Efficient Mechanisms without Money:
Randomization Won't Let You Escape from Dictatorships

Rohan Garg,[1]
Alexandros Psomas,[1]
[1] Purdue University
rohanvgarg@gmail.com, apsomas@cs.purdue.edu

Abstract

We study the fundamental mechanism design problem of allocating a set of items among additive agents, without monetary transfers. It is well known that the only deterministic mechanism that satisfies Pareto efficiency and truthfulness is the serial dictatorship. A central problem is whether randomized rules can simultaneously satisfy truthfulness, efficiency, and provide any non-trivial fairness guarantee. We settle this open problem in the negative, by showing that, even for the case of two agents, every Pareto efficient and truthful mechanism is a serial dictatorship.

1 Introduction

We study the fundamental mechanism design problem of allocating a set of items among a set of agents with additive utilities. In this elementary setting, the tension between truthfulness, fairness, and efficiency is well-understood. Achieving all three is generally impossible, see e.g. (Kojima 2009). There are randomized rules (or, equivalently, if items are divisible) that are envy-free and Pareto efficient, simultaneously (Varian 1973). And, even though envy-freeness is generally impossible for deterministic rules,[1] there is an efficient, deterministic rule that satisfies a compelling notion of fairness: envy-freeness up to one item (EF1) (Caragiannis et al. 2019). On the other hand, there is no deterministic and truthful mechanism that always outputs an EF1 allocation (Amanatidis et al. 2017), even for two additive agents. Randomization allows us to circumvent this negative result: randomly allocating each item is truthful and envy-free.

Truthfulness and efficiency can be simultaneously achieved via a serial dictatorship. Unfortunately, a dictatorship is typically the *only* deterministic, truthful and efficient mechanism, in a number of settings, see e.g. (Pápai 2000; Klaus and Miyagawa 2002; Ehlers and Klaus 2003). A central open problem is whether randomization can provide an escape from these negative results. Our research question is:

> Is there a (randomized) non-dictatorial, truthful and Pareto efficient mechanism for additive agents?

We prove that the answer is an emphatic "no"! The only truthful and Pareto efficient mechanism for two additive agents is a dictatorship, even when randomization is allowed (or, equivalently, when items are divisible).

Challenges. Structural results about truthful deterministic mechanisms typically start by fixing the outcome of the mechanism on a specific instance (e.g., "in instance I, item 1 goes to agent A, without loss of generality"), and then use truthfulness to argue about the outcomes in other, similar instances (e.g. "in instance I', item 1 should go to agent A, otherwise she would have an incentive to deviate in a way that results in instance I") to slowly propagate to every possible instance. Truthfulness is a constraint connecting the *utility* of an agent between instances, and there might be many different allocations of the items that yield the same utilities. Therefore, the main technical task is to argue about how utilities should map to allocations. This task is vastly simpler when allocations are deterministic. That is, if we've established that every instance with k items should give all items to agent A, there are only a few possibilities for an instance with $k + 1$ items: truthfulness implies strong lower bounds on the (previous) dictator's utility, and efficiency narrows down the possibilities even more, making even a case analysis tractable.

For randomized mechanisms this common proof structure immediately breaks down. E.g., for the case of a single item and two agents, ignoring the agents' reported values and flipping a biased coin that gives the item to agent A with probability α (and agent B with probability $1 - \alpha$) is truthful and efficient (as long as both agents have positive value). But, it is not clear what this choice implies for other instances: truthfulness gives some loose bounds on the utility of each agent, which, in turn, implies some loose constraints on the allocation space, but no insight on the mechanism's structure whatsoever. We bypass such issues by carefully selecting the instances for which truthfulness implies the strongest possible constraints on the allocation. Additive utilities imply additional, *linear* constraints for the agent's utilities, since the Pareto frontier is comprised of lines. In addition to restricting the utility space, these linear constraints significantly narrow down the *allocation space* for a given instance as well. Truthfulness can then be used to rule out many allocations, allowing us to make progress.

[1] Consider the case of a single item and two agents with strictly positive valuation.

1.1 Related Work

A long series of works characterize deterministic mechanisms in the absence of money, e.g. (Pápai 2000; Klaus and Miyagawa 2002; Ehlers and Klaus 2003; Nesterov 2017; Svensson 1999). Most of these works show that deterministic, truthful and efficient mechanisms are dictatorial, under various utility functions (unit-demand, additive, etc) and various combinations of other desirable axioms (e.g. anonymity, neutrality, non-bossiness, etc). (Amanatidis et al. 2017) completely characterize deterministic and truthful mechanisms for two additive agents. One immediate corollary of their characterization is that deterministic, truthful and Pareto efficient mechanisms are dictatorships.[2]. As we prove in this paper, this result persists even for randomized mechanisms. Closer to this work, (Kojima 2009; Aziz and Kasajima 2017) study randomized mechanisms. (Kojima 2009) shows that there exists no mechanism that is ordinally efficient, envy-free and weakly strategy-proof. (Aziz and Kasajima 2017) show that equal treatment of equals, efficiency, and strategy-proofness are incompatible. Here we show a much stronger impossibility result: the *only* strategy-proof and efficient mechanism is a dictatorship.

