

The Anarchy of Scheduling Without Money*

Yiannis Giannakopoulos[†] Elias Koutsoupias[‡] Maria Kyropoulou[§]

December 4, 2018

Abstract

We consider the scheduling problem on n strategic unrelated machines when no payments are allowed, under the objective of minimizing the makespan. We adopt the model introduced in [Koutsoupias 2014] where a machine is *bound* by her declarations in the sense that if she is assigned a particular job then she will have to execute it for an amount of time at least equal to the one she reported, even if her private, true processing capabilities are actually faster. We provide a (non-truthful) randomized algorithm whose pure *Price of Anarchy* is arbitrarily close to 1 for the case of a single task and close to n if it is applied independently to schedule many tasks, which is asymptotically optimal for the natural class of anonymous, task-independent algorithms. Previous work considers the constraint of truthfulness and proves a tight approximation ratio of $(n + 1)/2$ for one task which generalizes to $n(n + 1)/2$ for many tasks. Furthermore, we revisit the truthfulness case and reduce the latter approximation ratio for many tasks down to n , asymptotically matching the best known lower bound. This is done via a detour to the relaxed, fractional version of the problem, for which we are also able to provide an optimal approximation ratio of 1. Finally, we mention that all our algorithms achieve optimal ratios of 1 for the social welfare objective.

1 Introduction

We consider a variant of the scheduling problem proposed by Koutsoupias [18] where no payments are allowed and the machines are bound by their declarations. In particular, the goal is to allocate a set of tasks to strategic unrelated machines while minimizing the makespan. The time/cost needed by a machine to execute a task is private information of the machine. Each machine is rational and selfish, and will misreport her costs in an attempt to minimize her own overall running time, under the assumption that if she is allocated a task, she will execute it for at least the declared cost (more specifically, for the maximum among her true and reported execution times). We are interested in designing allocation protocols that do not use payments and the stable outcomes are not far from the non-strategic, centrally enforced optimum makespan.

The field of Mechanism Design [27] focuses on the implementation of desired outcomes. Given the strategic behaviour of the players who provide the input and a specific objective

*Supported by ERC Advanced Grant 321171 (ALGAME), EPSRC grant EP/M008118/1 and the Alexander von Humboldt Foundation with funds from the German Federal Ministry of Education and Research (BMBF). A preliminary version of this paper, not including all results, appeared in SAGT'16 [14].

A significant part of this work was carried out while Yiannis Giannakopoulos was at the University of Liverpool and Maria Kyropoulou was at the University of Oxford.

[†]Department of Mathematics, TU Munich. Email: yiannis.giannakopoulos@tum.de

[‡]Department of Computer Science, University of Oxford. Email: elias@cs.ox.ac.uk

[§]School of Computer Science and Electronic Engineering, University of Essex. Email: maria.kyropoulou@essex.ac.uk

function that measures the quality of the outcome, the challenge is to design mechanisms which are able to elicit a desired behaviour from the players, while at the same time optimizing that objective value. A primary designer goal that has been extensively studied is that of *truthfulness*, under the central solution concept of dominant strategies: a player should be able to optimize her own individual utility by reporting truthfully, no matter what strategies the other players follow. However, achieving this is not always compatible with maintaining a good objective value [15, 31]. The introduction of *payments* was suggested as a means towards achieving these goals, since a carefully designed payment scheme incentivizes the players to make truthful declarations. The goal now becomes to design such algorithms (termed *mechanisms*) which utilize monetary compensations in order to impose truthful behaviour while optimizing the objective function. A prominent positive result exists for utilitarian settings at which the objective function is the social welfare, where the well-known VCG mechanism [9, 16, 33] constitutes such an optimal mechanism. The study of the algorithmic aspect of mechanism design was initiated by Nisan and Ronen [26], and since then a significant body of work has been dedicated to optimization problems from the mechanism design point of view (see e.g. [4, 10, 21]).

There are many situations, though, where the use of payments might be considered unethical [27], illegal (e.g. organ donations) or even just impractical. For this reason researchers have started turning their attention to possible ways of achieving truthfulness without the use of payments. In such a setting, in order to circumvent Social Choice impossibility results (e.g. the seminal Gibbard-Satterthwaite [15, 31] theorem), domains with richer structure have to be considered. Procaccia and Tennenholtz [30] were the first to consider achieving truthfulness without using payments, by sacrificing the optimality of the solution and settling for just an approximation, in the context of facility location problems. This work was extended in [1, 23]. Similar questions have been considered in the context of inter-domain routing [22], in assignment problems [11], and in the setting of allocating items to two players (with the use of a certain artificial currency) [17]. Moreover, exact (as opposed to approximate) mechanism design without money has a rich history in the social choice literature, e.g. [24] characterizes functions that are truthfully implementable when the preferences of the agents are single-peaked.

Clearly, truthfulness is a property desired by every mechanism designer; if the mechanism can ensure that no player can benefit from misreporting, the designer knows what kind of player behaviour and outcome to expect. Moreover, the focus on truthful mechanisms has been largely motivated by the Revelation Principle stating that essentially every equilibrium state of a mechanism can be simulated by a truthful mechanism which achieves the same objective. However this is no longer possible in the variant we examine here, due to the fact that the players are bound by their declarations and thus do not have quasi-linear utilities. So, it is no longer without loss of generality if we restrict attention to truthful mechanisms. For mechanisms that are not truthful, *Price of Anarchy* (PoA) [19] analysis is the predominant, powerful tool for quantifying the potential suboptimality of the outcomes/equilibria; it measures the impact the lack of coordination (strategic behaviour) has on the solution quality, by comparing it to the optimal and non-strategic solution.

Scheduling is one of the most influential problems in Algorithmic Game Theory and has been studied extensively. In its most general form, the goal is to schedule m tasks to n parallel machines with arbitrary processing times, in order to minimize the makespan. In the front where payments are allowed, truthfulness comes at no extra cost given the strategic nature of the machines. Nisan and Ronen [26] first considered the mechanism design approach of the problem. They prove that the well known VCG mechanism achieves an n -approximation of the optimal makespan, while no truthful deterministic mechanism can achieve approximation ratio better than 2. The currently known best lower bound is 2.61 [20] while Ashlagi et al. [5] proved the tightness of the upper bound for anonymous mechanisms. With respect to

randomized (truthful in expectation) mechanisms as well as fractional ones, the best known bounds are $(n + 1)/2$ and $2 - 1/n$ [8, 25]. We note that the aforementioned lower bounds disregard computational feasibility and simply rely on the requirement for truthfulness.

In an attempt to get positive results when payments are not allowed in the scheduling context, Koutsoupias [18] first considered the plausible assumption that the machines are bound by their declarations. This was influenced by the notion of impositions, according to which a mechanism can restrict the set of reactions available to a player after the outcome is chosen. This notion appeared in [12, 28] and was applied in facility location as well as digital goods pricing settings. The notion of winner imposition fits within the framework of approximate mechanism design without payments. A more powerful framework that is also very much related to this assumption is the notion of verification that appears in [6, 26, 29]. The mechanisms in this context are allowed to use payments and simply give or deny payments to machines after they discover their true execution costs. In particular, the mechanism receives limited information about the players' types after observing the computed solution. Relevant works include [2, 7] where the scheduling problem of selfish tasks is considered again under the assumption that the players who control the tasks are bound by their declarations.

Our Results. In this work we adopt the model of [18]. For the case of scheduling a single task Koutsoupias [18] proved that the approximation ratio of any (randomized) truthful mechanism is at least $(n + 1)/2$ and gave a mechanism matching this bound, where n is the number of machines. When applied to many tasks, this mechanism immediately implies an $n(n + 1)/2$ approximation ratio for the makespan objective. In Section 3 we provide a (non-truthful) algorithm which performs considerably better than the best truthful mechanism; even the worst pure equilibrium/outcome of our algorithm achieves an optimal makespan, i.e. our algorithm has a pure PoA of 1. Pure Nash equilibria constitute a very strong solution concept, that in general need not even exist, and they are also very desirable from a mechanism designer's perspective. For our algorithm we actually prove both existence of pure Nash equilibria and almost optimal performance guarantees for *all* of them.¹ If we run this algorithm independently for each job, we get a task-independent and anonymous algorithm yielding a PoA of n for any number of tasks, which we show to be asymptotically optimal. Next, revisiting truthfulness, in Section 4 we also show that the mechanism inspired by the LP relaxation of the problem is provably truthful and achieves an optimal approximation ratio 1 for the fractional scheduling problem of divisible tasks while providing an n -approximation when interpreted as a randomized mechanism. This almost matches the lower bound of $(n + 1)/2$ for truthful mechanisms known from [18]. Finally, in Section 5 we briefly study the more optimistic objective of minimizing the makespan at the best possible equilibrium (instead of the worst one used in the Price of Anarchy metric) and show that the natural greedy algorithm achieves an optimal Price of Stability.

