Price of Anarchy in Sequencing Situations and the Impossibility to Coordinate

Herbert Hamers
Flip Klijn
Marco Slikker

August 2013

Barcelona GSE Working Paper Series

Working Paper nº 709

Price of Anarchy in Sequencing Situations and the Impossibility to Coordinate

Herbert Hamers* Flip Klijn† Marco Slikker ‡

August 29, 2013

Abstract

Scheduling jobs of decentralized decision makers that are in competition will usually lead to cost inefficiencies. This cost inefficiency is studied using the Price of Anarchy (PoA), i.e. the ratio between the worst Nash equilibrium cost and the cost attained at the centralized optimum. First, we provide a tight upperbound for the PoA that depends on the number of machines involved. Second, we show that it is impossible to design a scheduled-based coordinating mechanism in which a Nash equilibrium enforces the centralized or first best optimum. Finally, by simulations we illustrate that on average the PoA is relatively small with respect to the established tight upperbound.

keywords: Sequencing situations; outsourcing; first best solution; game theory; price of anarchy; coordinating mechanism

1 Introduction

Companies that produce advanced products in-house completely are becoming more and more scarce. Hence, apart from management of their own production facilities, companies have an increasing need to tightly control outsourced operations. Specialized suppliers may well be capable of providing high-quality parts that meet all product specifications. Interfering factors usually include the presence of multiple suitable suppliers, who all possibly serve other clients as well. One could think of multiple companies that produce cell phones and outsource the production of an essential chip to chip suppliers. In this paper we study such outsourcing decisions in a setting with multiple companies and multiple

*CentER and Department of Econometrics and Operations Research, Tilburg University.

†Institute for Economic Analysis (CSIC), Barcelona GSE, and CentER. The first draft of the paper was written while F. Klijn was visiting CentER and the Department of Econometrics and Operations Research, Tilburg University. He gratefully acknowledges the hospitality of Tilburg University and an extramural fellowship from CentER. Financial support from Plan Nacional I+D+i (ECO2011–29847) is also gratefully acknowledged.

‡Industrial Engineering and Innovation Sciences, Eindhoven University of Technology.

suppliers. In doing so, we acknowledge that outsourcing decisions are made locally by companies and that these companies are wary of long waiting times in the outsourcing process. We analyze the impact of decentralized decision making on performance vis--vis the centralized (first-best) solution and on the possibility to mitigate the difference. We model this situation by a sequencing problem with multiple decision makers, corresponding to the companies, who execute their (outsourcing) jobs on multiple machines, corresponding to the suppliers.

The importance of this type of problem in industrial and service organizations is evident and has been the driving force for significant progress in its analysis. Traditionally, the presence of multiple decision makers is neglected and the focus is on finding schedules that optimize a common goal. For an overview on sequencing problems with one decision maker and its applications we refer to Lawler *et al.* (1993) and Pinedo (2002).

The presence of multiple decision makers is not restricted to a setting with multiple companies, but appears for example in companies with multiple profit centers as well. Independently of the setting under consideration, decision makers can aim for coordination in favor of a common goal or each decision maker can focus on its own performance. Even though we take the second approach, the first approach has its relevance, for example in case binding agreements can be made. This cooperative approach has led to an established line of research that started with Curiel *et al.* (1989). They introduced one machine sequencing situations in which each job belongs to a different decision maker that has to be processed on a single machine. Each job has its own specific processing time and its costs are linear in completion time. Starting with a prescribed initial order, not necessarily optimal with respect to the weighted sum cost criterion, optimal orders are established for each coalition of players, which defines in a natural way the corresponding cooperative sequencing game. Several core stable allocation rules of the grand coalition are proposed. This model has been extended in different ways by considering restrictions on jobs such as ready times (Hamers *et al.* (1995)), different cost criteria (Curiel *et al.* (1994)), multiple machines such as job shop (van den Nouweland *et al.* (1992), multiple jobs owned by a player (Fernandez *et al.* (2008)), and multiple rearrangements (Slikker (2006)). This line of research is applied in industrial and service management by Cai and Vairaktarakis (2012) and Aydinliyim and Vairaktarakis (2010) who use this cooperative approach in a setting where coordination of outsourced operations plays a central role and by Hall and Liu (2010) in a setting where allocation and scheduling issues are combined in supply chains.

Coordinated decision making can have important benefits but it does not seem natural in all instances. For example, if binding agreements are not enforceable (or not allowed by competition law) or if the implementation of a coordinating policy is labor-intensive and therefore costly. The role of individual decision makers has, despite its evident importance, received increased attention in the last decade only.

One stream in this recent non-cooperative literature assumes complete information. More precisely, it is assumed that there is complete information about all the inputs of the scheduling problem. In this setting Bukchin and Hanany (2007) study the decentralization costs, i.e. the ratio between the best Nash equilibrium cost (or, equivalently, the

Nash equilibrium with lowest cost) and the cost attained at the centralized optimum, of a dispatching-sequencing problem. Each decision maker has a set of jobs that can be processed either in-house, which is less costly, or can be send to a subcontractor, which is more costly. They provide bounds for the decentralization costs for an arbitrary number of jobs and decision makers. Moreover, they introduce a scheduling-based coordinating mechanism such that the centralized optimum is obtained in a Nash equilibrium. Bukchin and Hanany (2011) consider a decentralized job shop scheduling system. Analyzing the bounds of the decentralization costs they propose a mechanism to reduce these costs. Vairaktarakis (2013) investigates scheduling situations in which a part of the workload can be subcontracted to a so-called third party. He develops pure Nash equilibria schedules for different production protocols. The Price of Anarchy (PoA) is the ratio between the worst Nash equilibrium cost (or, equivalently, the Nash equilibrium with highest cost) and the cost attained at the centralized optimum, and has been studied in parallel to the decentralized costs (dealing with the best Nash equilibrium). Koutsoupias and Papadimitriou (1999) introduce the PoA and show that the PoA is at most $\frac{3}{2}$ if there are two identical machines and the objective criterion is the makespan. Immorlica *et al.* (2009) consider general multiple machine scheduling situations with objective criterion the makespan. They provide mechanisms that minimize the PoA. Bounds for the PoA in scheduling situations with objective criterion the weighted completion time criterion are provided in Correa and Queyranne (2012) and Cole *et al.* (2013). For sequencing situations with the minsum objective, i.e. the sum of completion times, bounds for the PoA are provided in Hoeksma and Uetz (2012).

Another stream in the non-cooperative literature assumes there is incomplete information. In such a setting, decision makers have private information about the processing times of their jobs. An important goal in this literature is the design of mechanisms that induce the revelation of private information. Given the complete information focus of the current paper we only refer to Hain and Mitra (2004) who introduce a Vickrey-Clarke-Groves mechanism that enforces the decision makers to tell their true processing time and results in the centralized optimum, and to Heydenreich *et al.* (2006) who studies an online machine scheduling problem. Surveys that cover both complete and incomplete information literature are provided by Heydenreich *et al.* (2007) and Aydinliyim and Varaktarakis (2011).

In the multiple machine scheduling problem of the current paper, we focus on the comparison of the first-best outcome with a decentralized outcome as well. We do this by considering the price of anarchy which is an established measure in many fields (see, e.g., in serial cost sharing Moulin (2008), in congestion situations Roughgarden (2006). The multiple identical machine scheduling problems considered in this paper consist of a finite number of identical machines and a set of decision makers who each own a different set of jobs. Since the machines are identical, the processing times of the jobs are the same on each machine. A schedule is a production plan of the jobs that assigns each job to precisely one machine. The completion time of a job in a specific schedule indicates the cost of that job in this schedule. The cost of a player according to a schedule is attained using the minsum criterion. More precisely, the cost of a player is the sum of the completion times according to this schedule of the jobs that are owned by this player. Each player provides an allocation of his jobs to the machines independently of the other players. After this allocation on each machine the jobs are scheduled in order of shortest processing time. In the noncooperative game formulation the pure strategies of a player are allocations of his jobs to machines. The total cost, i.e. the sum of all player's cost, corresponding to a strategy profile is the sum of the player's cost that correspond to the schedule that is induced by the strategy profile. In this paper we identify a profile of mixed strategies that constitutes a worst Nash equilibrium. Exploiting the specific structure of this profile we show that an upperbound of the price of anarchy is $\frac{3m-1}{2m}$, where m denotes the number of machines. In fact, we prove that this upperbound is tight.

The identification of the possible gap between the performance of individual and coordinated decisions calls for an analysis to see whether this difference can be mitigated. In doing so, however, it would be very unnatural to consider all mechanisms. In extremis, a mechanism that penalizes all participants severely whenever they do not collectively end up with a schedule that is optimal from a centralized point of view would do the job, but does not seem a realistic mechanism. We therefore focus on mechanisms that satisfy several basic assumptions. More precisely. we assume there is no idle time between jobs and that the scheduling of a job does not depend on (the identity of) its owner. Additionally, we only require that all machines schedule similarly (machine anonymity) and that jobs are processed on the machines chosen by their owners (pre-schedule consistency). Imposing these natural constraints on possible coordinating mechanisms leads to an impossibility: there is no mechanism that satisfies these requirements and guarantees the existence of Nash equilibria that result in optimal schedules only. Hence, anarchy costs cannot be avoided via any "natural" mechanism.

Finally, we provide insight in the severity of the anarchy costs via a numerical experiment. In our simulations we find that the price of anarchy is non-negligible, identify instances close to the theoretical bound that are non-degenerate, and provide characteristics of the simulated distributions of the price of anarchy (mean, median, and 2.5th and 97.5th percentiles) that deviate significantly from both the theoretical upperbound derived in this paper and the lowerbound (which is 1). Insights into the impact of number of jobs, players, and machines are also provided.

