Games and Economic Behavior 110 (2018) 174–193

Contents lists available at ScienceDirect

Games and Economic Behavior

www.elsevier.com/locate/geb

Setting lower bounds on truthfulness

Ahuva Mu'alem [a,*], Michael Schapira [b]

[a] Computer Science Department, Holon Institute of Technology (HIT), Holon, Israel
[b] School of Computer Science and Engineering, Hebrew University, Jerusalem, Israel

ARTICLE INFO

Article history:
Received 31 July 2015
Available online 6 March 2018

JEL classification:
C60
D82

Keywords:
Algorithmic Mechanism Design
Scheduling
Bayesian Incentive-Compatible mechanisms
Truthfulness-in-expectation mechanisms
Randomized truthful mechanisms

ABSTRACT

This paper presents inapproximability results for paradigmatic multi-dimensional truthful mechanism design problems.
We first show a lower bound of $2 - \frac{1}{n}$ for the scheduling problem with n unrelated machines (formulated as a mechanism design problem in the seminal paper of Nisan and Ronen on Algorithmic Mechanism Design). Our lower bound applies to universally-truthful randomized mechanisms, regardless of any computational assumptions on the running time of these mechanisms. Moreover, it holds even for the wider class of truthfulness-in-expectation mechanisms.
We then turn to Bayesian settings and show a lower bound of 1.2 for Bayesian Incentive-Compatible (BIC) mechanisms. No lower bounds for truthful mechanisms in multi-dimensional settings which incorporate randomness were previously known.
Next, we define the workload-minimization problem in networks. We prove lower bounds for the inter-domain routing setting presented by Feigenbaum, Papadimitriou, Sami, and Shenker.
Finally, we prove lower bounds for Max–Min fairness, Min–Max fairness, and envy minimization.

1. Introduction

1.1. Inapproximability issues in Algorithmic Mechanism Design

Mechanism Design is a field of economic theory and game-theory that deals with protocols for optimizing global goals that require interaction with selfish players (Mas-Collel et al., 1995; Osborne and Rubinstein, 1994). *Algorithmic Mechanism Design* (Nisan and Ronen, 2001) combines economic perspectives (e.g., strategic behaviour of the players) with theoretical computer-science perspectives (such as computational-efficiency and approximability).

More formally, Algorithmic Mechanism Design attempts to solve problems of the following nature: given a finite set of *alternatives* $A = \{a, b, c, \ldots\}$, and a set of *strategic players* $N = \{1, \ldots, n\}$, each player i has a *valuation function* $v_i : A \to \mathbb{R}$ that is its own private information. The players are self-interested and only wish to maximize their own utility. The global goal is expressed by a *social choice function* f that assigns to every possible n-tuple of players' valuations $(v_1, \ldots, v_n)$ an alternative $a \in A$. Mechanisms are said to *truthfully implement* a social choice function f, if their outcome for every n-tuple

* Corresponding author.
E-mail addresses: ahumu@yahoo.com (A. Mu'alem), schapiram@huji.ac.il (M. Schapira).

https://doi.org/10.1016/j.geb.2018.02.001
0899-8256/

of players' valuations matches that of f, and if they enforce payments of the different players in a way that motivates truthful reporting of their valuations (no matter what the other players do).[1]

A canonical social choice function is the *utilitarian* function, which aims to maximize the *social welfare*, i.e., to find the alternative $a \in A$ for which the expression $\Sigma_i\ v_i(a)$ is maximized. Another canonical social choice function is the *Max–Min* function (based on the philosophical work of Rawls, 1971): For every n-tuple $(v_1, ..., v_n)$ of valuations, the Max–Min function assigns the alternative $a \in A$ that maximizes the expression $\min_i\ v_i(a)$. Intuitively, the Max–Min function chooses the alternative $a \in A$ in which the least satisfied player has the highest value.

In many computational and economic settings we may wish to implement a utilitarian social choice function in a truthful manner. In such cases we can rely upon a classic result of mechanism design which states that for every utilitarian social choice function there exists a mechanism that truthfully implements it – namely, a member of the celebrated family of *VCG mechanisms* (Vickrey, 1961; Clarke, 1971; Groves, 1973). However, no such general technique is known for non-utilitarian social choice functions such as revenue maximization in auctions (e.g., Fiat et al., 2002), minimizing the makespan in scheduling (e.g., Nisan and Ronen, 2001; Archer and Tardos, 2001; Andelman et al., 2007; Dhangwatnotai et al., 2011; Christodoulou and Kovács, 2013), fair allocation of resources (e.g., Bikhchandani et al., 2006; Bezáková and Dani, 2005; Lipton et al., 2004), etc. In fact, some non-utilitarian social-choice functions cannot be truthfully implemented (Bikhchandani et al., 2006; Nisan and Ronen, 2001). Hence, it is natural to ask how well non-utilitarian social choice functions can be *approximated* in a truthful manner.

In their seminal paper on Algorithmic Mechanism Design (Nisan and Ronen, 2001), Nisan and Ronen formulated the following natural scheduling problem as a mechanism design problem: The global goal is to minimize the makespan of the chosen schedule; i.e., to assign the tasks to the unrelated machines in a way that minimizes the latest finishing time. Obviously, the makespan-minimization social choice function is non-utilitarian and hence *may* not be truthfully implemented by any mechanism. Nisan and Ronen prove that not only is it impossible to minimize the makespan in a truthful manner, but that *any approximation strictly better than* 2 *cannot be achieved by a truthful deterministic mechanism*. Since a non-truthful $(1+\epsilon)$-approximation exists (Horowitz and Sahni, 1976) (assuming constant number of machines), this raises a natural question:

> Can optimality be achieved by incorporating randomness in a truthful manner?
>
> More formally, can near-optimal $(1+\epsilon)$-approximation *truthful* mechanisms *which incorporate randomness* exist for multi-dimensional non-utilitarian settings?

1.2. Our results

In this paper we present lower bounds on the approximability of truthful mechanisms. We obtain the first lower bounds for several canonical non-utilitarian multi-dimensional settings which incorporate randomness.

Section 3 proves several lower bounds for the scheduling problem studied by Nisan and Ronen. In particular, we prove that no universally-truthful randomized mechanism can achieve an approximation ratio better than $2 - \frac{1}{n}$. This nearly matches the known truthful upper bound of 1.58606 for the case in which there are only two machines (Chen et al., 2015). Surprisingly, this lower bound applies even for the substantially weaker notion of truthfulness for randomized mechanisms – truthfulness-in-expectation. Furthermore, we show a lower bound of 1.2 for Bayesian Incentive Compatible mechanisms (also known as Bayesian truthful mechanisms).

> Hence, truthful $(1+\epsilon)$-approximation with randomness is ruled out for the canonical unrelated machines problem (regardless of computational efficiency).

These are the first lower bounds for multi-dimensional settings which incorporate randomness. In fact, to the best of our knowledge these are the first lower bounds for universally truthful mechanisms, truthful-in-expectation mechanisms and Bayesian Incentive Compatible mechanisms in *multi-dimensional* settings in general.

In addition, we show how to prove lower bounds for the important class of *strongly-monotone* deterministic mechanisms. The strongly-monotone property (Lavi et al., 2003) is essentially similar to Arrow's Independence of Irrelevant Alternatives (IIA). Lavi et al. (2003) show that in several canonical domains this property can be assumed without loss of generality. This natural property says that the social choice between two alternatives depends only on the individual valuation difference between these two alternatives.[2] This is another step towards proving the long-standing conjecture of Ronen and Nisan that *no truthful deterministic mechanism can obtain an approximation ratio better than n.*

[1] It is well known (e.g., Mas-Collel et al., 1995) that, without loss of generality, we can limit ourselves to only considering direct-revelation truthful mechanisms. In such mechanisms participants are always rationally motivated to correctly report their private information.

[2] Together with decisiveness, strong-monotonicity essentially implies affine maximization in general combinatorial auctions domains and multi-unit domains (Lavi et al., 2003). In several discrete domains (such as unrestricted integer domains), strong-monotonicity is sufficient for truthful implementability, while weak-monotonicity is not (Mu'alem and Schapira, 2008). For a recent characterization of strongly-monotone scheduling mechanisms see Kovács and Vidali (2015).

In Section 4 we present another multi-dimensional non-utilitarian problem – minimizing the workload in communication networks. This problem arises naturally in the design of routing mechanisms. We study the approximability and inapproximability of this problem in the inter-domain routing setting presented by Feigenbaum et al. (2005).

Finally, in Section 5 we discuss three notions of non-utilitarian fairness – Max–Min fairness, Min–Max fairness, and envy-minimization. We highlight the connections between these notions and the problems studied in this paper and prove several general inapproximability results.

1.3. Related work

In a seminal paper Nisan and Ronen (2001) introduced the field of Algorithmic Mechanism Design. The main problem presented in Nisan and Ronen (2001) to illustrate the novelty of this new area of research was *scheduling with unrelated machines*. Nisan and Ronen explored the approximability of this non-utilitarian multi-dimensional problem and showed a lower bound of $2-\epsilon$ for truthful deterministic mechanisms. For this NP-hard scheduling problem there exist an FPTAS (Horowitz and Sahni, 1976) (assuming constant number of machines) and a polynomial-time 2-approximation algorithm (Karel Lenstra et al., 1990), that are both non-truthful. Additionally, this problem cannot be approximated in polynomial-time within a factor of less than $\frac{3}{2}$ (Karel Lenstra et al., 1990).

In recent years Algorithmic Mechanism Design has been the subject of extensive study (Nisan et al., 2007; Roughgarden, 2016). A substantial amount of this research has focused on single-dimensional settings (see e.g., Lehmann et al., 2002; Archer and Tardos, 2007; Mu'alem and Nisan, 2008; Fiat et al., 2002; Karlin et al., 2005). Nearly-optimal truthful mechanisms were designed for the single-dimensional problem of minimum makespan for scheduling tasks on related machines (Archer and Tardos, 2001; Andelman et al., 2007; Dhangwatnotai et al., 2011; Christodoulou and Kovács, 2013). The exploration of truthful mechanisms for multi-dimensional settings has, arguably, revolved mainly around the problem of welfare maximization in Multiple-Object auctions (Cramton et al., 2006; Nisan, 2015), that has gained the status of the paradigmatic problem of this field. As this is a utilitarian problem, it can be optimally and truthfully implemented by a VCG mechanism. However, it has been shown that the social welfare in combinatorial auctions cannot be maximized (or even closely approximated) in polynomial time (Lehmann et al., 2002; Nisan and Segal, 2006). As algorithmic mechanism design seeks time-efficient implementations, the main challenge faced by researchers is how to devise truthful polynomial-time mechanisms that approximately maximize the social welfare in combinatorial auctions (Lavi and Swamy, 2011; Dobzinski et al., 2010, 2012; Bartal et al., 2003; Holzman et al., 2004; Dobzinski and Dughmi, 2013).

There are few *inapproximability* results for truthful mechanisms. This is particularly true in multi-dimensional settings. Other than Nisan and Ronen's previously mentioned $2-\epsilon$ lower bound, the following inapproximability results are known for multi-dimensional combinatorial auctions:

- Lavi et al. (2003) proved that it is impossible to achieve an approximation factor better than 2 for a polynomial time *deterministic* truthful mechanism for a multi-unit auction between two players that always allocates all units.
- Dobzinski and Nisan (2010) proved inapproximability results of polynomial-time VCG mechanisms for multi-unit auctions.
- Dobzinski and Vondrák (2016) bound the power of polynomial-time universally-truthful randomized mechanisms in combinatorial auctions with submodular valuations.
- Several papers use VC dimensionality to prove inapproximability results of deterministic truthful mechanisms for combinatorial auctions (Daniely et al., 2015 and references therein).

We contribute to this ongoing research by presenting methods for deriving the first lower bounds for multi-dimensional non-utilitarian settings that apply to truthful mechanisms with randomness. Our lower bounds do not require any assumptions on the running-time of the mechanisms.

Our results rely heavily on the work of Bikhchandani et al. (2006). They characterize truthfulness in multi-dimensional settings by showing that any truthful deterministic mechanism must maintain a certain *weak-monotonicity* property. Using this characterization, Bikhchandani et al. (2006) showed that while welfare maximization can be truthfully implemented in combinatorial auctions, the Max–Min social choice function cannot be truthfully implemented, even in a very restricted type of combinatorial auctions. The weak-monotonicity property (and several of its extensions) will play a major role in our inapproximability proofs.

