Beyond Dominant Resource Fairness: Extensions, Limitations, and Indivisibilities

Citation

Parkes, David C., Ariel D. Procaccia, and Nisarg Shah. 2012. Beyond dominant resource fairness: Extensions, limitations, and indivisibilities. Proceedings of the 13th ACM Conference on Electronic Commerce (EC '12), June 4-8, 2012, Valencia, Spain, 808-825. New York, NY: ACM Press.

Published Version

doi:10.1145/2229012.2229075

Permanent link

<http://nrs.harvard.edu/urn-3:HUL.InstRepos:11956916>

Terms of Use

This article was downloaded from Harvard University's DASH repository, and is made available under the terms and conditions applicable to Open Access Policy Articles, as set forth at <http://nrs.harvard.edu/urn-3:HUL.InstRepos:dash.current.terms-of-use#OAP>

Share Your Story

The Harvard community has made this article openly available.
Please share how this access benefits you. [Submit a story](#).

[Accessibility](#)

Beyond Dominant Resource Fairness: Extensions, Limitations, and Indivisibilities

DAVID C. PARKES, Harvard University

ARIEL D. PROCACCIA, Carnegie Mellon University

NISARG SHAH, Carnegie Mellon University

We study the problem of allocating multiple resources to agents with heterogeneous demands. Technological advances such as cloud computing and data centers provide a new impetus for investigating this problem under the assumption that agents demand the resources in fixed proportions, known in economics as *Leontief preferences*. In a recent paper, Ghodsi et al. [2011] introduced the *dominant resource fairness (DRF)* mechanism, which was shown to possess highly desirable theoretical properties under Leontief preferences. We extend their results in three directions. First, we show that DRF generalizes to more expressive settings, and leverage a new technical framework to formally extend its guarantees. Second, we study the relation between social welfare and properties such as truthfulness; DRF performs poorly in terms of social welfare, but we show that this is an unavoidable shortcoming that is shared by every mechanism that satisfies one of three basic properties. Third, and most importantly, we study a realistic setting that involves indivisibilities. We chart the boundaries of the possible in this setting, contributing a new relaxed notion of fairness and providing both possibility and impossibility results.

Categories and Subject Descriptors: I.2.11 [Distributed Artificial Intelligence]: Multiagent Systems; J.4 [Computer Applications]: Social and Behavioral Sciences—*Economics*

General Terms: Economics, Theory

1. INTRODUCTION

Resource allocation is a fundamental issue in the design and implementation of computing systems, which are naturally constrained in terms of CPU time, memory, communication links, and other resources. We are interested in settings where these resources must be allocated to multiple *agents* with different requirements. Such situations arise, e.g., in operating systems (where the agents can be jobs) or in cloud computing and data centers (where the agents can be users, companies, or software programs representing them). To take one example, *federated clouds* [Rochwerger et al. 2009] involve multiple agents that contribute resources; the redistribution of these resources gives rise to delicate issues, including fairness as well as incentives for participation and revelation of private information, which must be carefully considered.

Despite the growing need for resource allocation policies that can address these requirements, state-of-the-art systems employ simple abstractions that fall short. For example, as pointed out by Ghodsi et al. [2011], Hadoop and Dryad—two of the most widely-used cluster computing frameworks—employ a *single resource* abstraction for resource allocation. Specifically, these frameworks partition the resources into bundles—known as *slots*—that contain fixed amounts of different resources. The slots

Author's addresses: D. C. Parkes, School of Engineering and Applied Sciences, Harvard University, parkes@eecs.harvard.edu; A. D. Procaccia and N. Shah, Computer Science Department, Carnegie Mellon University, {arielpro,nkshah}@cs.cmu.edu.

Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies show this notice on the first page or initial screen of a display along with the full citation. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, to republish, to post on servers, to redistribute to lists, or to use any component of this work in other works requires prior specific permission and/or a fee. Permissions may be requested from Publications Dept., ACM, Inc., 2 Penn Plaza, Suite 701, New York, NY 10121-0701 USA, fax +1 (212) 869-0481, or permissions@acm.org.

© 2012 ACM 1539-9087/2012/01-ARTA \$10.00

DOI 10.1145/0000000.0000000 <http://doi.acm.org/10.1145/0000000.0000000>

are then treated as the system's single resource type, and at this point the allocation can be handled using standard techniques that were developed by the systems community. However, in a realistic environment where agents have heterogeneous demands, the single resource abstraction inevitably leads to significant inefficiencies.

Ghodsi et al. [2011] suggest a compelling alternative. Their key insight is that even though agents may have heterogeneous demands for resources, their demands can be plausibly assumed to be highly structured, in maintaining a fixed proportion between resource types. For example, if an agent wishes to execute multiple instances of a job that requires 2 CPUs and 1 GB RAM, its demand for these two resources has a fixed ratio of 2. Given 5 CPUs and 1.8 GB RAM, the agent can run only 1.8 instances of its task (note that Ghodsi et al. allow divisible tasks) despite the additional CPU, hence the agent would be indifferent between this allocation and receiving only 3.6 CPUs and 1.8 GB RAM. Preferences over resource bundles that exhibit this proportional structure are known as *Leontief preferences* in the economics literature. There are some positive results on resource allocation under Leontief preferences [Nicolò 2004], but more often than not Leontief preferences are drawn upon for negative examples.

Leveraging this model, Ghodsi et al. [2011] put forward the *dominant resource fairness (DRF)* mechanism. Briefly, DRF allocates resources according to agents' proportional demands, in a way that equalizes the shares that agents receive of their most highly demanded resource (we elaborate in Section 2). Ghodsi et al. demonstrate that DRF satisfies a number of prominent desiderata, including: (i) *Pareto optimality (PO)*, a well-known criterion of economic efficiency; (ii) *sharing incentives (SI)*, in the sense that agents are at least as happy as they would be under an equal split of the resources; (iii) *envy-freeness (EF)*, meaning that agents do not wish to swap their allocated resources with other agents; (iv) *strategyproofness (SP)*, which requires that agents cannot gain by misreporting their demands. From the perspective of economics, Ghodsi et al. discovered that the *egalitarian equivalent* rule [Pazner and Schmeidler 1978] possesses excellent properties under Leontief preferences.

At the risk of sounding melodramatic, we believe that resource allocation in computing systems is one of the most important and real challenges at the intersection of economics and computer science today. In particular, we find it exciting that the theory of fair division [Moulin 2003], as exemplified by the egalitarian equivalent rule, provides an ideal, almost tailor-made framework for tackling modern technological challenges in data centers and cloud computing. Despite the significant step forward made by Ghodsi et al. [2011], there are still many key issues that need to be addressed. Is it possible to rigorously extend the DRF paradigm to more expressive settings where agents are weighted and may not demand every resource? How can we formally tackle settings where agents' demands are indivisible? How does DRF compare to alternative mechanisms when social welfare is a concern? We provide answers to these questions below.

Our results

Our first order of business is to rigorously extend some of the theoretical results of Ghodsi et al. [2011]. In Section 3, we consider a setting where agents can be weighted based on an initial endowment, and modify the definitions of SI and EF accordingly. We also address the case where agents may not require every resource. We extend DRF to tackle these extensions of the model, and formally establish that the extended version still satisfies PO, SI, EF, and in fact satisfies the stronger coalitional version of SP, *group strategyproofness (GSP)*. (See below for a discussion of closely related results.)

In Section 4, we explore the relation between the desirable properties mentioned above, and maximization of the social welfare—the sum of utilities of the agents. To study social welfare we must assume an *interpersonal comparison of utilities*, that is,

we need to exactly specify the utilities of agents for allocations, rather than focusing on ordinal preferences over allocations. Focusing on a natural, albeit specific, utility function, we observe that DRF can produce allocations that only provide roughly $1/m$ of the social welfare of the optimal allocation, where m is the number of resources. However, we demonstrate that this poor welfare property is necessary for any mechanism that satisfies *at least one* of the three properties SI, EF, and SP. While the proofs for SI and EF are straightforward, the proof for SP is our most technical. These results vindicate DRF by establishing that reasonable methods for the allocation of multiple resources are inherently incompatible with social welfare maximization. In other words, the poor performance of DRF in terms of social welfare does not stem from its implicit focus on egalitarian welfare.

We view Section 5 as the most important on both a conceptual and practical level. The definition of Leontief preferences, on which the above results hinge, implicitly requires that both resources and tasks be divisible, so an agent's utility can be represented simply by the *proportion* between its demand for different resources. However, in practice an agent's task would require a minimum, *indivisible* bundle of resources. For example, if an agent requires 2 CPUs and 1 GB RAM to run one instance of its task, allocating 1 CPU and 1/2 GB RAM would be no more preferred than allocating nothing at all. Indeed, the agent's utility function should be conceptualized as a step function that only increases when another *complete* bundle is available. This setting is novel; even its single-resource variant has not been studied, to the best of our knowledge.