Several papers in the aforementioned literature overview are in some way related to our paper, but they differ in at least one aspect different from our model. In many of the previously mentioned papers the objective criterion is the makespan whereas we consider the minsum. Moreover, in some papers the jobs have to be processed on all machines, whereas in our paper the jobs have to be processed on precisely one machine. Bukchin and Hanany (2007) and Hoeksma and Uetz (2012) are closest to the current paper. A key difference with the models of Bukchin and Hanany (2007) is found in the processing times of the jobs. They consider a single outsourcing machine on which the jobs have larger processing times than on their in-house machines. In our model, the processing times are all identical on each machine. Moreover, Bukchin and Hanany (2007) their focus is on decentralization costs whereas we analyze on the PoA. In contrast to our model, in their setting it is possible to find a coordinating mechanism. In the model of Hoeksma and Uetz (2012) the same cost criterion is considered. However, they assume that each decision maker owns exactly one job. In our setting, each decision maker may own a set of jobs. Moreover, they consider machines with different processing speed, whereas we assume that all machines have the same processing speed. Finally, they did not consider the possible design of a coordinating mechanism.

The paper is organized as follows. In Section 2, our model is formally introduced and some examples illustrate that there may not exist pure Nash equilibria or there may exist pure Nash equilibria that do not support the centralized optimum. A tight upperbound for the PoA is provided in Section 3. The impossibility theorem with respect to a scheduled-based coordinating mechanism is presented in Section 4. The behavior of the PoA is simulated in Section 5. Finally, Section 6 concludes.

2 Model

Let $M = \{1, \ldots, |M|\}$ be the finite set of (identical) **machines** with $|M| \geq 2$. Let N be the finite set of **agents**. Let J_i be the set of **jobs** owned by agent i. We assume that for all $i, j \in N$ with $i \neq j$, $J_i \cap J_j = \emptyset$. Let $J \equiv \cup_i J_i$. We assume all jobs have non–preemptive processing requirements. Each job $j \in J$ has processing time $p_j > 0$. To avoid degenerate situations that require cumbersome notation we assume that for $j, j' \in J$ with $j \neq j'$, $p_j \neq p_{j'}$. A **(scheduling) problem** is a quadruple $\Lambda = (M, N, (J_i)_{i \in N}, (p_j)_{j \in J})$. Whenever there is no possible confusion we omit the set of machines M from the specification of the scheduling problem.

A pre–schedule is an unordered assignment of the jobs to the machines. Formally, a (deterministic) **pre–schedule** is a function $\pi : J \to M$, where $\pi(j)$ indicates the machine on which job j is processed. Let Π be the set of pre–schedules. A schedule is an ordered assignment of the jobs to the machines. Formally, a (deterministic) **schedule** is a function $\sigma : J \to M \times \{1, \ldots, |J|\}$, where $\sigma(j) = (m, k)$ indicates that job j is scheduled in position k of machine m. We assume that on each machine there is no idle time between jobs nor before the first job. Given a schedule σ, job j's predecessors are the jobs $P(\sigma, j) = \{j' \in J : \sigma_1(j') = \sigma_1(j) \text{ and } \sigma_2(j') < \sigma_2(j)\}$. Then, job j's completion time can be written as

$$C_j(\sigma) = p_j + \left(\sum_{j' \in P(\sigma, j)} p_{j'} \right).$$

Each agent determines which of his jobs are processed on which machine. In other words, each agent i chooses $\pi_i : J_i \to M$, where $\pi_i(j)$ indicates the machine on which job j is processed. Then, the resulting pre–schedule is $\pi : J \to M$ with $\pi(j) = \pi_i(j)$ for each $i \in N$ and each job $j \in J_i$.

The central objective is to minimize the sum of completion times respecting the chosen pre–schedule. Let π be a pre–schedule. A **schedule $\boldsymbol{\sigma}$ respects pre–schedule $\boldsymbol{\pi}$** if for all $j \in J$, $\sigma_1(j) = \pi(j)$. A **schedule $\boldsymbol{\sigma^\pi}$ is $\boldsymbol{\pi}$–optimal** if it respects π and for any other schedule σ that respects π,

$$\sum_{j \in J} C_j(\sigma^\pi) \leq \sum_{j \in J} C_j(\sigma).$$

It is easy to see that σ^π is π–optimal if and only if for each machine m, jobs $\pi^{-1}(m)$ are scheduled in order of shortest processing time (SPT).

A **schedule $\boldsymbol{\sigma^*}$ is optimal** if for any other schedule σ,

$$\sum_{j \in J} C_j(\sigma^*) \leq \sum_{j \in J} C_j(\sigma).$$

The following algorithm can be used to find all optimal schedules.

Minimum Mean Flow Time[1] (MFT) algorithm. (Horowitz and Sahni, 1976)
For each machine m, set $l_m \equiv 0$. Set $J^* \equiv J$. As long as $J^* \neq \emptyset$, do Procedure.

BEGIN PROCEDURE.
Let $j^* \in J^*$ be such that $p_{j^*} > p_j$ for all $j \in J^*$. Let $m \in M$ be a machine with lowest l_m. Set $\pi^*(j^*) \equiv m$ and update $l_m \equiv l_m + 1$ and $J^* \equiv J^* \backslash \{j^*\}$.
END PROCEDURE.

Let σ^* be a π^*–optimal schedule.

Theorem 1. [Horowitz and Sahni, 1976] *A schedule is optimal if and only if it can be obtained from the minimum mean flow time (MFT) algorithm.*

We associate with each scheduling problem $(N, (J_i)_{i \in N}, (p_j)_{j \in J})$ a (non–cooperative) **scheduling game** $\Gamma = (N, (\Pi_i)_{i \in N}, (c_i)_{i \in N})$, which is explained next. For each $i \in N$, the set of **players** is given by N. The set of **(pure) strategies** of player i, denoted Π_i, is the collection of functions $\pi_i : J_i \to M$. With a slight abuse of notation, a strategy profile $\pi = (\pi_i)_{i \in N}$ straightforwardly induces a pre–schedule π. Since all processing times are distinct, π induces a **unique $\boldsymbol{\pi}$–optimal schedule $\boldsymbol{\sigma^\pi}$**. Player i's resulting "costs" are given by the sum of completion times of his jobs in σ^π. In other words, player i's **cost function** c_i is given by

$$c_i(\pi) \equiv \sum_{j \in J_i} C_j(\sigma^\pi).$$

Example 1.
Let $M = \{m_1, m_2\}$ and $N = \{1, 2\}$. Let $J_1 = \{a, c\}$ and $J_2 = \{b, d\}$. Suppose $(p_a, p_b, p_c, p_d) = (1, 2, 3, 4)$. Each player has 4 pure strategies: he can send both jobs to m_1, both jobs to m_2, or different jobs to different machines (two ways). Table 1 concisely depicts the scheduling game. Player 1 is the row player and each row indicates which jobs are sent to m_1 (the complement is sent to m_2). Player 2 is the column player and each column indicates which jobs are sent to m_1. Next, we illustrate that each pair

[1] Minimum mean flow time and minimum sum of completion times are equivalent objectives.

of numbers indicates the costs induced by the corresponding strategy–profile. Consider, for instance, the pair $(\{a\},\{b,d\})$, which corresponds with profile $\pi = (\pi_i)_{i=1,2}$ such that $\pi_1(a) = \pi_2(b) = \pi_2(d) = m_1$ and $\pi_1(c) = m_2$. Then, jobs a, b, and d end up together on m_1 and job c on m_2. The unique π–optimal schedule σ^π satisfies $\sigma^\pi(a) = (m_1, 1)$, $\sigma^\pi(b) = (m_1, 2)$, $\sigma^\pi(c) = (m_2, 1)$, and $\sigma^\pi(d) = (m_1, 3)$. Then, player 1's costs equal the sum of the completion times of his jobs a and c: $C_a(\sigma^\pi) + C_c(\sigma^\pi) = 1 + 3 = 4$. Similarly, player 2's costs equal the sum of the completion times of his jobs b and d: $C_b(\sigma^\pi) + C_d(\sigma^\pi) = (1 + 2) + (1 + 2 + 4) = 10$. Hence, in this case the costs of the two players are given by $(4, 10)$.

$1\backslash 2$	$\emptyset$	$\{b\}$	$\{d\}$	$\{b,d\}$
$\emptyset$	7,13	5,10	7,**7**	5,8
$\{a\}$	6,11	**4**,10	6,**7**	**4**,10
$\{c\}$	**4**,10	6,**7**	**4**,10	6,11
$\{a,c\}$	5,8	7,**7**	5,10	7,13

Table 1: Table of Example 1

The boldfaced numbers in Table 1 are related to the concept of Nash equilibrium, which is formally introduced next.

$\diamond$

Let $i \in N$. A **mixed strategy** $\tilde{\pi}_i$ of player i is a probability distribution over all pure strategies $\pi_i \in \Pi_i$. At mixed strategy $\tilde{\pi}_i$, let $Pr(\pi_i|\tilde{\pi}_i)$ be the probability assigned to pure strategy $\pi_i \in \Pi_i$. Let $\tilde{\pi} = (\tilde{\pi}_i)_{i\in N}$ be a profile of mixed strategies. For any deterministic pre–schedule $\pi \in \Pi$, let $Pr(\pi|\tilde{\pi})$ be the probability of π according to $\tilde{\pi}$. Then, player i's expected "costs" can be written as

$$\tilde{c}_i(\tilde{\pi}) \equiv \sum_{j\in J_i} \tilde{C}_j(\tilde{\pi}) = \sum_{j\in J_i} \left(\sum_{\pi\in\Pi} Pr(\pi|\tilde{\pi}) C_j(\sigma^\pi) \right),$$

where we denote the expected completion time of $j \in J$ by

$$\tilde{C}_j(\sigma^{\tilde{\pi}}) = \sum_{\pi\in\Pi} Pr(\pi|\tilde{\pi}) C_j(\sigma^\pi).$$

A profile of mixed strategies is a Nash equilibrium if no player has a profitable deviation. Formally, a profile of mixed strategies $\tilde{\pi}$ is a **(Nash) equilibrium** if there exists no player $i' \in N$ with a strategy $\tilde{\pi}'_{i'}$ such that

$$\tilde{c}_{i'}(\tilde{\pi}') < \tilde{c}_{i'}(\tilde{\pi}),$$

where $\tilde{\pi}' \equiv (\tilde{\pi}'_{i'}, (\tilde{\pi}_i)_{i\neq i'})$. Let $\mathcal{E}(\Gamma)$ be the set of Nash equilibria of game Γ.