Makespan in multi-parameter settings Nisan and Ronen (2001) present a truthful deterministic mechanism that obtains an n-approximation. They showed that no multi-dimensional truthful deterministic mechanism can achieve an approximation ratio strictly better than 2 (and also strengthened this lower bound to n for two specific classes of deterministic mechanisms whose payments satisfy some local properties). This lower bound has been improved and extended in a series of results. Christodoulou et al. (2009) showed a lower bound of $1+\sqrt{2}$ for 3 machines.[3] Koutsoupias and Vidali (2013) showed a lower

[3] Gamzu (2011) gives a considerably simpler alternative proof.

bound of $1+\phi \approx 2.618$ for truthful deterministic mechanisms with $n \to \infty$ machines (ϕ is the golden ratio). An optimal lower bound of n for *anonymous* truthful mechanisms is shown in Ashlagi et al. (2012).

For the case of two machines, Lehmann, Nisan, and Ronen exhibit a universally-truthful randomized mechanism that obtains an approximation of $\frac{7}{4}$ (Nisan and Ronen, 2001). This upper bound was later improved to 1.6737 (Pinyan and Yu, 2008) and then improved to 1.58606 (Chen et al., 2015). Dobzinski and Sundararajan show that every mechanism in the support of any universally-truthful randomized mechanism for two machines that obtains finite approximation ratio must be *task independent* (Dobzinski and Sundararajan, 2008). That is, the mechanism must assign each task separately from the others.[4]

For Bayesian settings, Daskalakis and Matthew Weinberg (2015) recently show that there is a polynomial time 2-approximately optimal mechanism for makespan minimization for unrelated machines. The approximation factor in this result is with respect to the optimal Bayesian Incentive Compatible mechanism (rather than the optimal algorithm for makespan minimization) and thus is not directly comparable with the lower bound presented in our paper.

Giannakopoulos and Kyropoulou (2015) show that the VCG mechanism achieves an approximation ratio of $O(\frac{\ln n}{\ln \ln n})$ when the processing times of the tasks are independent random variables, identical across machines. This essentially improves on the previously best known bound of $O(\frac{m}{n})$ given by Chawla et al. (2013).

For fractional settings (where the mechanism is allowed to split the task across several machines) Christodoulou et al. (2010) present a truthful task-independent $\frac{n+1}{2}$-approximation mechanism. To compliment this result they show a lower bound of $2-\frac{1}{n}$ for any fractional truthful mechanism. They give a lower bound of $\frac{n+1}{2}$ for task-independent fractional truthful mechanisms. Pinyan (2009) gives a lower bound of 1.5625 for scale-free universally-truthful randomized mechanism for two machines, where the allocation of tasks depends only on relative costs, not on scale. Lavi and Swamy (2009) and Yu (2009) show truthful mechanisms in a multi-dimensional scheduling special setting where the processing time of a task on each machine is either 'low' or 'high'.

In several interesting settings, truthful mechanisms are essentially equivalent to mechanisms that select envy-free allocations with the smallest supporting price vectors (Demange and Gale, 1985). A natural question to ask is whether envy-free pricing techniques can improve the current striking approximability and inapproximability bounds for truthful mechanisms. Mu'alem (2014) observed that the optimal envy-free mechanisms for two machines are not truthful.

In a follow-up work, Gamzu (2011) improved our truthful lower bound for minimizing the workload in inter-domain routing (from $\phi = \frac{1+\sqrt{5}}{2} \approx 1.618$ to 2) and our universally-truthful randomized lower bound (from $\frac{3+\sqrt{5}}{4} \approx 1.309$ to 2).

1.4. The organization of the paper

In Section 3 we study several lower bounds on truthfulness of *the scheduling problem with unrelated machines*. In Section 4 we study the problem of *workload-minimization in networks*. In Section 5 we study several notions of non-utilitarian fairness.

2. Preliminaries

We consider the standard mechanism design setting: There are n players, and a finite set of alternatives A. Each player $i \in [n]$ has a private valuation function $v_i \in V_i$ that assigns a non-negative real value to every $a \in A$ (the higher the value of the alternative the more desirable it is).

A (deterministic) mechanism $M(f, p)$ consists of a deterministic social choice function $f : V \to A$ (representing some global goal, e.g., makespan as defined below), and a payment function $p_i : V \to R$ for each player i. Each player i is simultaneously asked to report a valuation v_i (possibly deviating from his private valuation), and the mechanism then computes the outcome $f(v)$ and charges price $p_i(v)$ to player i (notice that payment might be negative or positive). The (quasi-linear) *utility* that player i derives by declaring valuation v'_i is $v_i(f(v'_i, v_{-i})) - p_i(v'_i, v_{-i})$ (assuming player i's private valuation function is v_i). Each player aims to maximize his own utility.

A mechanism $M(f, p)$ is called truthful if for every player i, for every $v_{-i} \in V_{-i}$, and for every $v_i, v'_i \in V_i$,

$$v_i(f(v_i, v_{-i})) - p_i(v_i, v_{-i}) \geq v_i(f(v'_i, v_{-i})) - p_i(v'_i, v_{-i}).$$

That is, a mechanism is truthful if no player can ever improve its utility by misreporting his private valuation to the mechanism (no matter what the other players do).

Remark 1. The above setting refers to value scenarios where players report their valuations to the mechanism and willing to maximize their values minus their payment to the mechanism. In what follows we will mainly consider cost scenarios where players report their costs (rather than valuations) to the mechanism and analogously are willing to minimize their costs minus the payment made to them by the mechanism. We will slightly abuse notation by letting v_i refer both to a valuation function, and to a cost function, but the meaning will be clear from the context.

[4] See Vidali (2009, 2011) for further characterizations of scheduling mechanisms.

The approximation ratio of a mechanism optimizing a global minimization goal (e.g., makespan) is defined to be the worst-case ratio (over all possible valuations $v \in V$) between the goal value of the chosen alternative and the optimal value.

A *universally-truthful randomized mechanism* is a probability distribution over truthful mechanisms. Formally, for every $v \in V$ the universally-truthful randomized mechanism produces a distribution $D(v)$ over deterministic truthful mechanisms and outputs a deterministic mechanism drawn from this distribution.

A randomized social choice function f is a function from n-tuples of players' valuations to probability distributions over the set of alternatives A. A randomized mechanism consists of a randomized social choice function $f : V \to A$, and a payment function $p_i : V \to R$ for each player i. A randomized mechanism (f, p) is *truthful-in-expectation* if for every player i, for every $v_{-i} \in V_{-i}$, and for every $v_i, v_i' \in V_i$,

$$E[v_i(f(v_i, v_{-i})) - p_i(v_i, v_{-i})] \geq E[v_i(f(v_i', v_{-i})) - p_i(v_i', v_{-i})].$$

Thus, if a mechanism is truthful-in-expectation then the *expected* utility of a player is maximized when he declares his true private valuation v_i (no matter what the other players do). Here, we assume *risk neutral* players aiming to maximize the expected difference between their true private valuation and their total payment (the expectation is taken over any randomness in the mechanism).

The approximation ratio of a mechanism with randomness optimizing a global minimization goal (e.g., makespan) is defined to be the worst-case ratio (over all possible valuations $v \in V$) between the *expected* goal value of the chosen alternative and the optimal value.

In a Bayesian (a.k.a. stochastic) setting the private valuations of each player i is drawn independently from D_i. The product distribution $D = D_1 \times \cdots \times D_n$ is assumed to be public knowledge. A mechanism $M(f, p)$ is Bayesian Incentive Compatible (given the public knowledge distribution D) if for all i

$$E_{v_{-i} \sim D_{-i}}[v_i(f(v_i, v_{-i})) - p_i(v_i, v_{-i})] \geq E_{v_{-i} \sim D_{-i}}\left[v_i(f(v_i', v_{-i})) - p_i(v_i', v_{-i})\right].$$

That is, player i's expected utility from reporting his true valuation v_i is no less than his expected utility from reporting a different valuation v_i' when others' true valuations are drawn from the product distribution D_{-i}. Here, we assume *risk neutral* players aiming to maximize the expected difference between their true private valuation and their total payment (the expectation is taken over the randomness in other players' valuations). Thus, if a mechanism is Bayesian Incentive Compatible then the *expected* utility of a player is maximized when he declares his true private valuation v_i (assuming all other players truthfully report their valuations).

The approximation ratio of a Bayesian mechanism optimizing a global minimization goal (e.g., makespan) is defined to be the ratio between the *expected* goal value of the chosen alternative and the *expected* optimal value, where the expectations are taken over D.

3. Scheduling problem

In this section we show how to derive lower bounds for the scheduling problem with n unrelated machines. Nisan and Ronen (2001) exhibited a truthful n-approximation deterministic mechanism for this problem. Their mechanism is a VCG mechanism, and can easily be shown to be strongly-monotone (see Subsection 3.2 for a formal definition of strong-monotonicity). This mechanism can be viewed as auctioning each task separately in a Vickrey auction.[5] It is straight forward to show that any VCG mechanism with a *deterministic* tie-breaking rule among alternatives is strongly-monotone. They also proved a lower bound of $2 - \epsilon$ for truthful deterministic mechanisms that applies even when there are only two machines and is tight for this case. However, Nisan and Ronen (2001) conjecture that their lower bound is not tight in general, and that *any truthful deterministic mechanism cannot obtain an approximation ratio better than n*.

For the case of two machines, Nisan and Ronen show that randomness helps get an approximation ratio better than 2. They present a universally-truthful randomized mechanism that has an approximation ratio of $\frac{7}{4}$. We generalize their result by designing a universally-truthful randomized mechanism that obtains an approximation ratio of $\frac{7n}{8}$ (see Appendix A.1). *Thus, we prove that randomness achieves better performances than the known truthful deterministic n upper bound for any number of machines*. In Subsection 3.2 we show ways of proving lower bounds for truthful deterministic mechanisms. Using these methods we provide a simple and shorter proof for Nisan and Ronen's $2 - \epsilon$ lower bound. Our proof (unlike the original) relies on exploiting the *weak-monotonicity* property defined in Bikhchandani et al. (2006). Subsection 3.2 also aid us in deriving a stronger lower bound for the important classes of *strongly-monotone* deterministic mechanisms. We note, that the mechanism in Nisan and Ronen (2001), which is the best currently known deterministic mechanism for the scheduling problem, is contained in this class. We prove that no approximation ratio better than n is possible for this class of strongly-monotone deterministic mechanisms (*thus making another step towards proving the long-standing conjecture of Nisan and Ronen, 2001*).

[5] Christodoulou et al. show that for two machines VCG is the unique mechanism achieving the optimal approximation of 2 (Christodoulou et al., 2008).

After discussing lower bounds for truthful deterministic mechanisms we turn our attention mechanisms with randomness. There are two possible definitions for the truthfulness of a randomized mechanism (Dobzinski et al., 2012; Nisan and Ronen, 2001). The first and stronger one is that of *universal truthfulness* that defines a truthful randomized mechanism as a probability distribution over truthful deterministic mechanisms. Thus, this definition requires that for *any* toss of the random coins made by the mechanism, the players still maximize their utility by reporting their true valuations. A considerably weaker definition of truthfulness is that of *truthfulness-in-expectation*. This definition only requires that players maximize their *expected* utility, where the expectation is over the random choices of the mechanism (but still for every behaviour of the other players). Unlike universally truthful mechanisms, truthful-in-expectation mechanisms only motivate *risk-neutral* bidders to act truthfully. Risk-averse bidders may benefit from strategic behaviour. In addition, truthful-in-expectation mechanisms induce truthful behaviour only as long as players have no information about the outcomes of the random coin flips before they need to act.

In Subsection 3.3 we prove the first lower bound on the approximability of universally-truthful randomized mechanisms in multi-dimensional settings. Namely, we show that any universally-truthful randomized mechanism for the scheduling problem cannot achieve an approximation ratio better than $2 - \frac{1}{n}$. This lower bound nearly matches the universally truthful 1.58606 upper bound for the case of two machines (Chen et al., 2015). To prove this lower bound, we apply Yao's powerful principle (Yao, 1977). Our proof for the $2 - \epsilon$ lower bound for deterministic mechanisms (in Subsection 3.2) serves as a building block in the proof of this lower bound.

In Subsection 3.4 we strengthen this result by proving that the same lower bound holds even when one is willing to settle for truthfulness-in-expectation. Our proof relies on some of the ideas that appear in the proof of the previous lower bound but takes a different approach. In particular, we generalize the weak-monotonicity requirement to fit the class of truthful randomized mechanisms, and explore the implications of this extended monotonicity on the probability distributions over allocations generated by such mechanisms.