One way to escape the aforementioned impossibility results is to focus on special cases. In recent work, (Halpern et al. 2020; Babaioff, Ezra, and Feige 2021) show that truthfulness (and even group-strategyproofness), Pareto efficiency, and approximate fairness are compatible when agents are additive and have binary values, or when agents have dichotomous marginals. Leontief valuations also allow for truthful, fair and efficient rules (Ghodsi et al. 2011; Friedman, Ghodsi, and Psomas 2014).

A line of work initiated by Procaccia and Tennenholtz strives to design truthful, fair, and approximately efficient mechanisms (Procaccia and Tennenholtz 2013). Approximately efficient mechanisms can bypass the limitations of efficient mechanisms by throwing away resources as a substitute for payments, a technique known as money-burning (Hartline and Roughgarden 2008; Fotakis et al. 2016; Friedman et al. 2019; Abebe et al. 2020). The works closest to ours are (Cole, Gkatzelis, and Goel 2013), which give truthful, fair and approximately efficient mechanisms, even for valuation functions more general than additive, and (Guo and Conitzer 2010; Han et al. 2011; Aziz et al. 2016) that study truthful and approximate optimal welfare maximization. One interpretation of our result is that, if truthfulness is a hard constraint, then money-burning is *necessary* to get any non-trivial fairness guarantees.

Finally, our setting is closely related to the problem of truthful cake-cutting; e.g., see (Mossel and Tamuz 2010; Maya and Nisan 2012; Chen et al. 2013; Tao 2021). We note that in the cake-cutting setting, there are no deterministic, truthful and proportional cake cutting mechanisms (Tao 2021), but, as opposed to what we prove in this paper, randomization does provide an escape (Mossel and Tamuz 2010).

[2]Other than the focus on deterministic mechanisms, one difference between our work and (Amanatidis et al. 2017) is that here we allow items' valuations to be zero. When valuations are always strictly positive, mechanisms that "reserve" a single item for an agent, and exchange it if there is a Pareto improvement, are truthful, slightly expanding the set of truthful and deterministic mechanisms beyond serial dictatorships;

2 Preliminaries

We consider the problem of allocating a set $\mathcal{M}$ of m items among 2 agents with additive utilities. To simplify notation, instead of explicitly arguing about indivisible items, randomized rules and expected utility, we will assume that the items are divisible; the two settings are equivalent (this equivalence is discussed in (Guo and Conitzer 2010; Aziz et al. 2016)). A fractional allocation $\mathbf{x}$ defines for each agent $i \in \{1,2\}$ and item $j \in \mathcal{M}$ the fraction x_j^i of the item that the agent will receive. A feasible allocation satisfies, for all items $j \in \mathcal{M}$, $\sum_{i \in \{1,2\}} x_j^i \leq 1$.

Agents are additive. Each agent $i \in \{1,2\}$ has a non-negative valuation v_j^i for receiving the entirety of item j, and has a utility $u_i(\mathbf{x}) = \sum_{j \in \mathcal{M}} x_j^i \cdot v_j^i$ for a fractional allocation $\mathbf{x}$. Without loss of generality, we assume that $v_j^i \in [0,1]$ for all items $j \in \mathcal{M}$ and agents $i \in \{1,2\}$.

An allocation $\mathbf{x}$ is Pareto efficient if there is no feasible allocation $\mathbf{x}'$ such that for all agents $i \in \{1,2\}$, $u_i(\mathbf{x}') \geq u_i(\mathbf{x})$, with at least one inequality being strict.

A mechanism elicits values v_j^i for each item $j \in \mathcal{M}$, for every player $i \in \{1,2\}$, and outputs a feasible allocation. We write $x_j^i(\vec{v}^1, \vec{v}^2)$ for the fraction of item j allocated to agent i when each agent $i \in \{1,2\}$ reports a vector of valuations $\vec{v}^i = \left(v_1^i, \ldots, v_m^i\right)$.

We focus on *efficient* mechanisms, i.e. those which always output a Pareto efficient allocation. We also focus on *truthful* mechanisms; a mechanism is truthful if agents cannot strictly improve their utility by misreporting their valuation. Formally, for every agent $i \in \{1,2\}$, every possible valuation vector $\vec{v}^i$, every possible valuation vector of the other agent $\vec{v}^{-i}$, and every possible report for agent i $\vec{b}^i$, we have $\sum_{j \in \mathcal{M}} v_j^i \cdot x_j^i(\vec{v}^i, \vec{v}^{-i}) \geq \sum_{j \in \mathcal{M}} v_j^i \cdot x_j^i(\vec{b}^i, \vec{v}^{-i})$.

3 Truthfulness and Efficiency imply a Dictatorship

In this section we prove our main result.

