Mechanism Design for Fair Division*

Richard Cole[†] Vasilis Gkatzelis[‡] Gagan Goel[§]

December 6, 2012

Abstract

We revisit the classic problem of fair division from a mechanism design perspective and provide an elegant truthful mechanism that yields surprisingly good approximation guarantees for the widely used solution of *Proportional Fairness*. This solution, which is closely related to Nash bargaining and the competitive equilibrium, is known to be not implementable in a truthful fashion, which has been its main drawback. To alleviate this issue, we propose a new mechanism, which we call the *Partial Allocation* mechanism, that discards a carefully chosen fraction of the allocated resources in order to incentivize the agents to be truthful in reporting their valuations.

For a multi-dimensional domain with an arbitrary number of agents and items, and for the very large class of homogeneous valuation functions, we prove that our mechanism provides *every agent* with at least a $1/e \approx 0.368$ fraction of her Proportionally Fair valuation. To the best of our knowledge, this is the first result that gives a constant factor approximation to every agent for the Proportionally Fair solution. To complement this result, we show that no truthful mechanism can guarantee more than 0.5 approximation, even for the restricted class of additive linear valuations. We also uncover a connection between the Partial Allocation mechanism and VCG-based mechanism design, which introduces a way to implement interesting truthful mechanisms in settings where monetary payments are not an option.

We also ask whether better approximation ratios are possible in more restricted settings. In particular, motivated by the massive privatization auction in the Czech republic in the early 90s we provide another mechanism for additive linear valuations that works really well when all the items are highly demanded.

1 Introduction

From inheritance and land dispute resolution to treaty negotiations and divorce settlements, the problem of fair division of diverse resources has troubled man since antiquity. Not surprisingly, it has now also found its way into the highly automated, large scale world of computing. As the leading internet companies guide the paradigm shift into cloud computing, more and more services that used to be run on isolated machines are being migrated to shared computing clusters. Moreoever, instead of just human beings bargaining or negotiating, one now also finds programmed strategic agents seeking resources. The goal of the resulting multiagent resource allocation problems [8] is to find solutions that are fair to the agents without introducing unnecessary inefficiencies.

One of the most challenging facets of this change is the need for higher quality incentive design in the form of protocols or mechanisms. As the peer-to-peer revolution has taught us, a proper set of incentives can make or break a system as the number of agents grows [30, Chapter 23]. We therefore revisit this classic fair division problem from a purely mechanism design approach, aiming to create simple and efficient mechanisms that are not susceptible to strategic manipulation by the participating agents; in particular, we want to design *truthful* mechanisms for fair division of heterogenous goods.

*This work was supported in part by NSF grants CCF-0830516 and CCF-1217989. The second author was an intern at Google when part of this work took place.

[†]Courant Institute, New York University, cole@cims.nyu.edu

[‡]Courant Institute, New York University, gkatz@cims.nyu.edu

[§]Google Research, New York, gagangoel@google.com

One distinguishing property of resource allocation protocols in computing is that, more often than not, they need to eschew monetary transfers completely. This is so because, for instance, agents could represent internal teams in an internet company which are competing for resources. This, of course, severely limits what the mechanism designer can achieve since the collection of payments is the most versatile method for designing truthful mechanisms. In light of this, essentially the only tool left for aligning the agents' incentives with the objectives of the system is what Hartline and Roughgarden referred to as "money burning" [19]. That is, the system can choose to intentionally degrade the quality of its services (in our case this will mean discarding resources) in order to influence the preferences of the agents. This degradation of service can often be interpreted as an implicit form of "payment", but since these payments do not correspond to actual trades, they are essentially burned or used for other purposes.

But even before dealing with the fact that the participating agents may behave strategically, one first needs to ask what is the right objective for fairness. This question alone has been the subject of long debates, in both social science and game theory, leading to a very rich literature. At the time of writing this paper, there are *five* academic books (Young, 1994 [36]; Brams and Taylor, 1996 [4]; Robertson and Webb, 1998 [34]; Moulin, 2003 [27]; Barbanel, 2004 [3]) written on the topic of fair division, providing an overview of various proposed solutions for fairness. In this paper we will be focusing on resources that are *divisible*; for such settings, the most attractive solution for efficient and fair allocation is the Proportionally Fair solution (PF). In brief, a PF allocation is a Pareto optimal allocation x^* which compares favorably to any other Pareto optimal allocation x in the sense that, when switching from x to x^*, the aggregate percentage gain in happiness of the agents outweigh the aggregate percentage loss. The notion of PF was first introduced in the seminal work of Frank Kelly [22] in the context of TCP congestion control. Since then it has become the de facto solution for bandwidth sharing in the networking community, and is in fact the *most widely* implemented solution in practice (for instance see, [2] and [21])[1]. The wide adoption of PF as the solution for fairness is not a fluke, but is grounded in the fact that PF is equivalent to the Nash bargaining solution [28], and to the Competitive Equilibria with Equal Incomes (CEEI) [35, 13] for a large class of valuation functions. Both Nash bargaining and the CEEI are well regarded solutions in microeconomics for bargaining and fairness.

A notable property of the PF solution is that it gives a good tradeoff between fairness and efficiency. One extreme notion of fairness is the Rawlsian notion of the egalitarian social welfare that aims to maximize the quality of service of the least satisfied agent irrespective of how much inefficiency this might be causing. On the other extreme, the utilitarian social welfare approach aims to maximize efficiency while disregarding how unsatisfied some agents might become. The PF allocation falls between these two extremes by providing a significant fairness guarantee without neglecting efficiency. As we showed in a recent work [10], for instances with just two players who have affine valuation functions the PF allocation has a social welfare of at least 0.933 times the optimal one.

Unfortunately, the PF allocation has one significant drawback: it cannot be implemented using truthful mechanisms without the use of payments; even for simple instances involving just two agents and two items, it is not difficult to show that no truthful mechanism can obtain a PF solution. This motivates the following natural question: can one design truthful mechanisms that yield a good approximation to the PF solution? Since our goal is to obtain a fair division, we seek a strong notion of approximation in which every agent gets a good approximation of her PF valuation. One of our main results is to give a truthful mechanism which guarantees that every agent will receive at least a $1/e$ fraction of her PF valuation for a very large class of valuation functions. We note that this is one of the very few positive results in multi-dimensional mechanism design without payments. We exhibit this hardness of achieving such truthful approximations by providing an almost matching negative result for a restricted class of valuations.

While a $1/e$ approximation factor is quite surprising for such a general setting, in some circumstances one would prefer to restrict the setting in order to achieve a ratio much closer to 1. Our final result concerns such a scenario, which is motivated by the real-world privatization auctions that took place in Czechoslovakia in the early 90s. At that time, the Czech government sought to privatize the state owned firms dating from

[1] We note that some of the earlier work on Proportional Fairness such as [22] and [23] have 2000+ and 3900+ citations respectively in google scholar, indicating the importance and usage of this solution.

the then recently ended communist era. The government's goal was two-fold — first, to distribute shares of these companies to their citizens in a fair manner, and second, to calculate the market prices of these companies so that the shares could be traded in the open market after the initial allocation. To this end, they ran an auction, as described in [1]. Citizens could choose to participate by buying 1000 vouchers at a cost of 1,000 Czech Crowns, about $35, a fifth of the average monthly salary. Over 90% of those eligible participated. These vouchers were then used to bid for shares in the available 1,491 firms. We believe that the PF allocation provides a very appropriate solution for this example, both to calculate a fair allocation and to compute market prices. Our second mechanism solves the problem of finding allocations very close to the PF allocation in a truthful fashion for such natural scenarios where there is high demand for each resource.

1.1 Our results

In this work we provide some surprising positive results for the problem of multi-dimensional mechanism design without payments. We focus on allocating divisible items and we use the widely accepted solution of proportional fairness as the benchmark regarding the valuation that each participating player deserves. In this setting, we undertake the design of truthful mechanisms that approximate this solution; we consider a strong notion of approximation, requiring that every player receives a good fraction of the valuation that she deserves according to the proportionally fair solution of the instance at hand.