In Subsection 3.5 we turn to the notion of Bayesian Incentive Compatible mechanisms, where players' valuations are drawn from a distribution that is public knowledge, and show a lower bound of 1.2 for deterministic Bayesian mechanisms.

3.1. The setting

We consider the standard unrelated strategic machines setting (Nisan and Ronen, 2001): There are m tasks that are to be assigned on n machines, where each task must be assigned to exactly one machine. Every machine i is a strategic player with an arbitrary valuation function $v_i : 2^{[m]} \to \mathbb{R}_+$ such that $v_i(\{j\})$ (or $v_i(j)$, for short) denotes the private *cost* of task $j \in [m]$ on machine i.

One can think of the private cost of task j on machine i as the time it takes i to complete j. For every subset of tasks $S \subseteq [m]$, $v_i(S) = \Sigma_{j \in S}\ v_i(j)$. That is, the total cost of a set of tasks on machine i is the sum of the costs of the individual tasks on that machine. Each machine wishes to minimize the cost of the set of tasks assigned to it minus the payment made to it by the mechanism. If the mechanism is truthful then machine i can never improve its utility by misreporting its private cost function v_i to the mechanism (no matter what the other machines do).

The set of alternatives A contains all possible allocations of tasks to the machines, where all tasks must be assigned, and each task is assigned to exactly one machine. The global goal is minimizing the makespan of the chosen allocation. I.e., find an allocation of tasks $a \in A$ to minimize the expression

$$\max\{v_1(a_1), \ldots, v_n(a_n)\}.$$

3.2. Lower bounds for truthful deterministic mechanisms

Bikhchandani et al. (2006) formally define the weak-monotonicity property for deterministic mechanisms: Consider an Algorithmic Mechanism Design setting with n strategic players. Before we present the formal definition we will require the following

Definition 1. Let M be a deterministic mechanism. Let $i \in [n]$ and let $v = (v_1, \ldots, v_n)$ be an n-tuple of players' valuations. Let v'_i be a valuation function. Denote by a the alternative that M outputs for v and by b the alternative that M outputs for (v'_i, v_{-i}). The mechanism M is said to be weakly monotone if for all such i, v, and v'_i it holds that:

$$v_i(a) + v'_i(b) \geq v'_i(a) + v_i(b).$$

Remark 2. This definition of weak-monotonicity is for value scenarios in which each player wishes to maximize the difference between his valuation and his total payment. In cost scenarios in which players have costs (such as the scheduling of unrelated machines problem, the workload minimization problem, and the Min–Max fairness considered in this paper) the inequality (in Definition 1) is in the other direction.

Bikhchandani et al. (2006) prove that any truthful deterministic mechanism must be weakly-monotone. For completeness, we present this simple proof.

Lemma 1. *Any truthful deterministic mechanism must be weakly-monotone.*

Proof. Let M be a truthful deterministic mechanism. Let $i \in [n]$ and let $v = (v_1, ..., v_n)$ be an n-tuple of players' valuations. Let v'_i be a valuation function. Denote by a the alternative that M outputs for v and by b the alternative that M outputs for (v'_i, v_{-i}). Consider player i. It is well known that the price a player is charged by the mechanism to ensure his truthfulness cannot depend on the player's report. Specifically, the payment of player i in a and b is a function of v_{-i} and of a and b, respectively. We denote by $p_i(v_{-i}, a)$ and by $p_i(v_{-i}, b)$ i's payment in a and b, respectively. It must hold that $v_i(a) - p_i(v_{-i}, a) \geq v_i(b) - p_i(v_{-i}, b)$ (for otherwise, if i's valuation function is v_i, he would have an incentive to declare his valuation to be v'_i). Similarly, $v'_i(b) - p_i(v_{-i}, b) \geq v'_i(a) - p_i(v_{-i}, a)$. By adding these two inequalities we reach the weak-monotonicity requirement. □

Relying on the weak-monotonicity property we provide an alternative proof for the $2 - \epsilon$ lower bound of Nisan and Ronen (2001) for the scheduling problem with unrelated machines. Our proof shows that any deterministic mechanism that achieves an approximation ratio better than 2 violates the weak-monotonicity property.

Theorem 1. *Any weakly-monotone mechanism cannot achieve an approximation ratio better than* 2.

Proof. Let ϵ be an arbitrarily small positive real number. Consider the scheduling problem with two machines and three tasks. For every machine $i = 1, 2$ we define two possible valuation functions v_i and v'_i:

$$v_i(t) = \begin{cases} 1 & t = i \text{ or } t = 3 \\ 100 & \text{otherwise} \end{cases}$$

$$v'_i(t) = \begin{cases} 0 & t = i \\ 1 + \epsilon & t = 3 \\ 100 & \text{otherwise.} \end{cases}$$

Let M be a deterministic, weakly-monotone, mechanism that achieves an approximation ratio better than 2. Then, when players 1 and 2 have the valuations v_1 and v_2 respectively, M must assign task 1 to player 1, task 2 to player 2, and can choose to which player to assign task 3 (because the optimal makespan is 2 and any other assignment results in a makespan of at least 100). W.l.o.g. assume that M assigns task 3 to player 2. Now, consider the instance with players' valuations (v'_1, v_2). Notice that the only task-allocation that guarantees an approximation ratio better than 2 is assigning tasks 1 and 2 to players 1 and 2 respectively, and task 3 to player 1. However, this turns out to be a violation of the weak-monotonicity requirement. Weak-monotonicity, in this case, dictates that for every player $i = 1, 2$ it must hold $v_i(a) + v'_i(b) \leq v'_i(a) + v_i(b)$. However, if we look at player 1 we find that $1 + (1 + \epsilon) = v_1(1) + v'_1(\{1, 2\}) > v'_1(1) + v_1(\{1, 2\}) = 2$. A contradiction. □

Lavi et al. (2003) present and study another property – strong-monotonicity. They show that in some canonical settings this property can be assumed without loss of generality. Strong-monotonicity is the strict version of weak-monotonicity (Definition 1). It says essentially that the social choice between two alternatives depend only on the individual valuation difference between these two alternatives.

Definition 2. Let M be a deterministic mechanism. Let $i \in [n]$ and let $v = (v_1, ..., v_n)$ be an n-tuple of players' valuations. Let v'_i be a valuation function. Denote by a the alternative that M outputs for v and by b the alternative that M outputs for (v'_i, v_{-i}). The mechanism M is said to be strongly monotone if for all such i, v, and v'_i it holds that: If $a \neq b$, then

$$v_i(a) + v'_i(b) > v'_i(a) + v_i(b).$$

Here, too, the inequality is in the other direction if players have costs rather than values. We prove that no member of the class of strongly-monotone mechanisms can obtain an approximation better than n for the scheduling problem (even for the case of zero/one valuations). The idea at the heart of our proof of Theorem 2 is an iterative use of the strong-monotonicity property to construct an instance of the problem for which the allocation generated by the mechanisms is very far from optimal.

Theorem 2. *Any strongly-monotone mechanism cannot obtain an approximation ratio better than n.*

Proof. Consider an instance of the scheduling problem with n machines and $m = n^2$ tasks. Let M be a deterministic mechanism for which the strong-monotonicity property holds. Let $\mathcal{I}$ be the instance of the scheduling problem in which every machine i has a valuation function v_i such that $v_i(j) = 1$ for all $j \in [m]$. Denote by $S = (S_1, ..., S_n)$ the allocation of tasks produced by M for the instance $\mathcal{I}$. It must be that there is some machine r such that $|S_r| \geq n$. Without loss of generality let $r = n$.

We will now create a new instance $\mathcal{I}^1$ by altering the valuation function of machine 1 to v'_1 while leaving all the other valuation functions unchanged (in case $S_1 = \emptyset$ we skip this part). That is, machine 1 will have the valuation function v'_1:

$$v'_1(t) = \begin{cases} 0 & t \in S_1 \\ 1 & t \notin S_1 \end{cases}$$

and every other machine $i \neq 1$ will have a valuation function v_i. Denote by $T = (T_1, ..., T_n)$ the allocation M generates for $\mathcal{I}^1$. The first step of the proof is showing that $S_1 = T_1$. This is guaranteed by the strong-monotonicity of M. Assume, by contradiction that $S_1 \neq T_1$. The strong-monotonicity property ensures that $v_1(S_1) + v'_1(T_1) < v_1(T_1) + v'_1(S_1)$. By assigning values we have:

$$|S_1| + |T_1 \setminus S_1| < |T_1| + 0.$$

Observe, that $|S_1| + |T_1 \setminus S_1| - |T_1| = |S_1 \setminus T_1|$, therefore:

$$|S_1 \setminus T_1| < 0.$$

A contradiction.

We shall now prove that not only does S_1 equal T_1, but in fact $S_i = T_i$ for every i. Since $S_1 = T_1$ it must be that $v_1(S_1) + v'_1(T_1) = v'_1(S_1) + v_1(T_1)$. However, the strong-monotonicity property (with respect to (v_1, v_{-1}) and (v'_1, v_{-1})) dictates that if this is true then $S = T$.

In an analogous manner we shall now turn the valuation function of machine 2 into v'_2 while keeping all the other valuation functions in $\mathcal{I}^1$ unchanged (in case $S_2 = \emptyset$ we skip this part). That is, the valuation function of machine 2 is changed into:

$$v'_2(t) = \begin{cases} 0 & t \in S_2 \\ 1 & t \notin S_2. \end{cases}$$

Similar arguments show that the allocation produced by the M for this new instance will remain S. We can now iteratively continue to change the valuation functions of machines $3, ..., n-1$ into $v'_3, ..., v'_{n-1}$ respectively, without changing the allocation the mechanism generates for these new instances. After performing this, we are left with an instance in which every machine $i \in [n-1]$ has the valuation function v'_i, and machine n has the valuation function v_n. We have shown that the allocation generated by M for this instance is S. Recall that $|S_n| \geq n$. Fix some $R \subseteq S_n$ such that $|R| = n$. We will now create a new instance $\mathcal{I}^n$ from the previous one by only altering the valuation function v_n into the following valuation function v'_n:

$$v'_n(t) = \begin{cases} 0 & t \in S_n \setminus R \\ 1 & \text{otherwise.} \end{cases}$$

By applying similar arguments to the ones used before, one can show that the allocation generated by M when given the instance $\mathcal{I}^n$ remains S. Observe that the finishing time of S for $\mathcal{I}^n$ is n because all the tasks in R are assigned to machine n. Also notice that the finishing time of the optimal allocation of tasks for $\mathcal{I}^n$ is precisely 1 (by assigning the i'th task in R to machine i, for $i = 1, ..., n$). The theorem follows. $\square$

3.3. A lower bound for universally-truthful randomized mechanisms

We now derive lower bounds for universally truthful mechanisms, based on Yao's principle (Yao, 1977). Consider a zero-sum game with two players. Let the row player's strategies be the various different instances of a specific problem, and let the column player's strategies be all the deterministic truthful mechanisms for solving that problem. Let entry g_{ij} in the matrix G depicting the game be the approximation ratio obtained by the algorithm of column j when given the instance of row i.

Recall that every universally-truthful randomized mechanism is a probability distribution over deterministic truthful mechanisms. The straight-forward approach for proving a lower bound for such mechanisms is to find an instance of the problem on which every such mechanism cannot achieve (in expectation) a certain approximation factor. By applying the well known Minimax Theorem to the game described above we get that an alternate and just as powerful way for setting lower bounds is to show that there is a probability distribution over instances on which any deterministic mechanism cannot obtain (in expectation) a certain approximation ratio.

We prove a $2 - \frac{1}{n}$ lower bound for universally truthful mechanisms for the scheduling problem. Our proof is based on finding a probability distribution over instances of the scheduling problem for which no deterministic truthful mechanism can provide an approximation ratio better than $2 - \frac{1}{n}$. To show this, we shall exploit the weak-monotonicity property of truthful deterministic mechanisms (as discussed in Subsection 3.2).

Theorem 3. *Any universally truthful mechanism cannot achieve an approximation ratio better than* $2 - \frac{1}{n}$.

Proof. Let ϵ be an arbitrarily small positive real number. Consider the scheduling problem with n machines and $m = n+1$ tasks. For every machine i we define two possible valuation functions:

$$v_i(t) = \begin{cases} 1 & t = i \text{ or } t = n+1 \\ \frac{4}{\epsilon} & \text{otherwise} \end{cases}$$

and

$$v'_i(t) = \begin{cases} 0 & t = i \\ 1+\epsilon & t = n+1 \\ \frac{4}{\epsilon} & \text{otherwise.} \end{cases}$$

Let $\mathcal{I}$ be the instance in which the valuation function of every machine i is v_i. For every j, let $\mathcal{I}^j$ be the instance in which every machine $i \neq j$ has the valuation function v_i, and machine j has the valuation function v'_j. We are now ready to define the probability distribution P over instances: instance $\mathcal{I}$ is assigned probability ϵ, and for every j instance $\mathcal{I}^j$ is picked with probability $\frac{1-\epsilon}{n}$.

