, all assignments in M_i respect preferences. Because each $\mu \in M_j$ is in a different equivalence class of $\widetilde{uEff}(\succ)$ and assigns the same number of people to acceptable shares, there must be an agent in each $[\mu \neq \emptyset]$ that is assigned the empty set in the other elements of M_j . It follows that each $\eta \in M_i$ is also in a different equivalence class in $\widetilde{uEff}(\succ)$. Because the mechanism is utilitarian efficient and the number of agents in M_i is the same as the number of agents in M_j , it follows that all assignments in M_i are in $uEff(\succ)$. Finally, as i reports his true preference and the mechanism chooses assignments from $\widetilde{uEff}(\succ)$ with uniform probability, it must be the case that $g(\succ)(M_i) = g(\succ)(M_j)$. Therefore, as $M_i \subset r(\succ_i)$, it must be the case that

$$\psi_{g(\succ)(i)} \succsim_i \psi_{g(\succ)(j)},$$

which proves that g is envy-free. □

7 Computational Aspects

A mechanism can be applied in a real-life problem only if we can carry out the computations that it requires in a “reasonable amount of time”, which of course depends on the application. Here we show that any utilitarian-efficient mechanism in our setting requires computations that in the worst case, run in time that is at least exponential in the number of agents. Worst-case analysis is the usual way of comparing running time of algorithms in computer science. However, we will also point out that many — maybe most — real life applications of such mechanisms should lead to running times that are at worst polynomial in the number of agents. In what follows, we consider the feasibility constraint on deterministic assignments μ to be $\mu_i \cap \mu_j = \emptyset$ for all $i \neq j$ in N .

The basis of our analysis will be to show the equivalence of the problem of computing an efficient assignment to the problem of finding an inclusion-maximal independent set in a graph. An *independent set* in a graph is a subset of vertices such that no two vertices are connected by an edge. An inclusion-maximal independent set is an independent set that is not strictly contained in another independent set

If we are to implement efficient mechanisms, we need to compute inclusion-maximal assignments, elements of $Eff(\succ)$. We now show that such computation can always be cast as an independent set problem, and that any independent set problem can be cast as an assignment problem.

Proposition 3. *For every assignment problem $AP = (N, X, \mathcal{F}, \succ)$ with $\succ \in 2Pref^n$, we can construct an undirected graph $G = (V, E)$ where every efficient assignment in AP corresponds to one and only one inclusion-maximal independent set in G and vice-versa. Conversely, for every undirected graph $G = (V, E)$, we can construct an assignment problem $AP = (N, X, \mathcal{F}, \succ)$ with $\succ \in 2Pref^n$ such that every inclusion-maximal independent set in G corresponds to one and only one efficient assignment in AP and vice-versa.*

Proof. Let $AP = (N, X, \mathcal{F}, \succ)$ with $\succ \in 2Pref^n$ be an assignment problem. For every agent $i \in N$, define $Z_i = \{i\} \times A(\succ_i)$, and let $V = \cup_{i \in N} Z_i$. Now, let E be the set of all pairs $(a, b) \in V$, $a = (i_a, r_a)$, $b = (i_b, r_b)$, such that either the individuals i_a and i_b coincide, or the desired shares r_a and r_b are not jointly feasible, that is, $r_a \cap r_b \neq \emptyset$. It is easy to see that there is a bijection between the feasible assignments of AP and the independent sets of $G = (V, E)$. Furthermore, it

es easy to check that the efficient assignments in AP correspond exactly to the inclusion maximal independent sets of G .

Conversely, let $G = (V, E)$ be a undirected graph. Define the following assignment problem $AP = (N, X, \mathcal{F}, \succ$ with $\succ \in 2\text{Pref}^n$: let $N = V$, $X = E$, $\mathcal{F} = 2^X$ and let the acceptable set $A(\succ_i)$ of every agent $i \in N$ contain only one share: the set of all edges for which i is one of the vertices. Then there is a bijection between independent sets of G and feasible assignments of AP and the inclusion-maximal independent sets of G correspond exactly to the efficient assignments of AP . \square

The proof of proposition 3 gives us a recipe for computing independent sets: cast the problem as an independent set problem and compute the desired independent sets. Software for computing independent sets (or cliques, which are the dual problem) are readily available¹⁰. Computing a inclusion maximal independent set can be done efficiently; computing all such sets is an NP-hard problem.