We observe that DRF performs poorly in this setting, but, more generally, strong impossibilities rear their head as one moves to a setting with indivisibilities. In particular, EF and PO are trivially incompatible; inspired by the fair division literature, we relax EF by considering the novel notion of *envy-freeness up to one bundle (EF1)*: an agent i prefers its own allocation to the allocation of another agent j , given that a single copy of the demanded bundle of i is removed from the allocation of j . Nevertheless, we show that SP is incompatible with PO and EF1 in a setting with indivisibilities, nor is it compatible with PO and SI. The main result of Section 5 is a possibility: we design a mechanism, SEQUENTIALMINMAX, which satisfies PO, SI, and EF1; given our impossibility results with respect to SP, this is the best combination of properties one could hope for. In general, the mechanism sequentially allocates bundles to minimize the maximum share of a resource that any agent has after allocation, but at each step only agents that satisfy certain conditions are eligible. We argue that this mechanism is practical for several reasons, in particular because it can be implemented dynamically (when agents arrive and depart, or change their demands over time).

Related work on the allocation of multiple resources

The work of Ghodsi et al. [2011] has quickly attracted significant attention from the algorithmic economics community. Consequently, during the process of preparing this paper, several independent (from our work and each other) related papers have become available. While none of these papers study social welfare maximization (our Section 4) or indivisibilities (our Section 5), there is a partial overlap with our results on extending DRF (Section 3). In the following, we discuss each of these papers in detail, starting from the original DRF paper.

Ghodsi et al. [2011] focus on the theoretical comparison of DRF to other potential mechanisms, and provide a systematic empirical study of the computational properties of DRF. Ghodsi et al. also mention that one can extend DRF to take into account weighted agents. However, the discussion of this extension is informal, and it is unclear whether DRF maintains its properties. In addition, although Ghodsi et al. do not assume that agents have strictly positive demands for every resource, this issue is addressed rather informally. Specifically, their proof of SP relies on a loose simulation

argument that, we believe, may be insufficient when some agents do not demand every resource. We introduce an alternative technical framework that allows us (in Section 3) to rigorously establish stronger properties (GSP instead of SP) while simultaneously tackling agent weights and zero demands on some resources.

A manuscript by Li and Xue [2011] provides a characterization of mechanisms that satisfy desirable properties under Leontief preferences. As a consequence of their results, one obtains a formal proof of GSP for certain mechanisms, even when agents are weighted. Curiously, their results do not capture DRF itself due to a technical assumption. More importantly, Li and Xue assume strictly positive demands; we argue in Section 3 that the case of zero demands is important, and our results suggest that it is also technically challenging.

Another manuscript by Friedman et al. [2011] explores the relations between resource allocation under Leontief preferences and bargaining theory. Among other results, they introduce a family of *weighted* DRF mechanisms. It is important to recognize that here the weights are not exogenously given, but are a way to induce different variations of DRF as mechanisms for allocating to *unweighted* agents. In addition, Friedman et al. explicitly assume strictly positive demands for simplicity. Under this assumption, they show that every weighted DRF mechanism that satisfies additional technical properties is GSP.

We mention two additional (already published!) related papers that are disjoint from ours in terms of results. Dolev et al. [2012] also study resource allocation under Leontief preferences. They consider an alternative fairness criterion, which they call *no justified complaints*. They compare allocations that satisfy this criterion with DRF allocations. Their main result is the existence of allocations that satisfy no justified complaints. Gutman and Nisan [2012] present polynomial time algorithms for computing allocations under a family of mechanisms that includes DRF. Moreover, they show that a competitive equilibrium (discussed in passing below) achieves the notion of fairness of Dolev et al. [2012], thereby leading to a polynomial time algorithm for computing allocations with this property. The two new mechanisms that we present are both polynomial time.

2. THE MODEL

The mathematical model is quite simple, and we therefore begin with an intuitive exposition based on an example from Ghodsi et al. [2011]; we also provide a different perspective on this example. Subsequently, we more rigorously introduce our notations and assumptions.

Intuition and an alternative interpretation of DRF

Consider a system with 9 CPUs, 18 GB RAM, and two agents. Agent 1 wishes to execute a (divisible) task with the demand vector $\langle 1 \text{ CPU}, 4 \text{ GB} \rangle$, and agent 2 has a (divisible) task that requires $\langle 3 \text{ CPU}, 1 \text{ GB} \rangle$. Note that each instance of the task of agent 1 demands $1/9$ of the total CPU and $2/9$ of the total RAM; the task of agent 2 requires $1/3$ of the total CPU and $1/18$ of the total RAM.

The *dominant resource fairness (DRF)* mechanism [Ghodsi et al. 2011] works as follows. The *dominant resource* of an agent is the resource for which the agent's task requires the largest fraction of total availability. In the example, the dominant resource of agent 1 is RAM, and the dominant resource of agent 2 is CPU. The DRF mechanism seeks to maximize the number of allocated tasks, under the constraint that the fractions of dominant resource that are allocated—called *dominant shares*—are equalized.

Returning to the running example, let y and z be the (possibly fractional) quantities of tasks allocated by DRF to agents 1 and 2, respectively; then overall $y + 3z$ CPUs and $4y + z$ GB are allocated, and these quantities are constrained by the availability

of resources. The dominant shares are $2y/9$ for agent 1 and $z/3$ for agent 2. The values of y and z can be computed as follows:

$$\begin{aligned} & \max (y, z) \\ & \text{subject to} \quad y + 3z \leq 9 \\ & \quad \quad \quad 4y + z \leq 18 \\ & \quad \quad \quad \frac{2y}{9} = \frac{z}{3} \end{aligned}$$

Due to the equality $2y/9 = z/3$ it is sufficient to maximize either y or z ; we maximize the pair (y, z) for consistency with [Ghodsi et al. 2011]. The solution is $y = 3$ and $z = 2$, i.e., agent 1 is allocated 3 CPUs and 12 GB RAM, and agent 2 is allocated 6 CPUs and 2 GB RAM.

We next introduce a novel, somewhat different way of thinking about DRF, which greatly simplifies the analysis of its properties. Let D_{ir} be the ratio between the demand of agent i for resource r , and the availability of that resource. In our example, when CPU is resource 1 and RAM is resource 2, $D_{11} = 1/9$, $D_{12} = 2/9$, $D_{21} = 1/3$, and $D_{22} = 1/18$. For all agents i and resources r , denote $d_{ir} = D_{ir}/(\max_{r'} D_{ir'})$; we refer to these demands as *normalized demands*. In the example, $d_{11} = 1/2$, $d_{12} = 1$, $d_{21} = 1$, $d_{22} = 1/6$.

We propose a linear program whose solution x is the dominant share of each agent. Observing the fixed proportion between an agent's demands for different resources, agent i is allocated an $(x \cdot d_{ir})$ -fraction of resource r .

$$\begin{aligned} & \max x \\ & \text{subject to} \quad \sum_i x \cdot d_{ir} \leq 1, \quad \forall r \end{aligned}$$

Allocations are bounded from above by 1 because the program allocates fractions of the total availability of resources. Clearly this linear program can be rewritten as:

$$x = \frac{1}{\max_r \sum_i d_{ir}}. \quad (1)$$

In the example, $x = 1/(1/2 + 1) = 2/3$. That is, the allocation to agent 1 is $1/3$ of the total CPU, and $2/3$ of the total RAM, which is equivalent to 3 CPUs and 12 GB RAM, as before. Similarly, agent 2 is allocated $2/3$ of the total CPU and $1/9$ of the total RAM, which is equivalent to 6 CPUs and 2 GB RAM, as before.

Rigorous model

Denote the set of agents by $N = \{1, \dots, n\}$, and the set of resources by R , $|R| = m$. As above, we denote the normalized demand vector of agent $i \in N$ by $\mathbf{d}_i = \langle d_{i1}, \dots, d_{im} \rangle$, where $0 \leq d_{ir} \leq 1$ for all $r \in R$. An allocation $\mathbf{A}$ allocates a fraction A_{ir} of resource r to agent i , subject to the feasibility condition $\sum_{i \in N} A_{ir} \leq 1$ for all $r \in R$. A resource allocation *mechanism* is a function that receives normalized demand vectors as input, and outputs an allocation.

Throughout the paper we assume that resources (e.g., CPU, RAM) are divisible. Up to (but excluding) Section 5, our model for preferences coincides with the domain of *Leontief preferences*. Let the utility of an agent for its allocation vector $\mathbf{A}_i$ be

$$u_i(\mathbf{A}_i) = \max\{y \in \mathbb{R}_+ : \forall r \in R, A_{ir} \geq y \cdot d_{ir}\}.$$

In words, the utility of an agent is the fraction of its dominant resource that it can actually use, given its proportional demands and its allocation of the various resources. Note that *this definition makes the implicit assumption that agents' tasks are divisible*; we relax this assumption in Section 5. Unless explicitly mentioned otherwise, we do

not rely on an interpersonal comparison of utilities. Put another way, with the exception of Section 4, an agent's utility function simply induces ordinal preferences over allocations, and its exact value is irrelevant.