We now need to show that any deterministic truthful mechanism M cannot achieve an approximation ratio better than $2 - \frac{1}{n}$ on P. Let T^j be the allocation of the $n+1$ tasks to the n machines in which every machine i gets task i, and machine j is also assigned task $n+1$. Observe, that T^j is the optimal allocation of tasks for the instance $\mathcal{I}^j$. Also observe, that while the finishing time of the allocation T^j for the instance $\mathcal{I}^j$ is $1+\epsilon$, the finishing time of any other allocation of tasks is at least 2. We shall denote the allocation that M outputs for the instance $\mathcal{I}$ by $M(\mathcal{I})$. Similarly, we shall denote the allocation that M outputs for the instance $\mathcal{I}^j$ by $M(\mathcal{I}^j)$ (for every $j \in [n]$). We will now examine two distinct cases: The case in which $M(\mathcal{I}) \neq T^r$ for any $r \in [n]$, and the case that $M(\mathcal{I}) = T^r$ for some $r \in [n]$.

Observe, that in the first case the finishing time is at least $\frac{4}{\epsilon}$ while the optimal finishing time is 2. Since instance $\mathcal{I}$ appears in P with probability ϵ we have that M's expected approximation ratio is at least

$$\frac{\frac{4}{\epsilon} \times \epsilon + (1+\epsilon) \times (1-\epsilon)}{2 \times \epsilon + (1+\epsilon) \times (1-\epsilon)} > 2 - \frac{1}{n}.$$

We are left with the case in which $M(\mathcal{I}) = T^r$ for some $r \in [n]$. Consider an instance $\mathcal{I}^j$ such that $j \neq r$. The following lemma states that M will not output the optimal allocation for $\mathcal{I}^j$ (that is, T^j).

Lemma 2. *If $M(\mathcal{I}) = T^r$ for some $r \in [n]$, then for every $j \neq r$ $M(\mathcal{I}^j) \neq T^j$.*

Proof. Let $j \neq r$. Let us assume by contradiction that $M(\mathcal{I}^j) = T^j$. The weak-monotonicity property dictates that $v_j(j) + v'_j(\{j, n+1\}) \leq v'_j(j) + v_j(\{j, n+1\})$. By assigning values we have that $1 + (1+\epsilon) \leq (1+1)$, and reach a contradiction. □

From Lemma 2 we learn that if $M(\mathcal{I}) = T^r$ (for some $r \in [n]$) then we have that $M(\mathcal{I}^j)$ (for every $j \neq r$) is an allocation that is not the optimal one (i.e., not T^j). In fact (as mentioned before), any allocation that M outputs given $\mathcal{I}^j$ will have a finishing time of at least 2, while the optimal allocation (T^j) has a finishing time of $1+\epsilon$. The expected approximation ratio of M for P is therefore at least

$$\frac{2 \cdot \frac{(n-1)\cdot(1-\epsilon)}{n} + (1+\epsilon) \cdot \frac{1-\epsilon}{n} + 2\epsilon}{(1+\epsilon)\cdot(1-\epsilon) + 2\epsilon} \geq \frac{(2-\frac{1}{n})\cdot(1-\epsilon)}{(1+\epsilon)\cdot(1-\epsilon)+2\epsilon}$$

Since $\lim_{\epsilon \to 0} \frac{(2-\frac{1}{n})\cdot(1-\epsilon)}{(1+\epsilon)\cdot(1-\epsilon)+2\epsilon} = 2 - \frac{1}{n}$, the theorem follows. □

3.4. A lower bound for mechanisms that are truthful-in-expectation

After handling the case of universally truthful mechanisms we now turn to the weaker notion of truthfulness-in-expectation. Any such mechanism can be regarded as a mechanism that for every instance of a problem produces a probability distribution over possible alternatives. In this subsection we consider risk neutral players with quasi-linear utility functions.

We start by generalizing the weak-monotonicity definition. We consider the expected value with respect to a distribution over the set of alternatives A:

Definition 3. Let v be a valuation function. We define the extended valuation function V_v as follows. For every probability distribution P over the set of alternatives A

$$V_v(P) = \Sigma_{a \in A} \, Pr_P[a] \times v(a).$$

Arguments similar to those of Lemma 1 show that randomized mechanisms that are truthful-in-expectations must be weakly monotone (given the new definition of the valuation functions). This *extended weak-monotonicity* is equivalent to the *monotonicity-in-expectation* property defined by Lavi and Swamy (2011).

Definition 4. Let M be a randomized mechanism. Let $i \in [n]$ and let $v = (v_1, ..., v_n)$ be an n-tuple of players' valuations. Let v'_i be a valuation function. Denote by P the distribution over alternatives that M outputs for v and by Q the distribution over alternatives that M outputs for (v'_i, v_{-i}). The mechanism M is said to be weakly monotone in the extended sense if for all such i, v, and v'_i it holds that: $V_{v_i}(P) + V_{v'_i}(Q) \geq V_{v'_i}(P) + V_{v_i}(Q)$.

Lemma 3 *(Lavi and Swamy, 2011). Any truthful-in-expectation mechanism must be weakly-monotone in the extended sense.*

We can exploit this extended definition of weak-monotonicity to prove inapproximability results. We show how this is done by strengthening our $2 - \frac{1}{n}$ lower bound for universally-truthful randomized mechanisms by showing that it applies even for the case of truthfulness-in-expectation.

A key element in the proof of Theorem 4 is the observation that instead of regarding a randomized mechanism for the scheduling problem as generating probability distributions over allocations of tasks, it can be regarded as generating, for each task, a probability distribution over the machines it is assigned to by the mechanism. This different view of a randomized mechanism for this specific problem, enables us to better analyse the contribution of each task to the expected makespan.

The main lemma in the proof of Theorem 4, namely Lemma 4, makes use of this fact together with the extended weak-monotonicity condition. Lemma 4 essentially proves that for two carefully chosen instances of the problem, the probability that a specific task is assigned to a specific machine in one of the instances, cannot be considerably higher than the probability it is assigned to the same machine in the other. Thus, we show that even though allocating this task to that machine in one of the instances leads to a good approximation, any truthful-in-expectation mechanism will fail to do so.

Theorem 4. *Any mechanism that is weakly-monotone in the extended sense cannot achieve an approximation ratio better than* $2 - \frac{1}{n}$.

Proof. Let ϵ be an arbitrarily small positive real number. Consider the scheduling problem with n machines and $m = n + 1$ tasks. For every machine i we define two possible valuation functions:

$$v_i(t) = \begin{cases} 1 & t = i \text{ or } t = n+1 \\ \frac{4}{\epsilon^2} & \text{otherwise} \end{cases} \tag{1}$$

and

$$v'_i(t) = \begin{cases} 0 & t = i \\ 1+\epsilon & t = n+1 \\ \frac{4}{\epsilon^2} & \text{otherwise.} \end{cases} \tag{2}$$

Let $\mathcal{I}$ be the instance in which the valuation function of every machine i is v_i. For every $j \in [n]$ let $\mathcal{I}^j$ be the instance in which every machine $i \neq j$ has the valuation function v_i, and machine j has the valuation function v'_j. Let T^j be the allocation of the $n+1$ tasks to the n machines in which every machine i gets task i, and machine j is also assigned task $n+1$.

Let M be a mechanism that is weakly-monotone in the extended sense. We shall denote by P the distribution over all possible allocations produced by M when given instance $\mathcal{I}$, and by P^j the distribution over all possible allocations M produces when given instance $\mathcal{I}^j$. Let R be some distribution over the possible allocations. Fix a machine i and a task t, we define $p_{i,t}(R)$ to be the probability that machine i gets task t given R. Formally, $p_{i,t}(R) = \Sigma_{a|t\in a_i} Pr_R[a]$. Observe that $V_{v_i}(R) = \Sigma_{t\in[m]}\ p_{i,t}(R)v_i(t)$ and $V_{v'_i}(R) = \Sigma_{t\in[m]}\ p_{i,t}(R)v'_i(t)$.

We are now ready to prove the theorem. In order to do so, we prove that for every mechanism M, as defined above, one can find an instance of the scheduling problem for which M fails to give an approximation ratio better than $2 - \frac{1}{n}$. Consider instance $\mathcal{I}$.

If $p_{i,i}(P) < 1 - \epsilon^2$ (for some $i \in [n]$) then machine i does not get task i with probability of at least ϵ^2. However, when machine i does not get task i, the finishing time of a schedule for $\mathcal{I}$ cannot be less than $\frac{4}{\epsilon^2}$, while the optimal finish time is 2. Therefore, with probability of at least ϵ^2 the makespan of the algorithm is at least $\frac{4}{\epsilon^2}$. If this is the case, the approximation ratio is at least 2 (and the theorem follows). Hence, from now on we will only deal with the case in which for every $i \in [n]$,

$$p_{i,i}(P) \geq 1 - \epsilon^2.$$

Let r be some machine such that $p_{r,n+1}(P) \leq \frac{1}{n}$ (recall that $\Sigma_{i\in[n]}\ p_{i,t}(P) = 1$ for every task $t \in [m]$). Intuitively, r is a machine that is hardly assigned task $n+1$ in P.

We will show that in this case we can choose the instance $\mathcal{I}^r$ to prove our lower bound. The main idea of the proof is showing that machine r will not be assigned task $n+1$ in P^r with probability that is significantly higher than the probability it was assigned the task in P.

Lemma 4. *Let r be some machine such that $p_{r,n+1}(P) \le \frac{1}{n}$. It holds that $p_{r,n+1}(P^r) \le \frac{1}{n} + \epsilon$.*

Proof. As M is weakly-monotone in the extended sense (and since this is a cost scenario) we have that $V_{v_r}(P) + V_{v'_r}(P^r) \le V_{v'_r}(P) + V_{v_r}(P^r)$. That is:

$$\Sigma_{t\in[m]}\ p_{r,t}(P)v_r(t) + \Sigma_{t\in[m]}\ p_{r,t}(P^r)v'_r(t) \le$$
$$\Sigma_{t\in[m]}\ p_{r,t}(P)v'_r(t) + \Sigma_{t\in[m]}\ p_{r,t}(P^r)v_r(t).$$

After subtracting identical terms from both sides we have:

$$p_{r,r}(P)v_r(r) + p_{r,n+1}(P)v_r(n+1) + p_{r,r}(P^r)v'_r(r) + p_{r,n+1}(P^r)v'_r(n+1) \le$$
$$p_{r,r}(P)v'_r(r) + p_{r,n+1}(P)v'_r(n+1) + p_{r,r}(P^r)v_r(r) + p_{r,n+1}(P^r)v_r(n+1)$$

By (1) and (2) we have:

$$p_{r,r}(P) + p_{r,n+1}(P) + p_{r,n+1}(P^r) \times (1+\epsilon) \le p_{r,n+1}(P) \times (1+\epsilon) + p_{r,r}(P^r) + p_{r,n+1}(P^r)$$

Therefore:

$$p_{r,r}(P) + p_{r,n+1}(P^r) \times \epsilon \le p_{r,n+1}(P) \times \epsilon + p_{r,r}(P^r)$$

Because $p_{r,r}(P) \ge 1-\epsilon^2$ and $p_{r,r}(P^r) \le 1$ we have:

$$(1-\epsilon^2) + p_{r,n+1}(P^r) \times \epsilon \le p_{r,n+1}(P) \times \epsilon + 1$$

Equivalently, $p_{r,n+1}(P^r) \le p_{r,n+1}(P) + \epsilon$. $\square$

We next show that M fails to provide an approximation ratio better than $2 - \frac{1}{n}$ for $\mathcal{I}^r$. The optimal allocation for $\mathcal{I}^r$ is T^r, which has a finishing time of $1+\epsilon$. Any other allocation has a finishing time of at least 2. However, with high probability T^r is not reached by M, since $p_{r,n+1}(P^r) \le \frac{1}{n} + \epsilon$, and since machine r gets task $n+1$ in T^r, we know the probability that M outputs T^r is at most $\frac{1}{n} + \epsilon$. The expected approximation ratio of M is therefore at least

$$\frac{(\frac{1}{n}+\epsilon) \times (1+\epsilon) + (1-(\frac{1}{n}+\epsilon)) \times 2}{1+\epsilon} \ge (2-\frac{1}{n}) \cdot \frac{1-\epsilon}{1+\epsilon}$$

Since $\lim_{\epsilon\to 0} \frac{1-\epsilon}{1+\epsilon} = 1$, the theorem follows. $\square$

3.5. A lower bound for Bayesian incentive compatible mechanisms

We now turn to the notion of Bayesian mechanisms, where players' valuations are drawn from a distribution that is public knowledge. We show a lower bound of 1.2 (even for two machines and three tasks). No lower bound for Bayesian Incentive Compatible mechanisms in multi-parameter settings was previously known. In what follows, we restrict our attention to deterministic mechanisms that deterministically output an allocation (in particular, for any given input each task will be always deterministically allocated to the same machine).

