

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

Additional Key Words and Phrases: Fair division, Resource allocation

ACM Reference Format:

David C. Parkes, Ariel D. Procaccia, and Nisarg Shah, 2013. Beyond Dominant Resource Fairness: Extensions, Limitations, and Indivisibilities. *ACM Trans. Econ. Comp.* V, N, Article A (June 2012), 22 pages.

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

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.

Author's addresses: D. C. Parkes, School of Engineering and Applied Sciences, Harvard University, parkes@eecs.harvard.edu; Ariel 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 1946-6227/2012/06-ARTA \$15.00

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

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 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. While it is reasonable to assume that the resources are divisible, in practice the tasks are often indivisible, that is, 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 in settings without money, 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 positive: 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. The mechanism sequentially allocates bundles to minimize the maximum share of a resource that any agent has after allocation. 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 the conference version of 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 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

We begin with an intuitive exposition based on an example from Ghodsi et al. [2011], and provide a different perspective on this example. Subsequently, we formulate a simple mathematical model, and 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 re-

quires 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:

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

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 [Ghods 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 .

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

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 such as group strategyproofness. We show that DRF can be mod-

ified 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 or the equivalent monetary contribution made by the agent for the resource). 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., $u_i(\mathbf{A}_i) \geq u_i(\langle w_{i1}, \dots, w_{im} \rangle)$. 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. We also note that setting $w_{ir} = 1/n$ for all agents $i \in N$ and resources $r \in R$ recovers the *unweighted case*.

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. Crucially, note that agent 1 ends up receiving all of its dominant resource while agents 2 and 3 receive 1/2 of their dominant resources. Thus, agents might receive unequal shares of their dominant resources at the end of such a multi-round DRF allocation in the presence of zero demands, unlike the single-round DRF allocation in the case of non-zero demands.

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 $\rho_i = w_{ir_i^*}/d_{ir_i^*}$ 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 fraction of resource $r \in R$ left unallocated at the beginning of round t (so $s_{r1} = 1$ for all $r \in R$). The fraction of resource r allocated to agent i in round t is denoted by A_{irt} . We remark that agent weights and demands do not change during the execution of the mechanism, but an agent's dominant resource might change depending on the surpluses of various resources available in each round.

ALGORITHM 1: EXTENDED DRF**Data:** Demands d , weights w **Result:** An allocation A $t \leftarrow 1$; $\forall r, s_{r1} \leftarrow 1$; $S_1 \leftarrow N$;**while** $S_t \neq \emptyset$ **do** $x_t \leftarrow \min_{r \in R} \left(\frac{s_{rt}}{\sum_{i \in S_t} \rho_i \cdot d_{ir}} \right)$; $\forall i \in S_t, r \in R, A_{irt} \leftarrow x_t \cdot \rho_i \cdot d_{ir}$; $\forall i \in N \setminus S_t, r \in R, A_{irt} \leftarrow 0$; $\forall r \in R, s_{r,t+1} \leftarrow s_{rt} - \sum_{i \in S_t} A_{irt}$; $t \leftarrow t + 1$; $S_t \leftarrow \{i \in S_{t-1} : \forall r \in R, d_{ir} > 0 \Rightarrow s_{rt} > 0\}$ /* S_t demand unsaturated resources */**end** $\forall i \in N, r \in R, A_{ir} \leftarrow \sum_{k=1}^{t-1} A_{irk}$

We present our mechanism EXTENDED DRF as Algorithm 1. It deviates from the unweighted mechanism in a couple of ways. Instead of using $x = 1/(\max_r \sum_i d_{ir})$ as in Equation (1), it defines x_t in round t using the ρ_i 's of the agents. Thus, each agent i is now allocated an $(x_t \cdot \rho_i \cdot d_{ir})$ -fraction of resource r in round t , and hence an $(x_t \cdot w_{ir^*})$ -fraction of its dominant resource, instead of being allocated an x -fraction of its dominant resource and other resources in proportion, as in the unweighted case.