An allocation $\mathbf{A}$ is called *non-wasteful* if for every agent $i \in N$ there exists $y \in \mathbb{R}_+$ such that for all $r \in R$, $A_{ir} = y \cdot d_{ir}$. Note that if $\mathbf{A}$ is a non-wasteful allocation then for all $i \in N$,

$$u_i(\mathbf{A}'_i) > u_i(\mathbf{A}_i) \Rightarrow \forall r \in R \text{ s.t. } d_{ir} > 0, A'_{ir} > A_{ir}. \quad (2)$$

Properties of mechanisms

DRF is intuitively fair as it equalizes dominant shares. To mathematically gauge its desirability we consider a set of formal properties that are deemed to be the most important by Ghodsi et al. [2011], and are also featured prominently in the work of Li and Xue [2011]. Ghodsi et al. established that these properties are all satisfied by DRF.

An allocation mechanism satisfies *sharing incentives (SI)* (also known in this context as *individual rationality*) if $u_i(\mathbf{A}_i) \geq u_i(1/n, \dots, 1/n)$ for all $i \in N$. In other words, an agent must receive an allocation that it values at least as much as the allocation that splits all resources equally between all agents. This property is appropriate when each agent contributes an equal amount of resources to the system (we generalize it below).

An allocation mechanism is *envy free (EF)* if $u_i(\mathbf{A}_i) \geq u_i(\mathbf{A}_j)$ for all agents $i, j \in N$. In words, each agent prefers its allocation to the allocation received by any other agent.

An allocation mechanism is *Pareto optimal (PO)* if it returns an allocation $\mathbf{A}$ such that for all allocations $\mathbf{A}'$, if there is an agent $i \in N$ such that $u_i(\mathbf{A}'_i) > u_i(\mathbf{A}_i)$ then there is an agent $j \in N$ such that $u_j(\mathbf{A}'_j) < u_j(\mathbf{A}_j)$, that is, there is no other allocation where all agents are at least as well off and at least one agent is strictly better off.

Finally, taking the game theoretic perspective, an agent's demand vector is its private information, and we would like to motivate agents to truthfully report their demands regardless of the reported demands of other agents. We say that an allocation mechanism is *strategyproof (SP)* if truth-telling is a dominant strategy. We define SP in words to avoid introducing additional notation that will not be used later on. Given a vector of demands for all agents $\mathbf{d}$, let $\mathbf{A}$ be the resulting allocation, and let $\mathbf{A}'$ be the allocation returned by the mechanism when other agents report the same demands $\mathbf{d}_{-i}$, and agent $i \in N$ reports $\mathbf{d}'_i$. Then under an SP mechanism $u_i(\mathbf{A}_i) \geq u_i(\mathbf{A}'_i)$, where u_i is the utility function induced by $\mathbf{d}_i$.

3. EXTENSIONS: WEIGHTS, ZERO DEMANDS, AND GROUP STRATEGYPROOFNESS

In this section we depart from the framework of Ghodsi et al. [2011] in three ways. First, we allow agents to be weighted, based on their contribution to the system. Second, we explicitly model the case where agents may have zero demands and therefore DRF needs to allocate in multiple rounds (these first two issues were informally considered by Ghodsi et al., as we discussed in Section 1). Third, we study stronger game theoretic properties. We show that DRF can be modified to address these realistic extensions, although the analysis of its game theoretic properties becomes more intricate.

For (contribution) weights, we assume that each agent $i \in N$ has a publicly known weight w_{ir} for resource $r \in R$, which reflects the agent's endowment of that resource (think of this as the amount of the resource contributed by the agent to the resource pool). We assume without loss of generality that for all resources $r \in R$, $\sum_{i \in N} w_{ir} = 1$. PO and SP are defined identically, but the weighted setting does require modifications to the notions of SI and EF, which are redefined as follows. SI now means that an agent

$i \in N$ receives as much value as it would get from the allocation that assigns it a w_{ir} -fraction of each resource $r \in R$, i.e., $u_i(\mathbf{A}_i) \geq u_i(\langle w_{i1}/w_{j1}, \dots, w_{im}/w_{jm} \rangle)$. EF requires that agent i does not envy agent j when the allocation of j is scaled by w_{ir}/w_{jr} on each resource r . Formally,

$$u_i(\mathbf{A}_i) \geq u_i(\langle (w_{i1}/w_{j1}) \cdot A_{j1}, \dots, (w_{im}/w_{jm}) \cdot A_{jm} \rangle).$$

An alternative definition for the weighted version of EF might simply require that in the special case where agents have a uniform weight $w_i = w_{ir}$ for all $r \in R$, agent i does not prefer the allocation of agent j scaled by w_i/w_j . However, note that the above definition is stronger, and, as we shall demonstrate, the stronger version is feasible.

Our second contribution in this section is explicitly modeling the case where agents may not demand every resource. At first glance the assumption of positive demands, which implies that DRF can be simulated via a single allocation according to Equation (1), seems very mild. After all, it may seem that the allocation when an agent has zero demand for a resource would be very similar to the allocation when that agent has almost zero demand for the same resource.

To see why this is not the case, let the normalized demand vector of agent 1 be $\langle 1, \epsilon \rangle$ (i.e., $d_{11} = 1$ and $d_{12} = \epsilon$), and let agents 2 and 3 have the demand vector $\langle \epsilon, 1 \rangle$. DRF allocates roughly $1/2$ of resource 1 to agent 1, and roughly $1/2$ of resource 2 to each of agents 2 and 3. If the ϵ demands are replaced by zero, the first-round allocation (using Equation (1)) would be similar. However, then a second round would take place; participation is restricted to agents that only demand non-saturated resources. In this second round agent 1 would receive the unallocated half of resource 1, hence agent 1 would ultimately be allocated all, instead of just half, of resource 1.

In dealing with both weights and zero demands, we extend DRF as follows. Let r_i^* be the weighted dominant resource of agent $i \in N$, defined by $r_i^* \in \arg \min_{r \in R: d_{ir} > 0} (w_{ir}/d_{ir})$. Informally, an agent's demand is now scaled by its weight on each resource. Next, we let $\rho_i = w_{ir^*}/d_{ir^*}$ be the ratio of weight to demand on the weighted dominant resource of agent $i \in N$; for now we assume, for ease of exposition, that $d_{ir} > 0$ implies $w_{ir} > 0$ (we explain how to drop this assumption later). The extended DRF mechanism proceeds in rounds; let s_{rt} be the surplus of resource $r \in R$ left unallocated at the beginning of round t (so $s_{r1} = 1$ for all $r \in R$). The allocation of resource r to agent i in round t is denoted A_{irt} .

ALGORITHM 1: EXTENDED DRF

Data: Demands $\mathbf{d}$, weights $\mathbf{w}$

Result: An allocation $\mathbf{A}$

$t \leftarrow 1$;

$\forall r, s_{r1} \leftarrow 1$;

$S_1 \leftarrow N$;

while $S_t \neq \emptyset$ **do**

$x_t \leftarrow \frac{s_{rt}}{\max_{r \in R} \sum_{i \in S_t} \rho_i \cdot d_{ir}}$;

$\forall i \in S_t, r \in R, A_{irt} \leftarrow x_t \cdot \rho_i \cdot d_{ir}$;

$\forall i \in N \setminus S_t, r \in R, A_{irt} \leftarrow 0$;

$\forall r \in R, s_{r,t+1} \leftarrow s_{rt} - \sum_{i \in S_t} A_{irt}$;

$t \leftarrow t + 1$;

$S_t \leftarrow \{i \in N : \forall r \in R, d_{ir} > 0 \Rightarrow s_{rt} > 0\}$ /* Agents demanding unsaturated resources */

end

$\forall i \in N, r \in R, A_{ir} \leftarrow \sum_{k=1}^{t-1} A_{irk}$

Our main result of this section is that EXTENDED DRF satisfies the four desirable properties introduced above. In fact, the mechanism satisfies an even stronger game-theoretic property known as *group strategyproofness (GSP)*: whenever a coalition of

agents misreports demands, there is a member of the coalition that does not strictly gain.

THEOREM 3.1. *EXTENDED DRF is PO, SI, EF, and GSP.*

It is trivial that the mechanism is PO, because resources are allocated in the correct proportions, i.e., it is non-wasteful, and moreover it allocates resources as long as players can derive value from them. Establishing SI is also a simple matter. Indeed, for every $r \in R$,

$$\sum_{i \in N} \rho_i \cdot d_{ir} \leq \sum_{i \in N: d_{ir} > 0} \left(\frac{w_{ir}}{d_{ir}} \cdot d_{ir} \right) = \sum_{i \in N: d_{ir} > 0} w_{ir} \leq 1$$

and therefore $x_1 \geq 1$. Thus, each agent i receives $x_1 \cdot \rho_i \cdot d_{ir_i}^* = x_1 \cdot w_{ir_i}^* \geq w_{ir_i}^*$ of resource r_i^* already in round 1. We need to show that each agent $i \in N$ values its allocation as much as the allocation $\langle w_{i1}, \dots, w_{im} \rangle$, and this now follows from the fact that the mechanism is non-wasteful and from Equation (2).