Theorem 1. *Every truthful and Pareto efficient mechanism for two additive agents and $m \geq 2$ items is a dictatorship. That is, one of the two agents deterministically receives all the items she has a positive value for.*

Proof. We will prove Theorem 1 via induction. Let $\mathcal{I}^k$ be the set of all possible instances where both agents report a value of zero for items $k+1$ through m, i.e. instances of the form

$$I = \begin{bmatrix} v_1^1 & v_2^1 & \dots & v_k^1 & 0 & \dots & 0 \\ v_1^2 & v_2^2 & \dots & v_k^2 & 0 & \dots & 0 \end{bmatrix},$$

where $v_j^i \geq 0$ (for $j \leq k$).

Fix an arbitrary truthful and efficient mechanism M^*. We will show that M^* is a dictatorship. Throughout this proof, we write $x_j^i(I)$ for the fraction of item j allocated to agent i by M^* in instance I, and $u_i(I)$ for agent i's utility.

Consider the following instance $I^* \in \mathcal{I}^1$:

$$I^* = \begin{bmatrix} 1 & 0 & \dots & 0 \\ 1 & 0 & \dots & 0 \end{bmatrix}.$$

Since M^* is efficient, we have that $x_1^1(I^*)+x_2^1(I^*) = 1$. Let $\alpha = x_1^1(I^*)$ (and therefore $x_2^1(I^*) = 1-\alpha$); without loss of generality, $\alpha \geq 1/2$. We will show that agent 1 is a dictator, i.e. she receives all items she has a (strictly) positive value for (implying that α has to, in fact, be equal to one). We will prove this statement via induction on k, starting with $k = 2$.

3.1 Induction Basis

Our first goal is to show that for all $I \in \mathcal{I}^2$, agent 1 receives all items she has a (strictly) positive value for. Recall that in instance I^* we had $x_1^1(I^*) = \alpha \geq 1/2$.

First, we show that all instances in $\mathcal{I}^1$ have exactly the same allocation.

Lemma 1. *For all $I \in \mathcal{I}^1$ such that $v_1^1, v_1^2 > 0$, $x_1^1(I) = \alpha$ and $x_1^2(I) = 1-\alpha$.*

Proof. Consider an instance $I_a \in \mathcal{I}^1$ of the form:

$$I_a = \begin{bmatrix} a & 0 & \dots & 0 \\ 1 & 0 & \dots & 0 \end{bmatrix}.$$

where $a > 0$. If $x_1^1(I_a) > x_1^1(I^*)$, then agent 1 has an incentive to deviate from I^* to I_a. If $x_1^1(I_a) < x_1^1(I^*)$, then agent 1 has an incentive to deviate from I_a to I^*. Therefore, $x_1^1(I_a) = x_1^1(I^*)$ (and therefore $x_1^2(I_a) = x_1^2(I^*)$). Now, consider an instance $I_{(a,b)} \in \mathcal{I}^1$ of the form:

$$I_{(a,b)} = \begin{bmatrix} a & 0 & \dots & 0 \\ b & 0 & \dots & 0 \end{bmatrix},$$

where $a, b > 0$. If $x_1^2(I_a) > x_1^2(I_{(a,b)})$, then agent 2 has an incentive to deviate from $I_{(a,b)}$ to I_a. If $x_1^2(I_a) < x_1^2(I_{(a,b)})$, then agent 2 has an incentive to deviate from I_a to $I_{(a,b)}$. Therefore, $x_1^2(I_a) = x_1^2(I_{(a,b)}) = \alpha$ (and therefore, $x_1^1(I_a) = x_1^1(I_{(a,b)}) = 1-\alpha$). The lemma follows. ☐

Next, consider an instance $I \in \mathcal{I}^2$, such that the second item is desired by exactly one of the two agents. Pareto efficiency implies that this item should go to the agent that values it. The next lemma shows that the allocation of the first item cannot deviate much from the allocation in $\mathcal{I}^1$ instances.

Lemma 2. *Consider instances $I^2_{(a,b,c)}, \hat{I}^2_{(a,b,c)} \in \mathcal{I}^2$:*

$$I^2_{(a,b,c)} = \begin{bmatrix} a & b & 0 & \dots \\ c & 0 & 0 & \dots \end{bmatrix}, \hat{I}^2_{(a,b,c)} = \begin{bmatrix} c & 0 & 0 & \dots \\ a & b & 0 & \dots \end{bmatrix}.$$

Then, $x_1^1(I^2_{(a,b,c)}) \in [\alpha - b, \alpha]$ and $x_1^2(\hat{I}^2_{(a,b,c)}) \in [1-\alpha-b, 1-\alpha]$. Furthermore, $x_j^i(I^2_{(a,b,c)}) = x_j^i(I^2_{(a',b',c')})$ and $x_j^i(\hat{I}^2_{(a,b,c)}) = x_j^i(\hat{I}^2_{(a',b',c')})$, for all $a, a', b, b', c, c' > 0$.