The reason why we consider mixed strategies is that not all games have a Nash equilibrium in pure strategies.

Example 2. (No Nash equilibrium in pure strategies.)[2]
Consider again the scheduling game discussed in Example 1. The boldfaced numbers in Table 1 indicate each player's *best responses* to the other player's strategies. For instance, player 1's strategy $\{c\}$ is the unique best response to player 2's strategy $\{d\}$ since any other strategy of player 1 yields higher costs for player 1: $\emptyset$ gives costs 7, $\{a\}$ gives costs 6, and $\{a, c\}$ gives costs 5, but $\{a\}$ gives costs 4.

It is easy to verify that for no strategy–profile each player plays a best response to the other player's strategy. Hence, there is no Nash equilibrium in pure strategies.

There is a unique Nash equilibrium in mixed strategies. To see this, note that for player 1 it is always strictly better to play $\{a\}$ than $\emptyset$, and it is always strictly better to play $\{c\}$ than $\{a, c\}$. Similarly, for player 2 it is always strictly better to play $\{b\}$ than $\emptyset$, and it is always strictly better to play $\{d\}$ than $\{b, d\}$. Therefore, in any Nash equilibrium, strategies $\emptyset$ (for both players), $\{a, c\}$ (for player 1), and $\{b, d\}$ (player 2) receive probability 0. Applying standard game–theoretic tools one can easily show that the strategy–profile in which each of the remaining strategies receives probability 1/2 constitutes the unique Nash equilibrium $\tilde{\pi}^-$ in mixed strategies. One easily verifies that the costs induced by the unique Nash equilibrium are $\tilde{c}_1(\tilde{\pi}^-) + \tilde{c}_2(\tilde{\pi}^-) = 13.5$, while the optimal costs are 13. $\diamond$

The next example shows that even if Nash equilibria in pure strategies exist, it is possible that none of them induces an optimal schedule.

Example 3. (Nash equilibria in pure strategies exist but induce sub–optimal schedules.)
Let $M = \{m_1, m_2\}$ and $N = \{1, 2\}$. Let $J_1 = \{a, e\}$ and $J_2 = \{b, c, d\}$. Let $(p_a, p_b, p_c, p_d, p_e) = (1, 2, 4, 6, 8)$. One easily verifies that $(\pi_i)_{i=1,2}$ with $\pi_1(a) = \pi_2(d) = m_1$ and $\pi_1(e) = \pi_2(b) = \pi_2(c) = m_2$ is a Nash equilibrium in pure strategies. The sum of completion times of the optimal π-respecting schedule σ^π is

$$[C_a(\sigma^\pi) + C_d(\sigma^\pi)] + [C_b(\sigma^\pi) + C_c(\sigma^\pi) + C_e(\sigma^\pi)] = [1 + 7] + [2 + 6 + 14] = 30.$$

Using MFT it follows that the sum of completion times of any optimal schedule equals 29. In Table 2 we give a profitable deviation for each profile of pure strategies that results in an optimal schedule. (We omit any schedule that is obtained from some listed schedule by switching all jobs on one machine with all jobs on the other machine.) For instance, consider the first row, which deals with profile $(\pi_i)_{i=1,2}$ such that $\pi_2(b) = \pi_2(d) = m_1$ and $\pi_1(a) = \pi_1(e) = \pi_2(c) = m_2$. The second column indicates the resulting π–optimal schedule σ^π. The third column gives a player i' that can profitably deviate. The fourth column provides the costs $c_{i'}(\pi)$ of player i' in the original profile π. The fifth column exhibits the pre–schedule π' after a deviation of player i'. The sixth column indicates the resulting π'–optimal schedule $\sigma^{\pi'}$. The last column provides the costs $c_{i'}(\pi')$ of player i' in profile π', showing that the deviation is indeed profitable (i.e. player i''s costs are lower).

[2]Examples 3 and 3 are not knife–edge in the sense that the processing times can slightly be varied without losing its features.

π	σ^{π}	i'	$c_{i'}(\pi)$	π'	$\sigma^{\pi'}$	$c_{i'}(\pi')$
$b,d \to m_1$ $a,c,e \to m_2$	$m_1 : b,d$ $m_2 : a,c,e$	1	$1+(1+4+8)$ $=14$	$\boldsymbol{a},b,d \to m_1$ $c,e \to m_2$	$m_1 : a,b,d$ $m_2 : c,e$	$1+(4+8)$ $=13$
$c,d \to m_1$ $a,b,e \to m_2$	$m_1 : c,d$ $m_2 : a,b,e$	1	$1+(1+2+8)$ $=12$	$\boldsymbol{a},c,d \to m_1$ $b,e \to m_2$	$m_1 : a,c,d$ $m_2 : b,e$	$1+(2+8)$ $=11$
$a,b,d \to m_1$ $c,e \to m_2$	$m_1 : a,b,d$ $m_2 : c,e$	2	$(1+2)+(1+2+6)+4$ $=16$	$a,d \to m_1$ $\boldsymbol{b},c,e \to m_2$	$m_1 : a,d$ $m_2 : b,c,e$	$(1+6)+2+(2+4)$ $=15$
$a,c,d \to m_1$ $b,e \to m_2$	$m_1 : a,c,d$ $m_2 : b,e$	2	$(1+4)+(1+4+6)+2$ $=18$	$a,d \to m_1$ $b,\boldsymbol{c},e \to m_2$	$m_1 : a,d$ $m_2 : b,c,e$	$(1+6)+2+(2+4)$ $=15$

Table 2: Table of Example 3

We conclude that the above scheduling problem induces a game that has Nash equilibria in pure strategies but they do not induce an optimal schedule. $\diamond$

Examples 2 and 3 show that Nash equilibria in pure strategies need not exist, and that if they do exist there still can be a performance loss with respect to the situation in which there would be a central authority.

3 Price of Anarchy

In this section, we first show that for any scheduling game there is a Nash equilibrium in mixed strategies. Then, we study the price of anarchy which is a worst-case measure of the inefficiency of selfish behavior. More precisely, the price of anarchy is the ratio between the highest costs across all Nash equilibria and the optimal costs. Formally, for a game Γ, the **price of anarchy** is defined as

$$PoA(\Gamma) = \frac{\max_{\tilde{\pi}\in\mathcal{E}(\Gamma)} \sum_{j\in J} \tilde{C}_j(\sigma^{\tilde{\pi}})}{\sum_{j\in J} C_j(\sigma^*)},$$

where σ^* is an optimal schedule.

We now show that for any scheduling game Γ, $\mathcal{E}(\Gamma) \neq \emptyset$. We do this by the construction of a Nash equilibrium in mixed strategies in two steps. In the first step, we find an optimal schedule for each individual player (assuming the other players are not present). In the second step, we define a mixed strategy for each player by randomly permuting his partition of jobs from the first step, i.e. switching partition elements between the machines (using a uniform distribution).

Let $(N, (J_i)_{i\in N}, (p_j)_{j\in J})$ be a scheduling problem. Let $i \in N$. Player i's **individual scheduling problem** $(\{i\}, J_i, (p_j)_{j\in J_i})$ is the scheduling problem obtained from the original scheduling problem by omitting all agents different from i and their corresponding jobs. Let σ^{i*} be an optimal schedule for player i's individual scheduling problem. For each $m \in M$, define $J_{im} \equiv \{j \in J_i : \sigma_1^{i*}(j) = m\}$. Let $\tilde{\pi}_i^-$ be the mixed strategy of player i that puts the same positive probability[3] on each pure strategy π_i^- such that for each

[3]If $|J_i| \geq |M| - 1$, then this probability is $\frac{1}{|M|!}$.

$m \in M$ and each pair of jobs $j, j' \in J_i$,

$$\pi_i^-(j) = \pi_i^-(j') \Longleftrightarrow j, j' \in J_{im}.$$

Let π_i^- be a pure strategy that receives positive probability in $\tilde{\pi}_i^-$. By construction of σ^{i*}, the π_i^-–optimal schedule $\sigma^{\pi_i^-}$ for $(\{i\}, J_i, (p_j)_{j \in J_i})$ is optimal for $(\{i\}, J_i, (p_j)_{j \in J_i})$. We call $\boldsymbol{\tilde{\pi}^- = (\tilde{\pi}_i^-)_{i \in N}}$ **the uniformly distributed profile (obtained from the individually optimal schedules $\boldsymbol{\sigma^{i*}}$)**.

Before we show that $\tilde{\pi}^-$ is a worst Nash equilibrium, we establish the following lemma, which is used in the proof of Theorem 2.