For EF, let $i, j \in N$; we argue that agent i does not envy j . Indeed, let t_i and t_j , respectively, be the last rounds in which these agents were allocated resources. If $t_j > t_i$, agent j does not demand some resource that i does and hence is not allocated a share of that resource. We can therefore assume that $t_j \leq t_i$. If $d_{jr_i}^* > 0$ then

$$A_{jr_i}^* = \left(\sum_{t=1}^{t_j} x_t \right) \cdot \rho_j \cdot d_{jr_i}^* \leq \left(\sum_{t=1}^{t_j} x_t \right) \cdot \frac{w_{jr_i}^*}{d_{jr_i}^*} \cdot d_{jr_i}^* = \left(\sum_{t=1}^{t_j} x_t \right) \cdot w_{jr_i}^*,$$

and otherwise $A_{jr_i}^* = 0$. Moreover,

$$A_{ir_i}^* = \left(\sum_{t=1}^{t_i} x_t \right) \cdot \rho_i \cdot d_{ir_i}^* = \left(\sum_{t=1}^{t_i} x_t \right) \cdot w_{ir_i}^* \geq \left(\sum_{t=1}^{t_j} x_t \right) \cdot w_{ir_i}^*.$$

Scaling $A_{jr_i}^*$ by $(w_{ir_i}^*/w_{jr_i}^*)$, we get at most $A_{ir_i}^*$; EF is then implied by Equation (2).

The challenge is to show that EXTENDED DRF is GSP. To gain some insight, let us concentrate first on a very special case: strictly positive demands (one round), no weights ($\rho_i = 1$ for all $i \in N$), and SP rather than GSP. For this, assume that agent $i \in N$ reports the demand vector $\mathbf{d}'_i$ instead of $\mathbf{d}_i$; this leads to the solution x' to Equation (1) instead of x , which induces the allocation $A'_{ir} = x' \cdot \mathbf{d}'_{ir}$. If $x' \leq x$ then agent i receives $x' \cdot \mathbf{d}'_{ir_i} \leq x \cdot 1 = x \cdot \mathbf{d}_{ir_i}$ of its dominant resource r_i , so its utility cannot increase by Equation (2). If $x' > x$, consider a resource r that was saturated when reporting $\mathbf{d}_i$. It holds that

$$A_{ir} = 1 - \sum_{j \neq i} A_{jr} = 1 - \sum_{j \neq i} x \cdot d_{jr} > 1 - \sum_{j \neq i} x' \cdot d_{jr} = 1 - \sum_{j \neq i} A'_{jr} \geq A'_{ir},$$

hence, once again, the utility of agent i cannot increase by Equation (2).

The next lemma extends the above argument to the case of possibly zero demands and weighted agents, and strengthens SP to get GSP, thereby completing the proof of Theorem 3.1.

LEMMA 3.2. *EXTENDED DRF is GSP.*

PROOF. Assume, without loss of generality, that the set of manipulating agents is $M \subseteq N$ and they report untruthful normalized demand vectors $\mathbf{d}'_M = \langle \mathbf{d}'_i \rangle_{i \in M}$ (which induce ratios ρ'_i); let $\mathbf{d}'$ be the collection of normalized demand vectors where $\mathbf{d}'_i = \mathbf{d}_i$ for all $i \in N \setminus M$.

Let t^* be the first round when a demanded resource of some $i \in M$ becomes saturated, under truthful or untruthful demands (i.e., the minimum of the two rounds). Let $isSaturated$ be a boolean variable that is true if and only if a demanded resource of a manipulator becomes saturated at the end of round t^* under the truthful reports, and define $isSaturated'$ similarly for the untruthful reports. As before, denote by A_{irt} (resp., A'_{irt}) the share of resource $r \in R$ allocated to agent $i \in N$ in round t , and let $A_{ir} = \sum_t A_{irt}$ (resp., A'_{ir}) be the total fraction of resource $r \in R$ allocated to agent $i \in N$ under $\mathbf{d}$ (resp., $\mathbf{d}'$). In addition, S_t (resp., S'_t) is the set of agents that demand only unsaturated resources in round t under $\mathbf{d}$ (resp., $\mathbf{d}'$).

CLAIM 3.3.

- (1) For all $t \leq t^*$ and $r \in R$ such that $d_{ir} = d'_{ir} = 0$ for all $i \in M$, $s_{rt} = s'_{rt}$.
- (2) For all $t \leq t^*$, $S_t = S'_t$.
- (3) For all $t < t^*$, $x_t = x'_t$.

PROOF. We prove the claim by induction on t . The base of the induction is trivial, as $s_{r1} = s'_{r1} = 1$ for all $r \in R$, and $S_1 = S'_1 = N$. We can also let $x_0 = x'_0$ as a formality.

For the induction step, consider round $t < t^*$. We assume statements (1) and (2) are true for all $t' \leq t$, and that statement (3) is true for all $t' < t$. We prove statement (3) for round t and then we prove statements (1) and (2) for round $t + 1$.

In round t , every agent $i \in M$ has no saturated resources under $\mathbf{d}$ and $\mathbf{d}'$, hence there are resources $r, r' \in R$ such that for all $i \in M$, $d_{ir} = d'_{ir} = 0$ and $d_{ir'} = d'_{ir'} = 0$, which become saturated at round t under the two demand vectors, i.e.,

$$x_t = \frac{s_{rt}}{\sum_{i \in S_t} \rho_i \cdot d_{ir}},$$

and

$$x'_t = \frac{s'_{r't}}{\sum_{i \in S'_t} \rho'_i \cdot d'_{ir'}}.$$

By the induction assumption, $s_{rt} = s'_{rt}$, $s_{r't} = s'_{r't}$, and $S_t = S'_t$. Moreover, $\sum_{i \in S_t} \rho_i \cdot d_{ir} = \sum_{i \in S'_t} \rho'_i \cdot d'_{ir}$ and $\sum_{i \in S_t} \rho_i \cdot d_{ir'} = \sum_{i \in S'_t} \rho'_i \cdot d'_{ir'}$, where last two equalities hold because every $i \in M$ does not demand either r or r' under $\mathbf{d}$ or $\mathbf{d}'$ (and hence the summations include zero terms for these agents). It follows that

$$x_t = \frac{s_{rt}}{\sum_{i \in S_t} \rho_i \cdot d_{ir}} = \frac{s'_{r't}}{\sum_{i \in S'_t} \rho'_i \cdot d'_{ir'}} \geq x'_t,$$

and similarly $x'_t \geq x_t$. We conclude that $x_t = x'_t$.

To establish statement (1) for round $t + 1$, let $r \in R$ such that $d_{ir} = d'_{ir} = 0$ for all $i \in M$. Using the induction assumption we conclude that

$$s_{r,t+1} = s_{rt} - x_t \sum_{i \in S_t} \rho_i d_{ir} = s'_{r't} - x'_t \sum_{i \in S'_t} \rho'_i d'_{ir'} = s'_{r',t+1}.$$

Finally, the assertion that $S_{t+1} = S'_{t+1}$ follows from the fact that all resources not demanded by members of the coalition have the same surplus under both $\mathbf{d}$ and $\mathbf{d}'$, and resources that are demanded by at least one member of the coalition are not saturated under $\mathbf{d}$ nor under $\mathbf{d}'$. $\square$

Having established Claim 3.3, we proceed with the lemma's proof by distinguishing between four cases.

Case 1: $x'_{t^} \geq x_{t^*}$ and $isSaturated$ is true.* Let r be a resource that is saturated under d in round t^* , and is demanded by some manipulator, i.e., there is $i \in M$ such that $d_{ir} > 0$. Note that such a resource exists by the definition of the variable $isSaturated$. Using Claim 3.3, we have that for every $t < t^*$ and every $i \in N \setminus M$,

$$A_{irt} = x_t \cdot \rho_i \cdot d_{ir} = x'_t \cdot \rho'_i \cdot d'_{ir} = A'_{irt},$$

and similarly $A_{irt^*} \leq A'_{irt^*}$ by the assumption that $x_{t^*} \leq x'_{t^*}$. It follows that for all $i \in N \setminus M$,

$$A_{ir} = \sum_{t=1}^{t^*} A_{irt} \leq \sum_{t=1}^{t^*} A'_{irt} \leq A'_{ir}.$$

Consider the set of agents $M' = \{i \in M \mid d_{ir} > 0\}$. From the definition of r , M' is nonempty. In addition, for all $i \in M \setminus M'$, $A_{ir} = 0 \leq A'_{ir}$. We conclude that for any $i \in N \setminus M'$, $A_{ir} \leq A'_{ir}$. It follows that

$$\sum_{i \in M'} A'_{ir} \leq 1 - \sum_{i \in N \setminus M'} A'_{ir} \leq 1 - \sum_{i \in N \setminus M'} A_{ir} = \sum_{i \in M'} A_{ir},$$

where the equality holds because r is saturated under d . In other words, the overall allocation of resource r to the set of manipulating agents in M' does not grow larger as a result of manipulation, so there must be $i \in M'$ such that $A'_{ir} \leq A_{ir}$. Since $d_{ir} > 0$, this implies by Equation (2) that the utility to agent i does not increase, as required.