The main contribution of this paper is the *Partial Allocation* mechanism. In Section 3 we analyze this mechanism and we prove that it is truthful and that it guarantees that every player will receive at least a $1/e$ fraction of her proportionally fair valuation. These results hold for the very general class of instances with players having arbitrary homogeneous valuation functions of degree one. This includes a wide range of well studied valuation functions, from additive linear and Leontief, to Constant Elasticity of Substitution and Cobb-Douglas [24]. We later extend these results to homogeneous valuations of any degree. To complement this positive result, we provide a negative result showing that no truthful mechanism can guarantee to every player an allocation with value greater than 0.5 of the value of the PF allocation, even if the mechanism is restricted to the class of additive linear valuations. In proving the truthfulness of the Partial Allocation mechanism we reveal a connection between the amount of resources that the mechanism discards and the payments in VCG mechanisms. As a result, we anticipate that this approach can have a significant impact on other problems in mechanism design without money. We have already verified this to be true for the problem of maximizing social welfare without payments for which a special two-agent version of the Partial Allocation mechanism allowed us to improve upon a setting for which mostly negative results were known [10].

In Section 4 we show that, restricting the set of possible instances to ones involving players with additive linear valuations[2] and items with high prices in the competitive equilibrium from equal incomes[3], will actually allow for the design of even more efficient and useful mechanisms. We present the *Strong Demand Matching* (SDM) mechanism, a truthful mechanism that performs increasingly well as the competitive equilibrium prices increase. More specifically, if p_j^* is the price of item j, then the approximation factor guaranteed by this mechanism is equal to $\min_j \left(p_j^* / \lceil p_j^* \rceil\right)$. It is interesting to note that scenarios such as the privatization auction mentioned above involve a number of bidders much larger than the number of items; as a rule, we expect this to lead to high prices and a very good approximation of the participants' PF utilities.

1.2 Related Work

Our setting is closely related to the large topic of fair division or cake-cutting [36, 4, 34, 27, 3], which has been studied since the 1940's, using the $[0, 1]$ interval as the standard representation of a cake. Each agent's preferences take the form of a valuation function over this interval, and then the valuations of unions of subintervals are additive. Note that the class of homogeneous valuation functions of degree one takes us beyond this standard cake-cutting model. Leontief valuations for example, allow for complementarities in the

[2] Note that our negative results imply that the restriction to additive linear valuations alone would not be enough to allow for significantly better approximation factors.

[3] The prices induced by the market equilibrium when all bidders have a unit of scrip money; also referred to as PF prices.

valuations, and then the valuations of unions of subintervals need not be additive. On the other hand, the additive linear valuations setting that we focus on in Section 4 is equivalent to cake-cutting with piecewise constant valuation functions over the $[0, 1]$ interval. Other common notions of fairness that have been studied in this literature are, proportionality[4], envy-freeness, and equitability [36, 4, 34, 27, 3].

Despite the extensive work on fair resource allocation, truthfulness considerations have not played a major role in this literature. Most results related to truthfulness were weakened by the assumption that each agent would be truthful in reporting her valuations unless this strategy was dominated. Very recent work [7, 26, 37, 25] studies truthful cake cutting variations using the standard notion of truthfulness according to which an agent need not be truthful unless doing so is a dominant strategy. Chen et al. [7] study truthful cake-cutting with agents having piecewise uniform valuations and they provide a polynomial-time mechanism that is truthful, proportional, and envy-free. They also design randomized mechanisms for more general families of valuation functions, while Mossel and Tamuz [26] prove the existence of truthful (in expectation) mechanisms satisfying proportionality in expectation for general valuations. Zivan et al. [37] aim to achieve envy-free Pareto optimal allocations of multiple divisible goods while reducing, but not eliminating, the agents' incentives to lie. The extent to which untruthfulness is reduced by their proposed mechanism is only evaluated empirically and depends critically on their assumption that the resource limitations are soft constraints. Very recent work by Maya and Nisan [25] provides evidence that truthfulness comes at a significant cost in terms of efficiency.

The recent papers of Guo and Conitzer [16] and of Han et al. [18] also consider the truthful allocation of multiple divisible goods; they focus on additive linear valuations and their goal is to maximize the social welfare (or efficiency) after scaling every player's reported valuations so that her total valuation for all items is 1. Guo and Conitzer [16] study two-agent instances, providing both upper and lower bounds for the achievable approximation; Han et al. [18] extend these results and also study the multiple agents setting. For problem instances that may involve an arbitrary number of items both papers provide negative results: no non-trivial approximation factor can be achieved by any truthful mechanism when the number of players is also unbounded. For the two-player case, after Guo and Conitzer [16] studied some classes of dictatorial mechanisms, Han et al. [18] showed that no dictatorial mechanism can guarantee more than the trivial 0.5 factor. Interestingly, we recently showed [10] that combining a special two-player version of the Partial Allocation mechanism with a dictatorial mechanism can actually beat this bound, achieving a 0.622 approximation.

The resource allocation literature has seen a resurgence of work studying fair and efficient allocation for Leontief valuations [15, 12, 32, 17]. These valuations exhibit perfect complements and they are considered to be natural valuation abstractions for computing settings where jobs need resources in fixed ratios. Ghodsi et al. [15] defined the notion of Dominant Resource Fairness (DRF), which is a generalization of the egalitarian social welfare to multiple types of resources. This solution has the advantage that it can be implemented truthfully for this specific class of valuations; as the authors acknowledge, the CEEI solution would be the preferred fair division mechanism in that setting as well, and its main drawback is the fact that it cannot be implemented truthfully. Parkes et al. [32] assessed DRF in terms of the resulting efficiency, showing that it performs poorly. Dolev et al. [12] proposed an alternate fairness criterion called Bottleneck Based Fairness, which Gutman and Nisan [17] subsequently showed is satisfied by the proportionally fair allocation. Gutman and Nisan [17] also posed the study of incentives related to this latter notion as an interesting open problem. Our results could potentially have significant impact on this line of work as we are providing a truthful way to approximate a solution which is recognized as a good benchmark. It would also be interesting to study the extent to which the Partial Allocation mechanism can outperform the existing ones in terms of efficiency.

Most of the papers mentioned above contribute to our understanding of the trade-offs between either truthfulness and fairness, or truthfulness and social welfare. Another direction that has been actively pursued is to understand and quantify the interplay between fairness and social welfare. Caragiannis et al. [6] measured the deterioration of the social welfare caused due to different fairness restrictions, the price of fairness. More recently, Cohler et al. [9] designed algorithms for computing allocations that (approximately)

[4]It is worth distinguishing the notion of PF from that of proportionality by noting that the latter is a much weaker notion, directly implied by the former.

maximize social welfare while satisfying envy-freeness.

Our results fit into the general agenda of approximate mechanism design without money, explicitly initiated by Procaccia et al. [33]. More interestingly, the underlying connection with VCG payments proposes a framework for designing truthful mechanisms without money and we anticipate that this might have a significant impact on this literature.

2 Preliminaries

Let M denote the set of m items and N the set of n bidders. Each item is divisible, meaning that it can be divided into arbitrarily small pieces and then allocated to different bidders. An allocation x of these items to the bidders defines the fraction x_{ij} of each item j that each bidder i will be receiving; let $\mathcal{F} = \{x \mid x_{ij} \geq 0 \text{ and } \sum_i x_{ij} \leq 1\}$ denote the set of feasible allocations. Each bidder is assigned a weight $b_i \geq 1$ which allows for interpersonal comparison of valuations, and can serve as priority in computing applications, as clout in bargaining applications, or as a budget for the market equilibrium interpretation of our results. We assume that b_i is defined by the mechanism as it cannot be truthfully elicited from the bidders. The preferences of each bidder $i \in N$ take the form of a valuation function $v_i(\cdot)$, that assigns nonnegative values to every allocation in $\mathcal{F}$. We assume that every player's valuation for a given allocation x only depends on the bundle of items that she will be receiving.

We will present our results assuming that the valuation functions are homogeneous of degree one, i.e. player i's valuation for an allocation $x' = f \cdot x$ satisfies $v_i(x') = f \cdot v_i(x)$, for any scalar $f > 0$. We later discuss how to extend these results to general homogeneous valuations of degree d for which $v_i(x') = f^d \cdot v_i(x)$. A couple of interesting examples of homogeneous valuations functions of degree one are additive linear valuations and Leontief valuations; according to the former, every player has a valuation v_{ij} for each item j and $v_i(x) = \sum_j x_{ij} v_{ij}$, and according to the latter, each player i's type corresponds to a set of values a_{ij}, one for each item, and $v_i(x) = \min_j \{x_{ij}/a_{ij}\}$. (i.e. player i desires the items in the ratio $a_{i1} : a_{i2} : \ldots : a_{im}$.)