2 Model and Notation

We have a set $N = \{1, 2, \dots, n\}$ of unrelated parallel machines and m tasks/jobs that need to be scheduled to these machines. Throughout the text we assume that matrix $\mathbf{t} \in \mathbb{R}_{\geq 0}^{n \times m}$ denotes the true execution times, i.e. $t_{i,j}$ is the time machine i needs to execute task j . This is private knowledge of each machine i . Let $\hat{\mathbf{t}}$ denote the corresponding (not necessarily true) *declarations* of the machines for these costs.

A (randomized) *allocation protocol* takes as input the machines' declarations $\hat{\mathbf{t}}$ and outputs

¹Clearly, pure equilibria are a refinement of mixed ones so, from the perspective of mixed Nash equilibrium analysis, our positive results can additionally be seen as an (almost) optimal Price of Stability guarantee (see also the discussion at the start of Section 5).

an allocation \mathbf{A} of tasks to machines where A_{ij} is a 0–1 random variable indicating whether or not machine i gets allocated task j and \mathbf{a} is the corresponding probability distribution of allocation, i.e. $a_{i,j} = \Pr[A_{i,j} = 1]$ where of course $\sum_{i=1}^n a_{i,j} = 1$ for any task j .

We assume that machines are *bound by their declarations*: if a machine i is allocated some task j , then she will execute the task for time $\max\{t_{i,j}, \hat{t}_{i,j}\}$. So, the expected cost/workload of machine i is defined as

$$C_i(\hat{\mathbf{t}}|\mathbf{t}_i) = \sum_{j=1}^m a_{i,j}(\hat{\mathbf{t}}) \max\{\hat{t}_{i,j}, t_{i,j}\}, \quad (1)$$

while the *makespan* is computed as the expected maximum execution time

$$\mathcal{M}(\hat{\mathbf{t}}|\mathbf{t}_i) = \mathbb{E}_{\mathbf{A} \sim \mathbf{a}(\hat{\mathbf{t}})} \left[\max_{i=1, \dots, n} \sum_{j=1}^m A_{i,j} \max\{\hat{t}_{i,j}, t_{i,j}\} \right].$$

To simplify notation, whenever the true execution times \mathbf{t} are clear from the context we will drop them and simply use $C_i(\hat{\mathbf{t}})$ and $\mathcal{M}(\hat{\mathbf{t}})$.

The allocation protocol is called *truthful*, or a truthful mechanism, if it does not give incentives to the machines to misreport their true execution costs. Formally, for every machine i and declarations matrix $\hat{\mathbf{t}}$,

$$C_i(\mathbf{t}_i, \hat{\mathbf{t}}_{-i}) \leq C_i(\hat{\mathbf{t}}),$$

where $(\mathbf{x}_i, \mathbf{y}_{-i})$ denotes the matrix of declarations where machine i has deviated to \mathbf{x}_i while all other machines report costs as in \mathbf{y} . The *approximation ratio* measures the performance of truthful mechanisms and is defined as the maximum ratio, over all instances, of the objective value (makespan) under that mechanism over the optimal objective value achievable by a centralized solution which ignores the truthfulness constraint.

If an allocation protocol is not truthful (we simply refer to it as an algorithm), we measure its performance by the quality of its Nash equilibria: the states from which no player has the incentive to unilaterally deviate. The *Price of Anarchy* (PoA) [19] is established as a meaningful benchmark and captures the maximum ratio, over all instances, of the objective value of the worst equilibrium over that of the optimal centralized solution that ignores the machines' incentives. For most part of this paper we restrict attention to pure Nash equilibria where the machines make deterministic reports about their execution costs, and we will from now on refer to them simply as *equilibria*. Then, the corresponding benchmark is called pure PoA. A more optimistic benchmark is the *Price of Stability* (PoS) [3, 32] which compares the objective value of the *best* equilibrium to the value of the optimal centralized solution.

The makespan objective is inherently different if we consider *divisible tasks*, i.e. fractional allocations. In that case, each machine is allocated a portion of each task by the protocol and the makespan is computed as the maximum of the execution times of the machines, namely

$$\mathcal{M}^f(\hat{\mathbf{t}}) = \max_{i=1, \dots, n} \sum_{j=1}^m \alpha_{i,j}(\hat{\mathbf{t}}) \max\{\hat{t}_{i,j}, t_{i,j}\}$$

where $\alpha_{i,j}(\hat{\mathbf{t}}) \in [0, 1]$ is the *fraction* of task j allocated to machine i . Again, it must be that $\sum_{i=1}^n \alpha_{i,j} = 1$ for any task j . Notice here that each fractional algorithm with allocation fractions α naturally gives rise to a corresponding randomized integral algorithm with allocation probabilities $\mathbf{a} = \alpha$, whose makespan is within a factor of n from the fractional one², i.e. for any cost vector \mathbf{t}

$$\mathcal{M}^f(\mathbf{t}) \leq \mathcal{M}(\mathbf{t}) \leq n \cdot \mathcal{M}^f(\mathbf{t}). \quad (2)$$

²This is due to the fact that for any random variables Y_1, Y_2, \dots, Y_n it is $\mathbb{E}[\max_i Y_i] \leq \mathbb{E}[\sum_i Y_i] = \sum_i \mathbb{E}[Y_i] \leq n \max_i \mathbb{E}[Y_i]$, and also $\max_i \mathbb{E}[Y_i] \leq \mathbb{E}[\max_i Y_i]$ due to the convexity of the max function.

	a_1	a_2
if $\hat{t}_1 = \hat{t}_2$	$\frac{1}{2}$	$\frac{1}{2}$
if $\hat{t}_1 < \hat{t}_2 < c \cdot \hat{t}_1$	$\frac{1}{L}$	$1 - \frac{1}{L}$
if $c \cdot \hat{t}_1 \leq \hat{t}_2$	$1 - \frac{1}{L} \frac{\hat{t}_1}{\hat{t}_2}$	$\frac{1}{L} \frac{\hat{t}_1}{\hat{t}_2}$

Figure 1: Algorithm $\mathcal{A}_{L,c}^{(2)}$ for scheduling a single task to two machines, parametrized by $L > 2$ and $c > 1$. The probability that machine $i = 1, 2$ gets the task is denoted by a_i , and \hat{t}_1, \hat{t}_2 are the reported execution times by the machines.

Except when clearly stated otherwise, in this paper we deal with the integral version of the scheduling problem.

Social Welfare. An alternative objective, very common in the Mechanism Design literature, is that of optimizing *social welfare*, i.e. minimizing the combined costs of all players: $\mathcal{W}(\hat{\mathbf{t}}) = \sum_{i=1}^n C_i(\hat{\mathbf{t}})$. It is not difficult to see² that the makespan and social welfare objectives are within a factor of n away, whatever the allocation algorithm \mathbf{a} and the input costs $\hat{\mathbf{t}}$ might be:

$$\mathcal{M}(\hat{\mathbf{t}}) \leq \mathcal{W}(\hat{\mathbf{t}}) \leq n \cdot \mathcal{M}(\hat{\mathbf{t}}). \quad (3)$$

Also notice that for the special case of a single task, since the job is eventually allocated entirely to some machine, the two objectives coincide no matter the number of machines n , i.e. $\mathcal{M}(\hat{\mathbf{t}}) = \mathcal{W}(\hat{\mathbf{t}})$. Because of that and the linearity of the social welfare with respect to the players' costs, it is easy to verify that all algorithms we present in this paper achieve optimal ratios of 1 for that objective, both with respect to equilibrium/PoA and truthfulness analysis (e.g. [Theorems 2](#) and [6](#)). We will not mention that explicitly again in the remaining of the paper and rather focus on makespan minimization, which is a more challenging objective for our scheduling problem.

3 Price of Anarchy

For clarity of exposition, we first describe our scheduling algorithm in the special case of just $n = 2$ machines (and one task) before presenting the algorithm for the general case of $n \geq 1$. Since we treat the case of only one task in this section, we use \hat{t}_i and t_i to denote the declared and true execution time of machine i , respectively, and use a_i to denote i 's allocation probability.

3.1 Warm Up: The Case of Two Machines

To simplify notation, throughout this section we will assume without loss of generality that $\hat{t}_1 \leq \hat{t}_2$, i.e. the input to our algorithm is sorted in nondecreasing order. Notice that the true execution times $\mathbf{t} = (t_1, t_2)$ do not have to preserve this ordering, since the highest bidding machine might very well in reality not have the fastest execution capabilities.

Our algorithm for the case of two machines, parameterized by two constants $L > 2$, $c > 1$, and denoted by $\mathcal{A}_{L,c}^{(2)}$, is defined by the allocation probabilities in [Fig. 1](#). Whenever parameter c is insignificant in a particular context³, we will just use $\mathcal{A}_L^{(2)}$.

³In such a case, as it is for example in the statement of [Theorem 1](#), one can simply pick e.g. $c = 1 + \frac{1}{L}$.