Case 2: $x'_{t^} \leq x_{t^*}$ and $isSaturated'$ is true.* Consider an agent $i \in M$ that demands a resource that becomes saturated in round t^* under d' . Let r_i^* be a weighted dominant resource of agent i . We have that $\rho_i = w_{ir_i^*}/d_{ir_i^*}$ and $\rho'_i \leq w_{ir_i^*}/d'_{ir_i^*}$. It follows that $\rho_i \cdot d_{ir_i^*} = w_{ir_i^*}$ and $\rho'_i \cdot d'_{ir_i^*} \leq w_{ir_i^*}$. From Claim 3.3, we know that for all $t < t^*$, $x_t = x'_t$, and our assumption is that $x'_{t^*} \leq x_{t^*}$. Therefore, for every $t \leq t^*$,

$$A'_{ir_i^*t} = x'_t \cdot \rho'_i \cdot d'_{ir_i^*t} \leq x'_t \cdot w_{ir_i^*} \leq x_t \cdot w_{ir_i^*} = x_t \cdot \rho_i \cdot d_{ir_i^*t} = A_{ir_i^*t}.$$

Next we notice that under d' , agent i is not allocated more resources after round t^* . We conclude that

$$A'_{ir_i^*} = \sum_{t=1}^{t^*} A'_{ir_i^*t} \leq \sum_{t=1}^{t^*} A_{ir_i^*t} \leq A_{ir_i^*}.$$

Thus, there exists an agent $i \in M$ who does not receive an increased share of weighted dominant resource under d' , and hence its utility under d' does not increase.

Case 3: $x'_{t^} > x_{t^*}$ and $isSaturated$ is false (and hence $isSaturated'$ is true).* We argue that this case is impossible. Let r be a resource that is saturated under d in round t^* . Since $isSaturated$ is false, it must hold that for all $i \in M$, $d_{ir} = 0$. Thus for all $t \leq t^*$ (using $S_t = S'_t$ by Claim 3.3), $\sum_{i \in S'_t} \rho'_i \cdot d'_{ir} \geq \sum_{i \in S_t} \rho_i \cdot d_{ir}$. Claim 3.3 further states that $x_t = x'_t$ for all $t < t^*$, and therefore

$$s_{rt^*} = 1 - \sum_{t=1}^{t^*-1} \left(x_t \cdot \sum_{i \in S_t} \rho_i \cdot d_{ir} \right) \geq 1 - \sum_{t=1}^{t^*-1} \left(x'_t \cdot \sum_{i \in S'_t} \rho'_i \cdot d'_{ir} \right) = s'_{rt^*}.$$

We conclude that

$$x_{t^*} = \frac{s_{rt^*}}{\sum_{i \in S_{t^*}} \rho_i \cdot d_{ir}} \geq \frac{s'_{rt^*}}{\sum_{i \in S'_{t^*}} \rho'_i \cdot d'_{ir}} \geq x'_{t^*},$$

Which contradicts our assumption that $x'_{t^*} > x_{t^*}$.

Case 4: $x'_{t^*} < x_{t^*}$ and *isSaturated'* is false (and hence *isSaturated* is true). This case is symmetric to case 3, by replacing the roles of d and d' .

It may not be immediately apparent that the above four cases are exhaustive, but note that it is never the case that both *isSaturated* and *isSaturated'* are false. Therefore, the three possible combinations of values for *isSaturated* and *isSaturated'* are covered by cases 1 and 3 when $x'_{t^*} > x_{t^*}$; by cases 2 and 4 when $x'_{t^*} < x_{t^*}$; and by cases 1 and 2 when $x'_{t^*} = x_{t^*}$. This completes the proof of Lemma 3.2. $\square$

While GSP is a strong game-theoretic axiom, an even stronger version has been studied in the literature. Under the stronger notion, it cannot be the case that the value of all manipulators is at least as high, and the value of at least one manipulator is strictly higher under the false reported demands. If one makes the assumption of strictly positive demands, it is easy to verify that EXTENDED DRF also satisfies the stronger notion (the allocation is made in a single round). However, it turns out that when there are multiple rounds, one agent can help another by causing a third agent to drop out early without losing value itself.

Recall that we have assumed for ease of exposition that $d_{ir} > 0$ implies $w_{ir} > 0$. We now briefly explain how to drop this assumption. Observe that if there exist $i \in N$ and $r \in R$ such that $d_{ir} > 0$ but $w_{ir} = 0$, agent i is not entitled to anything according to EF or SI, i.e., any allocation would satisfy these two properties with respect to i . Hence, we can initially remove such agents, and proceed as before with the remaining agents. We then add a second stage where the remaining resources are allocated to the agents that were initially removed, e.g., via unweighted DRF. Using Theorem 3.1, it is easy to verify that this two-stage mechanism is PO, EF, SI, and GSP.

4. LIMITATIONS: SP MECHANISMS AND WELFARE MAXIMIZATION

In Section 3 we have established that the DRF paradigm is robust to perturbations of the model, maintaining its many highly desirable properties. In this section we examine some of the limitations of DRF, and ask whether these limitations can be circumvented.

Throughout this section we assume that an agent's utility is exactly $u_i(A_i) = \max\{y \in \mathbb{R}_+ : \forall r \in R, A_{ir} \geq y \cdot d_{ir}\}$. In particular, if the mechanism is non-wasteful, an agent's utility is the allocated fraction of its dominant resource. In addition, we assume an interpersonal comparison of utilities; i.e., the utility function is not merely a formalism for comparing allocations ordinally. Therefore, given an allocation A , we define its (utilitarian) social welfare as $\sum_{i \in N} u_i(A_i)$. The results of this section, which are negative in nature, hold under this assumption specifically, but similar results hold under perturbations of the utility functions such as linear affine sums of utility. Because the negative results hold even when players are unweighted in terms of their resource contributions, we adopt the model and definitions of Section 2.

Our basic observation is that DRF may provide very low social welfare compared to a welfare-maximizing allocation. For example, consider a setting with m resources. For each resource r there is an agent i_r such that $d_{i_r, r} = 1$, and $d_{i_r, r'} = 0$ for all $r' \in R \setminus \{r\}$. Moreover, there is a large number of agents with normalized demands $d_{ir} = 1$ for all $r \in R$. The optimal allocation would give all of resource r to agent i_r , for a social welfare of m . In contrast, under DRF each agent would receive a $(1/(n - m + 1))$ -fraction of its dominant resource. As n grows larger, the social welfare of $n/(n - m + 1)$ approaches 1. The ratio between these two values is arbitrarily close to m . More formally, letting the *approximation ratio* of a mechanism be the worst-case ratio between the optimal

solution and the mechanism's solution, we say that for every $\delta > 0$, DRF cannot have an approximation ratio better than $m - \delta$ for the social welfare.

In some sense it is not surprising that DRF provides poor guarantees with respect to social welfare. After all, the DRF paradigm focuses on *egalitarian welfare* by equalizing the dominant shares across the agents, rather than aiming for utilitarian welfare. So, can other mechanisms do better?

Let us first examine mechanisms that satisfy SI. It is immediately apparent that the above example provides a similar lower bound. Indeed, any SI mechanism would have to allocate at least $1/n$ of each resource to each of the $n - m$ agents with an all-ones normalized demand vector. Intuitively, an SI mechanism must allocate almost ϵm of the various resources in order to obtain ϵ social welfare. That is, for any $\delta > 0$, any SI mechanism cannot have an approximation ratio better than $m - \delta$ for the social welfare. A similar (though slightly more elaborate) argument works for EF, by replacing the demand vectors of the m agents i_r with $d_{i_r, r} = 1$ (as before) and $d_{i_r, r'} = \epsilon$ for all $r' \in R \setminus \{r\}$ and an arbitrarily small $\epsilon > 0$.

These observations are disappointing, but not entirely unexpected. SI and EF are properties that place significant constraints on allocations. SP is an altogether different matter. *A priori* the constraints imposed by SP seem almost orthogonal to welfare maximization, and indeed in some settings SP mechanisms (even without the use of payments) provide optimal, or nearly optimal, social welfare [Procaccia and Tennenholtz 2009]. Nevertheless, the main result of this section, whose proof is omitted, is a similar lower bound for SP mechanisms.