An allocation $x^* \in \mathcal{F}$ is *Proportionally Fair* (PF) if, for any other allocation $x' \in \mathcal{F}$ the (weighted) aggregate proportional change to the valuations after replacing x^* with x' is not positive, i.e.:

$$\sum_{i \in N} \frac{b_i[v_i(x') - v_i(x^*)]}{v_i(x^*)} \leq 0. \tag{1}$$

This allocation rule is a strong refinement of Pareto efficiency, since Pareto efficiency only guarantees that if some player's proportional change is strictly positive, then there must be some player whose proportional change is negative. The Proportionally Fair solution can also be defined as an allocation $x \in \mathcal{F}$ that maximizes $\prod_i [v_i(x)]^{b_i}$, or equivalently (after taking a logarithm), that maximizes $\sum_i b_i \log v_i(x)$; we will refer to these two equivalent objectives as the PF objectives. Note that, although the PF allocation need not be unique for a given instance, it does provide unique achieved bidder valuations [14].

We note that the PF solution is equivalent to the Nash bargaining solution. John Nash in his seminal paper [28] considered an axiomatic approach to bargaining and gave four axioms that any bargaining solution must satisfy. He showed that these four axioms rendered a unique solution which is captured by a convex program; this convex program is equivalent to the one defined above for the PF solution. Another well-studied allocation rule which is equivalent to the PF allocation is the *Competitive Equilibrium*. Eisenberg [13] showed that if all agents have valuation functions that are quasi-concave and homogeneous of degree 1, then the competitive equilibrium is also captured by the same convex program as the one for the PF solution. The *Competitive Equilibrium with Equal Incomes* (CEEI) has been proposed as an allocation rule for fairness in microeconomics [35, 5, 31].

Given a valuation function reported from each bidder, we want to design mechanisms that output an allocation of items to bidders. We restrict ourselves to truthful mechanisms, i.e. mechanisms such that any false report from a bidder will never return her a more valuable allocation. Since proportional fairness cannot be implemented via truthful mechanisms, we will measure the performance of our mechanisms based on the extent to which they approximate this benchmark. More specifically, the approximation factor, or competitive factor of a mechanism will correspond to the minimum value of $\rho(\mathcal{I})$ across all relevant instances $\mathcal{I}$, where

$$\rho(\mathcal{I}) = \min_{i \in N} \left\{ \frac{v_i(x)}{v_i(x^*)} \right\},$$

and x, x^* are the allocation generated by the mechanism for instance $\mathcal{I}$ and the PF allocation of $\mathcal{I}$ respectively.

3 Partial Allocation Mechanism

In this section, we define the *Partial Allocation* (PA) mechanism as a novel way to allocate divisible items to bidders with homogeneous valuation functions of degree one. We subsequently prove that this non-dictatorial mechanism, not only achieves truthfulness, but also guarantees that every bidder will be receiving at least a $1/e$ fraction of the valuation that she deserves, as dictated by the PF solution. This mechanism depends on a subroutine that computes the PF allocation for the problem instance at hand; we therefore later study the running time of this subroutine, as well as the robustness of our results in case this subroutine returns only approximate solutions.

The PA mechanism elicits the valuation function $v_i(\cdot)$ from each player i and it computes the PF allocation x^* considering all the players' valuations. The final allocation x that the mechanism outputs, allocates to each player i only a fraction f_i of her PF bundle x_i^*, i.e. for every item j of which the PF allocation assigned to her a portion of size x_{ij}^*, the PA mechanism instead assigns to her a portion of size $f_i \cdot x_{ij}^*$, where $f_i \in [0,1]$ depends on the extent to which the presense of player i inconveniences the other players; the value of f_i may therefore vary across different players. The following steps give a more precise description of the mechanism.

1. Compute the PF allocation based on the reported bids.
2. For each player i, remove this player and compute the PF allocation x_{-i}^* that would arise in her absence.
3. Allocate to each player i a fraction f_i of everything that she receives according to x^*, where

$$f_i = \left(\frac{\prod_{i' \neq i} [v_{i'}(x^*)]^{b_{i'}}}{\prod_{i' \neq i} [v_{i'}(x_{-i}^*)]^{b_{i'}}} \right)^{1/b_i}. \tag{2}$$

Lemma 1. *The allocation x produced by the PA mechanism is feasible.*

Proof. Since the PF allocation x^* is feasible, to verify that the allocation produced by the PA mechanism is also feasible, it suffices to show is that $f_i \in [0,1]$ for every bidder i. The fact that $f_i \geq 0$ is clear since both the numerator and the denominator are non-negative. To show that $f_i \leq 1$, note that

$$x_{-i}^* = \arg\max_{x' \in \mathcal{F}} \left\{ \prod_{i' \neq i} v_{i'}(x') \right\}.$$

Since x^* remains a feasible allocation after removing bidder i (we may discard all her items), this implies

$$\prod_{i' \neq i} v_{i'}(x^*) \leq \prod_{i' \neq i} v_{i'}(x_{-i}^*)$$

. □

3.1 Truthfulness

We now show that, despite the fact that this mechanism is not dictatorial and does not use monetary payments, it is still in the best interest of every player to report her true valuation function, irrespective of what the other players do.

Theorem 1. *The PA mechanism is truthful.*

Proof. In order to prove this theorem, we approach the PA mechanism from the perspective of some arbitrary player i. Let $\bar{v}_{i'}(\cdot)$ denote the valuation function that each player $i' \neq i$ reports to the PA mechanism. We assume that the valuation functions reported by these players may differ from their true ones, $v_{i'}(\cdot)$. Player i is faced with the options of, either reporting her true valuation function $v_i(\cdot)$, or reporting some false valuation function $\bar{v}_i(\cdot)$. After every player has reported some valuation function, the PA mechanism computes the PF allocation with respect to these valuation functions; let $\bar{x}_t$ denote the PF allocation that arises if player i reports the truth and $\bar{x}_f$ otherwise. Finally, player i receives a fraction of what the computed PF allocation assigned to her, and how big or small this fraction will be depends on the computed PF allocation. Let f_i' denote the fraction of her allocation that player i will receive if $\bar{x}_t$ is the computed PF allocation and f_i'' otherwise. Since the players have homogeneous valuation functions of degree one, what we then need to show is that $f_i' v_i(\bar{x}_t) \geq f_i'' v_i(\bar{x}_f)$, or equivalently that

$$[f_i' v_i(\bar{x}_t)]^{b_i} \;\geq\; [f_i'' v_i(\bar{x}_f)]^{b_i}.$$

Note that the denominators of both fractions f_i' and f_i'', as given by Equation (2), will be the same since they are independent of the valuation function reported by player i. Our problem therefore reduces to proving that

$$[v_i(\bar{x}_t)]^{b_i} \cdot \prod_{i' \neq i} [\bar{v}_{i'}(\bar{x}_t)]^{b_{i'}} \;\geq\; [v_i(\bar{x}_f)]^{b_i} \cdot \prod_{i' \neq i} [\bar{v}_{i'}(\bar{x}_f)]^{b_{i'}}. \tag{3}$$

To verify that this inequality holds we use the fact that the PF allocation is the one that maximizes the product of the corresponding reported valuations. This means that

$$\bar{x}_t = \arg\max_{\bar{x} \in \mathcal{F}} \left\{ [v_i(\bar{x})]^{b_i} \cdot \prod_{i' \neq i} [\bar{v}_{i'}(\bar{x})]^{b_{i'}} \right\},$$

and since $\bar{x}_f \in \mathcal{F}$, this implies that Inequality (3) holds, and therefore reporting her true valuation function is a dominant strategy for every player i. □

The arguments used in the proof of Theorem 1 imply that, given the valuation functions reported by all the other players $i' \neq i$, player i can effectively choose any bundle that she wishes, but for each bundle the mechanism defines what fraction player i can keep. One can therefore think of the fraction of the bundle thrown away as a form of non-monetary "payment" that the player must suffer in exchange for that bundle, with different bundles having different payments. The fact that the PA mechanism is truthful implies that these payments, in the form of fractions, make the bundle allocated to her by allocation x^* the most desirable one. A closer look reveals a connection with the well known VCG-mechanism design approach: The valuation of player i for the output of the PA mechanism is $v_i(x) = f_i \cdot v_i(x^*)$, or

$$v_i(x) \;=\; \left(\frac{\prod_{i' \neq i} [v_{i'}(x^*)]^{b_{i'}}}{\prod_{i' \neq i} [v_{i'}(x^*_{-i})]^{b_{i'}}} \right)^{1/b_i} \cdot v_i(x^*). \tag{4}$$

The connection becomes apparent after considering the surrogate valuation $u_i(\cdot) = b_i \log v_i(\cdot)$ for each player i. After taking a logarithm on both sides of Equation (4) and then multiplying them by b_i, using surrogate valuation function $u_i(\cdot)$, yields

$$u_i(x) = u_i(x^*) - \left(\sum_{i' \neq i} u_{i'}(x^*_{-i}) \;-\; \sum_{i' \neq i} u_{i'}(x^*) \right).$$

This equation shows that the surrogate valuation of a player for the output of the PA mechanism equals her surrogate valuation for the PF allocation minus a payment which corresponds to the externalities with respect again to the surrogate valuations. The connection is complete if one notes that the PF allocation is actually maximizing the social welfare with respect to these surrogate valuations. Therefore, the fraction that is removed from each player's PF allocation can actually be interpreted as a VCG payment with respect to the surrogate functions.