The main result of this section is the following theorem, showing that by choosing parameter L arbitrarily high, the above algorithm can achieve an optimal Price of Anarchy:

Theorem 1. *For the case of one task and two machines, algorithm $\mathcal{A}_L^{(2)}$ has a (pure) Price of Anarchy of at most $1 + \frac{1}{L}$ (for any $L > 2$).*

We break down the proof of [Theorem 1](#) in distinct claims.

Claim 1. *At any equilibrium $\hat{\mathbf{t}}$ the ratio of the two bids must be at least c , i.e. $\hat{t}_2 \geq c \cdot \hat{t}_1$.*

Proof. Without loss assume $\hat{t}_1 \neq 0$, since otherwise the claim is trivially true. First, assume for a contradiction that $\hat{t}_1 < \hat{t}_2 < c \cdot \hat{t}_1$. Then the machine with larger report would have an incentive to deviate to bid $t'_2 = \max\{c\hat{t}_1, t_2\}$:

$$C_2(\hat{\mathbf{t}}) = \left(1 - \frac{1}{L}\right) \max\{\hat{t}_2, t_2\} > \frac{1}{L} \hat{t}_1 = \frac{\hat{t}_1}{L t'_2} \max\{t'_2, t_2\} = C_2(\hat{t}_1, t'_2)$$

where the inequality holds since $L > 2$ and the final two equalities hold because the deviating bid t'_2 equals $\max\{c\hat{t}_1, t_2\}$. Thus $\hat{\mathbf{t}} = (\hat{t}_1, \hat{t}_2)$ could not have been an equilibrium under the assumption that $\hat{t}_1 < \hat{t}_2 < c \cdot \hat{t}_1$.

A similar contradiction can be obtained for the remaining case of $\hat{t}_1 = \hat{t}_2$. In this case, both machines have an incentive to deviate to a bid t'_1 such that $\frac{\hat{t}_1}{c} < t'_1 < \hat{t}_1$, since

$$C_1(\hat{\mathbf{t}}) = \frac{1}{2} \max\{\hat{t}_1, t_1\} \geq \frac{1}{2} \max\{t'_1, t_1\} > \frac{1}{L} \max\{t'_1, t_1\} = C_1(t'_1, \hat{t}_2).$$

We can conclude that indeed $\hat{t}_2 \geq c\hat{t}_1$ at any equilibrium. □

Claim 2. *At any equilibrium $\hat{\mathbf{t}}$ the machine with the larger report will never have underbid, i.e. $\hat{t}_2 \geq t_2$.*

Proof. Assume for a contradiction that $\hat{t}_2 < t_2$. Then

$$C_2(\hat{\mathbf{t}}) = \frac{\hat{t}_1}{L \hat{t}_2} \max\{\hat{t}_2, t_2\} = \frac{\hat{t}_1}{L \hat{t}_2} t_2 > \frac{\hat{t}_1}{L t_2} t_2 = C_2(\hat{t}_1, t_2),$$

the first equality holding due to [Claim 1](#) and the last one because $t_2 > \hat{t}_2 \geq c\hat{t}_1$. □

Claim 3. *At any equilibrium $\hat{\mathbf{t}}$ the smaller bid is given by $\hat{t}_1 = \min\{t_1, \frac{\hat{t}_2}{c}\}$.*

Proof. Assume for a contradiction that $\hat{t}_1 \neq t'_1 = \min\{t_1, \frac{\hat{t}_2}{c}\}$. Then, we will show that the lower bidding machine would have an incentive to deviate from \hat{t}_1 to t'_1 .

Indeed, first consider the case when $\hat{t}_1 < t'_1$. Then

$$C_1(\hat{\mathbf{t}}) = \left(1 - \frac{\hat{t}_1}{L \hat{t}_2}\right) \max\{\hat{t}_1, t_1\} > \left(1 - \frac{t'_1}{L \hat{t}_2}\right) \max\{t'_1, t_1\} = C_1(t'_1, \hat{t}_2).$$

In the remaining case of $\hat{t}_1 > t'_1 = \min\{t_1, \frac{\hat{t}_2}{c}\}$, because of [Claim 1](#) it must be that $t'_1 = t_1 < \hat{t}_1 \leq \frac{\hat{t}_2}{c}$, thus

$$C_1(\hat{\mathbf{t}}) = \left(1 - \frac{\hat{t}_1}{L \hat{t}_2}\right) \max\{\hat{t}_1, t_1\} = \left(1 - \frac{\hat{t}_1}{L \hat{t}_2}\right) \hat{t}_1 > \left(1 - \frac{t_1}{L \hat{t}_2}\right) t_1 = C_1(t_1, \hat{t}_2).$$

where the inequality holds since $x \mapsto \left(1 - \frac{x}{y}\right)x$ is a strictly increasing function for $x \in [0, \frac{y}{2}]$, and indeed $t_1 < \hat{t}_1 < \hat{t}_2 < \frac{L \hat{t}_2}{2}$. □

Claim 4. *At any equilibrium $\hat{\mathbf{t}}$ bidding must preserve the relative order of the true execution times, i.e. $t_1 \leq t_2$.*

Proof. For a contradiction assume that $t_2 < t_1$, and first consider the case when $t_2 < \hat{t}_1$. If we pick $t'_2 \in (\frac{\hat{t}_1}{c}, \hat{t}_1)$ we have

$$C_2(\hat{\mathbf{t}}) = \frac{\hat{t}_1}{L\hat{t}_2} \max\{\hat{t}_2, t_2\} = \frac{\hat{t}_1}{L} > \frac{1}{L} \max\{t'_2, t_2\} = C_2(\hat{t}_1, t'_2),$$

meaning that the higher bidding machine would have an incentive to deviate from \hat{t}_2 to t'_2 .

For the remaining case of $\hat{t}_1 \leq t_2 < t_1$, first note that if $t_1 \leq \frac{\hat{t}_2}{c}$ then by [Claim 3](#) we would immediately derive that $\hat{t}_1 = t_1$, which is a contradiction. Hence, by [Claim 1](#) we can assume that $\hat{t}_1 = \frac{\hat{t}_2}{c} < t_1$. Then, if $\hat{t}_2 > t_1$ we have that

$$C_1(\hat{\mathbf{t}}) = \left(1 - \frac{\hat{t}_1}{L\hat{t}_2}\right) \max\{\hat{t}_1, t_1\} = \left(1 - \frac{\hat{t}_1}{L\hat{t}_2}\right) t_1 > \frac{1}{L} t_1 = C_1(t_1, \hat{t}_2),$$

the inequality holding because $\frac{\hat{t}_1}{\hat{t}_2} \frac{1}{L} \leq \frac{1}{L} < \frac{1}{2}$, and if $\hat{t}_2 \leq t_1$ then, in the same way, for $t'_1 = \max\{t_1, c\hat{t}_2\}$

$$C_1(\hat{\mathbf{t}}) > \frac{1}{L} t_1 \geq \frac{\hat{t}_2}{L} = \frac{\hat{t}_2}{Lt'_1} \max\{t'_1, t_1\} = C_1(t'_1, \hat{t}_2).$$

□

Proof of [Theorem 1](#). [Claims 1](#) to [4](#) imply that the makespan (and thus also the social cost since we have a single task) of any allocation at equilibrium can be bounded by

$$\begin{aligned} \mathcal{M}(\hat{\mathbf{t}}) &= \left(1 - \frac{\hat{t}_1}{L\hat{t}_2}\right) \max\{\hat{t}_1, t_1\} + \frac{\hat{t}_1}{L\hat{t}_2} \max\{\hat{t}_2, t_2\} \\ &\leq \max\{\hat{t}_1, t_1\} + \frac{1}{L} \hat{t}_1 \\ &\leq \left(1 + \frac{1}{L}\right) t_1, \end{aligned}$$

where t_1 is the optimal makespan.

Also, it is important to mention that it can be verified that there exists at least one (pure Nash) equilibrium, e.g. reporting $\hat{t}_1 = t_1$ and $\hat{t}_2 = \max\{Lc \cdot t_1, t_2\}$. □

3.2 The General Case

The algorithm for two machines (and a single task) can be naturally generalized to the case of any number of machines $n \geq 2$. We note that the essence of the techniques and the core ideas we presented in [Section 3.1](#) carry over to the general case. So, for clarity of exposition, we only give the definition of the algorithm here and the proof can be found in [Appendix A](#).

To present our algorithm $\mathcal{A}_{L,c}$ we first need to add some notation. We use \hat{t}_{\min} and \hat{t}_{\sec} to denote the smallest and second smallest declarations in $\hat{\mathbf{t}}$, and N_{\min} , N_{\sec} the corresponding sets of machine indices that make these declarations. (If $N = N_{\min}$, i.e. all machines make the same declaration we define $\hat{t}_{\sec} = \hat{t}_{\min}$). Also let $n_{\min} = |N_{\min}|$ and $n_{\sec} = |N_{\sec}|$.