A maximum independent set (or cardinality-maximal) is an independent set that has no less vertices than any other independent set. It is well known (see Kleinberg and Éva Tardos (2005)) that the problem of finding a maximum independent set or listing all inclusion-maximal independent sets is an NP-hard problem, that is, the worst-case running time is at least exponential in the number of vertices. Using the reduction provided in the proof of proposition 3, we conclude that to computing utilitarian-efficient assignments (those that assign the largest possible number of agents) is equivalent to computing cardinality-maximal independent sets.

Corollary 2. *Computing a utilitarian-efficient assignment is NP-hard: the worst-case running time is at least exponential in the number of agents.*

However, the worst case scenarios necessary to obtain exponential running times might not reflect many real-world problems. Independent sets can be computed in polynomial time for many special families of graphs, including interval graphs and planar graphs; also, for some families of “sparse” graphs, cliques¹¹ (the dual of independent sets) can be computed in polynomial time. See Chiba and Nishizeki (1985) for details.

We remark that the computational difficulties presented above have nothing to do with strategic issues: the problem is solely that of computing efficient and utilitarian-efficient assignments, even if agents report truthfully.

¹⁰For an incomplete list, check Skiena (2008), or the companion website <http://www.cs.sunysb.edu/~algorithm/>. A Python package called NetworkX, with all the necessary functionality, is developed at Los Alamos National Lab, <http://networkx.lanl.gov/>

¹¹A clique in a graph is a subset of vertices such that every two vertices are connected by an edge.

8 Robustness

So far, we have shown how to assign shares of a imperfectly divisible resource to agents that have dichotomous preferences in a way that is efficient (even utilitarian-efficient), strategyproof and envy-free. Such mechanisms are not very complicated, requiring agents only to report a subset of acceptable shares. Even so, as we argued in the previous section, such computations are potentially costly. Therefore, it is reasonable to want to use mechanisms that are no more complicated than the mechanism of proposition 2 for this kind of problem. This is especially the case if the mechanism designer only cares about assigning the largest number of agents to some acceptable share.

In real-world applications, it may be naive to assume that agents only care about obtaining an acceptable share or not. Even though that can be the only thing that the designer cares about, it may well be the case that agents have strict preferences between acceptable shares. It would be a very positive result if the designer could implement his objective of assigning the largest number of agents to some acceptable share by asking only each agent's acceptable shares (those that are better than nothing). In the next proposition, we show that this is impossible: indeed, any strategyproof mechanism for our assignment problem with dichotomous preferences gives incentives for some agents to report an acceptable set that is smaller than the true one. To make things simpler, in the following we assume that \mathcal{F} is a finite algebra of subsets of X .

Proposition 4. *If agents have two or more indifference classes in their acceptable sets, then a mechanism $g : 2\text{Pref}^n \rightarrow \Delta\mathcal{M}$ is strategyproof for all all possible utility functions that agents may have if and only if g is **monotonic** and **additive** in the following sense: for all \succsim_{-i} and for every \succsim_i^+ more flexible than \succsim_i we have respectively,*

$$g(\succsim_i^+, \succsim_{-i})(r(\succsim_i^+)) \geq g(\succsim_i, \succsim_{-i})(r(\succsim_i)),$$

and

$$g(\succsim_i^+, \succsim_{-i})(r(\succsim_i)) = g(\succsim_i, \succsim_{-i})(r(\succsim_i))$$

Proof. Suppose that we have a mechanism $g : 2\text{Pref}^n \rightarrow \Delta\mathcal{M}$ that is monotonic and additive as above. Then, for fixed preferences \succsim_{-i} agent i 's, mechanism g induces a probability distribution P_i over \mathcal{F} , such that $g(\succsim_i, \succsim_{-i})(r(\succsim_i)) = P_i(A(\succsim_i))$. It is easy to check that in that case, saying the truth is a dominant strategy for agent i , no matter which utility function $u_i : \mathcal{F} \rightarrow \mathbb{R}$ i has.