Lemma 1. *For any pure strategy π_1° of player 1, any $j \in J_1$, and any $j' \in J \backslash J_1$,*

$$\sum_{\pi \in \Pi:\, \pi(j') = \pi_1^\circ(j)} Pr(\pi | (\pi_1^\circ, (\tilde{\pi}_i^-)_{i \in N \backslash \{1\}})) = \frac{1}{|M|}. \tag{1}$$

Similarly, for any mixed strategies $(\tilde{\pi}_i^\circ)_{i \in N \backslash \{1\}}$ of players $N \backslash \{1\}$, any $j \in J_1$, and any $j' \in J \backslash J_1$,

$$\sum_{\pi \in \Pi:\, \pi(j') = \pi(j)} Pr(\pi | (\tilde{\pi}_1^-, (\tilde{\pi}_i^\circ)_{i \in N \backslash \{1\}})) = \frac{1}{|M|}. \tag{2}$$

Proof. From the construction of the mixed strategies $(\tilde{\pi}_i^-)_{i \in N \backslash \{1\}}$ of the players in $N \backslash \{1\}$ it follows that for any pure strategy of player 1 and any job of player 1, any job of any other player ends up on the same machine as player 1's job with probability $1/M$. This proves the first statement. The second statement follows from similar arguments. □

We now show that the uniformly distributed profile $\tilde{\pi}^- = (\tilde{\pi}_i^-)_{i \in N}$ is a Nash equilibrium. In fact, Theorem 2 states that $\tilde{\pi}^-$ is a particularly interesting Nash equilibrium: for each player, the expected costs in the constructed Nash equilibrium is higher than those in any other Nash equilibrium. In other words, the constructed Nash equilibrium is a worst Nash equilibrium for all players.

Theorem 2. *Let $\Gamma = (N, (\Pi)_{i \in N}, (c_i)_{i \in N})$ be a scheduling game. Then,*
(i) $\tilde{\pi}^- \in \mathcal{E}(\Gamma)$;
(ii) $\tilde{\pi} \in \mathcal{E}(\Gamma) \Longrightarrow \tilde{c}_i(\tilde{\pi}^-) \geq \tilde{c}_i(\tilde{\pi})$ for each $i \in N$.

Proof. Consider $\tilde{\pi}^-$. We first prove (i). Without loss of generality we show that player 1 has no profitable deviation. Let $\pi_1^- \in \Pi_1$ be such that $Pr(\pi_1^- | \tilde{\pi}_1^-) > 0$, i.e. player 1 plays π_1^- with positive probability at mixed strategy $\tilde{\pi}_1^-$. Let π_1' be a pure strategy of player 1.

Let $\tilde{\pi}' = (\pi_1', (\tilde{\pi}_i^-)_{i \in N \setminus \{1\}})$. Then,

$$
\begin{aligned}
\tilde{c}_1(\tilde{\pi}') &= \sum_{j \in J_1} \tilde{C}_j(\sigma^{\tilde{\pi}'}) \\
&= \sum_{j \in J_1} \left(\sum_{\pi \in \Pi} Pr(\pi|\tilde{\pi}') C_j(\sigma^{\pi}) \right) \\
&= \sum_{j \in J_1} \left(\sum_{\pi \in \Pi} Pr(\pi|\tilde{\pi}') \left(p_j + \sum_{j' \in P(\sigma^{\pi}, j)} p_{j'} \right) \right) \\
&= \sum_{j \in J_1} \left(p_j + \sum_{\pi \in \Pi} Pr(\pi|\tilde{\pi}') \sum_{j' \in P(\sigma^{\pi}, j)} p_{j'} \right) \\
&= \sum_{j \in J_1} \left(p_j + \sum_{\pi \in \Pi} Pr(\pi|\tilde{\pi}') \left(\sum_{j' \in J_1 \cap P(\sigma^{\pi}, j)} p_{j'} + \sum_{j' \in (J \setminus J_1) \cap P(\sigma^{\pi}, j)} p_{j'} \right) \right) \\
&= \sum_{j \in J_1} \left(p_j + \left(\sum_{\pi \in \Pi} Pr(\pi|\tilde{\pi}') \sum_{j' \in J_1 \cap P(\sigma^{\pi}, j)} p_{j'} \right) + \left(\sum_{\pi \in \Pi} Pr(\pi|\tilde{\pi}') \sum_{j' \in (J \setminus J_1) \cap P(\sigma^{\pi}, j)} p_{j'} \right) \right) \\
&= \sum_{j \in J_1} \left(p_j + \sum_{j' \in J_1 \cap P(\sigma^{\pi'_1}, j)} p_{j'} \right) + \sum_{j \in J_1} \left(\sum_{\pi \in \Pi} Pr(\pi|\tilde{\pi}') \sum_{j' \in (J \setminus J_1) \cap P(\sigma^{\pi}, j)} p_{j'} \right) \\
&= \sum_{j \in J_1} \left(p_j + \sum_{j' \in J_1 \cap P(\sigma^{\pi'_1}, j)} p_{j'} \right) + \sum_{j \in J_1} \left(\sum_{\pi \in \Pi} Pr(\pi|\tilde{\pi}') \sum_{i \in N \setminus \{1\}} \sum_{j' \in J_i \cap P(\sigma^{\pi}, j)} p_{j'} \right) \\
&= \sum_{j \in J_1} \left(p_j + \sum_{j' \in J_1 \cap P(\sigma^{\pi'_1}, j)} p_{j'} \right) + \sum_{j \in J_1} \sum_{i \in N \setminus \{1\}} \left(\sum_{\pi \in \Pi} Pr(\pi|\tilde{\pi}') \sum_{j' \in J_i \cap P(\sigma^{\pi}, j)} p_{j'} \right) \\
&= \sum_{j \in J_1} \left(p_j + \sum_{j' \in J_1 \cap P(\sigma^{\pi'_1}, j)} p_{j'} \right) + \sum_{j \in J_1} \sum_{i \in N \setminus \{1\}} \left(\sum_{\pi \in \Pi} Pr(\pi|\tilde{\pi}') \sum_{j' \in J_i :\, p_{j'} < p_j, \pi(j') = \pi(j)} p_{j'} \right) \\
&= \sum_{j \in J_1} \left(p_j + \sum_{j' \in J_1 \cap P(\sigma^{\pi'_1}, j)} p_{j'} \right) + \sum_{j \in J_1} \sum_{i \in N \setminus \{1\}} \left(\sum_{j' \in J_i :\, p_{j'} < p_j} \left(\sum_{\pi \in \Pi :\, \pi(j') = \pi(j)} Pr(\pi|\tilde{\pi}') \right) p_{j'} \right) \\
&= \sum_{j \in J_1} \left(p_j + \sum_{j' \in J_1 \cap P(\sigma^{\pi'_1}, j)} p_{j'} \right) + \sum_{j \in J_1} \sum_{i \in N \setminus \{1\}} \left(\sum_{j' \in J_i :\, p_{j'} < p_j} \left(\sum_{\pi \in \Pi :\, \pi(j') = \pi'_1(j)} Pr(\pi|\tilde{\pi}') \right) p_{j'} \right) \\
&\overset{(a)}{=} \sum_{j \in J_1} \left(p_j + \sum_{j' \in J_1 \cap P(\sigma^{\pi'_1}, j)} p_{j'} \right) + \sum_{j \in J_1} \left(\sum_{i \in N \setminus \{1\}} \sum_{j' \in J_i : p_{j'} < p_j} \frac{p_{j'}}{|M|} \right)
\end{aligned}
$$

$$\overset{(b)}{\geq} \sum_{j\in J_1}\left(p_j + \sum_{j'\in J_1\cap P(\sigma^{\pi_1^-},j)} p_{j'}\right) + \sum_{j\in J_1}\left(\sum_{i\in N\setminus\{1\}}\ \sum_{j'\in J_i:p_{j'}<p_j} \frac{p_{j'}}{|M|}\right)$$
$$\overset{(c)}{=} \tilde{c}_1(\pi_1^-, (\tilde{\pi}_i^-)_{i\in N\setminus\{1\}}).$$

Equality (a) follows from (1) with $\pi_1^\circ = \pi_1'$. Inequality (b) follows from the fact that by construction $\sigma^{\pi_1^-}$ is optimal for $(\{1\}, J_1, (p_j)_{j\in J_1})$. Equality (c) follows from arguments similar to those applied to establish all previous equalities and (1) with $\pi_1^\circ = \pi_1^-$.

Since $\tilde{c}_1$ is linear, it follows that

$$\begin{aligned}
\tilde{c}_1(\pi_1', (\tilde{\pi}_i^-)_{i\in N\setminus\{1\}}) &= \tilde{c}_1(\tilde{\pi}') \\
&\geq \sum_{\pi_1^-\in\Pi_1} Pr(\pi_1^-|\tilde{\pi}_1^-)\, \tilde{c}_1(\pi_1^-, (\tilde{\pi}_i^-)_{i\in N\setminus\{1\}}) \\
&= \tilde{c}_1(\tilde{\pi}_1^-, (\tilde{\pi}_i^-)_{i\in N\setminus\{1\}}) \\
&= \tilde{c}_1(\tilde{\pi}^-),
\end{aligned}$$

which shows that deviation π_1' is not profitable. Hence, $\tilde{\pi}^-$ is a Nash equilibrium. This completes the proof of (i).