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 $x_1 \cdot \rho_i \cdot d_{ir^*} = x_1 \cdot w_{ir^*} \geq w_{ir^*}$ 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 $d_{jr^*} > 0$ then

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

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 $A'_{ir} = x' \cdot d'_{ir}$. If $x' \leq x$ then agent i receives $x' \cdot d'_{ir} \leq x \cdot 1 = x \cdot d_{ir}$ 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 , no agent $i \in M$ has saturated resources under \mathbf{d} or \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 d or d' (and hence the summations include zero terms for these agents). It follows that

$$x_t = \frac{s_{rt}}{\sum_{i \in S_t} \rho_i \cdot d_{ir}} = \frac{s'_{rt}}{\sum_{i \in S'_t} \rho'_i \cdot d'_{ir}} \geq x'_t,$$

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

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

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 d and d' , and resources that are demanded by at least one member of the coalition are not saturated under d nor under 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 \mathbf{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 \mathbf{d}' , and hence its utility under \mathbf{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 \mathbf{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 \mathbf{d} and \mathbf{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. Under 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. To see this, consider the following setting with 5 agents with $w_{ir} = 1/5$ for each $i \in N$ and $r \in R$, and,

$$\mathbf{d}_1 = \langle 0, 1, 0 \rangle, \mathbf{d}_2 = \langle 1, 0, 0 \rangle, \mathbf{d}_3 = \langle 1, 0, 1/4 \rangle, \mathbf{d}_4 = \langle 0, 1, 1 \rangle, \mathbf{d}_5 = \langle 0, 1, 1 \rangle.$$

It can be seen that $x_1 = 1/3$. The surplus at the beginning of the second round is $s_{12} = 1/3$, $s_{22} = 0$, and $s_{32} = 1/4$. Resource 2 is saturated and $S_2 = \{2, 3\}$. In the second round, $x_2 = 1/6$. Therefore, the utilities of agents 1 and 2 in this setting are $u_1 = 1/3$ and $u_2 = 1/3 + 1/6 = 1/2$.

Now, assume that agents 1 and 2 collude, and report demand vectors $d'_1 = \{0, 1, 3/4\}$ and $d'_2 = d_2$. In this case $x'_1 = 1/3$ as before, but the surplus at the beginning of the second round is $s'_{12} = 1/3$, $s'_{22} = 0$, and $s'_{32} = 0$. Hence, both resources 2 and 3 are saturated and $S_2 = \{2\}$. In round 2 agent 2 receives the remaining surplus of resource 1. Hence, under the false demands, $u'_1 = 1/3$ and $u'_2 = 1/3 + 1/3 = 2/3$; agent 1 is not worse off and agent 2 is better off.

We conclude with two remarks. First, note that in each round, the mechanism exhausts a resource that at least one agent demanded. This gives an immediate upper bound of $\min(|N| = n, |R| = m)$ on the number of rounds. Since each round takes $O(n \cdot m)$ time to execute, it follows that the running time of the mechanism is polynomial in the number of agents and the number of resources.

Second, 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 $u_i(\mathbf{A}_i) = \max\{y \in \mathbb{R}_+ : \forall r \in R, A_{ir} \geq y \cdot d_{ir}\}$. 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 \mathbf{A} , we define its (utilitarian) social welfare as $\sum_{i \in N} u_i(\mathbf{A}_i)$. 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 less obstructive to welfare maximization than SI or EF, 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 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.*

PROOF. We first introduce some notation. Given an allocation \mathbf{A} , define the *index set* of \mathbf{A} as $I_{\mathbf{A}} = \{i \in N \mid \min_{r \in R} A_{ir} \leq 1/\sqrt{n}\}$. In words, $I_{\mathbf{A}}$ denotes the set of agents who have at most a $(1/\sqrt{n})$ -fraction of some resource allocated to them. We argue that for any allocation \mathbf{A} ,

$$|I_{\mathbf{A}}| \geq n - \sqrt{n}. \quad (3)$$

Indeed, for any $i \in N \setminus I_{\mathbf{A}}$ and $r \in R$, $A_{ir} > 1/\sqrt{n}$. Hence, there cannot be more than \sqrt{n} such agents, otherwise their total allocation for every resource would be more than 1.