We now prove the contrapositive of the converse. Suppose that the mechanism $g : 2\text{Pref}^n \rightarrow \Delta\mathcal{M}$ is monotonic, but strictly subadditive. That is, there are shares $A, B \in \mathcal{F}$, such that for some $i \in N$, $\succsim_{-i} \in 2\text{Pref}^{n-1}$ and \succsim_i^+ such that $A(\succsim_i^+) = \{A, B\}$, \succsim_i such that $A(\succsim_i) = \{A\}$,

$$g(\succsim_i^+, \succsim_{-i})(r(\succsim_i)) < g(\succsim_i, \succsim_{-i})(r(\succsim_i)).$$

That is, the chance that i obtains A is higher if i does not include B in the report. In that case, it is easy to see that an agent that only finds A and B as better than nothing, but has a utility for A that is sufficiently larger than the utility for B will have an incentive to report that only A is acceptable. Therefore, the mechanism is not strategyproof. \square

It is easy to see now that efficiency and strategyproofness (in the strong sense that it holds for all possible utility representations of non-dichotomous preferences) are not always compatible if we only ask agents to report their acceptable sets. A simple example is the following: suppose we have a monotonic and additive mechanism g and all agents but agent 1 report the empty set as the only acceptable set. Then efficiency requires that whenever agent 1 accepts something, he must be given that with probability one. In particular, if $A, B \in \mathcal{F}$ are disjoint, and 1 reports A as the only acceptable set, then efficiency demands that 1 gets A with probability one; if 1 reports B as the only acceptable set, then efficiency requires that 1 gets B with probability one. However, if 1 reports exactly A and B as acceptable, then additivity requires that 1 obtains A or B with probability one, which is absurd. It is easy to see that such problems would occur in less artificial settings. This remark yields the following corollary.

Corollary 3. *If agents have two or more indifference classes in their acceptable sets, then there is no mechanism $g : 2\text{Pref}^n \rightarrow \Delta\mathcal{M}$ that is strategyproof and efficient for all possible utilities that agents may have.*

9 Conclusion

In this paper we present alternatives for a mechanism designer that wants to simplify the complex process of assigning shares of a divisible resource by allowing agents to report only a set of acceptable shares.

Assuming preferences are dichotomous, we show that strategyproof assignment mechanisms are characterized by a monotonicity and a sub-additivity condition, which translate into the property that no agent would like to report a larger or a smaller acceptable set. We note that by requiring preferences to be dichotomous, we avoid impossibilities given by Kojima (2009) for more general preferences, and we provide a strategyproof, efficient and envy-free mechanism. However we show that, unlike the case of assignment of individual goods, the existence of a Lorenz dominant assignment is not guaranteed.

We also provide an analysis of how easy it is to run such mechanisms “in the real world”, by examining their computational complexity and the incentive compatibility of these mechanisms when agents have general preferences. Drawing on theorems about the computational

complexity of finding independent sets/cliques in graphs, we show that finding utilitarian efficient assignments in arbitrary problems is “hard”(NP-hard), but that this worst-case scenario might not reflect the difficulties of a real-world application. Drawing on the same literature, we remark that an efficient assignment, not necessarily utilitarian-efficient, can be computed in polynomial time. We also show that agents with general utility functions may want to report an acceptable set that is significantly smaller than $\{A \in \mathcal{F} : u(A) > u(\emptyset)\}$, and that there is no way to guarantee strategyproofness and efficiency when agents have strict preferences among acceptable shares (those better than nothing).

Note: this paper needs improvement in the last section, “Robustness”. A sharper characterization of the incompatibility of efficiency and strategyproofness must be provided. It also needs improvement in the section “Computational Aspects”: a better explanation of which preference profiles correspond to the families of graphs for which computations are easy must be provided.