Next, we prove (ii). Without loss of generality we show that player 1 has higher costs at $\tilde{\pi}^-$ than at any other Nash equilibrium. Let $\tilde{\pi}$ be a Nash equilibrium. Then,

$$\begin{aligned}
\tilde{c}_1(\tilde{\pi}) &\leq \tilde{c}_1(\tilde{\pi}_1^-, (\tilde{\pi}_i)_{i\in N\setminus\{1\}}) \\
&\overset{(d)}{=} \sum_{\pi_1^-\in\Pi_1} Pr(\pi_1^-|\tilde{\pi}_1^-)\left(\tilde{c}_1(\pi_1^-, (\tilde{\pi}_i)_{i\in N\setminus\{1\}})\right) \\
&= \sum_{\pi_1^-\in\Pi_1} Pr(\pi_1^-|\tilde{\pi}_1^-)\left(\sum_{j\in J_1}\left(p_j + \sum_{j'\in J_1\cap P(\sigma^{\pi_1^-},j)} p_{j'}\right)\right) + \\
&\quad \sum_{\pi_1^-\in\Pi_1} Pr(\pi_1^-|\tilde{\pi}_1^-)\left(\sum_{j\in J_1}\left(\sum_{i\in N\setminus\{1\}}\ \sum_{j'\in J_i:p_{j'}<p_j}\ \sum_{\pi\in\Pi:\pi(j')=\pi_1^-(j)} Pr(\pi|(\pi_1^-, (\tilde{\pi}_i)_{i\in N\setminus\{1\}}))p_{j'}\right)\right) \\
&= \left(\sum_{\pi_1^-\in\Pi_1} Pr(\pi_1^-|\tilde{\pi}_1^-)\sum_{j\in J_1}\left(p_j + \sum_{j'\in J_1\cap P(\sigma^{\pi_1^-},j)} p_{j'}\right)\right) + \\
&\quad \sum_{j\in J_1}\left(\sum_{i\in N\setminus\{1\}}\ \sum_{j'\in J_i:p_{j'}<p_j}\ \sum_{\pi\in\Pi:\pi(j')=\pi(j)} Pr(\pi|(\tilde{\pi}_1^-, (\tilde{\pi}_i)_{i\in N\setminus\{1\}}))p_{j'}\right)
\end{aligned}$$

$$\stackrel{(e)}{=} \left(\sum_{\pi_1^- \in \Pi_1} Pr(\pi_1^- | \tilde{\pi}_1^-) \sum_{j \in J_1} \left(p_j + \sum_{j' \in J_1 \cap P(\sigma^{\pi_1^-}, j)} p_{j'} \right) \right) +$$

$$\sum_{j \in J_1} \left(\sum_{i \in N \setminus \{1\}} \sum_{j' \in J_i : p_{j'} < p_j} \frac{p_{j'}}{|M|} \right)$$

$$= \sum_{\pi_1^- \in \Pi_1} Pr(\pi_1^- | \tilde{\pi}_1^-) \left(\sum_{j \in J_1} \left(p_j + \sum_{j' \in J_1 \cap P(\sigma^{\pi_1^-}, j)} p_{j'} \right) + \sum_{j \in J_1} \left(\sum_{i \in N \setminus \{1\}} \sum_{j' \in J_i : p_{j'} < p_j} \frac{p_{j'}}{|M|} \right) \right)$$

$$\stackrel{(f)}{=} \sum_{\pi_1^- \in \Pi_1} Pr(\pi_1^- | \tilde{\pi}_1^-) \left(\tilde{c}_1(\pi_1^-, (\tilde{\pi}_i^-)_{i \in N \setminus \{1\}}) \right)$$

$$\stackrel{(g)}{=} \tilde{c}_1(\tilde{\pi}^-).$$

Here, the inequality follows from the fact that $\tilde{\pi}$ is a Nash equilibrium. Equalities (d) and (g) follow from the linearity of $\tilde{c}_1$. Equality (e) follows from (2) with $\tilde{\pi}_i^\circ = \tilde{\pi}_i$, $i \in N \setminus \{1\}$. Equality (f) follows from (c). Therefore, for any player the expected costs at $\tilde{\pi}^-$ are higher than those at any other Nash equilibrium $\tilde{\pi}$. This completes the proof of (ii). $\square$

Remark 1. Any choice of individually optimal schedules $(\sigma^{i*})_{i \in N}$ induces a uniformly distributed profile $\tilde{\pi}^- = (\tilde{\pi}_i^-)_{i \in N}$ which by Theorem 2 is a Nash equilibrium of the scheduling game. Even though different choices of $(\sigma^{i*})_{i \in N}$ can induce different Nash equilibria, again by Theorem 2, for each player all these Nash equilibria yield the same associated expected costs. $\diamond$

Let $i \in N$. Denote $J_i = \{j_{i1}, \ldots, j_{in_i}\}$. We denote the processing time of job j_{il} by p_{il} and assume that $p_{i1} < \cdots < p_{in_i}$. In view of our objective to determine the price of anarchy, we henceforth conveniently assume that for each $i \in N$, σ^{i*} is such that for each $m \in M = \{1, \ldots, |M|\}$,

$$\boldsymbol{J_{im}} = \{j \in J_i : \sigma_1^{i*}(j) = m\} = \{j \in J_i \, : \, j = j_{il} \text{ with } (l-1)\text{mod}|M| = m-1\}.$$

As an illustration, suppose player 1 has 7 jobs, i.e. $J_1 = \{j_{1,1}, \ldots, j_{1,7}\}$. If there are $|M| = 3$ machines, then $J_{1,1} = \{j_{1,1}, j_{1,4}, j_{1,7}\}$ are the jobs assigned to machine 1, $J_{1,2} = \{j_{1,2}, j_{1,5}\}$ are the jobs assigned to machine 2, and $J_{1,3} = \{j_{1,3}, j_{1,6}\}$ are the jobs assigned to machine 3.

For $j \in J$, let $\boldsymbol{o(j)}$ denote the **owner of job $\boldsymbol{j}$**, i.e. $o(j) = i$ where $i \in N$ is such that $j \in J_i$. For $j \in J$, define

$$\begin{aligned} \boldsymbol{\lambda_j} &\equiv |\{j' \in J : p_{j'} > p_j\}| \text{ and} \\ \boldsymbol{\kappa_j} &\equiv |\{j' \in J : p_{j'} > p_j \text{ and } o(j') = o(j)\}|. \end{aligned}$$

The next lemma provides a convenient expression for the sum of expected costs induced by the worst Nash equilibria in terms of $(\kappa_j, \lambda_j)_{j \in J}$ and the processing times $(p_j)_{j \in J}$.

Lemma 2. *The sum of expected costs induced by a worst Nash equilibrium $\tilde{\pi}^-$ is given by*[4]

$$\sum_{j\in J} \tilde{C}_j(\tilde{\pi}^-) = \sum_{j\in J} p_j \left(1 + \frac{\lambda_j - \kappa_j}{|M|} + \left\lfloor \frac{\kappa_j}{|M|} \right\rfloor\right).$$

Proof. From equalities (f) and (g) in the proof of Theorem 2 it follows that

$$\begin{aligned}
\sum_{j\in J} \tilde{C}_j(\sigma^{\tilde{\pi}^-}) &= \sum_{i\in N} \tilde{c}_i(\tilde{\pi}^-) \\
&= \sum_{i\in N} \sum_{\pi_i^-\in\Pi_i} Pr(\pi_i^-|\tilde{\pi}_i^-) \left[\sum_{j\in J_i}\left(p_j + \sum_{j'\in J_i\cap P(\sigma^{\pi_i^-},j)} p_{j'}\right) + \sum_{j\in J_i}\left(\sum_{i'\in N\setminus\{i\}} \sum_{j'\in J_{i'}:p_{j'}<p_j} \frac{p_{j'}}{|M|}\right)\right] \\
&= \sum_{i\in N} \sum_{\pi_i^-\in\Pi_i} Pr(\pi_i^-|\tilde{\pi}_i^-) \left(\sum_{m\in M}\sum_{j\in J_{im}}\left(p_j + \sum_{j'\in J_i\cap P(\sigma^{\pi_i^-},j)} p_{j'}\right)\right) + \\
&\quad \sum_{i\in N} \sum_{\pi_i^-\in\Pi_i} Pr(\pi_i^-|\tilde{\pi}_i^-) \left(\sum_{j\in J_i}\sum_{i'\in N\setminus\{i\}} \sum_{j'\in J_{i'}:p_{j'}<p_j} \frac{p_{j'}}{|M|}\right) \\
&= \sum_{i\in N}\sum_{m\in M}\sum_{j\in J_{im}}\left(p_j + \sum_{j'\in J_{im}:p_{j'}<p_j} p_{j'}\right) + \sum_{i\in N}\sum_{j\in J_i}\sum_{i'\in N\setminus\{i\}}\sum_{j'\in J_{i'}:p_{j'}<p_j}\frac{p_{j'}}{|M|} \\
&= \sum_{i\in N}\sum_{m\in M}\sum_{j\in J_{im}}\left(p_j + \sum_{j'\in J_{im}:p_{j'}>p_j} p_j\right) + \sum_{i\in N}\sum_{j\in J_i}\sum_{i'\in N\setminus\{i\}}\sum_{j'\in J_{i'}:p_{j'}>p_j}\frac{p_j}{|M|} \\
&= \sum_{i\in N}\sum_{m\in M}\sum_{j\in J_{im}}\left(p_j(1 + \sum_{j'\in J_{im}:p_{j'}>p_j} 1)\right) + \sum_{i\in N}\sum_{j\in J_i}\frac{p_j}{|M|}\left(\sum_{i'\in N\setminus\{i\}}\sum_{j'\in J_{i'}:p_{j'}>p_j} 1\right) \\
&= \sum_{j\in J} p_j\left(1 + \left\lfloor\frac{\kappa_j}{|M|}\right\rfloor\right) + \sum_{j\in J}\frac{p_j}{|M|}(\lambda_j - \kappa_j) \\
&= \sum_{j\in J} p_j\left(1 + \frac{\lambda_j-\kappa_j}{|M|} + \left\lfloor\frac{\kappa_j}{|M|}\right\rfloor\right),
\end{aligned}$$

which proves the desired equality. □

The next lemma shows that for a given set of jobs the price of anarchy is maximal when all jobs are owned by different players. Let $\Lambda = (N, (J_i)_{i\in N}, (p_j)_{j\in J})$ be a scheduling problem. We define its associated **simple scheduling problem** by $\Lambda' \equiv (N', (J_i)_{i\in N'}, (p_j)_{j\in J})$ where N' is such that $|N'| = |J|$ and each agent in N' owns exactly one job in J, i.e. for each $i \in N'$, $|J_{i'}| = 1$.