Let $\mathbf{1} = (1, 1, \dots, 1)$ be the all-ones demand vector. Let ϵ_i be the demand vector that has 1 in i^{th} position and ϵ everywhere else. We are interested in settings where the agents in a subset $X \subseteq N$, $|X| = m$, have the demand vectors $\epsilon_1, \epsilon_2, \dots, \epsilon_m$ and the remaining agents have demand vectors $\mathbf{1}$. We show that if the resulting allocation \mathbf{A} is such that for all $i \in X$, $i \in I_{\mathbf{A}}$, then the theorem follows.

CLAIM 4.2. *Suppose that for any sufficiently large $n \in \mathbb{N}$, for any SP mechanism for n agents and m resources, and for any $\epsilon > 0$, there exists $X \subseteq N$ such that $|X| = m$ and when the agents in X report the demand vectors $\epsilon_1, \dots, \epsilon_m$, and the agents in $N \setminus X$ report the demand vector $\mathbf{1}$, the mechanism returns an allocation \mathbf{A} such that $X \subseteq I_{\mathbf{A}}$. Then for any $\delta > 0$, no SP mechanism for m resources can provide an approximation ratio smaller than $m - \delta$.*

PROOF. Given n , an SP mechanism, and ϵ , let $X \subseteq N$ as in the claim's statement. For all $i \in X$, $u_i(\mathbf{A}_i) \leq 1/(\sqrt{n} \cdot \epsilon)$. Clearly it also holds that

$$\sum_{i \in N \setminus X} u_i(\mathbf{A}_i) \leq 1.$$

Hence, the social welfare of an SP mechanism on this instance can be at most $m/(\sqrt{n} \cdot \epsilon) + 1$.

In contrast, maximum social welfare is obtained by dividing the resources among the m agents in X , with everyone getting equal shares of their dominant resources. Under this allocation, every agent $i \in X$ receives a $(1/(1 + (m - 1) \cdot \epsilon))$ -fraction of its dominant resource, which is its utility. Hence, the optimal social welfare is $m/(1 + (m - 1) \cdot \epsilon)$.

It follows that the approximation ratio of the SP mechanism cannot be smaller than

$$\frac{m}{1 + (m-1) \cdot \epsilon} \cdot \frac{1}{\frac{m}{\sqrt{n \cdot \epsilon}} + 1}. \quad (4)$$

To prove the claim it remains to show that for any $\delta > 0$, we can choose n and ϵ such that the expression in Equation (4) is greater than $m - \delta$. Indeed, choose

$$\epsilon < \frac{\delta}{3 \cdot m \cdot (m-1)}$$

and

$$n > \max \left(\left(\frac{3 \cdot m^2 \cdot (m-1)}{\delta} \right)^2, \left(\frac{3 \cdot m^2}{\epsilon \cdot \delta} \right)^2 \right).$$

Then $(m^2 \cdot (m-1))/\sqrt{n} < \delta/3$, $m \cdot (m-1) \cdot \epsilon < \delta/3$, and $m^2/(\sqrt{n} \cdot \epsilon) < \delta/3$. It follows that

$$\delta > \frac{m^2 \cdot (m-1)}{\sqrt{n}} + m \cdot (m-1) \cdot \epsilon + \frac{m^2}{\sqrt{n} \cdot \epsilon},$$

and in particular

$$1 + \frac{\delta}{m} > (1 + (m-1) \cdot \epsilon) \cdot \left(\frac{m}{\sqrt{n} \cdot \epsilon} + 1 \right).$$

Therefore, using the fact that $1/(1 - \frac{\delta}{m}) > 1 + \delta/m$,

$$\frac{m}{m-\delta} > (1 + (m-1) \cdot \epsilon) \cdot \left(\frac{m}{\sqrt{n} \cdot \epsilon} + 1 \right),$$

and finally

$$m - \delta < \frac{m}{1 + (m-1) \cdot \epsilon} \cdot \frac{1}{\frac{m}{\sqrt{n \cdot \epsilon}} + 1},$$

as required. \square (Claim 4.2)

Our goal is therefore to prove the hypothesis made in the statement of Claim 4.2. In the remainder of the proof ϵ is arbitrary and fixed, and we vary the number of agents n . Since we are dealing with an arbitrarily large n , we can assume that m divides n for ease of exposition. We start from a setting where the demand vector is $\mathbf{d}_i = \mathbf{1}$ for all $i \in N$. We partition the set of n agents into m buckets, B_1, B_2, \dots, B_m , each with n/m agents. For consistency we impose the following restriction: an agent in B_i can only change its demand from 1 to ϵ_i . In future references, we omit the unambiguous initial and final demands when we say that an agent changes its demand. We show not only that a set $X \subseteq N$ of m agents exists as in the statement of Claim 4.2, but one such set exists that consists of one agent from each bucket, such that after these m agents change their demands, they all belong to the index set of the resulting allocation.