Proof. For $I^2_{(a,b,c)}$, we have that item 2 must be allocated to agent 1, by Pareto efficiency. Consider the case where $a = 1$. Then, agent 1's utility when reporting honestly is $x_1^1(I^2_{(1,b,c)}) + b$. By Lemma 1, her utility when (dishonestly) reporting $v_2^1 = 0$ is exactly α, therefore, by truthfulness, $x_1^1(I^2_{(1,b,c)}) \geq \alpha - b$. Again by Lemma 1, if her true values are $v_1^1 = 1$ and $v_2^1 = 0$, she gets utility α, while reporting $v_2^1 = b$ gives her utility $x_1^1(I^2_{(1,b,c)})$, implying $\alpha \geq x_1^1(I^2_{(1,b,c)})$. Overall, we have that $\alpha \geq x_1^1(I^2_{(1,b,c)}) \geq \alpha - b$. Next, notice that if $x_1^1(I^2_{(a,b,c)}) \neq x_1^1(I^2_{(a',b',c)})$ for some $a, b, a', b', c > 0$, agent 1 has an obvious incentive to report the values that yield the higher allocation; a contradiction. Since $x_1^2(I^2_{(a,b,c)}) = 1 - x_1^1(I^2_{(a,b,c)})$, we have $x_j^i(I^2_{(a,b,c)}) = x_j^i(I^2_{(a',b',c)})$ for all $a, a', b, b', c > 0$. Finally, notice that if $x_j^i(I^2_{(a,b,c)}) \neq x_j^i(I^2_{(a,b,c')})$, for some $a, b, c, c' > 0$, then agent 2 has an incentive to misreport. The claims about $I^2_{(a,b,c)}$ follow.

Similarly, in instance $\hat{I}^2_{(a,b,c)}$ item 2 is allocated to agent 2, as long as $b > 0$. For $a = 1$ agent 2's (honest) utility is $x_1^1(\hat{I}^2_{(1,b,c)}) + b$, which is at least $1-\alpha$ (her allocation of item 1 when she reports $v_2^2 = 0$). Furthermore, $1-\alpha$ is at least $x_1^1(\hat{I}^2_{(1,b,c)})$, otherwise there is an incentive to misreport in the other direction. The second part of the lemma is identical to above. ☐

Notice that Lemma 2 implies that

$$\alpha \geq x_1^1(I^2_{(a,b,c)}) = x_1^1(I^2_{(a,b',c)}) \geq \alpha - b',$$

for all $b' > 0$. That is, the bounds on $x_1^1(I^2_{(a,b,c)})$ do not depend on a, b, c, and hold for all b'. The same observation can be made about $x_1^1(\hat{I}^2_{(a,b,c)})$. The following corollary is immediate:

Corollary 1. $x_1^1(I^2_{(a,b,c)}) = \alpha$ *and* $x_1^1(\hat{I}^2_{(a,b,c)}) = \alpha$*, for all* $a, b, c > 0$.

Fix some small $\epsilon \in (0,1)$. Our next lemma shows that the allocation in $I^2_{(\epsilon,1,1)}$ remains unchanged if agent 2 has value $v_2^2 = \epsilon$ instead of zero.

Lemma 3. *Consider the following instance $I_1 \in \mathcal{I}^2$:*

$$I_1 = \begin{bmatrix} \epsilon & 1 & 0 & \dots & 0 \\ 1 & \epsilon & 0 & \dots & 0 \end{bmatrix}.$$

Then, M^* satisfies $x_2^2(I_1) = 0$, and $x_1^2(I_1) = 1-\alpha$.

Proof. The Pareto frontier of I_1 can be seen in Figure 1, and consists of two lines: (1) $u_2 + \epsilon \cdot u_1 = 1 + \epsilon$, and (2) $\epsilon \cdot u_2 + u_1 = 1 + \epsilon$. Since M^* is Pareto efficient, the agents' utilities must satisfy at least one of two equations.

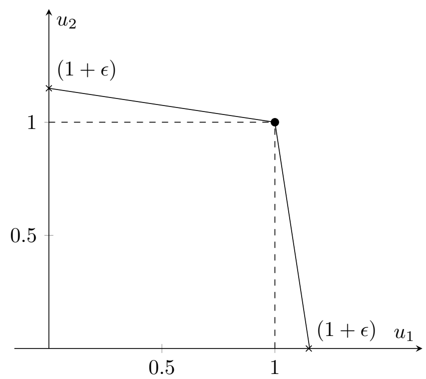

Figure 1: The Pareto frontier of instance I_1

.

First, consider the case that $u_2(I_1) + \epsilon \cdot u_1(I_1) = 1 + \epsilon$. Expanding these utilities, and using the fact that $x_j^2(I_1) = 1 - x_j^1(I_1)$, we have

$$\begin{aligned}1 - x_1^1(I_1) + \epsilon - \epsilon x_2^1(I_1) + \epsilon^2 x_1^1(I_1) + \epsilon x_2^1(I_1) &= 1 + \epsilon\\ -x_1^1(I_1) + \epsilon^2 x_1^1(I_1) &= 0\\ x_1^1(I_1)(\epsilon^2 - 1) &= 0,\end{aligned}$$

implying that $x_1^1(I_1) = 0$ (since $\epsilon \neq 1$). If this was the case, reporting $v_2^2 = \epsilon$ when the truth is $v_2^2 = 0$ would improve agent 2's utility from at most $1-\alpha$ (by Lemma 2) to 1. Since $\alpha > 0$, this is a contradiction.