[4]For $x \in \mathbb{R}$, $\lfloor x \rfloor$ denotes the largest integer n with $n \le x$.

Lemma 3. *Let Γ be the game associated with a scheduling problem Λ. Let Γ' be the game associated with the corresponding simple schedule problem Λ'. Then, $PoA(\Gamma) \leq PoA(\Gamma')$.*

Proof. Let σ^* be an optimal schedule (for both Λ and Λ'). Let $\tilde{\pi}^-$ be a worst Nash equilibrium for Λ, and let $\tilde{\tau}^-$ be a worst Nash equilibrium for Λ'. Then,

$$\begin{aligned}
\sum_{j \in J} \tilde{C}_j(\sigma^{\tilde{\pi}^-}) &= \sum_{j \in J} p_j \left(1 + \frac{\lambda_j - \kappa_j}{|M|} + \left\lfloor \frac{\kappa_j}{|M|} \right\rfloor\right) \\
&\leq \sum_{j \in J} p_j \left(1 + \frac{\lambda_j}{|M|}\right) \\
&= \sum_{j \in J} \tilde{C}_j(\sigma^{\tilde{\tau}^-}),
\end{aligned}$$

where the two equalities follow from Lemma 2. Therefore,

$$PoA(\Gamma) = \frac{\sum_{j \in J} \tilde{C}_j(\sigma^{\tilde{\pi}^-})}{\sum_{j \in J} C_j(\sigma^*)} \leq \frac{\sum_{j \in J} \tilde{C}_j(\sigma^{\tilde{\tau}^-})}{\sum_{j \in J} C_j(\sigma^*)} = PoA(\Gamma'),$$

which completes the proof. □

The following lemma provides a convenient expression for the optimal sum of costs (as if there were a central authority).

Lemma 4. *For any optimal schedule σ^*, the associated (optimal) sum of costs equals*

$$\sum_{j \in J} C_j(\sigma^*) = \sum_{j \in J} p_j \left(1 + \left\lfloor \frac{\lambda_j}{|M|} \right\rfloor\right).$$

Proof. Let π^* be the pre–schedule such that for each $j \in J$, $\pi^*(j) = m$ where $m = (|J| - \lambda_j - 1) \text{mod} |M| = m - 1$. One easily verifies that π^* can be obtained using the MFT algorithm. Therefore, the π^*–optimal schedule σ^* is optimal, and its associated sum of costs equals

$$\begin{aligned}
\sum_{j \in J} C_j(\sigma^*) &= \sum_{j \in J} \sum_{\substack{j' \in J: \\ p_{j'} > p_j, \\ \pi^*(j') = \pi^*(j)}} p_j \\
&= \sum_{j \in J} p_j \left|\{j' \in J : p_{j'} > p_j, \pi^*(j') = \pi^*(j)\}\right| \\
&= \sum_{j \in J} p_j \left(1 + \left\lfloor \frac{\lambda_j}{|M|} \right\rfloor\right).
\end{aligned}$$

□

Remark 2. Obviously, the optimal costs are independent of the owners of the jobs. Also, in case there is a unique player, the costs associated with any Nash equilibrium are optimal. Therefore, Lemma 4 can be obtained from Lemma 2 by assuming that there is a unique player (and hence, $\kappa_j = \lambda_j$ for all $j \in J$). $\diamond$

Finally, in the proof of our second main result we will use the following inequality.

Lemma 5. *For all integers $k \geq 0$,*

$$\frac{1 + k + \frac{|M|-1}{2|M|}}{1+k} \leq 1 + \frac{|M|-1}{2|M|}. \tag{3}$$

Proof. The proof is by induction on k. For $k = 0$, inequality (3) is in fact an equality. Suppose that (3) holds for some $k = k' \geq 0$. Then,

$$\frac{1 + (k'+1) + \frac{|M|-1}{2|M|}}{1 + (k'+1)} = \frac{1 + \left(1 + k' + \frac{|M|-1}{2|M|}\right)}{1 + (1+k')} \leq 1 + \frac{|M|-1}{2|M|},$$

where the inequality follows from the fact that for any $\beta, \delta > 0$ and any $\alpha, \gamma, \epsilon \in \mathbb{R}$, we have $[\ \frac{\alpha}{\beta}, \frac{\gamma}{\delta} \leq \epsilon \implies \frac{\alpha+\gamma}{\beta+\delta} \leq \epsilon\ ]$. Hence, (3) also holds for $k = k'+1$. $\square$

Now we can state and prove our second main result.

Theorem 3. *For the game Γ associated with a scheduling problem,*

$$PoA(\Gamma) \leq \frac{3|M|-1}{2|M|}.$$

Proof. Let $\Lambda = (N, (J_i)_{i \in N}, (p_j)_{j \in J})$ be a scheduling problem. Let Γ be the game associated with Λ. By Lemma 3, we may assume that for each $i \in N$, $|J_i| = 1$. With a slight abuse of notation, let $J = N = \{1, \ldots, n\}$. Without loss of generality we assume that $p_1 > \cdots > p_n$. Let $K = \lfloor \frac{|J|}{|M|} \rfloor$. We assume $K \neq \frac{|J|}{|M|}$ since the case $K = \frac{|J|}{|M|}$ follows from similar (but easier) arguments.

From Lemma 2, for the sum of expected costs induced by a worst Nash equilibrium $\tilde{\pi}^-$ we have

$$\begin{aligned}
\sum_{j \in J} \tilde{C}_j(\sigma^{\tilde{\pi}^-}) &= \sum_{j \in J} p_j \left(1 + \frac{\lambda_j - \kappa_j}{|M|} + \left\lfloor \frac{\kappa_j}{|M|} \right\rfloor\right) \\
&= \sum_{j \in J} p_j (1 + \frac{\lambda_j}{|M|}) \\
&= \sum_{k=0,\ldots,K-1} \sum_{l=1,\ldots,|M|} p_{(k|M|+l)} (1 + k + \frac{l-1}{|M|}) +
\end{aligned}$$

$$
\begin{aligned}
& \left(p_{(K|M|+1)}(1+K+\frac{0}{|M|})+\cdots+p_{|J|}(1+K+\frac{|J|-K|M|-1}{|M|})\right) \\
\leq\ & \sum_{k=0,\ldots,K-1}\ \sum_{l=1,\ldots,|M|} p_{(k|M|+l)}\,(1+k+\frac{|M|-1}{2|M|})\ + \\
& \left(p_{(K|M|+1)}(1+K+\frac{|M|-1}{2|M|})+\cdots+p_{|J|}(1+K+\frac{|M|-1}{2|M|})\right) \\
=\ & \sum_{k=0,\ldots,K-1}\left(1+k+\frac{|M|-1}{2|M|}\right)\left(\sum_{l=1,\ldots,|M|} p_{(k|M|+l)}\right)\ + \qquad (4)\\
& \left(1+K+\frac{|M|-1}{2|M|}\right)\left(p_{(K|M|+1)}+\cdots+p_{|J|}\right).
\end{aligned}
$$

Here, the inequality follows from the identity

$$
\sum_{l=1,\ldots,|M|}\frac{l-1}{|M|}=\frac{|M|-1}{2}
$$

and the redistribution of this sum in such a way that the jobs with longer (shorter) processing times have larger (smaller) coefficients on the right hand side than on the left hand side of the inequality. (For the case $k=K$, the sum of coefficients is even augmented.)

From Lemma 4, for the minimal sum of costs induced by an optimal schedule σ^* we have

$$
\begin{aligned}
\sum_{j\in J} C_j(\sigma^*) &= \sum_{j\in J} p_j\left(1+\left\lfloor\frac{\lambda_j}{|M|}\right\rfloor\right) \\
&= \sum_{k=0,\ldots,K-1}(1+k)\left(\sum_{l=1,\ldots,|M|} p_{(k|M|+l)}\right)+ \qquad (5)\\
&\quad (1+K)\left(p_{(K|M|+1)}+\cdots+p_{|J|}\right).
\end{aligned}
$$

Let $k=0,\ldots,K-1$. Then, from Lemma 5,

$$
\frac{\left(1+k+\frac{|M|-1}{2|M|}\right)\left(\sum_{l=1,\ldots,|M|} p_{(k|M|+l)}\right)}{(1+k)\left(\sum_{l=1,\ldots,|M|} p_{(k|M|+l)}\right)} \leq \left(1+\frac{|M|-1}{2|M|}\right). \qquad (6)
$$

Similarly,

$$
\frac{\left(1+K+\frac{|M|-1}{2|M|}\right)\left(p_{(K|M|+1)}+\cdots+p_{|J|}\right)}{(1+K)\left(p_{(K|M|+1)}+\cdots+p_{|J|}\right)} \leq \left(1+\frac{|M|-1}{2|M|}\right). \qquad (7)
$$

From (4), (5), (6), (7) and the fact that for any $\beta, \delta > 0$ and any $\alpha, \gamma, \epsilon \in \mathbb{R}$, [$\frac{\alpha}{\beta}, \frac{\gamma}{\delta} \leq \epsilon \implies \frac{\alpha+\gamma}{\beta+\delta} \leq \epsilon$] it follows that

$$PoA(\Gamma) = \frac{\sum_{j \in J} \tilde{C}_j(\sigma^{\tilde{\pi}^-})}{\sum_{j \in J} C_j(\sigma^*)} \leq \left(1 + \frac{|M|-1}{2|M|}\right) = \frac{3|M|-1}{2|M|},$$

which completes the proof. □

The following theorem shows that the bound in Theorem 3 is tight.