Call an m -tuple (i_1, \dots, i_m) *diverse* if $i_k \in B_k$ for every $k \in \{1, \dots, m\}$, that is, if it consists of one agent from each bucket. Let $T_{n,m}$ denote the set of all diverse m -tuples. Let $L_{n,m} \subseteq T_{n,m}$ be the set of diverse m -tuples such that when they all change their demands, at least one of them is not in the index set. We want to show that $|T_{n,m} \setminus L_{n,m}| > 0$. We prove the following claim.

CLAIM 4.3. *For any $n, m \in \mathbb{N}$, $|T_{n,m}| = (n/m)^m$ and $|L_{n,m}| \leq m \cdot (n/m)^{m-1} \cdot \sqrt{n}$.*

PROOF. Since $T_{n,m}$ is the set of all diverse m -tuples where one agent is selected from each of m buckets with each bucket containing n/m agents, it is easy to see that $|T_{n,m}| = (n/m)^m$. Recall that $L_{n,m}$ denotes the set of all diverse m -tuples such that when they change their demands, at least one of them is not in the index set. For any $k \in \{1, \dots, m\}$, define $L_{n,m}^k \subseteq L_{n,m}$ to be the set of m -tuples in $L_{n,m}$ such that when they change their demands, the agent from bucket B_k is not in the index set. Clearly $|L_{n,m}| \leq \sum_{k=1}^m |L_{n,m}^k|$.

Fix any $k \in \{1, \dots, m\}$. We upper-bound $|L_{n,m}^k|$ as follows. Take any tuple $t = (i_1, \dots, i_m) \in L_{n,m}^k$. Let $t_{-k} = (i_1, \dots, i_{k-1}, i_{k+1}, \dots, i_m)$. If all agents in t change their demands, agent i_k is not in the index set and gets more than a $(1/\sqrt{n})$ -fraction of every resource. Consider the case when only agents in t_{-k} change their demands. If agent i_k is in the index set, then it receives at most a $(1/\sqrt{n})$ -fraction of some resource. Thus, its utility under the demand vector $\mathbf{1}$ is at most $1/\sqrt{n}$ and it has strict incentive to misreport its demand vector (to ϵ_k) and be outside the index set. This is impossible since the mechanism is SP. Thus, i_k must not be in the index set when the agents in t_{-k} change their demands.

Now, for any fixed t_{-k} , the number of possible agents i_k that are not in the index set when the agents in t_{-k} change their demands is at most \sqrt{n} (from Equation (3)). Thus, for each possible t_{-k} , there are at most \sqrt{n} tuples in $L_{n,m}^k$. Since the number of ways of choosing t_{-k} is at most $(n/m)^{m-1}$, we have that $|L_{n,m}^k| \leq (n/m)^{m-1} \cdot \sqrt{n}$. Hence,

$$|L_{n,m}| \leq \sum_{k=1}^m |L_{n,m}^k| \leq m \cdot \left(\frac{n}{m}\right)^{m-1} \cdot \sqrt{n},$$

as required. \square (Claim 4.3)

Using Claim 4.3, we see that

$$|T_{n,m} \setminus L_{n,m}| \geq \left(\frac{n}{m}\right)^m - m \cdot \left(\frac{n}{m}\right)^{m-1} \cdot \sqrt{n} = \left(\frac{n}{m}\right)^{m-1} \cdot \sqrt{n} \cdot \left(\frac{\sqrt{n}}{m} - m\right).$$

Thus, for $n > m^4$, we have that $|T_{n,m} \setminus L_{n,m}| > 0$, as required. \square (Theorem 4.1)

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. Indeed, this is the case in 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

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.