3.2 Approximation

Before studying the approximation factor of the PA mechanism, we first show a lemma (the proof is deferred to the appendix) which will be useful for proving Theorem 2.

Lemma 2. *For any set of pairs (δ_i, β_i) with $\beta_i \geq 1$ and $\sum_i \beta_i \cdot \delta_i \leq b$ the following holds (where $B = \sum_i \beta_i$)*

$$\prod_i (1+\delta_i)^{\beta_i} \leq \left(1+\frac{b}{B}\right)^B .$$

Using this lemma we can now prove tight bounds for the approximation factor of the Partial Allocation mechanism. As we show in this proof, the approximation factor depends directly on the relative weights of the players. For simplicity in expressing the approximation factor, let b_{min} denote the smallest value of b_i across all bidders of an instance and let $\bar{B} = \left(\sum_{i\in N} b_i\right) - b_{min}$ be the sum of the b_i values of all the other bidders. Finally, let $\psi = \bar{B}/b_{min}$ denote the ratio of these two values.

Theorem 2. *The approximation factor of the PA mechanism for problem instances involving n players is exactly*

$$\left(1+\frac{1}{\psi}\right)^{-\psi} .$$

Proof. The PA mechanism allocates to each player i a fraction f_i of her PF allocation, and for the class of homogeneous valuation functions of degree one this means that the final valuation of player i will be $v_i(x) = f_i \cdot v_i(x^*)$. The approximation factor guaranteed by the mechanism is therefore equal to $\min_i\{f_i\}$. Without loss of generality, let player i be the one with the minimum value of f_i. In the PF allocation x^*_{-i} that the PA mechanism computes after removing player i, every other player i' experiences a value of $v_{i'}(x^*_{-i})$. Let $d_{i'}$ denote the proportional change between the valuation of player i' for allocation x^* and allocation x^*_{-i}, i.e.

$$v_{i'}(x^*_{-i}) = (1+d_{i'})v_{i'}(x^*).$$

Substituting for $v_{i'}(x^*_{-i})$ in Equation (2) yields:

$$f_i = \left(\frac{1}{\prod_{i'\neq i}(1+d_{i'})^{b_{i'}}}\right)^{1/b_i} . \tag{5}$$

Since x^* is a PF allocation, Inequality (1) implies that

$$\begin{aligned}
\sum_{i'\in N} \frac{b_{i'}[v_{i'}(x^*_{-i}) - v_{i'}(x^*)]}{v_{i'}(x^*)} &\leq 0 \iff \\
\sum_{i'\neq i} b_{i'}d_{i'} + \frac{b_i[v_i(x^*_{-i}) - v_i(x^*)]}{v_i(x^*)} &\leq 0 \iff \\
\sum_{i'\neq i} b_{i'}d_{i'} &\leq b_i. \qquad (6)
\end{aligned}$$

The last equivalence holds due to the fact that $v_i(x^*_{-i}) = 0$, since allocation x^*_{-i} clearly assigns nothing to player i.

Let $B_{-i} = \sum_{i' \neq i} b_{i'}$; using Inequality (6) and Lemma 2 (on substituting b_i for b, d'_i for δ_i, and B_{-i} for B), it follows from Equation (5) that

$$f_i \geq \left(1 + \frac{b_i}{B_{-i}}\right)^{-\frac{B_{-i}}{b_i}}. \tag{7}$$

To verify that this bound is tight, consider any instance of n players and one item. The PF solution dictates that each player should be receiving a fraction of the item proportional to the player's b_i value. The removal of a player i therefore leads to a proportional increase of exactly b_i/B_{-i} for all the other players. The PA mechanism therefore assigns to every player i a fraction of her PF allocation which is equal to the lower bound of Inequality (7). The player with the smallest b_i value receives the smallest fraction. □

The approximation factor of Theorem 2 implies that $f_i \geq 1/2$ for instances with two players having equal b_i values, and $f_i \geq 1/e$ even when ψ goes to infinity; we therefore get the following corollary.

Corollary 1. *For any problem instance the Partial Allocation mechanism yields an allocation x such that for every participating player i*

$$v_i(x) \;\geq\; \frac{1}{e} \cdot v_i(x^*).$$

To complement this approximation factor, we now provide a negative result showing that, even for the special case of additive linear valuations, no truthful mechanism can guarantee an approximation factor better than $\frac{n+1}{2n}$.

Theorem 3. *There is no truthful mechanism that can guarantee an approximation factor greater than $\frac{n+1}{2n}$ for all n-player problem instances, even if the valuations are restricted to being additive linear.*

Proof. For an arbitrary real value of $n > 1$, let $\rho = \frac{n+1}{2n}$, and assume that Q is a truthful resource allocation mechanism that guarantees a $(\rho + \epsilon)$ approximation for all n-player problem instances, where ϵ is a positive constant. This mechanism receives as input the bidders' valuations and it returns a valid (fractional) allocation of the items. We will define $n + 1$ different input instances for this mechanism, each of which will consist of n bidders and $m = (k + 1)n$ items, where $k > \frac{2}{\epsilon}$ will take very large values. In order to prove the theorem, we will then show that Q cannot simultaneously achieve this approximation guarantee for all these instances, leading to a contradiction. For simplicity we will refer to each bidder with a number from 1 to n, to each item with a number from 1 to $(k+1)n$, and to each problem instance with a number from 1 to $n+1$.

We start by defining the first n problem instances. For $i \leq n$, let problem instance i be as follows: Every bidder $i' \neq i$ has a valuation of $kn + 1$ for item i' and a valuation of 1 for every other item; bidder i has a valuation of 1 for all items. In other words, all bidders except bidder i have a strong preference for just one item, which is different for each one of them. The PF allocation for such additive linear valuations dictates that every bidder $i' \neq i$ is allocated only item i', while bidder i is allocated all the remaining $kn + 1$ items. Since Q achieves a $\rho + \epsilon$ approximation for this instance, then it needs to provide bidder i with an allocation which the bidder values at least at $(\rho + \epsilon)(kn + 1)$. In order to achieve this, mechanism Q can assign to this bidder fractions of the set M_{-i} of the $n - 1$ items that the PF solution allocates to the other bidders as well as fractions of the set M_i of the $kn + 1$ items that the PF allocation allocates to bidder i. Even if all of the $n - 1$ items of M_{-i} were fully allocated to bidder i, the mechanism would still need to assign to this bidder an allocation of value at least $(\rho + \epsilon)(kn + 1) - (n - 1)$ using items from M_i. Since $k > \frac{2}{\epsilon}$, $n - 1 < \frac{\epsilon}{2}(kn + 1)$, and therefore mechanism Q will need to allocate to bidder i a fractional assignment of items in M_i that the bidder values at least at $\left(\rho + \frac{\epsilon}{2}\right)(kn + 1)$. This implies that there must exist at least one item in M_i of which bidder i is allocated a fraction of size at least $\left(\rho + \frac{\epsilon}{2}\right)$. Since all the items in M_i are identical and the numbering of the items is arbitrary, we can, without loss of generality, assume that this item is item i. We have therefore shown that, for every instance $i \leq n$ mechanism Q will have to assign to bidder i at least $\left(\rho + \frac{\epsilon}{2}\right)$ of item i, and an allocation of items in M_i that guarantees her a valuation of at least $\left(\rho + \frac{\epsilon}{2}\right)(kn + 1)$.