THEOREM 4.1. *For any $\delta > 0$ there exists a sufficiently large number of agents n such that no SP mechanism can provide an approximation ratio smaller than $m - \delta$ for the social welfare.*

Despite their negative flavor, the results of this section can be seen as vindicating DRF as far as social welfare maximization is concerned. Indeed, asking for just one of the three properties (SI, EF, or SP) seems like a minimal requirement, but already leads to an approximation ratio that is as bad as the one provided by DRF itself (or EXTENDED DRF, for that matter).

5. INDIVISIBLE TASKS

So far, our theoretical analysis treated agents' tasks as divisible: if to run a task an agent needs 2 CPUs and 2 GB RAM, but it is allocated 1 CPU and 1 GB RAM, then it can run half a task. This assumption, which coincides with Leontief preferences, is the driving force behind the results of Section 3 as well as earlier and related results [Ghods et al. 2011; Li and Xue 2011; Friedman et al. 2011].

However, in practice agents' tasks would usually be indivisible, as is indeed the case with the implementation and simulations carried out by Ghods et al. [2011]. In other words, a more realistic domain model is a setting with divisible goods but indivisible tasks: an agent can only derive utility from tasks that are allocated enough resources to run.¹ Hence, informally, an agent's utility function is a step function that increases with every additional instance of its task that it can run. In this section we will *not* require interpersonal comparison of utilities; i.e., as in Section 3, a utility function is used only to induce *ordinal* preferences over allocations. However, for ease of exposition we do assume here that agents are unweighted (or, equivalently, all agents have equal weights).

¹The combination of divisible goods and indivisible tasks may lead to an additional, practical issue, in regard to resource fragmentation, which we discuss in Section 6.

More formally, each agent $i \in N$ now reports a demand bundle $\mathbf{b}_i$, where b_{ir} denotes the fraction of resource $r \in R$ that agent i requires to complete one instance of its task. Given an allocation $\mathbf{A}$, the utility function of agent $i \in N$ is given by

$$u_i(\mathbf{A}_i) = \max\{t \in \mathbb{N} \cup \{0\} : \forall r \in R, A_{ir} \geq t \cdot b_{ir}\},$$

that is, an agent's utility increases linearly with the number of *complete* instantiations of its task that it can run.

Our positive results focus, as before, on non-wasteful mechanisms that do not allocate resources to agents that cannot use them (our negative results do not make this assumption). For a non-wasteful mechanism we can simply denote by $x_i^{\mathbf{A}} = x_i$ the number of bundles allocated to agent $i \in N$ under the allocation $\mathbf{A}$, in which case the utility of agent i is $u_i(\mathbf{A}) = x_i^{\mathbf{A}}$. As a last piece of new notation, we define inequalities between two vectors as follows. We write $\mathbf{v} \geq \mathbf{w}$ if the inequality is satisfied pointwise. Furthermore, $\mathbf{v} > \mathbf{w}$ if $\mathbf{v} \geq \mathbf{w}$ and there is at least one coordinate k such that $v_k > w_k$.

Impossibility results

Observe that the definitions of PO, SI, EF, and SP are identical to the ones given in Section 2. Interestingly, DRF still satisfies SI, EF, and SP under the new “truncated” utilities. However, DRF is no longer PO. For example, say that there are two agents and one resource, and $b_{11} = 1/10$, $b_{21} = 2/5$. DRF would allocate $1/2$ of the resource to each agent, but allocating $3/5$ of the resource to agent 1 and $2/5$ of the resource to agent 2 would be better for agent 1 and equally good for agent 2. An implementation of DRF that allocates tasks sequentially to agents with currently minimum dominant share [Ghodsi et al. 2011] also suffers from several theoretical flaws when analyzed carefully, e.g., it does not satisfy SI.² Note that it is trivial to satisfy SI, EF, and SP without PO, by giving each agent exactly its proportional share of each resource.

The next few results imply that we cannot hope to tweak DRF to achieve all four desirable properties (PO, SI, EF, and SP) under indivisibilities.

THEOREM 5.1. *Under indivisibilities there is no mechanism that satisfies PO, SI, and SP.*

PROOF. Consider a setting with two agents and a single resource. Both agents have bundles $b_{11} = b_{21} = 1/2 + \epsilon$, for a small $\epsilon > 0$. Assume for contradiction that there exists a mechanism that satisfies PO, SI, and SP. By PO, the mechanism allocates one bundle to exactly one of the two agents, without loss of generality agent 1. Now suppose that agent 2 reports $b'_{21} = 1/2$; then any SI allocation must allocate at least $1/2$ of the resource to agent 2, hence the only allocation that is both PO and SI allocates all of the resource to agent 2. This is a contradiction to SP. $\square$

An even more fundamental issue is that PO and EF are trivially incompatible under indivisibilities. For example, when two agents have a task that requires all available resources, the only (two) Pareto optimal allocations give everything to one of the agents, but these allocations are not EF.

Fortunately there is precedent to studying relaxed notions of EF in related settings. A recent example is given by the work of Budish [2011] about the *combinatorial assignment* problem. Budish deals with a related resource allocation setting that models allocations of seats in university courses. Each resource (seats in course) has an integer availability, and agents have preferences over bundles of resources. Budish proposes an approximate version of the notion of *competitive equilibrium from equal*

²As pointed out by Ghodsi et al. [2011], this mechanism does satisfy a coarse version of an approximate EF notion that is discussed below.

incomes [Varian 1974].³ An interesting notion of approximate fairness satisfied by his solution is *envy bounded by a single good*: there exists some good in the allocation of agent $j \in N$ such that $i \in N$ does not envy j when that good is removed.⁴ This notion has a natural equivalent in our setting, which we formalize below.

Definition 5.2. A mechanism is *envy-free up to one bundle* (EF1) if for every vector of reported bundles $\mathbf{b}$ it outputs an allocation $\mathbf{A}$ such that for all $i, j \in N$, $u_i(\mathbf{A}_i) \geq u_i(\mathbf{A}_j - \mathbf{b}_i)$.

Subtraction between vectors is pointwise. In words, a mechanism is EF1 if each agent $i \in N$ does not envy agent $j \in N$ if one instance of the task of agent i is removed from the allocation of j . For non-wasteful mechanisms, this is equivalent to $x_j \mathbf{b}_j \not\geq (x_i + 2) \mathbf{b}_i$ for all $i, j \in N$. Unlike the strict notion of EF, we shall see that the relaxed notion is compatible with PO. However, the following straightforward result implies that the latter notion is incompatible with PO and SP.

THEOREM 5.3. *Under indivisibilities there is no mechanism that satisfies PO, EF1, and SP.*

PROOF. As above, consider a setting with two agents and one resource. Both agents have bundles $b_{11} = b_{21} = 1/3$. The only feasible allocations that are EF1 and PO allocate $2/3$ of the resource to one agent and $1/3$ to the other. Assume without loss of generality that agent 1 receives two bundles. If agent 2 reports $b'_{21} = 1/6$ then the only PO and EF1 allocation gives $2/3$ of the resource to agent 2 and $1/3$ to agent 1, violating SP. $\square$

To summarize, the combination SP+SI+EF1 is trivial, but if we add PO and insist on SP we immediately run into impossibilities with either SI or EF1. Unfortunately, SP seems to preclude reasonable mechanisms when tasks are indivisible. We therefore focus on the other three properties in the context of task indivisibilities, namely PO, SI, and EF1, which, as we shall see, induce a rich axiomatic framework and lead to the design of practical mechanisms.

Sequential Minmax: Achieving PO, SI, and EF1 simultaneously

To design a mechanism that is PO, SI, and EF1 under indivisible tasks, we first introduce a few notations. Given a non-wasteful allocation $\mathbf{A}$, let $\text{MaxDom}(\mathbf{A}) = \max_{i \in N} \max_{r \in R} A_{ir}$ be the maximum dominant share allocated to an agent, and similarly let $\text{MinDom}(\mathbf{A}) = \min_{i \in N} \max_{r \in R} A_{ir}$ be the minimum dominant share. We also let $\mathbf{A} \uparrow i$ be the allocation obtained by starting from the allocation $\mathbf{A}$ and giving agent $i \in N$ another bundle $\mathbf{b}_i$.

An important construction is what we call the *future envy (FE) graph* of an allocation, denoted $G^{\text{FE}}(\mathbf{A})$. Specifically, given an allocation $\mathbf{A}$, the set of vertices of $G^{\text{FE}}(\mathbf{A})$ is the set of agents N , and there is an edge from $j \in N$ to $i \in N$ if one of the following conditions holds:

- (1) Giving one bundle to i would make j envy i but giving one bundle to j would not make i envy j , i.e. $(x_i + 1) \cdot \mathbf{b}_i \geq (x_j + 1) \cdot \mathbf{b}_j$ but $(x_j + 1) \cdot \mathbf{b}_j \not\geq (x_i + 1) \cdot \mathbf{b}_i$.
- (2) Giving one bundle to either i or j would make the other agent envy it, i.e. $(x_i + 1) \cdot \mathbf{b}_i \geq (x_j + 1) \cdot \mathbf{b}_j$ and $(x_j + 1) \cdot \mathbf{b}_j \geq (x_i + 1) \cdot \mathbf{b}_i$, and thus $(x_i + 1) \cdot \mathbf{b}_i = (x_j + 1) \cdot \mathbf{b}_j$. In addition, we also require that $x_j < x_i$.