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).

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 [Ghods 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.

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

Fortunately there is precedent to studying relaxed notions of EF in related settings. A recent example is given by the work of Budish [2011] on 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 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 $x_j \mathbf{b}_j \not\leq (x_i + 2)\mathbf{b}_i$ 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

We remark that Theorems 5.1 and 5.3 can easily be extended to multiple resources by having every agent demand all the resources equally. 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. 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 .

We are now ready to present our mechanism, SEQUENTIALMINMAX, which is formally given as Algorithm 2. The mechanism sequentially allocates one bundle at each step to an agent that minimizes the maximum dominant share after allocation.

Our main result of this section is the following theorem.

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].

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** $M_k \leftarrow \{i \in T_k \mid \forall j \in T_k, \text{MaxDom}(\mathbf{A}^{k-1} \uparrow i) \leq \text{MaxDom}(\mathbf{A}^{k-1} \uparrow j)\};$ $i \leftarrow \text{any agent in } M_k;$ $\mathbf{A}^k \leftarrow \mathbf{A}^{k-1} \uparrow i;$ $T_{k+1} \leftarrow \{i \in T_k \mid \mathbf{A}^k \uparrow i \text{ is feasible}\};$ $k \leftarrow k + 1;$ **end****return** $\mathbf{A}^{k-1};$

THEOREM 5.4. 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 trivially from the mechanism itself since $T_k \neq \emptyset$ implies $M_k \neq \emptyset$. We establish SI and EF1 in the following two lemmas.

LEMMA 5.5. SEQUENTIALMINMAX *satisfies SI.*

PROOF. Let \mathbf{A}^{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}^{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}^{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}^{SI} . Note that \mathbf{A}^{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 next lemma shows that the mechanism maintains EF1 at every step.

LEMMA 5.6. SEQUENTIALMINMAX *satisfies EF1.*

PROOF. For consistency, we say that the mechanism has an iteration 0 at the end of which the allocation is $\mathbf{A}^0 = \mathbf{0}$. Note that the allocation $\mathbf{0}$ is EF1. Suppose for contradiction that the allocation returned by the mechanism at the end is not EF1.

Iteration k_1 : Let k_1 be the first iteration such that \mathbf{A}^{k_1} is not EF1. Thus, $k_1 > 0$. Let $i \in N$ be the agent that is allocated a bundle in iteration k_1 . There exists an agent $j \in N$ who envies i up to one bundle under \mathbf{A}^{k_1} , otherwise the allocation after iteration $k_1 - 1$ would not be EF1 as well, contradicting the fact that k_1 is the first such iteration. Since j envies i , it must hold that i has positive demand for every resource for which j has positive demand. That is,

$$\forall r \in R, b_{jr} > 0 \Rightarrow b_{ir} > 0. \quad (5)$$

Let x_i and x_j denote the number of bundles of i and j , respectively, in \mathbf{A}^{k_1} . It holds that $x_i \cdot \mathbf{b}_i \geq (x_j + 2) \cdot \mathbf{b}_j$, and therefore

$$\forall r \in R, x_i \cdot b_{ir} \geq (x_j + 2) \cdot b_{jr}. \quad (6)$$

For every $r \in R$ such that $b_{ir} > 0$, Equation (6) implies that $x_i \geq (x_j + 2) \cdot (b_{jr}/b_{ir})$. Let $\hat{r} \in \arg \max_{r: b_{ir} > 0} (b_{jr}/b_{ir})$. Then, we have

$$x_i \geq (x_j + 2) \cdot \frac{b_{j\hat{r}}}{b_{i\hat{r}}}. \quad (7)$$

Iteration k_2 : Consider the iteration k_2 in which j was allocated its x_j -th bundle. Since i was allocated a bundle in iteration k_1 , we have $k_2 \neq k_1$. Hence, $0 \leq k_2 < k_1$. Since allocations are monotonic, j has x_j bundles at the end of iteration k for every $k_2 \leq k \leq k_1$. If $k_2 = 0$, then j clearly does not envy i at the end of iteration k_2 . If $k_2 > 0$ and j envies i at the end of iteration k_2 , then j envied i up to one bundle at the end of iteration $k_2 - 1 < k_1$ since j had $x_j - 1$ bundles after iteration $k_2 - 1$. But this contradicts our assumption that k_1 is the first iteration such that \mathbf{A}^{k_1} is not EF1. Hence, it must be the case that j does not envy i at the end of iteration k_2 . However, j envies i (even up to one bundle) at the end of iteration k_1 .