References

- Atila Abdulkadiroglu and Tayfun Sonmez. Random serial dictatorship and the core from random endowments in house allocation problems. *Econometrica*, 66(3):689–701, 1998.
- Julius B. Barbanel. *The Geometry of Efficient Fair Division*. Cambridge University Press, 2005.
- Claude Berge. *The Theory of Graphs*. Dover Publications, 2001.
- Marcus Berliant, William Thomson, and Karl Dunz. On the fair division of a heterogeneous commodity. *Journal of Mathematical Economics*, 21(3):201–216, 1992.
- T. Bewley. Existence of equilibria in economies with infinitely many commodities. *Journal of Economic Theory*, 4(3):514–540, 1972.
- A. Bogomolnaia and H. Moulin. A new solution to the random assignment problem. *Journal of Economic Theory*, 100(2):295–328, 2001.
- A. Bogomolnaia and H. Moulin. Random matching under dichotomous preferences. *Econometrica*, 72(1):257–279, 2004.
- Steven J. Brams and Alan D. Taylor. *Fair Division: From Cake-Cutting to Dispute Resolution*. Cambridge University Press, 1996.
- Norishige Chiba and Takao Nishizeki. Arboricity and subgraph listing algorithms. *SIAM Journal on Computing*, 14(1):210–223, February 1985.

- Aanund Hylland and Richard Zeckhauser. The efficient allocation of individuals to positions. *The Journal of Political Economy*, 87(2):293–314, 1979.
- Jon Kleinberg and Éva Tardos. *Algorithm Design*. Addison Wesley, 2005.
- Fuhito Kojima. Random assignment of multiple indivisible objects. *Mathematical Social Sciences*, 57(1):134–142, 2009.
- Tjalling C. Koopmans and Martin Beckman. Assignment problems and the location of economic activities. *Econometrica*, 25(1):53–76, 1957.
- H. W. Kuhn. The hungarian method for the assignment problem. *Naval Research Logistics Quarterly*, 2(1-2):83–97, 1955.
- John Ledyard, Charles Noussair, and David Porter. The allocation of a shared resource within an organization. *Review of Economic Design*, 2(1):163–192, 1996.
- Mihai Manea. A constructive proof of the ordinal efficiency welfare theorem. *Journal of Economic Theory*, 141(1):276–281, 2008.
- Andrew McLennan. Ordinal efficiency and the polyhedral separating hyperplane theorem*1. *Journal of Economic Theory*, 105(2):435–449, 2002.
- Hervé Moulin. *Fair Division and Collective Welfare*. The MIT Press, 2004.
- Ingram Olkin and Albert W. Marshall. *Inequalities: Theory of Majorization and Its Applications*. Academic Press, 1980.
- Christos H. Papadimitriou, Kenneth Steiglitz, and Kenneth Steiglitz. *Combinatorial Optimization*. 1998.
- Jack Robertson and William Webb. *Cake Cutting Algorithms*. A K Peters Ltd, 1 edition, 1998.
- Alvin E. Roth. Incentive compatibility in a market with indivisible goods*1. *Economics Letters*, 9(2):127–132, 1982.
- Alvin E. Roth and Andrew Postlewaite. Weak versus strong domination in a market with indivisible goods. *Journal of Mathematical Economics*, 4(2):131–137, 1977.
- Mark A. Satterthwaite and Hugo Sonnenschein. Strategy-Proof allocation mechanisms at differentiable points. *The Review of Economic Studies*, 48(4):587–597, 1981.

Lloyd Shapley and Herbert Scarf. On cores and indivisibility. *Journal of Mathematical Economics*, 1(1):23—37, 1974.

Steven S. Skiena. *The Algorithm Design Manual*. Springer, 2nd edition, 2008.

H. Steinhaus. The problem of fair division. *Econometrica*, 16(1):101–104, 1948.

L. Zhou. On a conjecture by gale about one-sided matching problems. *Journal of Economic Theory*, 52(1):123–135, 1990.