³Specifically, he allows the market to clear with an error, and also slightly perturbs the initially equal budgets of agents.

⁴Related notions appear in earlier work, for example the approximate notion of equal treatment of equals of Moulin and Stong [2002] and an approximate notion of envy-freeness studied by Lipton et al. [2004].

Note that in the second case the bundles $\mathbf{b}_i$ and $\mathbf{b}_j$ are proportional, and because $x_j < x_i$, $\mathbf{b}_j = c \cdot \mathbf{b}_i$ for some $c > 1$.

LEMMA 5.4. *For any allocation $\mathbf{A}$, if the future envy graph $G^{\text{FE}}(\mathbf{A})$ is nonempty then it has a source vertex (with no incoming edges).*

PROOF. It is sufficient to show that the graph is finite, irreflexive, antisymmetric, and transitive. These properties imply that the graph is acyclic and has a source vertex. The first two properties are trivially satisfied.

For antisymmetry, if there is an edge from $i \in N$ to $j \in N$ and an edge from j to i then it must be the case that

$$(x_i + 1) \cdot \mathbf{b}_i = (x_j + 1) \cdot \mathbf{b}_j.$$

But then we must also have that $x_j > x_i$ and $x_i > x_j$, which is impossible.

For transitivity, note that if there are edges from j to i and k to j then

$$(x_i + 1) \cdot \mathbf{b}_i \geq (x_j + 1) \cdot \mathbf{b}_j \geq (x_k + 1) \cdot \mathbf{b}_k.$$

Assume for contradiction that there is no edge from k to i , then

$$(x_i + 1) \cdot \mathbf{b}_i = (x_j + 1) \cdot \mathbf{b}_j = (x_k + 1) \cdot \mathbf{b}_k$$

and $x_k \geq x_i$. However, then the existing edges imply that $x_k < x_j < x_i$, which contradicts $x_k \geq x_i$. $\square$

We are now ready to present our mechanism, **SEQUENTIALMINMAX**, which is formally given as Algorithm 2. The mechanism sequentially allocates a bundle at each step to a source vertex of the subgraph of the future envy graph that is restricted to agents that minimize the maximum dominant share after allocation.

ALGORITHM 2: SEQUENTIALMINMAX

Data: Bundles $\mathbf{b}$

Result: An allocation $\mathbf{A}$

$k \leftarrow 1; \mathbf{A}^0 \leftarrow \mathbf{0}; T_1 \leftarrow N;$

while $T_k \neq \emptyset$ **do**

$M_k \leftarrow \{i \in T_k \mid \forall j \in T_k, \text{MaxDom}(\mathbf{A}^{k-1} \uparrow i) \leq \text{MaxDom}(\mathbf{A}^{k-1} \uparrow j)\};$

$i \leftarrow \text{source of } G^{\text{FE}}(\mathbf{A}^{k-1}) \text{ restricted to } M_k;$

$\mathbf{A}^k \leftarrow \mathbf{A}^{k-1} \uparrow i;$

$k \leftarrow k + 1;$

$T_k \leftarrow \{i \in N \mid \mathbf{A}^k \uparrow i \text{ is feasible}\};$

end

return $\mathbf{A}^k;$

Our main result of this section is the following theorem.

THEOREM 5.5. *SEQUENTIALMINMAX satisfies PO, SI, and EF1.*

Note that the more intuitive alternative of maximizing the minimum dominant share does not achieve the same properties, nor do variations that consider dominant shares before rather than after allocation.

To establish PO we need to prove that we are allocating a bundle to some agent as long as there exists an agent to which we can allocate; this follows directly from Lemma 5.4. We establish SI and EF1 in the following two lemmas.

LEMMA 5.6. *SEQUENTIALMINMAX satisfies SI.*

PROOF. Let $\mathbf{A}^{\text{SI}}$ be the minimal SI allocation that is obtained by giving each agent $1/n$ of each resource, and then taking back resources that agents cannot use. We show that $\mathbf{A}^{\text{SI}} = \mathbf{A}^k$ for some k during the execution of SEQUENTIALMINMAX. This is sufficient because subsequent allocations can only increase the shares of resources that agents obtain.

We prove, by induction, that $\mathbf{A}^k \leq \mathbf{A}^{\text{SI}}$ (where allocations are treated as vectors and the inequality is pointwise) until it becomes equal to $\mathbf{A}^{\text{SI}}$ for the first time. Note that initially $\mathbf{A}^0 \leq \mathbf{A}^{\text{SI}}$. We will show that until $\mathbf{A}^k = \mathbf{A}^{\text{SI}}$, the mechanism would only allocate a bundle to an agent that is also allocated in $\mathbf{A}^{\text{SI}}$, and therefore after finitely many iterations $\mathbf{A}^k = \mathbf{A}^{\text{SI}}$. Formally, we want to prove that if $\mathbf{A}^k \leq \mathbf{A}^{\text{SI}}$ and $\mathbf{A}^k \neq \mathbf{A}^{\text{SI}}$ then $\mathbf{A}^{k+1} \leq \mathbf{A}^{\text{SI}}$.

Indeed, assume that $\mathbf{A}^k \leq \mathbf{A}^{\text{SI}}$ and $\mathbf{A}^k \neq \mathbf{A}^{\text{SI}}$. Let $i \in N$ be the agent that is given a bundle in iteration $k+1$ of the mechanism. Since $\mathbf{A}^k \neq \mathbf{A}^{\text{SI}}$, there exists an agent j who has a strictly smaller number of bundles in $\mathbf{A}^k$ than in $\mathbf{A}^{\text{SI}}$. Note that $\mathbf{A}^{\text{SI}}$ is a feasible allocation, hence the mechanism does have enough resources to allocate a bundle to j in iteration $k+1$, i.e., $j \in T_{k+1}$. It follows that

$$\text{MaxDom}(\mathbf{A}^{k+1}) = \text{MaxDom}(\mathbf{A}^k \uparrow i) \leq \text{MaxDom}(\mathbf{A}^k \uparrow j) \leq \text{MaxDom}(\mathbf{A}^{\text{SI}}) \leq \frac{1}{n}.$$

Therefore, $\mathbf{A}^{k+1} \leq \mathbf{A}^{\text{SI}}$. $\square$

The proof of the next lemma shows that the mechanism maintains EF1 at every step.

LEMMA 5.7. SEQUENTIALMINMAX satisfies EF1.

PROOF. Note that the allocation $\mathbf{0}$ is EF1. Suppose for contradiction that the allocation returned by the mechanism is not EF1, and let k be the minimum step where $\mathbf{A}^k$ is not EF1. Let $i \in N$ be the agent that is allocated a bundle in step k . Thus there exists an agent $j \in N$ who envies i up to one bundle under the allocation $\mathbf{A}^k$. Let x_i and x_j denote the number of bundles of i and j , respectively, in $\mathbf{A}^k$. It holds that $x_i \cdot \mathbf{b}_i \geq (x_j + 2) \cdot \mathbf{b}_j$, and therefore

$$\forall r \in R, x_i \cdot b_{ir} \geq (x_j + 2) \cdot b_{jr}. \quad (3)$$

It is easy to verify that $x_i \cdot \mathbf{b}_i \geq (x_j + 2) \cdot \mathbf{b}_j$ implies $(x_j + 1) \cdot \mathbf{b}_j \not\geq x_i \cdot \mathbf{b}_i$. Hence, there is an edge from j to i in $G^{\text{FE}}(\mathbf{A}^{k-1})$. In addition, it follows from Equation (3) and the fact that $i \in M_k$ that $j \in T_k$ implies $j \in M_k$, and in that case i would not be a source vertex. Since the mechanism allocated to i , it must be the case that $j \notin T_k$.