Theorem 4. *The bound for the price of anarchy in Theorem 3 is tight. That is, for any $\rho < \frac{3|M|-1}{2|M|}$ there is a scheduling problem such that for its associated game Γ, $PoA(\Gamma) \geq \rho$.*

Proof. Let $n = |M|$. Let $0 < p_1 < p_2 < \cdots < p_{n-1} < p_n$ be such that

$$\frac{p_1}{p_n}\left(\frac{3n-1}{2n}\right) \geq \rho. \tag{8}$$

Let $\Lambda = (N, (J_i)_{i \in N}, (p_j)_{j \in J})$ be the scheduling problem for which

- $|J| = |N| = |M|$;
- $J = \{1, \ldots, n\}$;
- for each $i \in N$, $|J_i| = 1$; and
- for each $j \in J$, the processing time of job j equals p_j.

From Lemma 2, for the sum of expected costs induced by a worst Nash equilibrium $\tilde{\pi}^-$ we have

$$\sum_{j \in J} \tilde{C}_j(\sigma^{\tilde{\pi}^-}) = \sum_{j=1,\ldots,n} p_j \left(1 + \frac{n-j}{n}\right). \tag{9}$$

Obviously, the minimal sum of costs induced by an optimal schedule σ^* equal

$$\sum_{j \in J} C_j(\sigma^*) = \sum_{j=1,\ldots,n} p_j. \tag{10}$$

From (9) and (10) it follows that for the game Γ associated with Λ we have

$$\begin{aligned}
PoA(\Gamma) &= \frac{\sum_{j\in J}\tilde{C}_j(\sigma^{\tilde{\pi}^-})}{\sum_{j\in J} C_j(\sigma^*)}\\
&= \frac{\sum_{j=1,\ldots,n} p_j\left(1+\frac{n-j}{n}\right)}{\sum_{j=1,\ldots,n} p_j}\\
&\geq \frac{\sum_{j=1,\ldots,n} p_j + \sum_{j=1,\ldots,n} p_j\left(\frac{n-j}{n}\right)}{np_n}\\
&\geq \frac{np_1 + p_1\sum_{j=1,\ldots,n}\left(\frac{n-j}{n}\right)}{np_n}\\
&= \frac{np_1 + p_1\left(\frac{n-1}{2}\right)}{np_n}\\
&= \frac{p_1\left(n+\frac{n-1}{2}\right)}{p_n n}\\
&= \frac{p_1}{p_n}\left(\frac{n+\frac{n-1}{2}}{n}\right)\\
&= \frac{p_1}{p_n}\left(\frac{3n-1}{2n}\right)\\
&\geq \rho,
\end{aligned}$$

where the last inequality follows from (8). $\square$

4 Mechanism Design

In this section we show that there is no "reasonable" mechanism such that optimal schedules can always be implemented in Nash equilibria. We will first introduce the properties that a "reasonable" mechanism should satisfy. Then, we will state and prove our impossibility result.

Let $\Lambda = (M, N, (J_i)_{i\in N}, (p_j)_{j\in J})$ be a scheduling problem. We denote the set of optimal schedules for Λ by $\Sigma^*(\Lambda)$. We associate with Λ a (non–cooperative) **scheduling game $\boldsymbol{\Gamma(\Lambda,\varphi)}$**, which is explained next. The set of **players** is N. For each $i \in N$, the set of **(pure) strategies** of player i, denoted Π_i, is the collection of functions $\pi_i : J_i \to M$. With a slight abuse of notation, a strategy profile $\pi = (\pi_i)_{i\in N}$ straightforwardly induces a pre–schedule π. The **mechanism** φ assigns each pre–schedule π to a schedule $\varphi^\pi = (\varphi_1^\pi, \varphi_2^\pi)$. Player i's resulting "costs" are given by the sum of completion times of his jobs in φ^π. In other words, player i's **cost function** c_i is given by

$$c_i(\pi) \equiv \sum_{j\in J_i} C_j(\varphi^\pi).$$

Remark 3. The definition of a mechanism as a function from the set of pre–schedules to the set of schedules involves two implicit assumptions. First, we assume **No Idle Time (NIT)**. More precisely, in view of our definition of schedule, we exclude from our analysis functions that can associate with a pre–schedule some assignment of jobs to machines that involves idle time. Second, we assume **Owner Anonymity (OA)**, i.e. for each pre–schedule the associated schedule does not depend on the owner of any of the involved jobs. $\diamond$

We consider that the following two additional properties should be satisfied by any "reasonable" mechanism φ. First, since all machines are identical, any reasonable mechanism should satisfy machine anonymity, which we explain next. Suppose a pre–schedule π' is obtained from some other pre–schedule π by permuting complete batches of jobs. Then, (a) the schedule $\varphi^{\pi'}$ is obtained from schedule φ^{π} by applying the same permutation to complete batches of jobs and (b) there is no change in the order in which the jobs are processed. Formally, φ satisfies

- **Machine Anonymity (MA)** if for any $\pi \in \Pi$, any permutation $\rho : M \to M$, any $j \in J$,

$$\varphi_1^{\rho\circ\pi}(j) = (\rho \circ \varphi_1^{\pi})(j) \text{ and} \tag{11}$$
$$\varphi_2^{\rho\circ\pi}(j) = \varphi_2^{\pi}(j), \tag{12}$$

  where $\circ$ denotes composition.

Second, any reasonable mechanism should associate with each pre–schedule a schedule that is consistent with the former. In other words, if a job is pre–scheduled on some machine, then the mechanism should schedule the job on that machine. Formally, φ satisfies

- **Pre–schedule Consistency (PC)** if for any $j \in J$ and any $\pi \in \Pi$, $\pi(j) = m \Longrightarrow \varphi_1^{\pi}(j) = m$.

Let $\mathcal{M}$ be the class of mechanisms that satisfy NIT, OA, MA, and PC. Note that mechanism σ from Section 2 (which assigns each pre–schedule π to its π–optimal schedule) satisfies NIT, OA, MA, and PC.

Let $i \in N$. A **mixed strategy** $\tilde{\pi}_i$ of player i is a probability distribution over all pure strategies $\pi_i \in \Pi_i$. At mixed strategy $\tilde{\pi}_i$, let $Pr(\pi_i|\tilde{\pi}_i)$ be the probability assigned to pure strategy $\pi_i \in \Pi_i$. Let $\tilde{\pi} = (\tilde{\pi}_i)_{i\in N}$ be a profile of mixed strategies. For any deterministic pre–schedule $\pi \in \Pi$, let $Pr(\pi|\tilde{\pi})$ be the probability of π according to $\tilde{\pi}$. Then, player i's expected "costs" can be written as

$$\tilde{c}_i(\tilde{\pi}) \equiv \sum_{j\in J_i} \tilde{C}_j(\varphi^{\tilde{\pi}}) = \sum_{j\in J_i}\left(\sum_{\pi\in\Pi} Pr(\pi|\tilde{\pi})C_j(\varphi^{\pi})\right),$$

where we denote the expected completion time of $j \in J$ by $\tilde{C}_j(\varphi^{\tilde{\pi}}) = \sum_{\pi\in\Pi} Pr(\pi|\tilde{\pi})C_j(\varphi^{\pi})$.

A profile of mixed strategies is a Nash equilibrium if no player has a profitable deviation. Formally, a profile of mixed strategies $\tilde{\pi}$ is a **(Nash) equilibrium** if there exists no player $i' \in N$ with a strategy $\tilde{\pi}'_{i'}$ such that

$$\tilde{c}_{i'}(\tilde{\pi}') < \tilde{c}_{i'}(\tilde{\pi}),$$

where $\tilde{\pi}' \equiv (\tilde{\pi}'_{i'}, (\tilde{\pi}_i)_{i \neq i'})$. Let $\mathcal{E}(\Gamma(\Lambda, \varphi))$ be the set of Nash equilibria of game $\Gamma(\Lambda, \varphi)$.

A mechanism $\varphi \in \mathcal{M}$ **implements optimal schedules** if for each scheduling problem $\Lambda = (M, N, (J_i)_{i \in N}, (p_j)_{j \in J})$, there exists some Nash equilibrium $\tilde{\pi} \in \mathcal{E}(\Gamma(\Lambda, \varphi))$ such that

$$\{\varphi^{\pi} : \pi \in \Pi,\ Pr(\pi|\tilde{\pi}) > 0\} \subseteq \Sigma^*(\Lambda). \tag{13}$$

Remark 4. Equivalently, a mechanism φ implements optimal schedules if for each scheduling problem $\Lambda = (M, N, (J_i)_{i \in N}, (p_j)_{j \in J})$, there exists some Nash equilibrium $\tilde{\pi} \in \mathcal{E}(\Gamma(\Lambda, \varphi))$ such that for some (or equivalently, all) $\sigma^* \in \Sigma^*(\Lambda)$,

$$\tilde{c}_i(\tilde{\pi}) = \sum_{j \in J} \tilde{C}_j(\varphi^{\tilde{\pi}}) = \sum_{j \in J} C_j(\sigma^*). \tag{14}$$

$\diamond$

We now state and prove our impossibility result.

Theorem 5. *There is no mechanism that satisfies NIT, OA, MA, and PC and that implements optimal schedules.*

Proof. Let $\varphi \in \mathcal{M}$. Let $\Lambda_{J_1} = (M, N, (J_1, J_2), (p_j)_{j \in J})$ be such that $M = \{m_1, m_2\}$, $N = \{1, 2\}$, $J = \{a, b, c\}$ with processing times $(p_a, p_b, p_c) = (1, 2, 3)$, and $J_1, J_2 \subseteq J$ such that $J_1 \cup J_2 = J$ and $J_1 \cap J_2 = \emptyset$. We will later specify J_1 (and thus J_2 as well), i.e. choose the owner of each job later in order to create a convenient scheduling game. Table 3 depicts the four optimal schedules, i.e. $\Sigma^*(\Lambda_{J_1}) = \{\alpha, \beta, \gamma, \delta\}$.

	optimal schedule
α	$m_1 : a, b$ $m_2 : c$
β	$m_1 : a, c$ $m_2 : b$
γ	$m_1 : c$ $m_2 : a, b$
δ	$m_1 : b$ $m_2 : a, c$

Table 3: Optimal schedules for $J = \{a, b, c\}$ with processing times $(p_a, p_b, p_c) = (1, 2, 3)$.