Therefore, the agents' utilities in I_1 must lie in the second line, i.e. $\epsilon \cdot u_2(I_1) + u_1(I_1) = 1 + \epsilon$. Expanding we get:

$$\begin{aligned}\epsilon - \epsilon x_1^1(I_1) + \epsilon^2 - \epsilon^2 x_2^1(I_1) + \epsilon x_1^1(I_1) + x_2^1(I_1)) &= 1 + \epsilon\\ (1-\epsilon^2) x_2^1(I_1) &= 1 - \epsilon^2,\end{aligned}$$

implying that $x_2^1(I_1) = 1$. Thus, $u_2(I_1) = x_1^2(I_1)$. Notice that if $x_1^2(I_1) \neq x_1^2(I^2_{\epsilon,1,1})$, agent 2 has an incentive to report whichever v_2^2 yields the higher allocation (out of $v_2^2 = 0$, giving $I^2_{(\epsilon,1,1)}$, or $v_2^2 = \epsilon$, giving I_1), which is a violation of truthfulness. Therefore, by Corollary 1, $x_1^2(I_1) = x_1^2(I^2_{(\epsilon,1,1)}) = 1 - \alpha$. □

Lemma 3 already shows that M^* slightly favors agent 1, at least in instance I_1. Notice that the proof of Lemma 3 itself never used any non-trivial facts about α, other than the fact it is at least 0. Therefore, the next natural step is to consider an instance symmetric to I_1 (since there cannot be two dictators, the same type of reasoning should lead to a contradiction). Specifically, consider

$$I_2 = \begin{bmatrix} 1 & \epsilon & 0 & \dots & 0 \\ \epsilon & 1 & 0 & \dots & 0 \end{bmatrix}.$$

We will argue that M^* favors agent 1 even in I_2, and in fact that $\alpha = 1$.

Lemma 4. *It holds that $\alpha = 1$.*

Proof. The Pareto frontier of I_2 is identical to I_1. This means that there are, again, two cases for the agents' utilities: (1) $u_2 + \epsilon \cdot u_1 = 1 + \epsilon$, or (2) $\epsilon \cdot u_2 + u_1 = 1 + \epsilon$.

Case 1. Assume that $u_2(I_2) + \epsilon \cdot u_1(I_2) = 1 + \epsilon$. Expanding and simplifying we get that $(1-\epsilon^2)x_2^2(I_2) = 1 - \epsilon^2$, i.e. $x_2^2(I_2) = 1$. Therefore, $u_1(I_2) = x_1^1(I_2)$. Furthermore, notice that if $x_1^1(I_2) \neq x_1^1(\hat{I}^2_{\epsilon,1,1})$, there is an incentive for agent 1 to report whichever v_2^1 yields the higher allocation, violating truthfulness. Therefore, by Corollary 1, $x_1^1(I_2) = x_1^1(\hat{I}^2_{\epsilon,1,1}) = \alpha$.

Next, consider the following instance:

$$\bar{I} = \begin{bmatrix} \epsilon & 1 & 0 & \dots & 0 \\ \epsilon & 1 & 0 & \dots & 0 \end{bmatrix}$$

Agent 1's truthfulness constraints from I_2 to $\bar{I}$ imply that $x_1^1(I_2) \geq x_1^1(\bar{I}) + \epsilon x_2^1(\bar{I}) = 1 + \epsilon - (x_1^2(\bar{I}) + \epsilon x_2^2(\bar{I}))$. Agent 2's truthfulness constraints from I_1 to $\bar{I}$ imply that $x_1^2(I_1) \geq x_1^2(\bar{I}) + \epsilon x_2^2(\bar{I})$. Therefore, M^* should satisfy

$$\begin{aligned}x_1^2(I_1) &\geq 1 + \epsilon - x_1^1(I_2)\\ 1 - \alpha &\geq 1 + \epsilon - \alpha\\ 0 &\geq \epsilon,\end{aligned}$$

where we used the facts that $x_1^1(I_2) = \alpha$, and $x_1^2(I_1) = 1-\alpha$ (Corollary 1). Since we chose $\epsilon > 0$, this is a contradiction. Therefore, $u_2(I_2) + \epsilon \cdot u_1(I_2) \neq 1 + \epsilon$.

Case 2 Since Case 1 is infeasible, it must be that $\epsilon \cdot u_2(I_2) + u_1(I_2) = 1 + \epsilon$. Expanding and simplifying gives

$$\begin{aligned}\epsilon^2 - \epsilon^2 x_1^1(I_2) + \epsilon - \epsilon x_2^1(I_2) + x_1^1(I_2) + \epsilon x_2^1(I_2) &= 1 + \epsilon\\ (1-\epsilon^2) x_1^1(I_2) &= 1 - \epsilon^2,\end{aligned}$$

which implies that $x_1^1(I_2) = 1$.