Consider the smallest k' such that j envies i in $\mathbf{A}^{k'+1}$ but not in $\mathbf{A}^{k'}$. It holds that $k' < k$. Let x'_i and x'_j be the number of bundles allocated to i and j , respectively, under $\mathbf{A}^{k'}$. Due to the monotonicity of the mechanism with respect to allocations we have that $x'_i < x_i$ and $x'_j \leq x_j$. In addition,

$$x'_i \cdot \mathbf{b}_i \not\geq (x'_j + 1) \cdot \mathbf{b}_j \quad \text{and} \quad (x'_i + 1) \cdot \mathbf{b}_i \geq (x'_j + 1) \cdot \mathbf{b}_j. \quad (4)$$

Let $r^* \in \arg \min_{r \in R: b_{jr} > 0} \{b_{ir}/b_{jr}\}$. From Equation (4) there exists $r \in R$ such that $x'_i \cdot b_{ir} < (x'_j + 1) \cdot b_{jr}$, and by rearranging we get that

$$\frac{b_{ir^*}}{b_{jr^*}} \leq \frac{b_{ir}}{b_{jr}} < \frac{x'_j + 1}{x'_i},$$

that is,

$$x'_i \cdot b_{ir^*} < (x'_j + 1) \cdot b_{jr^*}. \quad (5)$$

However, from Equation (3),

$$x_i \cdot b_{ir^*} \geq (x_j + 2) \cdot b_{jr^*}. \quad (6)$$

Taking the difference between Equations (6) and (5), we get

$$(x_i - x'_i) \cdot b_{ir^*} > (x_j - x'_j + 1) \cdot b_{jr^*} \geq b_{jr^*},$$

and hence for all $r \in R$ such that $b_{jr} > 0$,

$$\frac{b_{ir}}{b_{jr}} \geq \frac{b_{ir^*}}{b_{jr^*}} > \frac{1}{x_i - x'_i},$$

that is, $(x_i - x'_i) \cdot b_{ir} > b_{jr}$, with a weak inequality if $b_{jr} = 0$. Therefore, $(x_i - x'_i) \cdot \mathbf{b}_i > \mathbf{b}_j$.

Note that in iteration $k' + 1$ the available surplus of resources was at least $(x_i - x'_i) \cdot \mathbf{b}_i > \mathbf{b}_j$. Hence, $j \in T_{k'+1}$. Furthermore $i \in M_{k'+1}$ since the mechanism allocates to i . By Equation (4) it must also be the case that $j \in M_{k'+1}$.

For i to be the source of the subgraph of $G^{\text{FE}}(\mathbf{A}^{k'+1})$ restricted to $M_{k'+1}$ it must be the case that $(x'_i + 1) \cdot \mathbf{b}_i = (x'_j + 1) \cdot \mathbf{b}_j$ and that $\mathbf{b}_i \geq \mathbf{b}_j$. But using $\mathbf{b}_i \geq \mathbf{b}_j$ we see that in iteration k there were enough resources to allocate a bundle to j instead of i ; we have reached a contradiction to the assertion that $j \notin T_k$. $\square$

Crucially, as a special case of Theorem 5.5 we find that an allocation that is PO, SI, and EF1 always *exists* for indivisible tasks; this is not *a priori* clear, and in fact, we are not aware of a more direct proof of existence. Once we know such an allocation exists though, we can consider alternative methods of finding it. For example, the computation of such an allocation can be carried out using a mixed integer program (MIP). Using a MIP approach one can also optimize an objective subject to the three constraints; e.g., maximize the minimum dominant share to achieve the most egalitarian allocation among allocations that are PO, SI, and EF1.

Another potential approach is to use the properties of competitive equilibrium from equal incomes (CEEI). In the setting of Section 2 a competitive equilibrium exists and satisfies PO, SI, and EF [Ghods et al. 2011]. While indivisibilities preclude the existence of a CEEI (obviously, as PO and EF are incompatible), conceivably it may be possible to prove the existence of an approximate version that is EF1 using an approach similar to Budish [2011]. Even if such an approximate CEEI exists, its computation can be challenging [Othman et al. 2010].

In any case, as an algorithm for finding PO+SI+EF1 allocations, SEQUENTIALMIN-MAX has several major advantages over other potential approaches. First, its running time is $\mathcal{O}(n^2 m / b^*)$, where $b^* = \min_{i \in N} \max_{r \in R} b_{ir}$; i.e., $1/b^*$ is an upper bound on the number of bundles that can be allocated to an agent. Second, it can be implemented *dynamically*. Indeed, even in a realistic setting where agents change their demands over time, or arrive and depart, we can still carry out the policy of allocating to a source of the (evolving) FE graph in a way that minimizes the maximum dominant share. Specifically, at the beginning of step t we can compute the sets T_k and M_k regardless of the currently present agents and their current demands; a source vertex i of $G^{\text{FE}}(\mathbf{A}^{k-1})$ restricted to M_k will always exist by Lemma 5.4, and hence the allocation $\mathbf{A}^{k-1} \uparrow i$ will always be well defined.

6. DISCUSSION

This paper enhances our understanding of resource allocation in settings with multiple resources. Nevertheless, several key challenges remain unresolved.

We have introduced a model of indivisible tasks and demonstrated that this departure from the previous model requires a rethinking of desirable properties and mechanisms. However, we did make the assumption that there is one pool of divisible

resources. As discussed in passing by Ghodsi et al. [2011], in practice clusters consist of many small machines. This issue of *resource fragmentation* may complicate matters further, and requires careful attention.

In addition, while working on this paper we spent some time studying a setting where each agent has multiple divisible tasks. We focused on a specific utility function: an agent's utility is the sum of dominant shares of its tasks, under the optimal division of its allocated resources between different tasks. Such a utility function can be motivated by routing settings where resources are links and tasks correspond to different paths from a source to a sink. It is possible to show that a competitive equilibrium exists in this setting, and the resulting allocation satisfies PO, SI, and EF. However, we were unable to settle the existence of a mechanism that satisfies PO, SI, EF, and SP, even under restricted utility functions.

Finally, our analysis is restricted to a static setting where demands are revealed and a single allocation is made. Above we touched on the issue of dynamic implementation, and indeed in practice agents may arrive and depart, or change their demands over time (cf. [Parkes 2007] for analogous settings with money). What theoretical guarantees would we look for in such a setting? In particular, what would be the appropriate notion of fairness in a dynamic resource allocation setting? These exciting questions suggest fundamentally new directions for the study of resource allocation.

REFERENCES

- BUDISH, E. 2011. The combinatorial assignment problem: Approximate competitive equilibrium from equal incomes. *Journal of Political Economy*. Forthcoming.
- DOLEV, D., FEITELSON, D. G., HALPERN, J. Y., KUPFERMAN, R., AND LINIAL, N. 2012. No justified complaints: On fair sharing of multiple resources. In *Proceedings of the 3rd Innovations in Theoretical Computer Science Conference (ITCS)*. Forthcoming.
- FRIEDMAN, E. J., GHODSI, A., SHENKER, S., AND STOICA, I. 2011. Strategyproofness, Leontief economies and the Kalai-Smorodinsky solution. Manuscript.
- GHODSI, A., ZAHARIA, M., HINDMAN, B., KONWINSKI, A., SHENKER, S., AND STOICA, I. 2011. Dominant resource fairness: Fair allocation of multiple resource types. In *Proceedings of the 8th USENIX Conference on Networked Systems Design and Implementation (NSDI)*. 24–37.
- GUTMAN, A. AND NISAN, N. 2012. Fair allocation without trade. In *Proceedings of the 11th International Joint Conference on Autonomous Agents and Multi-Agent Systems (AAMAS)*. Forthcoming.
- LI, J. AND XUE, J. 2011. Egalitarian division under Leontief preferences. Manuscript.
- LIPTON, R. J., MARKAKIS, E., MOSSEL, E., AND SABERI, A. 2004. On approximately fair allocations of indivisible goods. In *Proceedings of the 6th ACM Conference on Electronic Commerce (EC)*. 125–131.
- MOULIN, H. 2003. *Fair Division and Collective Welfare*. MIT Press.
- MOULIN, H. AND STONG, R. 2002. Fair queuing and other probabilistic allocation methods. *Mathematics of Operations Research* 27, 1, 1–30.
- NICOLÒ, A. 2004. Efficiency and truthfulness with Leontief preferences. A note on two-agent, two-good economies. *Review of Economic Design* 8, 4, 373–382.
- OTHMAN, A., SANDHOLM, T., AND BUDISH, E. 2010. Finding approximate competitive equilibria: Efficient and fair course allocation. In *Proceedings of the 9th International Joint Conference on Autonomous Agents and Multi-Agent Systems (AAMAS)*. 873–880.
- PARKES, D. C. 2007. Online mechanisms. In *Algorithmic Game Theory*, N. Nisan, T. Roughgarden, E. Tardos, and V. Vazirani, Eds. Cambridge University Press, Chapter 16.
- PAZNER, E. AND SCHMEIDLER, D. 1978. Egalitarian equivalent allocations: A new concept of economic equity. *Quarterly Journal of Economics* 92, 4, 671–687.
- PROCACCIA, A. D. AND TENNENHOLTZ, M. 2009. Approximate mechanism design without money. In *Proceedings of the 10th ACM Conference on Electronic Commerce (EC)*. 177–186.
- ROCHWERGER, B., BREITGAND, D., LEVY, E., GALIS, A., NAGIN, K., LLORENTE, I., MONTERO, R., WOLFSHAL, Y., ELMROTH, E., CÁCERES, J., BEN-YEHUDA, M., EMMERICH, W., AND GALÁN, F. 2009. The RESERVOIR model and architecture for open federated cloud computing. *IBM Journal of Research and Development* 53, 4.
- VARIAN, H. 1974. Equity, envy and efficiency. *Journal of Economic Theory* 9, 63–91.