Theorem 5. *Any Bayesian Incentive Compatible deterministic mechanism cannot achieve an approximation ratio strictly better than* 1.2.

Proof. Let ϵ be an arbitrarily small positive real number. Consider a setting with two machines and three tasks (the generalization for $n > 2$ is straightforward). We define a product distribution with two possible *equally likely* valuation functions v_i and v'_i for every machine $i = 1, 2$ (notice that the processing times of the tasks are not identical across machines):

$$v_i(t) = \begin{cases} 1 & t = i \text{ or } t = 3 \\ \frac{4}{\epsilon} & \text{otherwise} \end{cases}$$

and

$$v_i'(t) = \begin{cases} 0 & t = i \\ 1+\epsilon & t = 3 \\ \frac{4}{\epsilon} & \text{otherwise.} \end{cases}$$

Assume by contradiction that there exists a *deterministic* Bayesian Incentive Compatible mechanism M with an expected approximation ratio $1.2 - \delta$ for some $\delta > 0$. We shall denote the allocation that M outputs for the instance $\mathcal{I}$ by $M(\mathcal{I})$.

Let T^j be the allocation of the tasks to machines in which task i is assigned to machine i, and task 3 is assigned to machine j (where $i, j \in \{1,\ 2\}$). Notice that $M(v_1, v_2), M(v_1', v_2') \in \{T^1, T^2\}$ since for any allocation that gives the first task to the second machine or the second task to the first machine we have

$$\frac{\frac{1}{4} \cdot \frac{4}{\epsilon} + \frac{3}{4} \cdot (1+\epsilon)}{\frac{1}{4} \cdot 2 + \frac{3}{4} \cdot (1+\epsilon)} > 1.2 - \delta.$$

Now, T^1, T^2 are optimal for the instances (v_1', v_2), (v_1, v_2'), respectively. Furthermore, we must have that $M(v_1', v_2) = T^1, M(v_1, v_2') = T^2$ under the assumption that an expected approximation ratio strictly better than 1.2 can be achieved. More formally, observe that the finishing time of any other allocation of tasks is at least 2, and clearly for a small enough ϵ we have

$$\frac{\frac{2}{4} \cdot 2 + \frac{2}{4} \cdot (1+\epsilon)}{\frac{1}{4} \cdot 2 + \frac{3}{4} \cdot (1+\epsilon)} > 1.2 - \delta.$$

Without loss of generality assume that $M(v_1, v_2) = T^2$. By symmetry it is enough to consider two cases: $M(v_1', v_2') = T^1$ and $M(v_1', v_2') = T^2$.

In the first case we have that $M(v_1, v_2) = M(v_1, v_2') = T^2, M(v_1', v_2) = M(v_1', v_2') = T^1$. Recall that the probability that the valuation function of machine 2 is v_2 (or v_2') is $\frac{1}{2}$. By Bayesian incentive compatibility (with respect to player 1) we have that (for the case where the valuation function of player 1 is v_1):

$$\frac{1}{2}(p_1(v_1, v_2) - v_1(T^2)) + \frac{1}{2}(p_1(v_1, v_2') - v_1(T^2)) \geq \frac{1}{2}(p_1(v_1', v_2) - v_1(T^1)) + \frac{1}{2}(p_1(v_1', v_2') - v_1(T^1)).$$

By Bayesian incentive compatibility (with respect to player 1) we have that (for the case where the valuation function of machine 1 is v_1'):

$$\frac{1}{2}(p_1(v_1', v_2) - v_1'(T^1)) + \frac{1}{2}(p_1(v_1', v_2') - v_1'(T^1)) \geq \frac{1}{2}(p_1(v_1, v_2) - v_1'(T^2)) + \frac{1}{2}(p_1(v_1, v_2') - v_1'(T^2)).$$

By adding up these two inequalities and cancelling out the payments from both sides we have that $v_1(T^1) + v_1'(T^2) \geq v_1(T^2) + v_1'(T^1)$. A contradiction (since $2 + 0 < 1 + (1+\epsilon)$).

We are left with the case in which $M(v_1, v_2) = M(v_1, v_2') = M(v_1', v_2') = T^2,\ M(v_1', v_2) = T^1$. By Bayesian incentive compatibility we have that

$$\frac{1}{2}(p_1(v_1, v_2) - v_1(T^2)) + \frac{1}{2}(p_1(v_1, v_2') - v_1(T^2)) \geq \frac{1}{2}(p_1(v_1', v_2) - v_1(T^1)) + \frac{1}{2}(p_1(v_1', v_2') - v_1(T^2)),$$

and

$$\frac{1}{2}(p_1(v_1', v_2) - v_1'(T^1)) + \frac{1}{2}(p_1(v_1', v_2') - v_1'(T^2)) \geq \frac{1}{2}(p_1(v_1, v_2) - v_1'(T^2)) + \frac{1}{2}(p_1(v_1, v_2') - v_1'(T^2)).$$

Equivalently,

$$\frac{1}{2}(p_1(v_1, v_2) - v_1(T^2)) + \frac{1}{2}p_1(v_1, v_2') \geq \frac{1}{2}(p_1(v_1', v_2) - v_1(T^1)) + \frac{1}{2}p_1(v_1', v_2'),$$

and

$$\frac{1}{2}(p_1(v_1', v_2) - v_1'(T^1)) + \frac{1}{2}p_1(v_1', v_2') \geq \frac{1}{2}(p_1(v_1, v_2) - v_1'(T^2)) + \frac{1}{2}p_1(v_1, v_2').$$

By adding up these two inequalities and rearranging we have that $v_1(T^1) + v_1'(T^2) \geq v_1(T^2) + v_1'(T^1)$. A contradiction (since $2 + 0 < 1 + (1+\epsilon)$). $\square$

4. Workload minimization in inter-domain routing

In this section, we study another non-utilitarian multi-dimensional problem – *workload minimization in inter-domain routing*. Feigenbaum, Papadimitriou, Sami, and Shenker formulated the inter-domain routing problem as a distributed mechanism design problem (inspired by the extensive literature on the real-life problem of inter-domain routing in the Internet) (Feigenbaum et al., 2005). Several works that study their model and its extensions have been published (Feigenbaum et al., 2006, 2007, 2011). All these works deal with the realization of utilitarian social-choice functions, and focus on the efficient and distributed design of VCG mechanisms.

Workload minimization is a problem that arises naturally in the design of routing protocols, as we wish that no single Autonomous System (AS) will be overloaded with work. It can easily be shown that any VCG mechanism performs very poorly with respect to workload minimization. Thus, while optimally minimizing the total cost, or maximizing the social welfare, the known truthful mechanisms for this problem can result in workloads that are very far from optimal (in which one AS is burdened by the traffic sent by all other ASes). We initiate the study of truthful workload minimization in inter-domain routing by presenting constant lower bounds that apply to any truthful mechanism (deterministic and randomized).

4.1. The setting

We are given a *directed graph* $G = \langle N, L\rangle$ (called the *AS graph*) in which the set of nodes N corresponds to the Autonomous Systems (ASes) of which the Internet is comprised. The set N consists of a *destination* node d, and n *source* nodes (see Example 1 below). The set of edges L corresponds to communication links between the ASes. Each source node $i \in [n]$ is a strategic player. The number of packets (intensity of traffic) originating in source node i and destined for d is denoted by t_i. The directed graph G and number of packets t_is are public knowledge.

Let $neighbours(i)$ be all the ASes that are directly linked from i in the AS graph. Each source node i has a private cost function c_i: $neighbours(i) \to \mathbb{R}_+$ that specifies the per-packet cost incurred by this node for carrying traffic. This cost function represents the additional load imposed on the internal AS network when sending a packet from i to an adjacent AS. In the formulation of the problem in Feigenbaum et al. (2005), a node does not incur a cost for packets that originate in that node. However, since we are interested in workload minimization, this is not the case in our formulation. Additionally, as we are interested in proving lower bounds we can restrict our attention to the model in which the number of packets t_is are public knowledge. If the mechanism is truthful then player i can never improve its utility by misreporting his private cost function c_i to the mechanism (no matter what the other players do). In the *single-dimensional version of this problem* an AS i incurs the *same* per-packet private cost c_i for sending traffic to each of its neighbours.

The set of alternatives A contains all possible *route allocation* that form a confluent tree to the destination d. I.e., no source node is allowed to transfer traffic to two adjacent nodes. We seek truthful mechanisms that output routing trees in which the *workload* imposed on the busiest source node is minimized. Formally, let N_i^T be the set of all nodes whose paths in the routing tree T go through node i. Let $s^T(i)$ be the subsequent node i transfers traffic to in the routing tree T. The global goal is to minimize the expression

$$max_i\ \Sigma_{j\in N_i^T}\ t_j \times c_i(s^T(i))$$

over all possible routing trees T. Each source node i is a strategic player who wishes to minimize his workload ($= \Sigma_{j\in N_i^T}\ t_j \times c_i(s^T(i))$) minus the payment made to him by the mechanism. If the mechanism is truthful then player i can never improve its utility by misreporting his private cost function c_i to the mechanism (no matter what the other players do).

Remark 3. It is easy to verify that the single-dimensional related machine scheduling makespan minimization problem (Archer and Tardos, 2001) is a special important case of the single-dimensional workload minimization problem (furthermore, from any c-approximation truthful mechanism for workload minimization one can construct a c-approximation truthful mechanism for makespan minimization). However, the problem of multi-dimensional workload minimization is *strategically* different from the multi-dimensional unrelated machine scheduling makespan minimization problem that we study in Section 3 (e.g., since t_i is a public information and the cost c_i is a function of the direct successor nodes rather than the direct predecessor nodes).

Example 1. Consider the routing instance in Fig. 1. Each source node has a single packet it wishes to send to the destination. The number beside every directed link (u, v) in the figure represents the cost u incurs for transferring a packet to v. Consider the routing tree in which both y and z send packets through x, and x forwards packets directly to d. The workload on x is 3 in this routing tree. In the optimal routing tree all source nodes chose to send their packets directly to d. This routing tree has a maximal workload of $1 + \epsilon$.

4.2. Approximability of the single-dimensional case

It is easy to show (via a simple reduction from Partition) that even the single-dimensional version of the workload-minimization problem is NP-hard. However, is it at all possible to *optimally* solve this problem in a truthful manner? The

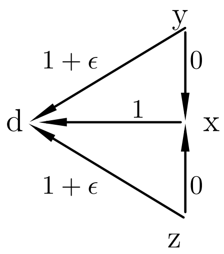

Fig. 1. Routing Instance for Example 1.

answer to this question is yes. However, the worst-case running time is exponential (we note that this is also the case in the single-dimensional version of the machine-scheduling problem studied by Archer and Tardos, 2001).

Lemma 5. *There exists a truthful, deterministic, exponential-time mechanism that always finds a workload-minimizing route allocation in the single-dimensional case.*

Proof. The mechanism M simply goes over all possible route allocations and outputs the *optimal* one with respect to workload-minimization. As in Archer and Tardos (2001), our truthful mechanism outputs the lexicographically-minimal optimal route allocation. Specifically, let a and b be two distinct optimal route allocations (if *two* such allocations exist). Let $a_1, \ldots, a_n$ be a *decreasing* order of the number of packets that go through each of different nodes in a (using a deterministic tie-breaking rule). Similarly, let $b_1, \ldots, b_n$ be *decreasing* order of the number of packets that go through each of different nodes in b. Let $j \in [n]$ be the first index such that $a_j \neq b_j$ or $j = n$ if no such index exists (notice that the sorted order of the coordinates of a might be identical to the sorted order of the coordinates of b, even if $a \neq b$). The mechanism will choose a if $a_j < b_j$, b if $b_j < a_j$, and otherwise according to a predefined deterministic tie-breaking rule.