We now define problem instance $n + 1$, in which every bidder i has a valuation of $kn + 1$ for item i and a valuation of 1 for all other items. The PF solution for this instance would allocate to each bidder i all of item i, as well as k items from the set $\{n+1, ..., (k+1)n\}$ (or more generally, fractions of these items that add up to k). Clearly, every bidder i can unilaterally misreport her valuation leading to problem instance i instead of this instance so, in order to maintain truthfulness, mechanism Q will have to provide every bidder i of problem instance $n+1$ with at least the value that such a deviation would provide her with. One can quickly verify that, even if mechanism Q when faced with problem instance i provided bidder i with no more than a $\left(\rho + \frac{\epsilon}{2}\right)$ fraction of item i, still such a deviation would provide bidder i with a valuation of at least

$$\left(\rho + \frac{\epsilon}{2}\right)(kn+1) \;+\; \left(\rho + \frac{\epsilon}{2}\right)kn \;\geq\; \left(\rho + \frac{\epsilon}{2}\right)2kn.$$

The first term of the left hand side comes from the fraction of item i that the bidder receives and the second term comes from the average fraction of the remaining items. If we substitute $\rho = \frac{n+1}{2n}$, we get that the truthfulness of Q implies that every bidder i of problem instance $n+1$ will have to receive an allocation of value at least

$$\left(\frac{n+1}{2n} + \frac{\epsilon}{2}\right)2kn \;=\; kn + k + \epsilon kn.$$

For any given constant value of ϵ though, since $k > \frac{2}{\epsilon}$ and $n > 1$, every bidder will need to be assigned an allocation that she values at more than $kn+k+2$, which is greater than the valuation of $kn+k+1$ that the player receives in the PF solution. This is obviously a contradiction since the PF solution is Pareto efficient and there cannot exist any other allocation for which all bidders receive a strictly greater valuation. □

Theorem 3 implies that, even if all the players have equal b_i values, no truthful mechanism can guarantee greater than 3/4 approximation even for instances with just two bidders, and this bound drops further as the number of bidders increases, finally converging to 1/2. To complement the statement of Corollary 1, we therefore get the following corollary.

Corollary 2. *No truthful mechanism can guarantee that, for every problem instance, it will yield an allocation x such that for every participating player i*

$$v_i(x) \;>\; \frac{1}{2} \cdot v_i(x^*).$$

3.3 Running Time and Robustness

The PA mechanism has reduced the problem of truthfully implementing a constant factor approximation of the PF allocation to computing exact PF allocations for several different problem instances, as this is the only subroutine that the mechanism calls. If the valuation functions of the players are affine, then there is a polynomial time algorithm to compute the exact PF allocation [11, 20].

We now show that, even if the PF solution can be only approximately computed in polynomial time, our truthfulness and approximation related statements are robust with respect to such approximations (all the proofs of this subsection are deferred to the appendix). More specifically, we assume that the PA mechanism uses a polynomial time algorithm that computes a feasible allocation $\widetilde{x}$ instead of x^* such that

$$\left[\prod_i [v_i(\widetilde{x})]^{b_i}\right]^{1/B} \quad\geq\quad \left[(1-\epsilon)\prod_i [v_i(x^*)]^{b_i}\right]^{1/B}, \quad \text{where } B = \sum_{i=1}^{n} b_i.$$

Using this algorithm, the PA mechanism can be adapted as follows:

1. Compute the approximate PF allocation $\widetilde{x}$ based on the reported bids.
2. For each player i, remove this player and compute the approximate PF allocation $\widetilde{x}_{-i}$ that would arise in her absence.

3. Allocate to each player i a fraction $\widetilde{f}_i$ of everything that she receives according to $\widetilde{x}$ where

$$\widetilde{f}_i = \min\left\{1\ ,\ \left(\frac{\prod_{i'\neq i}[v_{i'}(\widetilde{x})]^{b_{i'}}}{\prod_{i'\neq i}[v_{i'}(\widetilde{x}_{-i})]^{b_{i'}}}\right)^{1/b_i}\right\}. \tag{8}$$

For this adapted version of the PA mechanism to remain feasible, we need to make sure that $\widetilde{f}_i$ remains less or equal to 1. Even if, for some reason, the allocation $\widetilde{x}_{-i}$ computed by the approximation algorithm does not satisfy this property, the adapted mechanism will then choose $\widetilde{f}_i = 1$ instead.

We start by showing two lemmas verifying that this adapted version of the PA mechanism is robust both with respect to the approximation factor it guarantees and with respect to the truthfulness guarantee.

Lemma 3. *The approximation factor of the adapted PA mechanism for problem instances involving n players is at least*

$$(1-\epsilon)\left(1+\frac{1}{\psi}\right)^{-\psi}.$$

Lemma 4. *If a player misreports her preferences to the adapted PA mechanism, she may increase her valuation by at most a factor $(1-\epsilon)^{-2}$.*

Finally, we show that if the valuation functions are, for example, concave and homogeneous of degree one, then a feasible approximate PF allocation can indeed be computed in polynomial time.

Lemma 5. *For concave homogeneous utility functions of degree one, there exists an algorithm that computes a feasible allocation $\widetilde{x}$ in time polynomial in $\log 1/\epsilon$ and the problem size, such that*

$$\prod_i [v_i(\widetilde{x})]^{b_i} \ \geq\ (1-\epsilon)\prod_i [v_i(x^*)]^{b_i}.$$

3.4 Extension to General Homogeneous Valuations

We can actually extend most of the results that we have shown for homogeneous valuation functions of degree one to any valuation function that can be expressed as $v_i(f\cdot x) = g_i(f)\cdot v_i(x)$, where $g_i(\cdot)$ is some increasing invertible function; for homogeneous valuation functions of degree d, this function is $g_i(f) = f^d$. If this function is known for each bidder, we can then adapt the PA mechanism as follows: instead of allocating to bidder i a fraction f_i of her allocation according to x^* as defined in Equation (2), we instead allocate to this bidder a fraction $g_i^{-1}(f_i)$, where $g_i^{-1}(\cdot)$ is the inverse function of $g_i(\cdot)$. If, for example, some bidder has a homogeneous valuation function of degree d, then allocating her a fraction $f_i^{1/d}$ of her PF allocation has the desired effect and both truthfulness and the same approximation factor guarantees still hold. The idea behind this transformation is that all that we need in order to achieve truthfulness and the approximation factor is to be able to discard some fraction of a bidder's allocation knowing exactly what fraction of her valuation this will correspond to.

4 Additive Linear Valuations

The results of the previous section show that one can guarantee a good constant factor approximation for a very large class of valuations of the agents. In this section we provide an improved mechanism for a very natural and well motivated class of instances in which the players' valuations are additive linear. We make an assumption that every item is well demanded in the sense that the PF price (or equivalently competitive equilibrium price) of every item is large when the budget of every player is fixed to one unit of scrip money[5].

[5] Remark: Our mechanism does not make this assumption, but the approximation guarantees are much better once you make this assumption.

Also we restrict to the case when all players are treated equally. The motivating example for this class of instances is dealing with problems such as the one that arose with the Czech privatization auctions [1].

Let p^* be the PF prices. Our main theorem in this section is the following:

Theorem 4. *For additive linear valuations, when all bidders are treated equally, i.e., their budget b_i is equal to 1, there exists a mechanism that achieves an approximation factor of* $\min_j \left\{p_j^*/\lceil p_j^* \rceil\right\}$.

Note that if $k = \min_j p_j^*$, this is an approximation factor of at least $k/(k+1)$.

We now describe our mechanism which we call the *Strong Demand Matching* mechanism (SDM). Informally speaking, SDM starts by giving every bidder a unit amount of *scrip* money. It then aims to discover *minimal* item prices such that the demand of each bidder at these prices can be satisfied using (a fraction of) just one item. Surprisingly, doing so renders the mechanism truthful and gives a good approximation for the case when every item is well demanded. Next we define our mechanism in more detail.

Let p_j denote the price of item e_j, and let the *bang per buck* that Bidder i gets from item e_j equal v_{ij}/p_j. We say that item e_j is an MBB item of Bidder i if Bidder i gets the maximum bang per buck from that item[6]. For a given price vector p, let the demand graph $D(p)$ be a bipartite graph with bidders on one side and items on the other, such that an edge between Bidder i and item e_j exists if and only if e_j is an MBB item of Bidder i. We call $c_j = \lfloor p_j \rfloor$ the *capacity* of item e_j when its price is p_j, and we say an assignment of bidders to items is *valid* if it matches each bidder to at most one item and no item e_j is matched to more than c_j bidders. Given a valid assignment A, we say an item e_j is *reachable* from Bidder i if there exists an alternating path $(i, j_1, i_1, j_2, i_2, \cdots, j_k, i_k, j)$ in the graph $D(p)$ such that edges $(i_1, j_1), \cdots, (i_k, j_k)$ lie in the assignment A. Finally, let $d(R)$ be the collection of bidders with all their MBB items in set R.