Our main algorithm $\mathcal{A}_{L,c}$ for the case of one task and n machines, parameterized by $L > 2(n-1)$, $c > 1$, is defined by the allocation probabilities a_i for each machine $i \in N$ given in [Fig. 2](#).

As the following theorem suggests, by picking a high enough value for L the above algorithm can achieve an optimal performance under equilibrium:

	$i \in N_{\min}$	$i \in N_{\sec}$	$i \in N \setminus (N_{\min} \cup N_{\sec})$
if $\hat{t}_{\min} = \hat{t}_{\sec}$	$\frac{1}{n}$	$\frac{1}{n}$	$\frac{1}{n}$
if $\hat{t}_{\min} < \hat{t}_{\sec} < c \cdot \hat{t}_{\min}$	$\left(\frac{1}{L}\right) / n_{\min}$	$\left(1 - \frac{1}{L}\right) / n_{\sec}$	0
if $\hat{t}_{\sec} \geq c \cdot \hat{t}_{\min}$	$\left(1 - \sum_{k \in N \setminus N_{\min}} \frac{\hat{t}_{\min}}{L \cdot \hat{t}_k}\right) / n_{\min}$	$\frac{\hat{t}_{\min}}{L \cdot \hat{t}_i}$	$\frac{\hat{t}_{\min}}{L \cdot \hat{t}_i}$

Figure 2: Algorithm $\mathcal{A}_{L,c}$ for scheduling a single task to $n \geq 2$ machines, parametrized by $L > 2(n-1)$ and $c > 1$. The first and second highest reported execution times by the machines are denoted by \hat{t}_{\min} and \hat{t}_{\sec} respectively, while N_{\min} , N_{\sec} denote the corresponding sets of machine indices, and n_{\min} , n_{\sec} their cardinalities.

Theorem 2. *For the problem of scheduling one task without payments to $n \geq 2$ machines, algorithm \mathcal{A}_L has a (pure) Price of Anarchy of at most $1 + \frac{n-1}{L}$ (for any $L > 2(n-1)$).*

Multiple Tasks. It is not difficult to extend our single-task algorithm and the result of [Theorem 2](#) to get a task-independent, anonymous algorithm with a pure PoA of n for any number of tasks $m \geq 1$: simply run \mathcal{A}_L *independently* for each job. Then, the equilibria of the extended setting correspond exactly to players not having an incentive to deviate for any task/round, and the approximation ratio of $1 + \frac{n-1}{L}$ with respect to the minimum cost $\min_i t_{i,j}$ at every such round $j = 1, \dots, m$, guarantees optimality with respect to the social welfare and thus provides indeed a worst-case n -approximation for the makespan objective (see [Eq. \(3\)](#)). The following result shows that this is in fact (asymptotically) optimal for the natural class of anonymous⁴, task-independent algorithms:

Theorem 3. *No task-independent, anonymous algorithm for scheduling without payments on n machines can have a (pure) Price of Anarchy better than $\frac{n}{2} - o(1)$.*

Proof. Fix an algorithm that allocates tasks by running an anonymous single-task algorithm \mathcal{A}_j , independently for each task j . Each algorithm \mathcal{A}_j takes as input a declared cost vector $\hat{\mathbf{t}} = (\hat{t}_1, \dots, \hat{t}_n)^\top$ by the machines, where \hat{t}_i is the report of machine i for the task, $i = 1, \dots, n$. Of course, there is also an underlying *true* cost vector \mathbf{t} . We will study cost vectors \mathbf{t}^j under which the corresponding task j has a high execution time $M > n^3$ for all machines except two of them, namely $i = 1$ and $i = j$. Formally, for any $j = 1, \dots, n$ define:

$$t_i^j = \begin{cases} 1, & \text{if } i = 1 \text{ or } i = j, \\ M, & \text{otherwise.} \end{cases}$$

Let $\hat{\mathbf{t}}^j$ be a (pure) Nash equilibrium strategy vector under algorithm \mathcal{A}_j , when the true execution costs are \mathbf{t}^j .

First we notice that we can safely assume that, under any such true instance \mathbf{t}^j , the probability of \mathcal{A}_j allocating the task to one of the fast machines, at equilibrium, has to be at least $a_1(\hat{\mathbf{t}}^j) + a_j(\hat{\mathbf{t}}^j) \geq 1 - \frac{1}{n^2}$. Otherwise the expected makespan would be at least

⁴A scheduling algorithm is called *anonymous* if it does not discriminate among the players (with respect to their identities). Formally, any permutation of the players results in the same permutation to the output/allocation of the algorithm, i.e., for any permutation π of the set of player indices $\{1, 2, \dots, n\}$ and any input \mathbf{t} , it is $a(\pi \cdot \mathbf{t}) = \pi \cdot a(\mathbf{t})$, where the dot product $\pi \cdot \mathbf{t}$ denotes the permutation of the rows of matrix \mathbf{t} according to π .

$(1 - \frac{1}{n^2}) \cdot 1 + \frac{1}{n^2} \cdot M > n + 1 - \frac{1}{n^2} > \frac{n}{2}$ (while the optimal one is 1, by assigning the task to machine $i = 1$) thus easily proving the Price of Anarchy lower bound of the theorem. So, at equilibrium $\hat{\mathbf{t}}^j$, one of machines $i = 1$ and $i = k$ has to receive the job with probability at least $\frac{1}{2} (1 - \frac{1}{n^2}) = \frac{1}{2} - \frac{1}{2n^2}$. Assume that the machine with this property is $i = 1$; otherwise, due to the anonymity of algorithm \mathcal{A}_j , it is not difficult to see that the vector that results from $\hat{\mathbf{t}}^j$ by swapping its entries between positions $i = 1$ and $i = j$, and keeping all other machines fixed, is also a Nash equilibrium and furthermore has the desired property.

Now we are ready to construct the bad instance for our task-independent, anonymous algorithm. Consider $m = n$ jobs, job $j = 1, 2, \dots, n$ having true execution time vector \mathbf{t}^j . By the previous analysis for the single-task algorithms \mathcal{A}_j , $\hat{\mathbf{t}} = (\hat{t}^1, \hat{t}^2, \dots, \hat{t}^n)$ is a Nash equilibrium for our multi-task algorithm, and the probability that machine $i = 1$ is the one allocated task j is at least $\frac{1}{2} - \frac{1}{2n^2}$, for every step j . So, by utilizing a standard Chernoff bound⁵ it is straightforward to see that the probability that machine $i = 1$ gets more than $(1 - o(1)) \cdot n (\frac{1}{2} - \frac{1}{2n^2}) = (1 - o(1)) (\frac{n}{2} - o(1))$ jobs in total is at least $1 - o(1)$, resulting to an expected makespan of at least $\frac{n}{2} - o(1)$. The expected optimal makespan is 1: just allocate task j to machine j , for all $j = 1, 2, \dots, n$. \square

Although anonymity is definitely a very natural assumption in the scheduling domain (see e.g. the discussion in [5]), the above results do not rule out the possibility of an improved Price of Anarchy if the requirements of Theorem 3 are relaxed, and as a matter of fact we find the design of such algorithms, that will take into consideration non-trivial correlations between tasks and among machines, to be one of the most interesting and challenging directions for future work. However keep in mind that, as Theorem 9 in Appendix B.1 demonstrates, just relaxing the anonymity constraint would not suffice in order to hope for constant-factor performance.

4 Truthful Mechanisms

In this section we turn our attention to truthful algorithms for many tasks and provide a mechanism that achieves approximation ratio n , almost matching the $(n + 1)/2$ known lower bound on truthfulness [18]. The best known ratio before our work was $n(n + 1)/2$, achieved by running the algorithm of Koutsoupias [18] independently for each task. Unfortunately this guarantee turns out to be tight for the particular algorithm (see Appendix B.2 for a bad instance), thus here we have to devise more involved, *non* task-independent mechanisms.

4.1 The LP mechanism

It is a known fact that the LP relaxation of a problem can be a useful tool for designing mechanisms (both randomized and fractional). What the LP mechanism essentially does is that it forces all machines to have equal expected (minimum) cost. This aligns the goal of the designer with the objective of the machines, thus making it easier for truthfulness to be guaranteed; similar in spirit is the EQUAL COST mechanism of [13]. We recall that the LP relaxation for the scheduling problem is as shown in Fig. 3.

We denote an *optimal* solution⁶ to the above LP by $\alpha^{\text{LP}}(\mathbf{t})$, and the optimal objective value by $\mu_{\mathbf{t}}^{\text{LP}}$ (dropping the LP superscript whenever this is clear from the context). The vector $\alpha^{\text{LP}}(\mathbf{t})$

⁵ Use the tail inequality $\Pr[X \leq (1 - \delta)\mu] \leq e^{-\frac{\delta^2}{2}\mu}$ where X is the number of tasks that are allocated to machine $i = 1$, $\mu = n (\frac{1}{2} - \frac{1}{2n^2})$, and choose $\delta \leq n^{-\frac{1}{3}}$.