We next show the truthfulness of the mechanism. It is well known that a mechanism is truthful in a single-dimensional setting such as ours *if and only if* it is weakly-monotone (Archer and Tardos, 2001). Let a be the route allocation that M outputs when the per-packet cost of i is c_i, and the per-packet costs of the other nodes are $c_{-i} = c_1, \ldots, c_{i-1}, c_{i+1}, \ldots, c_n$. Let b be the route allocation that M outputs when the per-packet cost of i is c_i', and the per-packet costs of the other nodes are c_{-i}. Weak-monotonicity states that if $c_i < c_i'$ then $k_i \geq k_i'$, where k_i and k_i' are the number of packets that go through i in a and b, respectively (and this is true for every node i, for every tuple c_{-i}, and for every $c_i \neq c_i'$).

Fix a node i. Assume, by contradiction, that there are $c_i < c_i'$, and c_{-i} such that $k_i < k_i'$. Let $a_1, \ldots, a_n$ and $b_1, \ldots, b_n$ be defined as before. Let $c = (c_i, c_{-i})$, $c' = (c_i', c_{-i})$. We shall use the notation $\mathrm{W}(a, c)$ to denote the maximum workload of a player in the route allocation a when the per-packet costs are $c = (c_1, c_2, \ldots, c_n)$ (that is, $\mathrm{W}(a, c) = max_i \ \Sigma_{j \in N_i^T} \ t_j \times c_i(s^a(i))$). Similarly, $\mathrm{W}(b, c)$ denotes the maximum workload of a player in the route allocation b when the per-packet costs are $c = (c_1, c_2, \ldots, c_n)$. The terms $\mathrm{W}(a, c')$ and $\mathrm{W}(b, c')$ are defined analogously. By the optimality of M it is immediate to verify that:

$$\mathrm{W}(a,c) \leq \mathrm{W}(b,c) \leq \mathrm{W}(b,c') \leq \mathrm{W}(a,c') \tag{3}$$

$$\mathrm{W}(a,c') = \max\{\mathrm{W}(a,c),\ k_i c_i'\} \tag{4}$$

$$k_i' c_i' \leq \mathrm{W}(b,c') \tag{5}$$

By (4) it suffices to consider two cases:

Case 1: $\mathrm{W}(a, c') = \mathrm{W}(a, c)$. From (3) we have

$$\mathrm{W}(a,c) = \mathrm{W}(b,c) = \mathrm{W}(b,c') = \mathrm{W}(a,c').$$

Now, $\mathrm{W}(a, c) = \mathrm{W}(b, c)$ dictates that a comes before b in the lexicographic order (since a is the optimal route allocation at c), and $\mathrm{W}(b, c') = \mathrm{W}(a, c')$ dictates that b comes before a in the lexicographic order (since b is the optimal route allocation at c'), a contradiction.

Case 2: $\mathrm{W}(a, c') = k_i c_i'$. From (3) and (5) we have $k_i' c_i' \leq \mathrm{W}(b, c') \leq \mathrm{W}(a, c') = k_i c_i'$, contradicting our assumption that $k_i' > k_i$. □

4.3. Approximability of the multi-dimensional case

Feigenbaum et al. (2005) present a truthful polynomial-time VCG mechanism that always outputs the *cost-minimizing tree* (a tree that minimizes the total sum of costs incurred for the packets sent to d). We begin our discussion on the multi-dimensional version of the workload minimization problem by showing that this VCG mechanism obtains an n-approximation for the multi-dimensional version of our problem (and hence also for the single-dimensional version) in polynomial time.

Theorem 6. *There is a truthful polynomial-time deterministic n-approximation mechanism for the workload minimization problem in inter-domain routing.*

Proof. We prove that any mechanism that minimizes the total cost provides an n-approximation to the minimal workload. Hence, the mechanism of Feigenbaum et al. (2005) obtains the required approximation ratio.

Suppose the per packet costs are $c = (c_1, \ldots, c_n)$. Denote by T the corresponding cost-minimizing routing-tree and by T^* the corresponding workload-minimizing routing-tree. Let $C(T, c)$ and $C(T^*, c)$ be the total costs of T and T^*, respectively. Recall that $W(T^*, c)$ is the value of the optimal solution for the workload-minimization problem. The result now follows immediately from $C(T^*, c) \leq n \cdot W(T^*, c)$ and $C(T, c) \geq W(T^*, c)$. □

Unfortunately, it can be shown that any mechanism that minimizes the total cost (and in particular the mechanism in Feigenbaum et al. (2005)) cannot obtain a better approximation ratio.

Theorem 7. *Any mechanism that minimizes the total cost of the routing tree cannot achieve an approximation ratio strictly better than n for the workload minimization problem in inter-domain routing.*

Proof. Recall the routing instance in Fig. 1. Observe, that any total-cost minimizing mechanism would choose the routing tree in which both y and z send packets through x, and x forwards packets directly to d. This means that the workload on x is 3. However, if all nodes chose to send their packets directly to d we would reach a maximum workload of $1 + \epsilon$. Clearly, the example in Fig. 1 can be generalized to n source nodes. Notice also that a similar example can be used to show that the theorem holds for singe-dimensional problems. The idea is to replace the link (y, d) with the links $(y, y'), (y', d)$ and the link (z, d) with $(z, z'), (z', d)$ while $c_y(y, y') = c_z(z, z') = 0$, and $c_{y'} = c_{z'} = 1 + \epsilon$ assuming x, y, z have each a single packet to send, and y', z' have no packets to send to the destination. □

Therefore, there exists a trade-off between the goal of minimizing the total-cost and the goal of minimizing the workload. It would be interesting to construct a truthful mechanism that optimizes (or at least closely approximates) the minimal workload. We present two negative results for this problem in Appendix A.2. In a follow-up work, Gamzu (2011) improved our truthful lower bound for minimizing the workload in inter-domain routing from $\frac{1+\sqrt{5}}{2} \approx 1.618$ to 2, and our universally-truthful randomized lower bound from $\frac{3+\sqrt{5}}{4} \approx 1.309$ to 2.

5. Non-utilitarian fairness

Utilitarian functions represent the *overall* satisfaction of the players, as they maximize the sum of players' values. This notion of *fairness* is but one of several that have been considered (explicitly and implicitly) in mathematical, economic and computational literature. A well known example of *non-utilitarian* fairness is the *cake-cutting problem*, presented by the Polish school of mathematicians in the 1950's (Steinhaus, 1948). Fair allocations of indivisible items have also been studied (Lavi et al., 2003; Lipton et al., 2004).

In this section, we discuss three general notions of *non-utilitarian* fairness – Max–Min fairness, Min–Max fairness, and envy-minimization. We show that Max–Min fairness is inapproximable within *any ratio*, even for extremely restricted special cases. In sharp contrast, we show that Min–Max fairness (which is a generalization of both the scheduling and workload-minimization problems considered in this paper) can always be truthfully approximated via a simple VCG mechanism. Finally, we prove a lower bound for the envy-minimization problem.

5.1. The setting

In this setting we have m indivisible items and n strategic players. Each player i is defined by a private valuation function $v_i : 2^{[m]} \to \mathbb{R}_+$. We assume that $v_i(\emptyset) = 0$ (free disposal), and $v_i(S) \leq v_i(T)$ (monotonicity) for every i and every two sets of items $S, T \subseteq [m]$ such that $S \subseteq T$.

Each player wishes to maximize the value of the set of items assigned to it minus its payment to the mechanism. If the mechanism is truthful then player i can never improve its utility by misreporting his private valuation function v_i to the mechanism (no matter what the other players do). The set of alternatives A contains all possible allocations of items to the players, where all items must be assigned, and each item is assigned to exactly one player. The global goal of each of our three notions of *non-utilitarian* fairness – Max–Min fairness, Min–Max fairness, and envy-minimization – will be defined below.

5.2. Max–Min fairness

The global goal The Max–Min social choice function is concerned with maximizing the value of the least satisfied player. Formally, for every n-tuple of v_i valuations the Max–Min function assigns the alternative $a \in A$ that maximizes the expression $\min_i v_i(a)$.

The Max–Min fairness allocation problem is based on the philosophical work of Rawls (1971). Lavi et al. (2003) proved that Max–Min fairness in allocations of indivisible items cannot be *optimally* implemented in a truthful manner. Non-truthful algorithms for this problem were designed (Bansal and Sviridenko, 2006), as well as algorithms that settle for restricted notions of truthfulness (Bezáková and Dani, 2005; Golovin, 2005). We prove that no truthful deterministic mechanism can obtain *any* approximation ratio to the Max–Min fairness value. We prove this lower bound even for the case of 2 players and 2 items.

Theorem 8. *No truthful deterministic mechanism can obtain* any *approximation to the Max–Min fairness value in the allocation of indivisible items. This holds even for the case of* 2 *players and* 2 *items.*

Proof. Let $c > 1$ and let ϵ be an arbitrarily small positive real number. Consider an instance with two players $i = 1, 2$ and two goods g_a, g_b. Each player i has an additive valuation function. Let

$$\begin{aligned} &v_1(g_a) = 2, && v_1(g_b) = \tfrac{1}{c} \\ &v_2(g_a) = 4 - \epsilon, && v_2(g_b) = 1 + \epsilon. \end{aligned}$$

Note, that the optimal allocation assigns g_a to player 1 and g_b to player 2, thus obtaining a Max–Min value of $1 + \epsilon$. Also note, that this allocation will also be chosen by any c-approximation mechanism.

We alter the valuation of player 2 into v'_2 such that $v'_2(g_a) = \frac{1}{c}$, $v'_2(g_b) = \frac{1}{c^2} - \epsilon$. The optimal Max–Min value is now $\frac{1}{c}$. Observe that any c-approximation mechanism must assign g_b to player 1 and g_a to player 2. However, if this happens we have that:

$$(1+\epsilon) + \frac{1}{c} = v_2(g_b) + v'_2(g_a) < v_2(g_a) + v'_2(g_b) = (4-\epsilon) + \frac{1}{c^2} - \epsilon.$$

This violates weak-monotonicity, and so no truthful c-approximation mechanism exists. Since this is true for any $c > 1$ the theorem follows. □

5.3. Min–Max fairness

The global goal Min–Max fairness can be thought of as the dual notion of Max–Min fairness. It is relevant in settings in which each player incurs a cost for every chosen alternative.[6] The Min–Max social choice function is concerned with minimizing the cost incurred by the least satisfied player. Formally, for every n-tuple of v_i valuations the Max–Min function assigns the alternative $a \in A$ that minimizes the expression $\max_i \; v_i(a)$.

Observe, that both the scheduling problem and the workload-minimization problem discussed in this paper, are in fact special cases of this notion of fairness. Studying Max–Min fairness in this more abstract setting enables us to state this simple observation – any Min–Max social-choice function can be truthfully approximated within a factor of n (recall that n is the number of players) by a simple VCG mechanism. Since the best currently known approximation-mechanisms for both scheduling and workload-minimization are VCG-based, this result can be viewed as a generalization of both.

Theorem 9. *Let f be a Min–Max social choice function. Then, there exists a truthful deterministic mechanism that for every n-tuple of valuations $v_1, \ldots, v_n$ outputs an alternative $a \in A$ such that* $\max_i \; v_i(a)$ *is an n-approximation to the value of the solution f outputs for these valuations.*

Proof. Let $v_1, \ldots, v_n$ be the valuation function of the players. Let b be the allocation that f outputs for $v_1, \ldots, v_n$. Consider the VCG mechanism that minimizes the total cost the players incur. The truthfulness of this mechanism is guaranteed by the VCG technique. Let $a \in A$ be the allocation that this mechanism outputs. The result now follows immediately from $\max_i \; v_i(a) \le \Sigma_i \; v_i(a)$ and the fact that $\max_i \; v_i(b) \ge \frac{1}{n} \cdot \Sigma_i \; v_i(a)$. □

5.4. Envy-minimization

The global goal Lipton et al. (2004) presented the problem of envy-minimization for indivisible items. They consider truthful mechanisms for this problem. An *envy-minimizing allocation* of items is a partition of the m items into disjoint sets $S_1, \ldots, S_n$ (player i is assigned S_i) that minimizes the expression $\max_{i,j} \{v_i(S_j) - v_i(S_i), 0\}$ (over all possible allocations). Intuitively, we wish to minimize the maximal envy a player might feel by comparing his value for a set of items given to another player to the value he assigns the items allocated to him.

They prove several approximability results for this problem. The parameter considered in Lipton et al. (2004) is the *maximal marginal utility*.

[6] Here each player wishes to minimize the cost of the set of items assigned to it minus the payment made to it by the mechanism.