Iteration k_3 : Take the smallest k_3 such that $k_2 < k_3 \leq k_1$ and j envies i at the end of iteration k_3 . Clearly, the mechanism must have allocated one bundle to i in iteration k_3 . Let x'_i be the number of bundles allocated to i at the end of iteration k_3 (hence $x'_i \geq 1$). Note that j must have exactly x_j bundles at the end of iteration k_3 since $k_2 < k_3 \leq k_1$. Further, since j does not envy i at the end of iteration $k_3 - 1$, there exists $r_1 \in R$ such that $(x'_i - 1) \cdot b_{ir_1} < (x_j + 1) \cdot b_{jr_1}$. First, this implies that $b_{jr_1} > 0$. Hence, we also have $b_{ir_1} > 0$ from Equation (5). Thus, we have $x'_i - 1 < (x_j + 1) \cdot b_{jr_1}/b_{ir_1}$. Since $\hat{r} \in \arg \max_{r: b_{ir} > 0} (b_{jr}/b_{ir})$, we have

$$x'_i - 1 < (x_j + 1) \cdot \frac{b_{j\hat{r}}}{b_{i\hat{r}}}. \quad (8)$$

Subtracting Equation (8) from Equation (7), we obtain $x_i - x'_i + 1 \geq b_{j\hat{r}}/b_{i\hat{r}}$. Thus, $x_i - x'_i + 1 \geq b_{jr}/b_{ir}$ for every $r \in R$ such that $b_{ir} > 0$. This implies that $(x_i - x'_i + 1) \cdot b_{ir} \geq b_{jr}$ for every $r \in R$ such that $b_{ir} > 0$. Moreover, if $b_{ir} = 0$ then we have $b_{jr} = 0$ from Equation (5), and $(x_i - x'_i + 1) \cdot b_{ir} = b_{jr} = 0$. Thus, for every $r \in R$, $(x_i - x'_i + 1) \cdot b_{ir} \geq b_{jr}$, i.e., $(x_i - x'_i + 1) \cdot \mathbf{b}_i \geq \mathbf{b}_j$. Note that at least $(x_i - x'_i + 1) \cdot \mathbf{b}_i$ resources were available at the beginning of iteration k_3 since the mechanism allocated $(x_i - x'_i + 1)$ bundles to i from iteration k_3 until k_1 . In particular, the mechanism had enough resources to allocate a bundle to j at the beginning of iteration k_3 , i.e., $j \in T_{k_3}$.

Next, since j envies i at the end of iteration k_3 , we have

$$\forall r \in R, x'_i \cdot b_{ir} \geq (x_j + 1) \cdot b_{jr}. \quad (9)$$

Let r_i^* and r_j^* be dominant resources of agents i and j respectively. Then Equation (9) implies that

$$x'_i \cdot b_{ir_i^*} \geq x'_i \cdot b_{ir_j^*} \geq (x_j + 1) \cdot b_{jr_j^*}. \quad (10)$$

Since the mechanism allocated a bundle to i in iteration k_3 , it must be the case that $i \in M_{k_3}$. Using our conclusion that $j \in T_{k_3}$, we have $x'_i \cdot b_{ir_i^*} \leq (x_j + 1) \cdot b_{jr_j^*}$. Therefore, Equation (10) holds with equalities. It follows that for every $r \in R$ such that $b_{ir} > 0$,

$$\frac{b_{jr}}{b_{ir}} \leq \frac{x'_i}{x_j + 1} = \frac{b_{jr_j^*}}{b_{ir_j^*}}, \quad (11)$$

where the first transition is due to Equation (9). Note that $b_{jr_j^*} > 0$, hence Equation (5) implies $b_{ir_j^*} > 0$. From Equation (11), we can see that $x'_i \cdot b_{ir_j^*} < (x_j + 2) \cdot b_{jr_j^*}$. Thus, j does not envy i up to one bundle at the end of iteration k_3 . This implies that $k_3 < k_1$. Since i was allocated a bundle in iteration k_1 , we also have $x'_i < x_i$.