⁶Notice that although $\mu_{\mathbf{t}}^{\text{LP}}$ is unique, there might be various allocation fractions $\alpha_{i,j}$ that give rise to the optimal makespan $\mu_{\mathbf{t}}^{\text{LP}}$, in which case we can choose an arbitrary one for $\alpha_{\mathbf{t}}^{\text{LP}}$, e.g. take the lexicographically smaller.

$\text{minimize } \mu$ $\forall j : \quad \sum_{i=1}^n \alpha_{i,j} = 1 \quad (\text{each task is allocated entirely})$ $\forall i : \quad \mu - \sum_{j=1}^m \alpha_{i,j} t_{i,j} \geq 0 \quad (\text{the cost of each machine does not exceed makespan})$ $\forall i, j : \quad \alpha_{i,j} \geq 0 \quad (\text{the allocation probabilities are non-negative})$

Figure 3: The LP relaxation for the scheduling problem. Our LP mechanism is defined by an optimal solution $\alpha_{i,j}^{\text{LP}}(\mathbf{t})$ to this program.

can be straightforwardly interpreted as allocation probabilities or allocation fractions giving rise to a randomized or a fractional mechanism, respectively. We refer to the corresponding mechanisms as the LP randomized and the LP fractional mechanism. In [Theorem 4](#) we show that both mechanisms are truthful, hence, we can think of $\mu_{\mathbf{t}}^{\text{LP}}$ as corresponding to the maximum (expected) cost/workload perceived by any machine.

It is a simple observation that in an optimal solution the workload must be fully balanced among all machines and that μ^{LP} can only increase when all execution times increase, i.e. $\mu_{\mathbf{t}}^{\text{LP}} \leq \mu_{\mathbf{t}'}^{\text{LP}}$ for $\mathbf{t} \leq \mathbf{t}'$ (pointwise).

Note that the proof of [Theorem 4](#) is identical in both cases where the α corresponds to fractions or allocation probabilities. Hence, the result holds for both the LP randomized and the LP fractional mechanisms. Note that a strategy is (weakly) dominant for a machine if the strategy will result in a cost at most equal to the cost of any other strategy, for every choice of strategies/declarations of the other machines.

Theorem 4. *Under the LP (fractional or randomized) mechanism, truthfully reporting the execution times is a (weakly) dominant strategy for every machine.*

Proof. Recall that \mathbf{t} and $\hat{\mathbf{t}}$ denote the true and (some) declared execution times, respectively, of the machines. Fix some machine i and define matrix \mathbf{t}^{\max_i} as follows: row i , $\mathbf{t}_i^{\max_i}$, is the vector of point-wise maxima between true and declared times for machine i , that is

$$\mathbf{t}_i^{\max_i} = (\max\{\hat{t}_{i,1}, t_{i,1}\}, \max\{\hat{t}_{i,2}, t_{i,2}\}, \dots, \max\{\hat{t}_{i,m}, t_{i,m}\}),$$

while every other row $k \neq i$ is $\mathbf{t}_k^{\max_i} = \hat{\mathbf{t}}_k$, i.e. $\mathbf{t}^{\max_i} = (\mathbf{t}_i^{\max_i}, \hat{\mathbf{t}}_{-i})$. Seen as a matrix of declarations, \mathbf{t}^{\max_i} corresponds to machine i 's deviation from $\hat{\mathbf{t}}$ to $\mathbf{t}_i^{\max_i}$, while each other machine makes declarations according to $\hat{\mathbf{t}}$. Then we can derive the following:

$$\sum_{j=1}^m \alpha_{k,j}(\hat{\mathbf{t}}) t_{k,j}^{\max_i} = \sum_{j=1}^m \alpha_{k,j}(\hat{\mathbf{t}}) \hat{t}_{k,j} = \mu_{\hat{\mathbf{t}}} = \sum_{j=1}^m \alpha_{i,j}(\hat{\mathbf{t}}) \hat{t}_{i,j} \leq \sum_{j=1}^m \alpha_{i,j}(\hat{\mathbf{t}}) t_{i,j}^{\max_i}$$

and thus from the optimality of the LP solutions it must be that

$$\mu_{\mathbf{t}^{\max_i}} \leq \max_{l=1, \dots, n} \left\{ \sum_{j=1}^m \alpha_{l,j}(\hat{\mathbf{t}}) t_{l,j}^{\max_i} \right\} = \sum_{j=1}^m \alpha_{i,j}(\hat{\mathbf{t}}) t_{i,j}^{\max_i}.$$

Bringing everything together and taking into consideration that $(\mathbf{t}_i, \hat{\mathbf{t}}_{-i}) \leq \mathbf{t}^{\max_i}$ we get

$$C_i(\hat{\mathbf{t}}) = \sum_{j=1}^m \alpha_{i,j}(\hat{\mathbf{t}}) \max\{\hat{t}_{i,j}, t_{i,j}\} \geq \mu_{\mathbf{t}^{\max_i}} \geq \mu_{(\mathbf{t}_i, \hat{\mathbf{t}}_{-i})} = \sum_{j=1}^m \alpha_{i,j}(\mathbf{t}_i, \hat{\mathbf{t}}_{-i}) t_{i,j} = C_i(\mathbf{t}_i, \hat{\mathbf{t}}_{-i}),$$

which shows that indeed, whatever the declarations of the other machines $\hat{\mathbf{t}}_{-i}$, machine i is always (weakly) better off by truthfully reporting \mathbf{t}_i . \square

[Theorem 4](#) gives rise to the following two results.

Theorem 5. *The LP fractional mechanism has approximation ratio 1 for the fractional scheduling problem without money, for any number of machines and tasks.*

As discussed in [Section 2](#), by (2) we know that the above performance guarantee can deteriorate at most by a factor of n when we use the fractions as allocation probabilities for the integral case:

Theorem 6. *The LP randomized mechanism has approximation ratio at most n for (integrally) scheduling any number of tasks to n machines without money.*

4.2 The Proportional Mechanism

In this section we briefly consider the proportional mechanism which allocates to each machine i a $t_i^{-1} / \sum_{k=1}^n t_k^{-1}$ fraction of the task or probability of getting the task, respectively, depending on whether we consider the randomized or the fractional variant. In [\[18\]](#) it was shown that this algorithm is truthful and that its approximation ratio for randomized allocations of a single task is n . With the following theorem we wish to stress the difference between fractional and (randomized) integral allocations. The theorem is about the fractional case and proves the optimality of the proportional mechanism for scheduling one task without payments.

Theorem 7. *The proportional mechanism has an optimal approximation ratio of 1 for the fractional scheduling problem of a single task. For m tasks the approximation ratio increases to at least m .*

Proof. First consider the case of a single task. Under the proportional mechanism each machine i is allocated an $\alpha_i = t_i^{-1} / \sum_{k=1}^n t_k^{-1}$ fraction of the task and executes it for time t_i . Hence, all machines have the same execution time (makespan) equal to

$$\mathcal{M}^f = \frac{1}{\sum_i t_i^{-1}}.$$

In the optimal fractional allocation, all machines will have the same execution cost, otherwise we could remove an ε fraction of the task from a machine with high makespan and allocate it to a machine with small workload, hence reducing the makespan. This implies that for the optimal allocation fractions $\{\alpha_i^*\}_{i=1}^n$ it holds that $t_1 \alpha_1^* = t_2 \alpha_2^* = \dots = t_n \alpha_n^*$. Thus, for any machine i it is $\alpha_i^* = \frac{t_1}{t_i} \alpha_1^*$, and by the fact that $\sum_i \alpha_i^* = 1$ we get that $\sum_i t_i^{-1} = \frac{1}{\alpha_1^* t_1}$ and so

$$\mathcal{M}^f = \alpha_1^* t_1 = \text{OPT}^f.$$

However the above result does not generalize to the case of many tasks, where the proportional mechanism is run independently for each task. First, it is easy to see that this independence preserves truthfulness. Secondly, it also preserves the optimality of the proportional mechanism with respect to social welfare. But regarding makespan, we will prove a lower bound of m . Consider the following instance with m tasks and m machines. Every machine i can execute all tasks in time 1, except from the i -th task that can be run very quickly in time M^{-1} , where $M \gg 1$. Formally,

$$t_{ij} = \begin{cases} M^{-1}, & j = i \\ 1, & j \neq i. \end{cases}$$

Then, the proportional mechanism computes allocation fractions

$$\alpha_{ij} = \begin{cases} \frac{M}{M+m-1}, & j = i \\ \frac{1}{M+m-1}, & j \neq i, \end{cases}$$

which results to a makespan of

$$\sum_{j=1}^m \alpha_{ij} t_{ij} = M^{-1} \frac{M}{M+m-1} + (m-1) \frac{1}{M+m-1} = \frac{m}{M+m-1}.$$

On the other hand, the allocation that assigns each task j to its fastest machine j , i.e. $\alpha_{ij} = 1$ for $i = j$, results to a makespan of M^{-1} , giving an approximation ratio lower bound of $\frac{Mm}{M+m-1} \rightarrow m$ as $M \rightarrow \infty$. \square

5 Price of Stability and Mixed Equilibria

In this section we attempt a more optimistic approach regarding the problem of scheduling without payments. The Price of Stability can be used to measure the best possible performance of a given mechanism in practice, i.e., it can provide some indication on what is its potential power or, in other words, what is the best we can hope for given a specific mechanism. There are settings where this can be very useful, e.g. if the context allows the mechanism to suggest a specific strategy profile to the players, and the players then decide if it is in their best interest to follow the suggested strategies.