The SDM mechanism initializes all item prices to $p_j = 1$ and iterates as follows:

1. Find a valid assignment that maximizes the number of matched bidders.
   If all the bidders are matched, conclude with Step 3.
2. Let U be the set of bidders who are not matched in Step 1. Let R be the set of all items reachable from bidders in the set U.
   Raise the price of each item e_j in R from p_j to $r \cdot p_j$,
   where $r \geq 1$ is the minimum value for which one of the following events takes place:
   (a) The price of an item in R reaches an integral value. If this happens, repeat Step 1.
   (b) For some bidder $b_i \in d(R)$, her set of MBB items increases, causing R to grow:
      i. If for each item e_j added to R, the number of bidders matched to it equals c_j,
      continue with Step 2.
      ii. If some item e_j added to R has c_j greater than the number of bidders matched to it, continue with Step 1.
3. Every bidder matched to some item e_j is allocated a fraction $1/p_j$ of that item.

It remains to explain how to carry out Step 2. Set R can be found using a breadth-first-search like algorithm. To determine when (a) is reached, we just need to know the smallest $\lceil p_j \rceil / p_j$ ratio over all items whose price is being increased. For (b), we need to calculate, for each bidder in $d(R)$, the ratio of the *bang per buck* for her MBB items and for the items outside the set R.

Running time If $c(R) = \sum_{j \in R} c_j$ denotes the total capacity in R, it is not difficult to see that if U is non-empty, $|d(R)| > c(R)$. Note that each time either event (a) or event (b)-ii occurs, $c(R)$ increases by at least 1, and thus, using the alternating path from a bidder in the set U to the corresponding item, we can increase the number of matched bidders by at least 1; this means that this can occur at most n times. The

[6]Note that for each bidder there could be multiple MBB items.

only other events are the unions resulting from (b)-i. There can be at most $\min(n, m)$ of these, and they are followed by either Step (a) or (b)-ii. Thus there are $O(n * \min(n, m))$ iterations of Step (b)-i and $O(n)$ iterations of Steps 1 and (b)-ii.

Correctness Let p^* be the PF prices and let q be the prices computed by the algorithm.

Lemma 6. *Let $f = \max_j \lceil p_j^* \rceil / p_j^*$. Then $q \le fp^*$.*

Proof. First note that at prices fp^*, the MBB items for each bidder are the same as at prices p^*. It is not difficult to see that at prices fp^* every bidder can be allocated to exactly one item from among her MBB items such that the number of bidders allocated to an item e_j is less than or equal to fp_j^*. To show this, consider a PF allocation. Form the following graph on items and bidders — add an edge between a bidder and an item if a portion of this item is assigned to this bidder in the PF solution. If there exists a cycle in this graph, one can remove an edge in this cycle by reallocating along the cycle while maintaining the valuation of every bidder. Hence there is a PF allocation in which this graph is a forest. Now for a given tree, root it at an arbitrary bidder. For each Bidder b in this tree, assign it to one of its child items, if any, and otherwise to its parent. The result is that for each item e_j, at most $\lceil p_j \rceil$ bidders will be assigned to it.

Now, suppose that some $q_j > fp_j^*$. Consider the first time t at which some price q_i starts to increase from fp_i^*. Let S be the set of items e_i whose price is currently fp_i^*. We will show that no item in set S will be part of the set R, and hence the prices of these items will not increase. Let T_q and T_{fp^*} be the sets of bidders who have edges to some item in the set S at the current prices and at prices fp^*, respectively. Clearly $T_q \subseteq T_{fp^*}$. Also suppose a bidder $b \in T_{fp^*}$ has an edge to an item outside set S at prices fp^*; this means that $b \notin T_q$ at the current prices as b will strictly prefer the item outside the set S. Thus if $b \in T_q$, this implies that b has no edges outside the set S at prices fp^*, and so b was allocated to some item in set S at prices fp^*. Since we know that at prices fp^* all the bidders can be allocated to some item, this implies that $|T_q| \le c(S)$. Thus even at current prices, all the bidders in T_q can be allocated to items in set S, and hence no item in set S can be part of the set R. □

Truthfulness We argue by contradiction. Suppose that some Bidder b were not truthful in the above algorithm, which we name algorithm $\mathcal{A}$. First, we consider an alternate algorithm $\mathcal{A}'$, and show that it is a dominant strategy for b to be truthful in $\mathcal{A}'$. Then we show that $\mathcal{A}$ and $\mathcal{A}'$ produce the same outcomes, and consequently b should also be truthful in $\mathcal{A}$. $\mathcal{A}'$ proceeds as follows. It begins by running algorithm $\mathcal{A}$ but with b absent (the first run), yielding prices p'. Then it runs $\mathcal{A}$ on all n bidders, but starting from prices p' (the second run).

Lemma 7. *b is truthful in algorithm $\mathcal{A}'$.*

Proof. Suppose the second run of the algorithm ends when b can be matched using an alternating path that ends at an item e. Suppose at this point, the price of an item $e_j \in R$ is p_j. It is easy to see that bidder b has no incentives to lie to obtain an item that is not in R as the prices of these items is completely defined by other bidders. Suppose that by lying, Bidder b is able to get an item $e_j \in R$ at a price $p'_j < p_j$. Suppose that this happens with an alternating path that starts at b and ends at some item e'. Now if this path existed in the truthful scenario when the price of item e_j reaches p'_j, then even in the truthful scenario b would have been matched when the price of item e_j is p'_j. Thus this path doesn't exist in the truthful scenario when the price of item e_j is p'_j. But why does this path not exist in the truthful scenario? It must be the case that some item on this path has a higher price than the price in the lying scenario. A higher price on this item means that in the truthful scenario item e_j would have been matched via some alternating path ending at e' before the price of item e_j reached p'_j, a contradiction. □

Lemma 8. *$\mathcal{A}$ and $\mathcal{A}'$ have the same outcome.*

The proof of this Lemma is similar to that of Lemma 6, and is deferred to the appendix.

Corollary 3. *$\mathcal{A}$ is truthful.*

We can now conclude with the proof of Theorem 4.

Proof of Theorem 4. If a bidder is allocated a portion of item e_j, she receives a $1/q_j$ fraction. But by Lemma 6, $1/q_j \geq 1/(fp_j^*) \geq 1/\lceil p_j^* \rceil$. In value, her PF allocation equals a $1/p_j^*$ fraction of item e_j. Thus she achieves an approximation factor of $p_j^*/\lceil p_j^* \rceil$. The result follows on minimizing over all bidders. □

5 Conclusion

Our work was motivated by the fact that no incentive compatible mechanisms were known for the natural and widely used fairness concept of Proportional Fairness. In hindsight our work provides several new contributions. Firstly, the class of bidder valuation functions for which our results apply is surpringly large and it contains several well studied functions; previous truthful mechanisms for fairness were studied for much more restricted classes of valuation functions. Secondly, to the best of our knowledge, this is first work that defines and gives guarantees for a strong notion of approximation for fairness, where one desires to approximate the valuation of every buyer. Lastly, our Partial Allocation Mechanism can be seen as a framework for designing truthful mechanims without money. This mechanism can be generalized further by restricting the range of the outcomes (similar to maximal-in-range mechanisms when one can use money). We believe that this generalization is a powerful one, and might allow for new solutions to other mechanism design problems without money. We plan to explore this in our future research.

Acknowledgement: The second author would like to thank Vasilis Syrgkanis for his help in clarifying the connection between the Partial Allocation mechanism and VCG payments.

References

[1] R. Aggarwal and J.T. Harper, *Equity valuation in the czech voucher privatization auctions*, Financial Management (Winter 2000), 77–100.

[2] M. Andrews, L. Qian, and A. L. Stolyar, *Optimal utility based multi-user throughput allocation subject to throughput constraints*, INFOCOM, 2005, pp. 2415–2424.

[3] J.B. Barbanel, *The geometry of efficient fair division*, Cambridge University Press, 2004.

[4] S. Brams and A. Taylor, *Fair division: from cake cutting to dispute resolution*, Cambridge University Press, Cambridge, 1996.

[5] E. Budish, *The combinatorial assignment problem: approximate competitive equilibrium from equal incomes*, BQGT '10, ACM, 2010, pp. 74:1–74:1.