Definition 5. The maximal marginal utility α is defined as follows:

$$\alpha(v_1, \ldots, v_n) = \max_{i\in[n],\ j\in[m],\ S\subseteq[m]} v_i(S \cup \{j\}) - v_i(S).$$

That is, α is the maximal value by which the value of a player increases when one item is added to his bundle.

Lipton et al. (2004) exhibit a universally-truthful randomized mechanism that has an envy of at most $O(\sqrt{\alpha}n^{\frac{1}{2}+\epsilon})$ w.h.p. for large values of n. They show that no truthful mechanism can guarantee an optimal solution with respect to envy minimization. We strengthen this lower bound by showing that no truthful deterministic mechanism can guarantee an allocation that has an envy value within α from optimal.

Theorem 10. *Any truthful deterministic mechanism cannot obtain an approximation ratio better than α for envy minimization.*

Proof. Assume for contradiction that there exists a truthful deterministic mechanism M that α-approximates the envy minimization problem. Consider an instance with 2 players and 3 items. Each player $i = 1, 2$ has the same additive valuation function v_i that assigns any of the single items a value of 1. Observe, that $\alpha(v_1, v_2) = 1$. Notice, that the minimal envy for this instance is 1.[7] Hence, if M assigns all items to one of the players the envy is precisely $3 > 1 = \alpha(v_1, v_2) \cdot 1$. Therefore, we can assume w.l.o.g. that the α-approximation mechanism M allocates items 1, 2 to player 1 and item 3 to player 2 (call this partition (S_1, S_2)).

Let ϵ be an arbitrarily small positive real number. We now change the valuation function of player 1 into the following additive valuation:

$$v_1'(j) = \begin{cases} 1+\epsilon & j = 1, 2 \\ \epsilon & j = 3. \end{cases}$$

Now $\alpha(v_1', v_2) = 1 + \epsilon$. Also observe that the minimal envy for this new instance is ϵ (e.g., assign item 1 to player 1 and items 2, 3 to player 2). However, it is easy to verify that weak-monotonicity dictates that the allocation remains the same even after the alteration of the valuation of player 1 (since for any partition $(T_1, T_2) \neq (S_1, S_2)$ of the items we have that $v_1'(T_1) - v_1(T_1) < v_1'(S_1) - v_1(S_1) = 2\epsilon$). Therefore, we end up with an allocation in which the envy is $1 > (1+\epsilon) \cdot \epsilon = \alpha(v_1', v_2) \cdot \epsilon$, a contradiction. $\square$

Acknowledgments

We thank Ron Lavi, Noam Nisan, Chaitanya Swamy, Amir Ronen and anonymous referees for many helpful suggestions.

Appendix A

A.1. A universally-truthful randomized approximation mechanism for the scheduling problem with unrelated machines

Nisan and Ronen (2001) present a truthful deterministic mechanism that obtains an n-approximation. For the case of 2 machines, they exhibit a universally-truthful randomized mechanism that obtains an approximation of $\frac{7}{4}$ (this bound was later improved to 1.58606 by Chen et al., 2015). We generalize this result by presenting a universally-truthful randomized mechanism that obtains an approximation-ratio of $\frac{7n}{8}$. We now turn to the description of our mechanism for n machines:

Input: An n-tuple of valuations $(v_1, \ldots, v_n)$ and m tasks.

Output: An allocation $T = T_1, \ldots, T_n$ of tasks, and payments $p_1, \ldots, p_n$ such that T has a makespan value which is a $\frac{7n}{8}$-approximation to the optimal makespan value, and the payments induce truthfulness.

The Mechanism:

1. For every machine i let $T_i \leftarrow \emptyset$ and $p_i \leftarrow 0$.
2. Partition the set of machines into two sets $S_1 \leftarrow \{1, \ldots, \frac{n}{2}\}$ and $S_2 \leftarrow \{\frac{n}{2}+1, \ldots, n\}$.
3. For each task $j = 1, \ldots, m$ perform the following actions (assuming that min and argmin break ties arbitrarily):
   - Let $v^1 \leftarrow \min_{i\in S_1} v_i(j)$, and let $i^1 \leftarrow \text{argmin}_{i\in S_1} v_i(j)$.
   - Let $v'^1 \leftarrow \min_{i\in S_1\setminus\{i^1\}} v_i(j)$.
   - Let $v^2 \leftarrow \min_{i\in S_2} v_i(j)$, and let $i^2 \leftarrow \text{argmin}_{i\in S_2} v_i(j)$.
   - Let $v'^2 \leftarrow \min_{i\in S_2\setminus\{i^2\}} v_i(j)$.
   - Randomly and uniformly choose a value $R \in \{0, 1\}$.

[7] The *envy* of $S_1, \ldots, S_n$ equals $\max_{i,j} \{v_i(S_j) - v_i(S_i), 0\}$.

- If $R = 0$ and $v^1 \leq \frac{4}{3}v^2$ set $T_{i^1} \leftarrow T_{i^1} \bigcup \{j\}$ and set $p_{i^1} \leftarrow p_{i^1} + \min\{v'^1, \frac{4}{3}v^2\}$.
- If $R = 0$ and $v^1 > \frac{4}{3}v^2$ set $T_{i^2} \leftarrow T_{i^2} \bigcup \{j\}$ and set $p_{i^2} \leftarrow p_{i^2} + \min\{v'^2, \frac{3}{4}v^1\}$.
- If $R = 1$ and $v^2 \leq \frac{4}{3}v^1$ set $T_{i^2} \leftarrow T_{i^2} \bigcup \{j\}$ and set $p_{i^2} \leftarrow p_{i^2} + \min\{v'^2, \frac{4}{3}v^1\}$.
- If $R = 1$ and $v^2 > \frac{4}{3}v^1$ set $T_{i^1} \leftarrow T_{i^1} \bigcup \{j\}$ and set $p_{i^1} \leftarrow p_{i^1} + \min\{v'^1, \frac{3}{4}v^2\}$.

4. Allocate each machine i the tasks in T_i, and pay it a sum of p_i.

Remark 4. If n cannot be divided by 2 simply add the extra machine to either S_1 or S_2.

Theorem 11. *There exists a universally-truthful randomized mechanism for the scheduling problem that obtains an approximation ratio of $\frac{7n}{8}$.*

Proof. We prove the theorem for the case that n can be divided by 2. The proof for the other case is similar. Our proof relies on the proof of Nisan and Ronen (2001). Observe, that the utility of each machine after the algorithm finishes is the sum of its utilities for the different tasks. Hence, it is sufficient to prove that for each individual task a machine has no incentive to lie. As in Nisan and Ronen (2001), this is guaranteed because the allocation of each task is in fact a weighted VCG mechanism (see Nisan and Ronen, 2001 for further explanations), which is known to be truthful. Hence, this mechanism is universally truthful.

We now need to prove that the approximation ratio guaranteed by the mechanism is indeed $\frac{7n}{8}$. Let $\mathcal{I}$ be an instance of the scheduling problem with m tasks, and with n machines that have the valuation functions $v_1, \ldots, v_n$. We define an instance $\mathcal{I}'$ of scheduling problem with m tasks, and with 2 machines that have the valuation function v'_1, v'_2, in the following way: $v'_1(j) = min_{i \in S_1}\ v_i(j)$ for all $j \in [m]$. Similarly, $v'_2(j) = min_{i \in S_2}\ v_i(j)$ for all $j \in [m]$. We denote by $M(\mathcal{I})$ and by $M(\mathcal{I}')$ the makespan values our mechanism generates for $\mathcal{I}$ and $\mathcal{I}'$ respectively. We denote by $OPT(\mathcal{I})$ and by $OPT(\mathcal{I}')$ the optimal makespan values for $\mathcal{I}$ and $\mathcal{I}'$ respectively.

First, notice that $M(\mathcal{I}) \leq M(\mathcal{I}')$. This is because applying our mechanism to $\mathcal{I}'$ results in the same makespan value as applying it to $\mathcal{I}$ in the worst-case scenario in which tasks are always assigned to the same machines in S_1 and in S_2. It also holds that $M(\mathcal{I}') \leq \frac{7}{4}OPT(\mathcal{I}')$ because in the case that there are only two machines our mechanism is precisely that of Nisan and Ronen (2001), which guarantees a $\frac{7}{4}$ approximation ratio. We now have that $M(\mathcal{I}) \leq \frac{7}{4}OPT(\mathcal{I}')$. All that is left to show is that $OPT(\mathcal{I}') \leq \frac{n}{2}OPT(\mathcal{I})$. Consider the optimal allocation of tasks for $\mathcal{I}$. By giving all tasks assigned to machines in S_1 to machine 1 in $\mathcal{I}'$, and allocating all tasks assigned to machines in S_2 to machine 2 in $\mathcal{I}'$, we end up with a makespan value for $\mathcal{I}'$ that is at most $\frac{n}{2}OPT(\mathcal{I})$. The theorem follows. $\square$

A.2. Lower bounds for minimizing the workload in inter-domain routing

Theorem 12. *No truthful deterministic mechanism for minimizing the workload in inter-domain routing can obtain an approximation ratio better than $\frac{1+\sqrt{5}}{2} \approx 1.618$.*

Proof. This proof is similar to the proof of Theorem 1. To prove the lower bound consider the instances of the workload-minimization problem with 3 source nodes x, y, z depicted in Figs. 2 and 3.

Each source node has a single packet it wishes to send to the destination. The number beside every directed link (u, u') in these figures represents the cost u incurs for transferring a packet to u'. Denote the instance in Fig. 2 by $\mathcal{I}$ and the instance in Fig. 3 by $\mathcal{I}'$. Observe that only the cost function of node x is different in $\mathcal{I}$ and $\mathcal{I}'$. We denote the cost function of x in $\mathcal{I}$ by c_x and his cost function in $\mathcal{I}'$ by c'_x.

Assume, by contradiction, that M is a truthful deterministic mechanism that obtains an approximation ratio better than $\phi = \frac{1+\sqrt{5}}{2}$. Observe, that for the instance $\mathcal{I}$ M must direct the traffic originating in node x through node y (otherwise this contradicts the fact that M obtains an approximation ratio better than ϕ). Similarly, for the instance $\mathcal{I}'$ M must direct the traffic originating in node x through node z. However, this violates the weak-monotonicity of M as $1 + \phi = c_x((x, y)) + c'_x((x, z)) > c_x((x, y)) + c'_x((x, z)) = \phi^2 - \epsilon$ (since ϕ is the golden ratio). $\square$

Theorem 13. *No universally-truthful randomized mechanism for minimizing the workload in inter-domain routing can obtain an approximation ratio better than $\frac{3+\sqrt{5}}{4} \approx 1.309$.*

Proof. We define $\mathcal{I}$ and $\mathcal{I}'$ as in the proof of Theorem 12. Consider the uniform distribution over $\mathcal{I}$ and $\mathcal{I}'$. Let M be a truthful deterministic mechanism. As shown in the proof of Theorem 12, M cannot achieve an approximation better than $\frac{1+\sqrt{5}}{2}$ on both $\mathcal{I}$ and $\mathcal{I}'$ due to its weak-monotonicity. Therefore, the expected approximation of M is at least $\frac{1}{2} \times 1 + \frac{1}{2} \times \frac{1+\sqrt{5}}{2} \approx 1.309$. $\square$

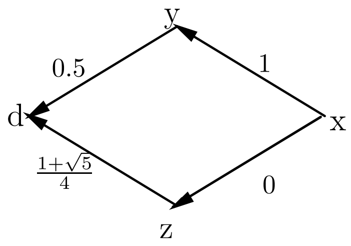

Fig. 2. The Instance $\mathcal{I}$ for the proof of Theorem 12.

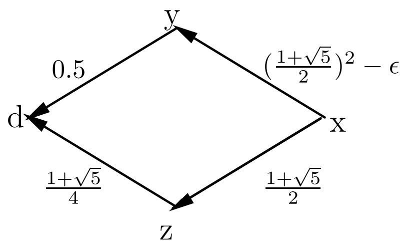

Fig. 3. The Instance $\mathcal{I}$' for the proof of Theorem 12.

References

Andelman, Nir, Azar, Yossi, Sorani, Motti, 2007. Truthful approximation mechanisms for scheduling selfish related machines. Theory Comput. Syst. 40 (4), 423–436.

Archer, Aaron, Tardos, Eva, 2001. Truthful mechanisms for one-parameter agents. In: IEEE Symposium on Foundations of Computer Science, pp. 482–491.

Archer, Aaron, Tardos, Éva, 2007. Frugal path mechanisms. ACM Trans. Algorithms 3 (1), 3.