We consider the benchmark of the *best* (mixed Nash) equilibrium and prove that the following, most natural greedy algorithm can achieve optimality: allocate each task independently to the machine declaring the minimum cost (breaking ties arbitrarily).

Theorem 8. *The Price of Stability of the Greedy algorithm is 1 for scheduling without money any number of tasks to any number of machines.*

Proof. We will prove the stronger statement that every feasible (integral) allocation of m tasks to n machines can arise at some (mixed) Nash equilibrium of the Greedy mechanism, from which the theorem immediately follows. First we observe the following, not difficult to prove fact:

Fact. Fix some task and a nonnegative constant T . If all but one machine i^* play the mixed strategy of declaring (independently) a value $x \in [T, \infty)$ with probability distribution $F(x) = 1 - \sqrt[n]{T/x}$, and the remaining machine i^* declares deterministically any value $x^* > T$, then i^* 's declared cost will be the minimum among all declarations for the task with probability $(1 - F(x^*))^{n-1}$. Thus, under the Greedy mechanism, machine i^* will incur an expected cost of at least $x^*(1 - F(x^*))^{n-1} = T$ for executing the particular task.

Consider any true instance \mathbf{t} of the scheduling without payments problem and fix a particular allocation A of tasks to machines. For each task j let $i_j^* \in N$ denote the machine that j is assigned to under allocation A . Consider now the following strategy profile: machine i_j^* truthfully declares her cost for task j , i.e. $\hat{t}_{i_j^*,j} = t_{i_j^*,j} \equiv T_j$ deterministically, while all other machines $k \neq i_j^*$ each play (independently) a mixed strategy of declaring for task j a higher value of $\hat{t}_{k,j} = x > T_j$ with cumulative distribution $F_j(x) = 1 - \sqrt[n]{T_j/x}$. We argue that this constitutes a Nash equilibrium for the Greedy mechanism. Indeed, under these declarations the Greedy algorithm allocates each task j to machine i_j^* for a cost (of executing this task) of T_j . But by the previous fact, machine i_j^* cannot avoid incurring at least the same cost if misreporting any higher execution time $x^* > T_j$, while if she underbids $x^* < T_j$ she will obviously still get the task for effectively the same cost of $\max\{x^*, t_{i_j^*,j}\} = T_j$. \square

Acknowledgements. We thank the anonymous reviewers for their careful reading of our manuscript and their valuable feedback.

References

- [1] N. Alon, M. Feldman, A. D. Procaccia, and M. Tennenholtz. Strategyproof approximation of the minimax on networks. *Mathematics of Operations Research*, 35(3):513–526, 2010.
- [2] E. Angel, E. Bampis, F. Pascual, and A.-A. Tchetgnia. On truthfulness and approximation for scheduling selfish tasks. *Journal of Scheduling*, 12(5):437–445, 2009.
- [3] E. Anshelevich, A. Dasgupta, J. Kleinberg, É. Tardos, T. Wexler, and T. Roughgarden. The price of stability for network design with fair cost allocation. *SIAM J. on Computing*, 38(4):1602–1623, 2008.
- [4] A. Archer and É. Tardos. Truthful mechanisms for one-parameter agents. In *FOCS*, pages 482–491, 2001.
- [5] I. Ashlagi, S. Dobzinski, and R. Lavi. Optimal lower bounds for anonymous scheduling mechanisms. *Mathematics of Operations Research*, 37(2):244–258, 2012.
- [6] V. Auletta, R. De Prisco, P. Penna, and G. Persiano. The power of verification for one-parameter agents. In *ICALP*, pages 171–182, 2004.
- [7] G. Christodoulou, L. Gourvès, and F. Pascual. Scheduling selfish tasks: About the performance of truthful algorithms. In *COCOON*, pages 187–197, 2007.
- [8] G. Christodoulou, E. Koutsoupias, and A. Kovács. Mechanism design for fractional scheduling on unrelated machines. *ACM Trans. Algorithms*, 6(2):38:1–38:18, 2010.
- [9] E. H. Clarke. Multipart pricing of public goods. *Public Choice*, 11(1):17–33, 1971.
- [10] P. Dhangwatnotai, S. Dobzinski, S. Dughmi, and T. Roughgarden. Truthful approximation schemes for single-parameter agents. *SIAM J. Comput.*, 40(3):915–933, 2011.
- [11] S. Dughmi and A. Ghosh. Truthful assignment without money. In *EC*, pages 325–334, 2010.
- [12] D. Fotakis and C. Tzamos. Winner-imposing strategyproof mechanisms for multiple facility location games. In *WINE*, pages 234–245, 2010.
- [13] D. Fotakis and C. Tzamos. Strategyproof facility location for concave cost functions. In *EC*, pages 435–452, 2013.
- [14] Y. Giannakopoulos, E. Koutsoupias, and M. Kyropoulou. The anarchy of scheduling without money. In *9th International Symposium on Algorithmic Game Theory (SAGT)*, pages 302–314, 2016.
- [15] A. Gibbard. Manipulation of voting schemes: A general result. *Econometrica*, 41(4):587–601, 1973.
- [16] T. Groves. Incentives in Teams. *Econometrica*, 41(4):617–31, 1973.
- [17] M. Guo and V. Conitzer. Strategy-proof allocation of multiple items between two agents without payments or priors. In *AAMAS*, pages 881–888, 2010.

- [18] E. Koutsoupias. Scheduling without payments. *Theory of Computing Systems*, 54(3): 375–387, 2014.
- [19] E. Koutsoupias and C. Papadimitriou. Worst-case equilibria. *Computer Science Review*, 3 (2):65–69, 2009.
- [20] E. Koutsoupias and A. Vidali. A lower bound of $1+\varphi$ for truthful scheduling mechanisms. *Algorithmica*, 66(1):211–223, 2013.
- [21] R. Lavi and C. Swamy. Truthful and near-optimal mechanism design via linear programming. *J. ACM*, 58(6):25:1–25:24, 2011.
- [22] H. Levin, M. Schapira, and A. Zohar. Interdomain routing and games. In *STOC*, pages 57–66, 2008.
- [23] P. Lu, X. Sun, Y. Wang, and Z. A. Zhu. Asymptotically optimal strategy-proof mechanisms for two-facility games. In *EC*, pages 315–324, 2010.
- [24] H. Moulin. On strategy-proofness and single peakedness. *Public Choice*, 35(4):437–455, 1980.
- [25] A. Mu’alem and M. Schapira. Setting lower bounds on truthfulness: Extended abstract. In *SODA*, pages 1143–1152, 2007.
- [26] N. Nisan and A. Ronen. Algorithmic mechanism design. *Games and Economic Behavior*, 35(1/2):166–196, 2001.
- [27] N. Nisan, T. Roughgarden, E. Tardos, and V. V. Vazirani. *Algorithmic Game Theory*. Cambridge University Press, New York, NY, USA, 2007.
- [28] K. Nissim, R. Smorodinsky, and M. Tennenholtz. Approximately optimal mechanism design via differential privacy. In *ITCS*, pages 203–213, 2012.
- [29] P. Penna and C. Ventre. Optimal collusion-resistant mechanisms with verification. *Games and Economic Behavior*, 86:491 – 509, 2014.
- [30] A. D. Procaccia and M. Tennenholtz. Approximate mechanism design without money. In *EC*, pages 177–186, 2009.
- [31] M. A. Satterthwaite. Strategy-proofness and Arrow’s conditions: Existence and correspondence theorems for voting procedures and social welfare functions. *Journal of Economic Theory*, 10(2):187 – 217, 1975.
- [32] A. S. Schulz and N. S. Moses. On the performance of user equilibria in traffic networks. In *SODA*, pages 86–87, 2003.
- [33] W. Vickrey. Counterspeculation, auctions, and competitive sealed tenders. *Journal of Finance*, 16(1):8–37, 1961.

A Proof of Theorem 2

Analogously to the proof of [Theorem 1](#), we break down our exposition into distinct claims.