[6] I. Caragiannis, C. Kaklamanis, P. Kanellopoulos, and M. Kyropoulou, *The efficiency of fair division*, Theory Comput. Syst. **50** (2012), no. 4, 589–610.

[7] Y. Chen, J.K. Lai, D.C. Parkes, and A.D. Procaccia, *Truth, justice, and cake cutting*, AAAI, 2010.

[8] Y. Chevaleyre, P.E. Dunne, U. Endriss, J. Lang, M. Lemaître, N. Maudet, J.A. Padget, S. Phelps, J.A. Rodríguez-Aguilar, and P. Sousa, *Issues in multiagent resource allocation*, Informatica (Slovenia) **30** (2006), no. 1, 3–31.

[9] Y.J. Cohler, J.K. Lai, D.C. Parkes, and A.D. Procaccia, *Optimal envy-free cake cutting*, AAAI, 2011.

[10] R. Cole, V. Gkatzelis, and G. Goel, *Positive results for mechanism design without money*, Submitted.

[11] N.R. Devanur, C.H. Papadimitriou, A. Saberi, and V.V. Vazirani, *Market equilibrium via a primal-dual algorithm for a convex program*, JACM. **55** (2008), no. 5.

[12] D. Dolev, D.G. Feitelson, J.Y. Halpern, R. Kupferman, and N. Linial, *No justified complaints: on fair sharing of multiple resources*, ITCS, 2012, pp. 68–75.

[13] E. Eisenberg, *Aggregation of utility functions*, Management Science **7** (1961), no. 4, 337–350.

[14] E. Eisenberg and D. Gale, *Consensus of subjective probabilities: The pari-mutuel method*, Ann. Math. Stat. **30** (1959), 165–168.

[15] A. Ghodsi, M. Zaharia, B. Hindman, A. Konwinski, S. Shenker, and I. Stoica, *Dominant resource fairness: fair allocation of multiple resource types*, NSDI'11, 2011, pp. 24–24.

[16] M. Guo and V. Conitzer, *Strategy-proof allocation of multiple items between two agents without payments or priors*, AAMAS, 2010, pp. 881–888.

[17] A. Gutman and N. Nisan, *Fair allocation without trade.*

[18] L. Han, C. Su, L. Tang, and H. Zhang, *On strategy-proof allocation without payments or priors*, WINE, 2011, pp. 182–193.

[19] J.D. Hartline and T. Roughgarden, *Optimal mechanism design and money burning*, STOC, 2008, pp. 75–84.

[20] K. Jain and V.V. Vazirani, *Eisenberg-gale markets: algorithms and structural properties*, STOC, 2007, pp. 364–373.

[21] F. Kelly, *Proportional fairness*, www.statslab.cam.ac.uk/~frank/pf/.

[22] F. P. Kelly, *Charging and rate control for elastic traffic*, European Transactions on Telecommunications **8** (1997), 33–37.

[23] F. P. Kelly, A. K. Maulloo, and D. K. H. Tan, *Rate control in communication networks: shadow prices, proportional fairness and stability*, Journal of Operational Research Society **49** (1998), 237–252.

[24] A. Mas-Colell, M. Whinston, and J. Green, *Microeconomic theory*, Oxford, 1995.

[25] A. Maya and N. Nisan, *Incentive compatible two player cake cutting*, WINE, 2012.

[26] E. Mossel and O. Tamuz, *Truthful fair division*, SAGT, 2010, pp. 288–299.

[27] H. Moulin, *Fair division and collective welfare*, The MIT Press, 2003.

[28] J. Nash, *The bargaining problem*, Econometrica **18** (1950), no. 2, 155–162.

[29] A. Nemirovski, *Advances in convex optimization: conic programming*, Proceedings oh the International Congress of Mathematicians, 2006, pp. 413–444.

[30] N. Nisan, T. Roughgarden, É. Tardos, and V.V. Vazirani, *Algorithmic game theory*, Cambridge University Press, New York, NY, USA, 2007.

[31] A. Othman, T. Sandholm, and E. Budish, *Finding approximate competitive equilibria: efficient and fair course allocation*, AAMAS '10, 2010, pp. 873–880.

[32] D.C. Parkes, A.D. Procaccia, and N. Shah, *Beyond dominant resource fairness: extensions, limitations, and indivisibilities*, ACM Conference on Electronic Commerce, 2012, pp. 808–825.

[33] A.D. Procaccia and M. Tennenholtz, *Approximate mechanism design without money*, ACM Conference on Electronic Commerce, 2009, pp. 177–186.

[34] J.M. Robertson and W.A. Webb, *Cake-cutting algorithms - be fair if you can*, A K Peters, 1998.

[35] H.R. Varian, *Equity, envy, and efficiency*, Journal of Economic Theory **9** (1974), no. 1, 63–91.

[36] H.P. Young, *Equity*, Princeton University Press, 1995.

[37] R. Zivan, M. Dudík, S. Okamoto, and K.P. Sycara, *Reducing untruthful manipulation in envy-free pareto optimal resource allocation*, IAT, 2010, pp. 391–398.

A Omitted Proofs

In what follows, we provide some of the proofs that are missing from the main section of the paper.

Proof of Lemma 2. We first prove that this lemma is true for any number k of pairs when $\beta_i = 1$ for every pair. For this special case we need to show that, if $\sum_{i=1}^{k} \delta_i \le b$, then

$$\prod_{i=1}^{k} (1+\delta_i) \le \left(1+\frac{b}{k}\right)^k.$$

Let $\bar{\delta}_i$ denote the values that actually maximize the left hand side of this inequality and $\Delta_{k'} = \sum_{i=1}^{k'} \bar{\delta}_i$ denote the sum of these values up to $\bar{\delta}_{k'}$. Note that it suffices to show that $\bar{\delta}_i = b/k$ for all i since we have

$$\prod_{i=1}^{k} (1+\delta_i) \le \prod_{i=1}^{k} (1+\bar{\delta}_i),$$

and replacing $\bar{\delta}_i$ with b/k yields the inequality that we want to prove.

To prove that $\bar{\delta}_i = b/k$ we first prove that for any $k' \le k$ and any $i \le k'$ we get $\bar{\delta}_i = \Delta_{k'}/k'$; we prove this fact by induction on k': For the basis step ($k'=2$) we show that $\bar{\delta}_1 = \Delta_2/2$. For any given value of Δ_2 we know that any choice of δ_1 will yield

$$\prod_{i=1}^{2} (1+\delta_i) \;=\; (1+\delta_1)(1+\Delta_2-\delta_1).$$

Taking the partial derivative with respect to δ_1 readily shows that this is maximized when $\delta_1 = \Delta_2/2$, thus $\bar{\delta}_1 = \Delta_2/2$. For the inductive step we assume that $\bar{\delta}_i = \Delta_{k'-1}/(k'-1)$ for all $i \le k'-1$. This implies that for any given value of $\Delta_{k'}$, given a choice of $\delta_{k'}$ the remaining product is maximized if the following holds

$$\prod_{i=1}^{k'} (1+\delta_i) = \left(1+\frac{\Delta_{k'}-\delta_{k'}}{k'-1}\right)^{k'-1} (1+\delta_{k'}).$$

Once again, taking the partial derivative of this last formula with respect to $\delta_{k'}$ for any given $\Delta_{k'}$ shows that this is maximized when $\delta_{k'} = \Delta_{k'}/k'$. This of course implies that $\Delta_{k'-1} = \frac{k'-1}{k'}\Delta_{k'}$ so $\bar{\delta}_i = \Delta_{k'}/k'$ for all $i \le k'$.

This property of the $\bar{\delta}_i$ that we just proved, along with the fact that $\Delta_k \le b$ implies

$$\prod_{i=1}^{k} (1+\delta_i) \;\le\; \left(1+\frac{\Delta_k}{k}\right)^k \;\le\; \left(1+\frac{b}{k}\right)^k.$$

We now use what we proved above in order to prove the lemma for any rational δ_i using a proof by contradiction. Assume that there exists a multiset $\mathcal{A}$ of pairs (δ_i, β_i) with $\beta_i \ge 1$ and $\sum_i \beta_i \cdot \delta_i \le b$ such that

$$\prod_i (1+\delta_i)^{\beta_i} > \left(1+\frac{b}{B}\right)^B, \tag{9}$$

where $B = \sum_i \beta_i$. Let M be an arbitrarily large value such that $\beta_i' = M\beta_i$ is a natural number for all i. Also, let $b' = Mb$. Then $\sum_i \beta_i' \cdot \delta_i \leq b'$, and $B' = M \cdot B = \sum_i \beta_i'$. Raising both sides of Inequality 9 to the power of M yields

$$\prod_i (1+\delta_i)^{\beta_i'} > \left(1 + \frac{b'}{B'}\right)^{B'}.$$

To verify that this is a contradiction, we create a multiset to which, for any pair (δ_i, β_i) of multiset $\mathcal{A}$, we add β_i' pairs $(\delta_i, 1)$. This multiset contradicts what we showed above for the special case of pairs with $\beta_i = 1$.