Finally, consider the truthfulness constraints of agent 1 from $\hat{I}^2_{(\epsilon,1,1)}$ to I_2, recalling that

$$\hat{I}^2_{(\epsilon,1,1)} = \begin{bmatrix} 1 & 0 & 0 & \dots \\ \epsilon & 1 & 0 & \dots \end{bmatrix}.$$

We have $u_1(\hat{I}^2_{(\epsilon,1,1)}) = x_1^1(\hat{I}^2_{(\epsilon,1,1)}) = \alpha \geq x_1^1(I_2) = 1$. □

To conclude the proof of the induction basis, it suffices to show that $\alpha = 1$ implies that agent 1 receives all items she has a positive value for, in every instance in $\mathcal{I}^2$.

First, consider instances of the form

$$I_{(a,b)} = \begin{bmatrix} a & b & 0 & \dots \\ 1 & \epsilon & 0 & \dots \end{bmatrix}.$$

By Lemmas 3 and 4 we have that agent 1 can always get any items she wants (from the first two) if she simply reports $v_1^1 = \epsilon$ and $v_2^1 = 1$, instead of a and b. Therefore,

$x_1^1(I_{(a,b)}) = x_2^1(I_{(a,b)}) = 1$, if $a, b > 0$. Then, if there was an instance $I \in \mathcal{I}^2$ of the form

$$I = \begin{bmatrix} a & b & 0 & \dots \\ c & d & 0 & \dots \end{bmatrix},$$

such that $x_1^2(I) > 0$ or $x_2^2(I) > 0$ (with the corresponding a, b also strictly positive), then agent 2 would have an incentive to deviate (from $I_{(a,b)}$); a contradiction.

This concludes the proof of the induction basis.

3.2 Induction step

Assume that for all $I \in \mathcal{I}^k$ we have that agent 1 receives all items she has a positive value for. Our goal in this section is to show that the same is true for all instances in $\mathcal{I}^{k+1}$.

First, fix an arbitrary $\epsilon \in (0, \sqrt{1/k})$ and a vector $\vec{a} = (a_1, \dots, a_k)$.

Consider the following instance:

$$S_{(\vec{a},\epsilon)} = \begin{bmatrix} 1 & 1 & \dots & 1 & \epsilon & 0\dots \\ a_1 & a_2 & \dots & a_k & 0 & 0\dots \end{bmatrix}.$$

By Pareto efficiency, $x_{k+1}^1(S_{(\vec{a},\epsilon)}) = 1$. And, since agent 1 can get all of the first k items by reporting $v_{k+1}^1 = 0$, truthfulness implies that $\sum_{i=1}^k x_i^1(S_{(\vec{a},\epsilon)}) + \epsilon \geq k$.

Similarly, it holds that $\sum_{i=1}^k x_i^1(S_{(\vec{a},\delta)}) \geq k - \delta$, for every $\delta \in (0, 1)$. If $\sum_{i=1}^k x_i^1(S_{(\vec{a},\delta)}) \neq \sum_{i=1}^k x_i^1(S_{(\vec{a},\epsilon)})$, agent 1 has an incentive to report v_{k+1}^1 (either ϵ or δ) in a way that maximizes her allocation; a contradiction.

Therefore, $k \geq \sum_{i=1}^k x_i^1(S_{(\vec{a},\epsilon)}) = \sum_{i=1}^k x_i^1(S_{(\vec{a},\delta)}) \geq k - \delta$, for all $\delta > 0$. Thus, $\sum_{i=1}^k x_i^1(S_{(\vec{a},\epsilon)}) = k$.

Next, consider the following instance $S \in \mathcal{I}^{k+1}$:

$$S = \begin{bmatrix} \epsilon & \epsilon & \dots & \epsilon & 1 & 0\dots \\ 1 & 1 & \dots & 1 & 0 & 0\dots \end{bmatrix}.$$

First, notice that $x_{k+1}^1(S) = 1$, by Pareto efficiency. Furthermore, since agent 1 can deviate to $S_{(\vec{1},\epsilon)}$, $\sum_{i=1}^k x_i^1(S)$ has to be at least $\sum_{i=1}^k x_i^1(S_{(\vec{1},\epsilon)}) = k$. And, since the allocation of the first k items is at most k, we overall have $\sum_{i=1}^k x_i^1(S) = k$. That is, agent 1 gets all items in S.

The next lemma shows that this situation remains unchanged when v_{k+1}^2 slightly increases.

Lemma 5. *Consider the instance $S_1 \in \mathcal{I}^{k+1}$*

$$S_1 = \begin{bmatrix} \epsilon & \epsilon & \dots & \epsilon & 1 & 0\dots \\ 1 & 1 & \dots & 1 & \epsilon & 0\dots \end{bmatrix}.$$

Then, agent 1 receives all items in S_1, i.e. $x_i^1(S_1) = 1$, for all $i = 1, \dots, k$.