Iteration k_4 : Consider the iteration k_4 , $k_3 < k_4 \leq k_1$, in which i is allocated its $(x'_i + 1)$ -th bundle. From Equation (6), we know that $x_i \cdot b_{ir_j^*} \geq (x_j + 2) \cdot b_{jr_j^*}$. Subtracting $x'_i \cdot b_{ir_j^*} = (x_j + 1) \cdot b_{jr_j^*}$ (equality holds due to Equation (11)), we get $(x_i - x'_i) \cdot b_{ir_j^*} \geq b_{jr_j^*}$. Again using Equation (11), we conclude that for every $r \in R$ such that $b_{ir} > 0$,

$$x_i - x'_i \geq \frac{b_{jr_j^*}}{b_{ir_j^*}} \geq \frac{b_{jr}}{b_{ir}}.$$

Thus, $(x_i - x'_i) \cdot b_{ir} \geq b_{jr}$ if $b_{ir} > 0$. Also, if $b_{ir} = 0$ then Equation (5) implies $b_{jr} = 0$ and we have $(x_i - x'_i) \cdot b_{ir} = b_{jr} = 0$. Hence, $(x_i - x'_i) \cdot \mathbf{b}_i \geq \mathbf{b}_j$. Note that at least $(x_i - x'_i) \cdot \mathbf{b}_i$ resources were available at the beginning of iteration k_4 since $x_i - x'_i$ bundles were allocated to i from iteration k_4 till k_1 . Hence, there were enough resources to allocate a bundle to j at the beginning of iteration k_4 , i.e., $j \in T_{k_4}$. Equation (10) implies that $(x_j + 1) \cdot b_{jr_j^*} < (x'_i + 1) \cdot b_{ir_i^*} \leq \text{MaxDom}(A^{k_4})$. Note that this is still not a contradiction because allocating to j may not decrease $\text{MaxDom}(A^{k_4})$.

Iteration k_5 : Consider the first iteration k_5 , $k_5 \leq k_4$, such that $\text{MaxDom}(A^{k_5}) > (x_j + 1) \cdot b_{jr_j^*}$. Hence, we have

$$\text{MaxDom}(A^{k_5-1}) \leq (x_j + 1) \cdot b_{jr_j^*}. \quad (12)$$

Let l be the agent that is allocated a bundle in iteration k_5 . Thus, $A^{k_5} = A^{k_5-1} \uparrow l$. As j has at most x_j bundles at the end of iteration $k_5 - 1$, Equation (12) implies $\text{MaxDom}(A^{k_5-1} \uparrow j) \leq (x_j + 1) \cdot b_{jr_j^*} < \text{MaxDom}(A^{k_5}) = \text{MaxDom}(A^{k_5-1} \uparrow l)$. Also, $j \in T_{k_4}$ and $k_5 \leq k_4$, hence $j \in T_{k_5}$. It follows that $l \notin M_{k_5}$, which is a contradiction since the mechanism allocated a bundle to l in iteration k_5 . Thus, our original assumption that there exists an iteration where the allocation is not EF1 is false. \square

Crucially, as a special case of Theorem 5.4 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 [Ghodsi 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, SEQUENTIALMINMAX has several major advantages over other potential approaches. First, since the mechanism allocates a bundle to some agent in each execution of the while loop, it can be seen that the mechanism terminates in time $\mathcal{O}(n^2m/b^*)$, 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 an agent minimizing the maximum dominant share. Specifically, at the beginning of iteration t we can compute the sets T_k and M_k regardless of the currently present agents and their current demands.

6. DISCUSSION

This paper enhances our understanding of resource allocation in settings with multiple resources. First, we assumed that tasks are divisible and showed that EXTENDEDDRF is PO, SI, EF, and GSP even when the agents have initial endowments and may not demand every resource. Second, we observed that no SI or EF mechanism can provide an approximation ratio smaller than m for the social welfare, and proved that the same property holds when one requires SP. Third, under indivisible tasks, we showed that SEQUENTIALMINMAX is PO, SI, and EF1. Nevertheless, several key challenges remain unresolved.