Extending the result to real valued δ_i just requires approximating the δ_i closely enough with rational valued terms. Specifically, let $\delta_i = \delta_i' + \epsilon_i$, where $\epsilon_i \geq 0$ and δ_i' is rational. Then $\sum_i \delta_i' \beta_i \leq b$, and by the result for rational δ,

$$\prod_i (1+\delta_i')^{\beta_i} \leq \left(1 + \frac{b}{B}\right)^{B}.$$

But then

$$\begin{aligned}
\prod_i (1+\delta_i)^{\beta_i} &\leq \prod_i (1+\delta_i'+\epsilon_i)^{\beta_i} \\
&\leq \prod_i \left[(1+\delta_i')\left(1+\frac{\epsilon_i}{1+\delta_i'}\right)\right]^{\beta_i} \\
&\leq \left(1+\frac{b}{B}\right)^{B} \prod_i \left(1+\frac{\epsilon_i}{1+\delta_i'}\right)^{\beta_i}.
\end{aligned}$$

As the ϵ_i can be chosen to be arbitrarily small, it follows that

$$\prod_i (1+\delta_i)^{\beta_i} \leq \left(1 + \frac{b}{B}\right)^{B}$$

for real valued δ_i too. □

Proof of Lemma 3. For any given approximate PF allocation $\widetilde{x}$, one can quickly verify that the valuation of bidder i for her final allocation only decreases as the value of $\prod_{i' \neq i} [v_{i'}(\widetilde{x}_{-i})]^{b_{i'}}$ increases. We can therefore assume that the approximation factor is minimized when the denominator of Equation (8) takes on its maximum value, i.e. $\widetilde{x}_{-i} = x^*_{-i}$. This implies that the fraction in this equation will always be less than or equal to 1, and the valuation of bidder i will therefore equal

$$\begin{aligned}
\widetilde{f}_i \cdot v_i(\widetilde{x}) &\geq \left(\frac{\prod_{i'} [v_{i'}(\widetilde{x})]^{b_{i'}}}{\prod_{i' \neq i} [v_{i'}(x^*_{-i})]^{b_{i'}}}\right)^{1/b_i} \\
&\geq (1-\epsilon) \left(\frac{\prod_{i'} [v_{i'}(x^*)]^{b_{i'}}}{\prod_{i' \neq i} [v_{i'}(x^*_{-i})]^{b_{i'}}}\right)^{1/b_i} \\
&= (1-\epsilon) f_i \cdot v_i(x^*).
\end{aligned}$$

The first inequality holds because the right hand side is minimized when $\widetilde{x}_{-i} = x^*_{-i}$, and the second inequality holds because $\widetilde{x}$ is defined to be an allocation that approximates x^*.

□

Proof of Lemma 4. In the proof of the previous lemma we showed that, if bidder i is truthful, then her valuation in the final allocation produced by the adapted PA mechanism will always be at least $(1-\epsilon)$ times the valuation $f_i \cdot v_i(x^*)$ that she would receive if all the PF allocations could be computed optimally rather than approximately. We now show that her valuation cannot be more than $(1-\epsilon)^{-1}$ times greater than

$f_i \cdot v_i(x^*)$, even if she misreports her preferences. Upon proving this statement, the theorem follows from the fact that, even if bidder i being truthful results in the worst possible approximation for this bidder, still any lie can increase her valuation by a factor of at most $(1-\epsilon)^{-2}$.

For any allocation x' we know that $\prod_{i'} [v_{i'}(x')]^{b_{i'}} \leq \prod_{i'} [v_{i'}(x^*)]^{b_{i'}}$, by definition of PF. Also, any allocation x'_{-i} that the approximation algorithm may compute instead of x^*_{-i} will satisfy $\prod_{i' \neq i} [v_{i'}(x')]^{b_{i'}} \geq (1-\epsilon) \prod_{i' \neq i} [v_{i'}(x^*)]^{b_{i'}}$. Using Equation (8) we can therefore infer that no matter what the computed allocations x' and x'_{-i} are, bidder i will experience a valuation of at most

$$\begin{aligned} \left(\frac{\prod_{i'} [v_{i'}(x')]^{b_{i'}}}{\prod_{i' \neq i} [v_{i'}(x'_{-i})]^{b_{i'}}} \right)^{1/b_i} &\leq \left(\frac{\prod_{i'} [v_{i'}(x^*)]^{b_{i'}}}{\prod_{i' \neq i} [v_{i'}(x'_{-i})]^{b_{i'}}} \right)^{1/b_i} \\ &\leq (1-\epsilon) \left(\frac{\prod_{i'} [v_{i'}(x^*)]^{b_{i'}}}{\prod_{i' \neq i} [v_{i'}(x^*_{-i})]^{b_{i'}}} \right)^{1/b_i} \\ &\leq (1-\epsilon)\ f_i \cdot v_i(x^*). \end{aligned}$$

□

Proof of Lemma 5. As the valuation functions are all concave and homogeneous of degree one, so is the following product,

$$\left(\prod_i [v_i(x)]^{b_i} \right)^{1/B} \tag{10}$$

and it has the same optima as the PF objective. Consequently the above optimization is an instance of convex programming with linear constraints, which can be solved approximately in polynomial time. More precisely, an approximation with an additive error of ϵ to the optimal product of the valuations can be found in time polynomial in the problem instance size and $\log(1/\epsilon)$ [29]. In addition, the approximation is a feasible allocation.

We normalize the individual valuations to have a value 1 for an allocation of everything. If $B = \sum_i b_i$ is the sum of the bidders' weights then, at the optimum, bidder i has valuation at least b_i/B. To verify that this is true, just note that the sum of the prices of all goods in the competitive equilibrium with equal incomes will be B and every bidder will have a budget of b_i. Since the bidder will spend all her budgets on the items she values the most for the prices at hand, her valuation for her bundle will have to be at least b_i/B. This implies that the optimum product valuation is at least $\prod_i (b_i/B)^{b_i/B} \geq \min_i b_i/B$; this can be approximated to within an additive factor $\epsilon \cdot \min_i b_i/B$ in time polynomial in $\log 1/\epsilon + \log B$, and this is an approximation to within a multiplicative factor of $1-\epsilon$.

□

Proof of Lemma 8. Suppose, for a contradiction, that $\mathcal{A}$ and $\mathcal{A}'$ have different outcomes. Without loss of generality suppose that some price is higher in the outcome of $\mathcal{A}'$ (if not, switch the roles of $\mathcal{A}$ and $\mathcal{A}'$).

In the run of the mechanism $\mathcal{A}'$, consider a time when no item has exceeded the final price given by the mechanism $\mathcal{A}$ but some items have reached that price. Let S be the set of items at this point that have prices equal to their final prices in $\mathcal{A}$. We will show that no item in set S will be part of the set R from this time onwards in the mechanism $\mathcal{A}'$, and hence the prices of these items will not increase in $\mathcal{A}'$. Let T' and T be the set of bidders who have edges to some item in the set S at the current prices in $\mathcal{A}'$, and at the final prices in $\mathcal{A}$, respectively. Clearly, $T' \subseteq T$. Also suppose a bidder $b \in T$ has an edge to an item outside set S at the final prices of $\mathcal{A}$; this means that $b \notin T'$ at the current prices as b will strictly prefer the item outside the set S. Thus if $b \in T'$, this implies that b has no edges to an item outside the set S at the final prices given by $\mathcal{A}$, and so in $\mathcal{A}$, b was allocated to some item in set S. Since we know that at the final prices given by $\mathcal{A}$, all the bidders can be allocated to some item, this implies that $|T'| \leq c(S)$. Thus, even at current prices, all the bidders in T' can be allocated to items in the set S, and hence no item in set S can be part of the set R. Thus no item can have higher final price in $\mathcal{A}'$ than in $\mathcal{A}$. A contradiction. □