Proof. Assume that M^* allocates a (strictly) positive fraction of one of the first k items to agent 2 in S_1. However, since agent 2 gets nothing in S, this would imply a violation of truthfulness, since agent 2 would prefer to report $v_{k+1}^2 = \epsilon$ when the truth is $v_{k+1}^2 = 0$. Therefore, $\sum_{i=1}^k x_i^2(S_1) = 0$ and $\sum_{i=1}^k x_i^1(S_1) = k$. Also, we immediately have that $u_2(S_1) = \epsilon x_{k+1}^2(S_1) \leq \epsilon$

Next, we argue that $k\epsilon u_2(S_1) + u_1(S_1) = 1 + k\epsilon$.

First, notice that

$$\begin{aligned} k\epsilon u_2(S_1) + u_1(S_1) &= k\epsilon\left(\epsilon x_{k+1}^2(S_1)\right) + \epsilon k + x_{k+1}^1(S_1) \\ &= 1 + \epsilon k + x_{k+1}^2(S_1)(k\epsilon^2 - 1) \\ &\leq 1 + \epsilon k, \end{aligned}$$

where we used the fact that $k\epsilon^2 - 1 < 0$, since we picked $\epsilon < \sqrt{1/k}$. Second, it is feasible to get utility points $(1, k)$ and $(1 + k\epsilon, 0)$: the first point is achievable by allocating the first k items to agent 2 and the last item to agent 1, and the second point is achievable by allocating all items to agent 1. Third, the line that connects $(1, k)$ and $(1 + k\epsilon, 0)$ is exactly $k\epsilon u_2 + u_1 = 1 + k\epsilon$, i.e. all the utility points on this line are feasible. Combining the first and third observations, we have that allocations that yield utilities such that $k\epsilon u_2 + u_1 = 1 + k\epsilon$, are Pareto efficient. Finally, since M^* is efficient and $u_2(S_1) \leq \epsilon$, it must be that M^* selects such at point, i.e. $k\epsilon u_2(S_1) + u_1(S_1) = 1 + k\epsilon$.

Then, notice that

$$\begin{gathered} k\epsilon u_2(S_1) + u_1(S_1) = 1 + k\epsilon \\ 1 + \epsilon k + x_{k+1}^2(S_1)(k\epsilon^2 - 1) = 1 + k\epsilon \\ x_{k+1}^2(S_1)(k\epsilon^2 - 1) = 0, \end{gathered}$$

implying that $x_{k+1}^2(S_1) = 0$. The lemma follows. □

Given Lemma 5 it is straightforward to conclude the proof of Theorem 1.

First, consider any possible valuation $\vec{a} = (a_1, \dots, a_{k+1})$ for agent 1, and instances of the form:

$$S_{\vec{a}} = \begin{bmatrix} a_1 & a_2 & \dots & a_k & a_{k+1} & 0\dots \\ 1 & 1 & \dots & 1 & \epsilon & 0\dots \end{bmatrix}.$$

Agent 1 should receive all items she has a positive value for, otherwise she will deviate to S_1 (and receive all items).

Finally, consider any possible valuation $\vec{b} = (b_1, \dots, b_{k+1})$ for agent 2, and instances of the form:

$$S_{(\vec{a},\vec{b})} = \begin{bmatrix} a_1 & a_2 & \dots & a_k & a_{k+1} & 0\dots \\ b_1 & b_2 & \dots & b_k & b_{k+1} & 0\dots \end{bmatrix}.$$

If agent 2 gets a (strictly) positive fraction of an item j such that $a_j > 0$, this would give an incentive to deviate from $S_{\vec{a}}$, where she receives nothing, to $S_{(\vec{a},\vec{b})}$; a contradiction.

This concludes the proof of Theorem 1. □

4 Conclusion

In this paper we prove that the only truthful and Pareto efficient mechanism for allocating items to additive agents is the serial dictatorship, even among the set of randomized mechanisms. Therefore, as opposed to the closely-related problem of truthful cake-cutting (Mossel and Tamuz 2010; Tao 2021), randomization does not provide a means for escaping strong impossibility results.

In the context of fair division, our results imply that if we want to be truthful and fair (for any reasonable definition of fairness) we *must* be inefficient. This gives a concrete, technical justification for studying truthful and approximately efficient no-money mechanisms à la (Cole, Gkatzelis, and Goel 2013; Guo and Conitzer 2010; Han et al. 2011; Friedman et al. 2019; Abebe et al. 2020).

Another interesting research direction is whether efficiency, fairness are compatible with weaker notions of truthfulness, e.g. similar to (Mennle and Seuken 2014, 2016; Tao 2021).

References

Abebe, R.; Cole, R.; Gkatzelis, V.; and Hartline, J. D. 2020. A truthful cardinal mechanism for one-sided matching. In *Proceedings of the Fourteenth Annual ACM-SIAM Symposium on Discrete Algorithms*, 2096–2113. SIAM.

Amanatidis, G.; Birmpas, G.; Christodoulou, G.; and Markakis, E. 2017. Truthful allocation mechanisms without payments: Characterization and implications on fairness. In *Proceedings of the 2017 ACM Conference on Economics and Computation*, 545–562.