Claim 5. *At any equilibrium $\hat{\mathbf{t}}$ it holds that $\hat{t}_{\text{sec}} \geq c \cdot \hat{t}_{\text{min}}$, i.e., the minimum declaration differs from the remaining declarations by at least a factor of c .*

Proof. Consider an equilibrium $\hat{\mathbf{t}}$ and assume for a contradiction that $\hat{t}_{\text{min}} < \hat{t}_{\text{sec}} < c \cdot \hat{t}_{\text{min}}$ at an equilibrium. Consider a machine $i \in N_{\text{sec}}$ and her deviation to $\hat{t}'_i = \max\{c \cdot \hat{t}_i, t_i\}$. It holds that

$$C_i(\hat{\mathbf{t}}) = \left(1 - \frac{1}{L}\right) \frac{\max\{\hat{t}_{\text{sec}}, t_i\}}{n_{\text{sec}}} > \frac{1}{L} \hat{t}_{\text{min}} = \frac{\hat{t}_{\text{min}}}{L \cdot \hat{t}'_i} \hat{t}'_i = \frac{\hat{t}_{\text{min}}}{L \cdot \hat{t}'_i} \max\{\hat{t}'_i, t_i\} \geq C_i(\hat{t}'_i, \hat{\mathbf{t}}_{-i})$$

where the first inequality holds since $\hat{t}_{\text{sec}} > \hat{t}_{\text{min}}$, and $L > 2(n-1) \geq n_{\text{sec}} + 1$ by assumption, and the second inequality holds because i 's cost will be either $\frac{\hat{t}_{\text{min}}}{L \cdot \hat{t}'_i} \hat{t}'_i$ or 0 depending on whether there exists an other declaration in $(\hat{t}_{\text{min}}, c \cdot \hat{t}_{\text{min}})$. If $\hat{t}_{\text{min}} = \hat{t}_{\text{sec}}$, i.e. all machines declare \hat{t}_{min} , consider any machine i and let $\frac{\hat{t}_{\text{min}}}{c} < \hat{t}'_i < \hat{t}_{\text{min}}$. Then

$$C_i(\hat{\mathbf{t}}) = \frac{1}{n} \max\{\hat{t}_{\text{min}}, t_i\} > \frac{1}{L} \max\{\hat{t}_{\text{min}}, t_i\} \geq \frac{1}{L} \max\{\hat{t}'_i, t_i\} = C_i(\hat{t}'_i, \hat{\mathbf{t}}_{-i})$$

where the first inequality holds since $L > 2(n-1) \geq n$ for $n \geq 2$. We can conclude that, at any equilibrium of the game, \hat{t}_{sec} has to be at least as large as $c \cdot \hat{t}_{\text{min}}$. An immediate consequence is that all non-minimal declarations are at least equal to $c \cdot \hat{t}_{\text{min}}$. \square

Claim 6. *At any equilibrium $\hat{\mathbf{t}}$ it holds that $\hat{t}_i \geq t_i$ for any machine $i \in N \setminus N_{\text{min}}$, i.e. any machine whose declaration is not the minimum one, declares a value at least equal to its true execution cost t_i .*

Proof. Consider an equilibrium $\hat{\mathbf{t}}$ and assume for a contradiction that $\hat{t}_i < t_i$ for some machine $i \notin N_{\text{min}}$. Recall that from [Claim 5](#) we have that $\hat{t}_i \geq c \cdot \hat{t}_{\text{min}}$. But deviating from \hat{t}_i to t_i is always beneficial for this machine, as

$$C_i(\hat{\mathbf{t}}) = \frac{\hat{t}_{\text{min}}}{L \cdot \hat{t}_i} \max\{\hat{t}_i, t_i\} = \frac{\hat{t}_{\text{min}}}{L \cdot \hat{t}_i} t_i > \frac{\hat{t}_{\text{min}}}{L \cdot t_i} t_i = C_i(t_i, \hat{\mathbf{t}}_{-i}),$$

which proves our claim. \square

Claim 7. *At any equilibrium $\hat{\mathbf{t}}$ it holds that $n_{\text{min}} = 1$, i.e. only one machine makes the minimum declaration.*

Proof. Consider an equilibrium $\hat{\mathbf{t}}$ and assume for a contradiction that $n_{\text{min}} > 1$. Let i be any machine such that $i \in N_{\text{min}}$. It holds that

$$\begin{aligned} C_i(\hat{\mathbf{t}}) &= \left(1 - \sum_{k \in N \setminus N_{\text{min}}} \frac{\hat{t}_{\text{min}}}{L \cdot \hat{t}_k}\right) \frac{\max\{\hat{t}_{\text{min}}, t_i\}}{n_{\text{min}}} \\ &> \left(1 - \sum_{k \in N \setminus N_{\text{min}}} \frac{\hat{t}_{\text{min}}}{L \cdot \hat{t}_{\text{min}}}\right) \frac{\max\{\hat{t}_{\text{min}}, t_i\}}{n_{\text{min}}} \\ &= \left(1 - \frac{n - n_{\text{min}}}{L}\right) \frac{\max\{\hat{t}_{\text{min}}, t_i\}}{n_{\text{min}}} \end{aligned}$$

$$\begin{aligned}
&> \frac{1}{L} \max\{\hat{t}_{\min}, t_i\} \\
&\geq \frac{\hat{t}_{\min}}{L} \\
&= \frac{\hat{t}_{\min}}{L \cdot \max\{\hat{t}_{\text{sec}}, t_i\}} \max\{\hat{t}_{\text{sec}}, t_i\} \\
&= C_i(\max\{\hat{t}_{\text{sec}}, t_i\}, \hat{\mathbf{t}}_{-i})
\end{aligned}$$

where the first inequality holds since $\hat{t}_k > \hat{t}_{\min}$ for all $k \in N \setminus N_{\min}$, the second inequality holds since $L > 2(n-1) \geq n$ by assumption and the final equality holds since there would be at least one other declaration \hat{t}_{\min} and since by [Claim 5](#) we have that $\hat{t}_{\text{sec}} \geq c \cdot \hat{t}_{\min}$. \square

Claim 8. *At any equilibrium $\hat{\mathbf{t}}$ it holds that $\hat{t}_{\min} = \min\{t_i, \frac{\hat{t}_{\text{sec}}}{c}\}$ for machine $i = N_{\min}$, i.e., if i is the machine that makes the minimum declaration then that declaration is equal to $\min\{t_i, \frac{\hat{t}_{\text{sec}}}{c}\}$.*

Proof. Consider an equilibrium $\hat{\mathbf{t}}$ and note that i is well defined from [Claim 7](#). Assume for a contradiction that $\hat{t}_{\min} < \min\{t_i, \frac{\hat{t}_{\text{sec}}}{c}\}$. We have

$$\begin{aligned}
C_i(\hat{\mathbf{t}}) &= \left(1 - \sum_{k \in N \setminus N_{\min}} \frac{\hat{t}_{\min}}{L \cdot \hat{t}_k}\right) \frac{\max\{\hat{t}_{\min}, t_i\}}{n_{\min}} \\
&= \left(1 - \sum_{k \in N \setminus N_{\min}} \frac{\hat{t}_{\min}}{L \cdot \hat{t}_k}\right) t_i \\
&> \left(1 - \sum_{k \in N \setminus N_{\min}} \frac{\min\{t_i, \frac{\hat{t}_{\text{sec}}}{c}\}}{L \cdot \hat{t}_k}\right) t_i \\
&= C_i\left(\min\left\{t_i, \frac{\hat{t}_{\text{sec}}}{c}\right\}, \hat{\mathbf{t}}_{-i}\right),
\end{aligned}$$

where we have used $n_{\min} = 1$ proved in [Claim 7](#). Now, it suffices to show that $\hat{\mathbf{t}}$ is not an equilibrium if $\hat{t}_{\min} > \min\{t_i, \frac{\hat{t}_{\text{sec}}}{c}\}$. We know from [Claim 5](#), that at any equilibrium $\hat{t}_{\min} \leq \frac{\hat{t}_{\text{sec}}}{c}$, so we only need to consider the case where $t_i < \hat{t}_{\min} \leq \frac{\hat{t}_{\text{sec}}}{c}$. Then

$$\begin{aligned}
C_i(\hat{\mathbf{t}}|\mathbf{t}) &= \left(1 - \sum_{k \in N \setminus N_{\min}} \frac{\hat{t}_{\min}}{L \cdot \hat{t}_k}\right) \frac{\max\{\hat{t}_{\min}, t_i\}}{n_{\min}} \\
&= \left(1 - \sum_{k \in N \setminus N_{\min}} \frac{\hat{t}_{\min}}{L \cdot \hat{t}_k}\right) \hat{t}_{\min} \\
&> \left(1 - \sum_{j \in N \setminus N_{\min}} \frac{t_j}{L \cdot \hat{t}_k}\right) t_i \\
&= C_i(t_i, \hat{\mathbf{t}}_{-i}|\mathbf{t}),
\end{aligned} \tag{4}$$

where the inequality holds since $2\hat{t}_{\min} \sum_{k \in N \setminus N_{\min}} \frac{1}{L \cdot \hat{t}_k} < 2\hat{t}_{\min} \frac{n-1}{L \hat{t}_{\min}} < \frac{2(n-1)}{L} < 1$ and the derivative of (4) with respect to $\hat{t}_{\min} < \min_k t_k$ is positive for $L > 2(n-1)$. \square