Ashlagi, Itai, Dobzinski, Shahar, Lavi, Ron, 2012. Optimal lower bounds for anonymous scheduling mechanisms. Math. Oper. Res. 37 (2), 244–258.

Bansal, Nikhil, Sviridenko, Maxim, 2006. The Santa Claus problem. In: Proceedings of the 38th Annual ACM Symposium on Theory of Computing. Seattle, WA, USA, May 21–23, 2006, pp. 31–40.

Bartal, Yair, Gonen, Rica, Nisan, Noam, 2003. Incentive compatible multi unit combinatorial auctions. In: Proceedings of the 9th Conference on Theoretical Aspects of Rationality and Knowledge. TARK-2003, Bloomington, Indiana, USA, June 20–22, 2003, pp. 72–87.

Bezáková, Ivona, Dani, Varsha, 2005. Allocating indivisible goods. ACM SIGecom Exch. 5 (3), 11–18.

Bikhchandani, Sushil, Chatterji, Shurojit, Lavi, Ron, Mu'alem, Ahuva, Nisan, Noam, Sen, Arunava, 2006. Weak monotonicity characterizes deterministic dominant strategy implementation. Econometrica 74 (4), 1109–1132.

Chawla, Shuchi, Hartline, Jason D., Malec, David L., Sivan, Balasubramanian, 2013. Prior-independent mechanisms for scheduling. In: Symposium on Theory of Computing Conference. STOC 2013, Palo Alto, CA, USA, June 1–4, 2013, pp. 51–60.

Chen, Xujin, Du, Donglei, Zuluaga, Luis Fernando, 2015. Copula-based randomized mechanisms for truthful scheduling on two unrelated machines. Theory Comput. Syst. 57 (3), 753–781.

Christodoulou, George, Koutsoupias, Elias, Kovács, Annamária, 2010. Mechanism design for fractional scheduling on unrelated machines. ACM Trans. Algorithms 6 (2), 38.

Christodoulou, George, Koutsoupias, Elias, Vidali, Angelina, 2008. A characterization of 2-player mechanisms for scheduling. In: Algorithms – ESA 2008, 16th Annual European Symposium, Proceedings. Karlsruhe, Germany, September 15–17, 2008, pp. 297–307.

Christodoulou, George, Koutsoupias, Elias, Vidali, Angelina, 2009. A lower bound for scheduling mechanisms. Algorithmica 55 (4), 729–740.

Christodoulou, George, Kovács, Annamária, 2013. A deterministic truthful PTAS for scheduling related machines. SIAM J. Comput. 42 (4), 1572–1595.

Clarke, Edward H., 1971. Multipart pricing of public goods. Public Choice 11, 17–33.

Cramton, Peter, Shoham, Yoav, Steinberg, Richard (Eds.), 2006. Combinatorial Auctions. MIT Press.

Daniely, Amit, Schapira, Michael, Shahaf, Gal, 2015. Inapproximability of truthful mechanisms via generalizations of the VC dimension. In: Proceedings of the Forty-Seventh Annual ACM on Symposium on Theory of Computing. STOC 2015, Portland, OR, USA, June 14–17, 2015, pp. 401–408.

Daskalakis, Constantinos, Matthew Weinberg, S., 2015. Bayesian truthful mechanisms for job scheduling from bi-criterion approximation algorithms. In: Proceedings of the Twenty-Sixth Annual ACM-SIAM Symposium on Discrete Algorithms. SODA 2015, San Diego, CA, USA, January 4–6, 2015, pp. 1934–1952.

Demange, Gabrielle, Gale, David, 1985. The strategy structure of two-sided matching markets. Econometrica 53 (4), 873–888.

Dhangwatnotai, Peerapong, Dobzinski, Shahar, Dughmi, Shaddin, Roughgarden, Tim, 2011. Truthful approximation schemes for single-parameter agents. SIAM J. Comput. 40 (3), 915–933.

Dobzinski, Shahar, Dughmi, Shaddin, 2013. On the power of randomization in algorithmic mechanism design. SIAM J. Comput. 42 (6), 2287–2304.

Dobzinski, Shahar, Nisan, Noam, 2010. Mechanisms for multi-unit auctions. J. Artif. Intell. Res. (JAIR) 37, 85–98.

Dobzinski, Shahar, Nisan, Noam, Schapira, Michael, 2010. Approximation algorithms for combinatorial auctions with complement-free bidders. Math. Oper. Res. 35 (1), 1–13.

Dobzinski, Shahar, Nisan, Noam, Schapira, Michael, 2012. Truthful randomized mechanisms for combinatorial auctions. J. Comput. Syst. Sci. 78 (1), 15–25.

Dobzinski, Shahar, Sundararajan, Mukund, 2008. On characterizations of truthful mechanisms for combinatorial auctions and scheduling. In: Proceedings 9th ACM Conference on Electronic Commerce. EC-2008, Chicago, IL, USA, June 8–12, 2008, pp. 38–47.

Dobzinski, Shahar, Vondrák, Jan, 2016. Impossibility results for truthful combinatorial auctions with submodular valuations. J. ACM 63 (1), 5.

Feigenbaum, Joan, Karger, David R., Mirrokni, Vahab S., Sami, Rahul, 2007. Subjective-cost policy routing. Theor. Comput. Sci. 378 (2), 175–189.

Feigenbaum, Joan, Papadimitriou, Christos, Sami, Rahul, Shenker, Scott, 2005. A BGP-based mechanism for lowest-cost routing. Distrib. Comput. 18, 61–72.

Feigenbaum, Joan, Ramachandran, Vijay, Schapira, Michael, 2011. Incentive-compatible interdomain routing. Distrib. Comput. 23 (5–6), 301–319.

Feigenbaum, Joan, Sami, Rahul, Shenker, Scott, 2006. Mechanism design for policy routing. Distrib. Comput. 18, 293–305.

Fiat, Amos, Goldberg, Andrew V., Hartline, Jason D., Karlin, Anna R., 2002. Competitive generalized auctions. In: Proceedings on 34th Annual ACM Symposium on Theory of Computing. Montréal, Québec, Canada, May 19–21, 2002, pp. 72–81.

Gamzu, Iftah, 2011. Improved lower bounds for non-utilitarian truthfulness. Theor. Comput. Sci. 412 (7), 626–632.
Giannakopoulos, Yiannis, Kyropoulou, Maria, 2015. The VCG mechanism for Bayesian scheduling. In: Web and Internet Economics – 11th International Conference, Proceedings. WINE 2015, Amsterdam, the Netherlands, December 9–12, 2015, pp. 343–356.
Golovin, Daniel, 2005. Max–Min Fair Allocation of Indivisible Goods. Technical Report CMU-CS-05-144. Carnegie Mellon University.
Groves, Theodore, 1973. Incentives in teams. Econometrica 41 (4), 617–631.
Holzman, Ron, Kfir-Dahav, Noa, Monderer, Dov, Tennenholtz, Moshe, 2004. Bundling equilibrium in combinatorial auctions. Games Econ. Behav. 47, 104–123.
Horowitz, Ellis, Sahni, Sartaj, 1976. Exact and approximate algorithms for scheduling nonidentical processors. J. ACM 23 (2), 317–327.
Karlin, Anna R., Kempe, David, Tamir, Tami, 2005. Beyond VCG: frugality of truthful mechanisms. In: 46th Annual IEEE Symposium on Foundations of Computer Science, Proceedings. FOCS 2005, Pittsburgh, PA, USA, October 23–25, 2005, pp. 615–626.
Koutsoupias, Elias, Vidali, Angelina, 2013. A lower bound of $1+\varphi$ for truthful scheduling mechanisms. Algorithmica 66 (1), 211–223.
Kovács, Annamária, Vidali, Angelina, 2015. A characterization of n-player strongly monotone scheduling mechanisms. In: Proceedings of the Twenty-Fourth International Joint Conference on Artificial Intelligence. IJCAI 2015, Buenos Aires, Argentina, July 25–31, 2015, pp. 568–574.
Lavi, Ron, Mu'alem, Ahuva, Nisan, Noam, 2003. Towards a characterization of truthful combinatorial auctions. In: 44th Symposium on Foundations of Computer Science, Proceedings. FOCS 2003, Cambridge, MA, USA, October 11–14, 2003, pp. 574–583.
Lavi, Ron, Swamy, Chaitanya, 2009. Truthful mechanism design for multidimensional scheduling via cycle monotonicity. Games Econ. Behav. 67 (1), 99–124. Special Section of Games and Economic Behavior. Dedicated to the 8th ACM Conference on Electronic Commerce.
Lavi, Ron, Swamy, Chaitanya, 2011. Truthful and near-optimal mechanism design via linear programming. J. ACM 58 (6), 25.
Lehmann, Daniel, O'Callaghan, Liadan, Shoham, Yoav, 2002. Truth revelation in approximately efficient combinatorial auctions. J. ACM 49 (5), 577–602.
Karel Lenstra, Jan, Shmoys, David B., Tardos, Éva, 1990. Approximation algorithms for scheduling unrelated parallel machines. Math. Program. 46, 259–271.
Lipton, Richard J., Markakis, Evangelos, Mossel, Elchanan, Saberi, Amin, 2004. On approximately fair allocations of indivisible goods. In: Proceedings 5th ACM Conference on Electronic Commerce. EC-2004, New York, NY, USA, May 17–20, 2004, pp. 125–131.
Pinyan, Lu., 2009. On 2-player randomized mechanisms for scheduling. In: Internet and Network Economics, 5th International Workshop, Proceedings. WINE 2009, Rome, Italy, December 14–18, 2009, pp. 30–41.
Pinyan, Lu, Yu, Changyuan, 2008. An improved randomized truthful mechanism for scheduling unrelated machines. In: 25th Annual Symposium on Theoretical Aspects of Computer Science, Proceedings. STACS 2008, Bordeaux, France, February 21–23, 2008, pp. 527–538.
Mas-Collel, Andreu, Whinston, Michael D., Green, Jerry R., 1995. Microeconomic Theory. Oxford University Press.
Mu'alem, Ahuva, 2014. Fair by design: multidimensional envy-free mechanisms. Games Econ. Behav. 88, 29–46.
Mu'alem, Ahuva, Nisan, Noam, 2008. Truthful approximation mechanisms for restricted combinatorial auctions. Games Econ. Behav. 64, 612–631.
Mu'alem, Ahuva, Schapira, Michael, 2008. Mechanism design over discrete domains. In: Proceedings 9th ACM Conference on Electronic Commerce. EC-2008, Chicago, IL, USA, June 8–12, 2008, pp. 31–37.
Nisan, Noam, 2015. Algorithmic mechanism design: through the lens of multiunit auctions. In: Handbook of Game Theory with Economic Applications, vol. 4. Elsevier, pp. 477–515. Chapter 9.
Nisan, Noam, Ronen, Amir, 2001. Algorithmic mechanism design. Games Econ. Behav. 35, 166–196.
Nisan, Noam, Rougarden, Tim, Tardos, Eva, Vazirani, Vijay V. (Eds.), 2007. Algorithmic Game Theory. Cambridge University Press.
Nisan, Noam, Segal, Ilya, 2006. The communication requirements of efficient allocations and supporting prices. J. Econ. Theory 129 (1), 192–224.
Osborne, Martin J., Rubinstein, Ariel, 1994. A Course in Game Theory. MIT Press.
Rawls, John, 1971. A Theory of Justice. Belknap Press of Harvard University Press, Cambridge, MA.
Roughgarden, Tim, 2016. Twenty Lectures on Algorithmic Game Theory. Cambridge University Press.
Steinhaus, Hugo, 1948. The problem of fair division. Econometrica 16 (1), 101–104.
Vickrey, William, 1961. Counterspeculation, auctions and competitive sealed tenders. J. Finance 16 (1), 8–37.
Vidali, Angelina, 2009. The geometry of truthfulness. In: Internet and Network Economics, 5th International Workshop, Proceedings. WINE 2009, Rome, Italy, December 14–18, 2009, pp. 340–350.
Vidali, Angelina, 2011. Extending characterizations of truthful mechanisms from subdomains to domains. In: Internet and Network Economics – 7th International Workshop, Proceedings. WINE 2011, Singapore, December 11–14, 2011, pp. 408–414.
Yao, Andrew Chi-Chih, 1977. Probabilistic computations: toward a unified measure of complexity (extended abstract). In: 18th Annual Symposium on Foundations of Computer Science. FOCS, Providence, Rhode Island, USA, 31 October–1 November 1977, pp. 222–227.
Yu, Changyuan, 2009. Truthful mechanisms for two-range-values variant of unrelated scheduling. Theor. Comput. Sci. 410 (21–23), 2196–2206.