Aziz, H.; Filos-Ratsikas, A.; Chen, J.; Mackenzie, S.; and Mattei, N. 2016. Egalitarianism of Random Assignment Mechanisms. In *Proceedings of the 2016 International Conference on Autonomous Agents & Multiagent Systems*, 1267–1268.

Aziz, H.; and Kasajima, Y. 2017. Impossibilities for probabilistic assignment. *Social Choice and Welfare*, 49(2): 255–275.

Babaioff, M.; Ezra, T.; and Feige, U. 2021. Fair and Truthful Mechanisms for Dichotomous Valuations. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 35, 5119–5126.

Caragiannis, I.; Kurokawa, D.; Moulin, H.; Procaccia, A. D.; Shah, N.; and Wang, J. 2019. The unreasonable fairness of maximum Nash welfare. *ACM Transactions on Economics and Computation (TEAC)*, 7(3): 1–32.

Chen, Y.; Lai, J. K.; Parkes, D. C.; and Procaccia, A. D. 2013. Truth, justice, and cake cutting. *Games and Economic Behavior*, 77(1): 284–297.

Cole, R.; Gkatzelis, V.; and Goel, G. 2013. Mechanism design for fair division: allocating divisible items without payments. In *Proceedings of the fourteenth ACM conference on Electronic commerce*, 251–268.

Ehlers, L.; and Klaus, B. 2003. Coalitional strategy-proof and resource-monotonic solutions for multiple assignment problems. *Social Choice and Welfare*, 21(2): 265–280.

Fotakis, D.; Tsipras, D.; Tzamos, C.; and Zampetakis, E. 2016. Efficient money burning in general domains. *Theory of Computing Systems*, 59(4): 619–640.

Friedman, E.; Ghodsi, A.; and Psomas, C.-A. 2014. Strategyproof allocation of discrete jobs on multiple machines. In *Proceedings of the fifteenth ACM conference on Economics and computation*, 529–546.

Friedman, E. J.; Gkatzelis, V.; Psomas, C.-A.; and Shenker, S. 2019. Fair and efficient memory sharing: Confronting free riders. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 33, 1965–1972.

Ghodsi, A.; Zaharia, M.; Hindman, B.; Konwinski, A.; Shenker, S.; and Stoica, I. 2011. Dominant Resource Fairness: Fair Allocation of Multiple Resource Types. In *Nsdi*, volume 11, 24–24.

Guo, M.; and Conitzer, V. 2010. Strategy-proof allocation of multiple items between two agents without payments or priors. In *AAMAS*, 881–888. Citeseer.

Halpern, D.; Procaccia, A. D.; Psomas, A.; and Shah, N. 2020. Fair division with binary valuations: One rule to rule them all. In *International Conference on Web and Internet Economics*, 370–383. Springer.

Han, L.; Su, C.; Tang, L.; and Zhang, H. 2011. On strategy-proof allocation without payments or priors. In *International Workshop on Internet and Network Economics*, 182–193. Springer.

Hartline, J. D.; and Roughgarden, T. 2008. Optimal mechanism design and money burning. In *Proceedings of the fortieth annual ACM symposium on Theory of computing*, 75–84.

Klaus, B.; and Miyagawa, E. 2002. Strategy-proofness, solidarity, and consistency for multiple assignment problems. *International Journal of Game Theory*, 30(3): 421–435.

Kojima, F. 2009. Random assignment of multiple indivisible objects. *Mathematical Social Sciences*, 57(1): 134–142.

Maya, A.; and Nisan, N. 2012. Incentive compatible two player cake cutting. In *International Workshop on Internet and Network Economics*, 170–183. Springer.

Mennle, T.; and Seuken, S. 2014. An axiomatic approach to characterizing and relaxing strategyproofness of one-sided matching mechanisms. In *Proceedings of the fifteenth ACM conference on Economics and computation*, 37–38.

Mennle, T.; and Seuken, S. 2016. The Pareto frontier for random mechanisms. In *Proceedings of the 2016 ACM Conference on Economics and Computation*, 769–769.

Mossel, E.; and Tamuz, O. 2010. Truthful fair division. In *International Symposium on Algorithmic Game Theory*, 288–299. Springer.

Nesterov, A. S. 2017. Fairness and efficiency in strategy-proof object allocation mechanisms. *Journal of Economic Theory*, 170: 145–168.

Pápai, S. 2000. Strategyproof multiple assignment using quotas. *Review of Economic Design*, 5(1): 91–105.

Procaccia, A. D.; and Tennenholtz, M. 2013. Approximate mechanism design without money. *ACM Transactions on Economics and Computation (TEAC)*, 1(4): 1–26.

Svensson, L.-G. 1999. Strategy-proof allocation of indivisible goods. *Social Choice and Welfare*, 16(4): 557–567.

Tao, B. 2021. On Existence of Truthful Fair Cake Cutting Mechanisms. *arXiv preprint arXiv:2104.07387*.

Varian, H. R. 1973. Equity, envy, and efficiency.