Our model of Section 5 relaxes the assumption that tasks are divisible, but we do 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; Zou et al. 2010] 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? Some answers are given by Kash et al. [2013], but many fundamental questions remain open.

7. ACKNOWLEDGMENTS

We thank Ian Kash for helpful discussions, in particular for suggesting a more elegant proof for Theorem 4.1. Procaccia and Shah are supported in part by NSF under award CCF-1215883, and by a gift from the CMU-MSR Center for Computational Thinking.

REFERENCES

- E. Budish. 2011. The combinatorial assignment problem: Approximate competitive equilibrium from equal incomes. *Journal of Political Economy* 119, 6 (2011), 1061–1103.

- D. Dolev, D. G. Feitelson, J. Y. Halpern, R. Kupferman, and N. Linial. 2012. No justified complaints: On fair sharing of multiple resources. In *Proceedings of the 3rd Innovations in Theoretical Computer Science Conference (ITCS)*. 68–75.
- E. J. Friedman, A. Ghodsi, S. Shenker, and I. Stoica. 2011. Strategyproofness, Leontief Economies and the Kalai-Smorodinsky Solution. Manuscript. (2011).
- A. Ghodsi, M. Zaharia, B. Hindman, A. Konwinski, S. Shenker, and I. Stoica. 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.
- A. Gutman and N. Nisan. 2012. Fair allocation without trade. In *Proceedings of the 11th International Joint Conference on Autonomous Agents and Multi-Agent Systems (AAMAS)*. 719–728.
- I. Kash, A. D. Procaccia, and N. Shah. 2013. No Agent Left Behind: Dynamic Fair Division of Multiple Resources. In *Proceedings of the 12th International Joint Conference on Autonomous Agents and Multi-Agent Systems (AAMAS)*. Forthcoming.
- J. Li and J. Xue. 2011. Egalitarian division under Leontief preferences. *Economic Theory* (2011), 1–26.
- R. J. Lipton, E. Markakis, E. Mossel, and A. Saberi. 2004. On approximately fair allocations of indivisible goods. In *Proceedings of the 6th ACM Conference on Electronic Commerce (EC)*. 125–131.
- H. Moulin. 2003. *Fair Division and Collective Welfare*. MIT Press.
- H. Moulin and R. Stong. 2002. Fair Queuing and Other Probabilistic Allocation Methods. *Mathematics of Operations Research* 27, 1 (2002), 1–30.
- A. Nicolò. 2004. Efficiency and truthfulness with Leontief preferences. A note on two-agent, two-good economies. *Review of Economic Design* 8, 4 (2004), 373–382.
- A. Othman, T. Sandholm, and E. Budish. 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.
- D. C. Parkes. 2007. Online Mechanisms. In *Algorithmic Game Theory*, N. Nisan, T. Roughgarden, É. Tardos, and V. Vazirani (Eds.). Cambridge University Press, Chapter 16.
- E. Pazner and D. Schmeidler. 1978. Egalitarian equivalent allocations: A new concept of economic equity. *Quarterly Journal of Economics* 92, 4 (1978), 671–687.
- A. D. Procaccia and M. Tennenholtz. 2009. Approximate mechanism design without money. In *Proceedings of the 10th ACM Conference on Electronic Commerce (EC)*. 177–186.
- B. Rochwerger, D. Breitgand, E. Levy, A. Galis, K. Nagin, I. Llorente, R. Montero, Y. Wolfsthal, E. Elmroth, J. Cáceres, M. Ben-Yehuda, W. Emmerich, and F. Galán. 2009. The RESERVOIR Model and Architecture for Open Federated Cloud Computing. *IBM Journal of Research and Development* 53, 4 (2009).
- H. Varian. 1974. Equity, envy and efficiency. *Journal of Economic Theory* 9 (1974), 63–91.
- J. Zou, S. Gujar, and D. C. Parkes. 2010. Tolerable Manipulability in Dynamic Assignment without Money. In *Proceedings of the 24th AAAI Conference on Artificial Intelligence (AAAI)*. 947–952.