Claim 9. *At any equilibrium $\hat{\mathbf{t}}$ it holds that $t_i \leq t_k$ for machine $i = N_{\min}$ and any machine $k \in N$, i.e., the true execution cost of the machine that makes the minimum declaration is indeed at most equal to the true execution cost of any other machine.*

Proof. Consider an equilibrium $\hat{\mathbf{t}}$ and assume for a contradiction that there exists some machine $k \in N \setminus N_{\min}$ such that $t_k < t_i$. We consider cases depending on the relative order of \hat{t}_{\min} , t_i , \hat{t}_i , and t_k . Assume that $t_k < \hat{t}_{\min}$ and let $\frac{\hat{t}_{\min}}{c} < \hat{t}'_k < \hat{t}_{\min}$. Then

$$C_k(\hat{\mathbf{t}}) = \frac{\hat{t}_{\min}}{L \cdot \hat{t}_k} \max\{\hat{t}_k, t_k\} = \frac{\hat{t}_{\min}}{L} > \frac{1}{L} \max\{\hat{t}'_k, t_k\} = C_i(\hat{t}'_i, \hat{\mathbf{t}}_{-i})$$

where the second equality holds by [Claim 6](#). This implies that at any equilibrium $\hat{\mathbf{t}}$ it holds that $t_k \geq \hat{t}_{\min}$. Recall that from [Claim 8](#) we know that $\hat{t}_{\min} = \min\left\{t_i, \frac{\hat{t}_{\sec}}{c}\right\}$, so if $t_i \leq \hat{t}_{\sec}/c$ we immediately get the desired inequality $t_k \geq t_i$. It remains to check the case $\frac{\hat{t}_{\sec}}{c} = \hat{t}_{\min} \leq t_k < t_i$. It holds that

$$\begin{aligned} C_i(\hat{\mathbf{t}}) &= \left(1 - \sum_{k \in N \setminus N_{\min}} \frac{\hat{t}_{\min}}{L \cdot \hat{t}_k}\right) \frac{\max\{\hat{t}_{\min}, t_i\}}{n_{\min}} \\ &= \left(1 - \sum_{k \in N \setminus N_{\min}} \frac{\hat{t}_{\min}}{L \cdot \hat{t}_k}\right) t_i \\ &> \left(1 - \frac{n-1}{L}\right) t_i \\ &> \frac{1}{L} t_i \end{aligned} \tag{5}$$

where the last inequality holds since $L > 2(n-1) \geq n$ by assumption. Note that if $t_i < \hat{t}_{\sec}$ then the right hand side of (5) is equal to the cost the machine would incur by declaring her true cost, i.e., $C_i(t_i, \hat{\mathbf{t}}_{-i} | \mathbf{t})$. Otherwise, if $t_i \geq \hat{t}_{\sec}$ then we get from (5) and for $\hat{t}'_i \geq \max\{t_i, c \cdot \hat{t}_{\sec}\}$ that

$$C_i(\hat{\mathbf{t}}) > \frac{1}{L} t_i \geq \frac{\hat{t}_{\sec}}{L \cdot \hat{t}'_i} \hat{t}'_i = \frac{\hat{t}_{\sec}}{L \cdot \hat{t}'_i} \max\{\hat{t}'_i, t_i\} \geq C_i(\hat{t}'_i, \hat{\mathbf{t}}_{-i}).$$

In each of the cases examined above we conclude that a beneficial deviation exists for the corresponding machine which is a contradiction to the original assumption that the starting configuration was at equilibrium. \square

Proof of [Theorem 2](#). We now use the above claims to bound the makespan at any equilibrium $\hat{\mathbf{t}}$ of the game. We denote by t_{\min} the minimum true execution cost of any machine (note that from [Claim 9](#) we know that the machine who makes the minimum declaration \hat{t}_{\min} indeed achieves the minimum true cost t_{\min}). Then, the expected makespan is

$$\begin{aligned} \mathcal{M}(\hat{\mathbf{t}}) &= \left(1 - \sum_{k \in N \setminus N_{\min}} \frac{\hat{t}_{\min}}{L \cdot \hat{t}_k}\right) \frac{\max\{\hat{t}_{\min}, t_{\min}\}}{n_{\min}} n_{\min} + \sum_{k \in N \setminus N_{\min}} \frac{\hat{t}_{\min}}{L \cdot \hat{t}_k} \max\{\hat{t}_k, t_k\} \\ &\leq \max\{\hat{t}_{\min}, t_{\min}\} + \frac{\hat{t}_{\min}}{L} (n-1) \\ &\leq t_{\min} + \frac{t_{\min}(n-1)}{L} \\ &\leq \left(1 + \frac{n-1}{L}\right) t_{\min}, \end{aligned}$$

where \hat{t}_{\min} is the optimal makespan. For the first equality we use [Claim 5](#), for the first inequality we use [Claims 6](#) and [7](#), and for the second inequality we use [Claims 8](#) and [9](#).

It can be easily verified that there exists at least one equilibrium $\hat{\mathbf{t}}$, e.g. choose $\hat{t}_{i^*} = t_{i^*} = t_{\min}$ for a machine $i^* \in \arg\min_i t_i$ that has a minimum *true* execution cost and $\hat{t}_k = \max\{Lc \cdot t_{\min}, t_k\}$ for all other machines $k \in N \setminus \{i\}$. \square

B Lower Bounds

B.1 Task-Independent Algorithms

Theorem 9. *No task-independent algorithm for scheduling without payments on n machines can have a Price of Anarchy better than $\frac{1}{2}\sqrt{n} - o(1)$.*

Proof. The bad instance here is simpler than the one in the proof of [Theorem 3](#): consider $m = n$ tasks, for all of which machine $i = 1$ has true execution cost 1 and all remaining machines are way slower with costs $M = \sqrt{n}$.

We can assume that at every single-task step there has to be an equilibrium where the fast machine $i = 1$ gets the job with probability at least $p = \frac{1}{2}$, otherwise the expected makespan of that single-shot algorithm would be at least $p \cdot 1 + (1 - p) \cdot M = \frac{1}{2}\sqrt{n} + \frac{1}{2}$, resulting indeed in the desired Price of Anarchy lower bound since the optimal cost is 1 (by allocating the task to machine $i = 1$).

Therefore, like in the proof of [Theorem 3](#), it is not difficult to deploy a Chernoff bound and see that the probability of machine $i = 1$ receiving a $p(1 - o(1)) = \frac{1}{2} - o(1)$ fraction of the overall number of jobs is at least $1 - o(1)$, thus resulting in an expected makespan of at least $\frac{n}{2} - o(1)$. The optimal allocation has a cost of $M = \sqrt{n}$, simply allocating each task j to machine j . This gives indeed a Price of Anarchy lower bound of $\frac{\frac{n}{2} - o(1)}{\sqrt{n}} = \frac{1}{2}\sqrt{n} - o(1)$. \square

B.2 The Mechanism of Koutsoupias [\[18\]](#)

Consider an instance of n machines and n tasks, machine i having (true) execution cost of 1 for task i and $M > 1$ (to be determined later) for all other tasks. Formally,

$$t_{i,j} = \begin{cases} 1, & j = i, \\ M, & j \neq i. \end{cases}$$

Consider running the (truthful) mechanism of Koutsoupias [\[18\]](#) independently for each task. Then (see [\[18, Eq. \(1\)\]](#)) the probability of a specific task j getting assigned to its unique “fast” machine (that is machine $i = j$ having execution cost 1) can be computed as

$$p = \int_0^1 \left(1 - \frac{y}{M}\right)^{n-1} dy = \frac{M}{n} \left[1 - \left(1 - \frac{1}{M}\right)^n\right].$$

Thus, the probability of at least one task being allocated to a “slow” machine is $1 - p^n$, since all tasks are being assigned independently. At such an event, the resulting makespan would be at least M (there would be at least one machine executing a task of duration M). So, the total expected makespan is at least $(1 - p^n)M$, which by selecting $M = n^3$ becomes

$$(1 - p^n)M = n^3 - n^{2n+3} \left[1 - \left(1 - \frac{1}{n^3}\right)^n\right]^n$$

and it is a matter of straightforward calculations to verify that for any number of machines $n \geq 2$ this is lower bounded by

$$(1 - p^n)M \geq \frac{n(n+1)}{2 + \frac{3}{n}} = \frac{n(n+1)}{2 + o(1)}.$$

Taking into consideration that the optimal makespan is 1 (by allocating each task j to machine j), this gives indeed a lower bound of $\frac{n(n+1)}{2+o(1)}$ to the expected makespan that asymptotically matches the mechanism’s upper bound of $\frac{n(n+1)}{2}$ given in [\[18\]](#).