Since mechanism φ satisfies MA and PC, it follows that for any 2 jobs $x, y \in \{a, b, c\}$, $x \neq y$, if they are the only jobs sent to some machine, then they are processed on that machine and their order is uniquely determined by φ, i.e. independently of the identity of the machine and the identity of the owners of any of the jobs. Therefore, we can write

$[x, y]$ or $[y, x]$ to indicate that order. We first distinguish between the two orders of jobs b and c. The analysis of [b,c] is subsequently split dependent on the order between jobs a and b, while the analysis of [c,b] is subsequently split dependent on the order between jobs a and c. So, we distinguish among the following four cases.

Case I: φ yields $[b, c]$ and $[a, b]$.

Let $J_1 = \{a, c\}$ and $J_2 = \{b\}$. Assume there exists some Nash equilibrium $\tilde{\pi} \in \mathcal{E}(\Gamma(\Lambda_{\{a,c\}}, \varphi))$ with

$$\{\alpha, \beta\} \subseteq \{\varphi^{\pi} : \pi \in \Pi \text{ and } Pr(\pi|\tilde{\pi}) > 0\}. \tag{15}$$

In α, job b is processed on machine m_1 and in β it is processed on machine m_2. Similarly, in α, job c is processed on machine m_2 and in β it is processed on machine m_1. From (15) and the fact that φ satisfies PC it follows that player 2 (player 1) sends his job b (his job c) to either machine with positive probability. Since jobs b and c are owned by different players, it follows that with positive probability the two jobs end up being processed on the same machine. Hence, (13) is violated. Similarly, assuming that any of the sets $\{\alpha, \gamma\}$, $\{\delta, \beta\}$, or $\{\delta, \gamma\}$ is a subset of $\{\varphi^{\pi} : \pi \in \Pi \text{ and } Pr(\pi|\tilde{\pi}) > 0\}$ leads to a violation of (13). Therefore, for each Nash equilibrium $\tilde{\pi} \in \mathcal{E}(\Gamma(\Lambda_{\{a,c\}}, \varphi))$,

$$\{\alpha, \beta\}, \{\alpha, \gamma\}, \{\delta, \beta\}, \{\delta, \gamma\} \not\subseteq \{\varphi^{\pi} : \pi \in \Pi \text{ and } Pr(\pi|\tilde{\pi}) > 0\}. \tag{16}$$

Now assume that there exists some Nash equilibrium $\tilde{\pi} \in \mathcal{E}(\Gamma(\Lambda_{\{a,c\}}, \varphi))$ with

$$\{\varphi^{\pi} : \pi \in \Pi \text{ and } Pr(\pi|\tilde{\pi}) > 0\} \subseteq \{\alpha, \delta\}.$$

Then, player 2 uses the pure strategy of sending his job b to m_1. Suppose that in fact $\{\varphi^{\pi} : \pi \in \Pi \text{ and } Pr(\pi|\tilde{\pi}) > 0\} = \{\alpha\}$. Then, player 1's mixed strategy is in fact the pure strategy of sending job a to m_1 and job c to m_2. Since mechanism φ yields $[b, c]$, player 2 would be better off by sending his job to machine m_2 instead. However, this contradicts $\tilde{\pi} \in \mathcal{E}(\Gamma(\Lambda_{\{a,c\}}, \varphi))$. Hence, $\delta \in \{\varphi^{\pi} : \pi \in \Pi \text{ and } Pr(\pi|\tilde{\pi}) > 0\}$. As δ results with positive probability we conclude that φ has to yield $[a, c]$. So, with positive probability, say $q_\delta > 0$, player 1's costs equal $(1+(1+3)) = 5$ as ϕ yields $[a, c]$,, and with probability $1-q_\delta$, player 1's costs equal $1 + 3 = 4$. But then, since φ yields $[a, b]$, player 1 would be better of by playing the pure strategy that consists of sending job a to m_1 and job c to m_2: it would give (with probability 1) the lower costs $1 + 3 = 4$. This contradicts $\tilde{\pi} \in \mathcal{E}(\Gamma(\Lambda_{\{a,c\}}, \varphi))$. Similarly, assuming that $\{\varphi^{\pi} : \pi \in \Pi \text{ and } Pr(\pi|\tilde{\pi}) > 0\} \subseteq \{\beta, \gamma\}$ leads to a contradiction with $\tilde{\pi} \in \mathcal{E}(\Gamma(\Lambda_{\{a,c\}}, \varphi))$. Therefore, for each Nash equilibrium $\tilde{\pi} \in \mathcal{E}(\Gamma(\Lambda_{\{a,c\}}, \varphi))$,

$$\{\varphi^{\pi} : \pi \in \Pi \text{ and } Pr(\pi|\tilde{\pi}) > 0\} \not\subseteq \{\alpha, \delta\}, \{\beta, \gamma\}. \tag{17}$$

We now prove that for each Nash equilibrium $\tilde{\pi} \in \mathcal{E}(\Gamma(\Lambda_{\{a,c\}}, \varphi))$, and any subset $\Sigma \subseteq \Sigma^*(\Lambda_{\{a,c\}})$,

$$\{\varphi^{\pi} : \pi \in \Pi,\ Pr(\pi|\tilde{\pi}) > 0\} \not\subseteq \Sigma. \tag{18}$$

If $|\Sigma| \geq 3, 4$, then (18) follows from (16). If $|\Sigma| = 1$, then (18) follows from (17). If $|\Sigma| = 2$, then (18) follows from (16) together with (17). From (18) it immediately follows that for $\Lambda = \Lambda_{\{a,c\}}$ we have that each Nash equilibrium $\tilde{\pi} \in \mathcal{E}(\Gamma(\Lambda, \varphi))$ does not satisfy (13).

Case II: φ yields $[b, c]$ and $[b, a]$.

Let $J_1 = \{b, c\}$ and $J_2 = \{a\}$. Let $\Lambda = \Lambda_{\{b,c\}}$. Assume there exists some Nash equilibrium $\tilde{\pi} \in \mathcal{E}(\Gamma(\Lambda, \varphi))$ that satisfies (13). Since φ yields $[b, a]$,

$$\{\varphi^{\pi} : \pi \in \Pi \text{ and } Pr(\pi|\tilde{\pi}) > 0\} \subseteq \{\beta, \delta\}. \tag{19}$$

As β or δ results we conclude that φ has to yield $[a, c]$. Suppose $\{\varphi^{\pi} : \pi \in \Pi \text{ and } Pr(\pi|\tilde{\pi}) > 0\} = \{\beta, \delta\}$. In β, job a is processed on machine m_1 and in δ it is processed on machine m_2. Similarly, in β, job b is processed on machine m_2 and in δ it is processed on machine m_1. From (19) and the fact that φ satisfies PC it follows that player 1 (player 2) sends his job b (his job a) to either machine with positive probability. Since jobs a and b are owned by different players, it follows that with positive probability the two jobs end up being processed on the same machine. Hence, (19) is violated. Therefore, either $\{\beta\} = \{\varphi^{\pi} : \pi \in \Pi \text{ and } Pr(\pi|\tilde{\pi}) > 0\}$ or $\{\delta\} = \{\varphi^{\pi} : \pi \in \Pi \text{ and } Pr(\pi|\tilde{\pi}) > 0\}$.

Assume, without loss of generality, that $\{\beta\} = \{\varphi^{\pi} : \pi \in \Pi \text{ and } Pr(\pi|\tilde{\pi}) > 0\}$. Then, both players use a pure strategy and player 1's costs equal $2 + (1 + 3) = 6$ as φ yields $[a, c]$. However, player 1 would be better off by playing the strategy that consists of sending job b to m_1 and job c to m_2 as it would yield the lower costs of $2 + 3 = 5$. This contradicts $\tilde{\pi} \in \mathcal{E}(\Gamma(\Lambda_{\{b,c\}}, \varphi))$. Therefore, for $\Lambda = \Lambda_{\{b,c\}}$ we have that each Nash equilibrium $\tilde{\pi} \in \mathcal{E}(\Gamma(\Lambda, \varphi))$ does not satisfy (13).

Case III: φ yields $[c, b]$ and $[a, c]$.

Let $J_1 = \{a, b\}$ and $J_2 = \{c\}$. Let $\Lambda = \Lambda_{\{a,b\}}$. Applying the same arguments as in the first part of Case I it follows again that for each Nash equilibrium $\tilde{\pi} \in \mathcal{E}(\Gamma(\Lambda, \varphi))$,

$$\{\alpha, \beta\}, \{\alpha, \gamma\}, \{\delta, \beta\}, \{\delta, \gamma\} \nsubseteq \{\varphi^{\pi} : \pi \in \Pi \text{ and } Pr(\pi|\tilde{\pi}) > 0\}. \tag{20}$$

Now assume that there exists some Nash equilibrium $\tilde{\pi} \in \mathcal{E}(\Gamma(\Lambda_{\{a,b\}}, \varphi))$ with

$$\{\varphi^{\pi} : \pi \in \Pi \text{ and } Pr(\pi|\tilde{\pi}) > 0\} \subseteq \{\alpha, \delta\}. \tag{21}$$

Then, player 2 uses the pure strategy of sending his job c to m_2. Suppose that in fact $\{\varphi^{\pi} : \pi \in \Pi \text{ and } Pr(\pi|\tilde{\pi}) > 0\} = \{\delta\}$. Then, player 1's mixed strategy is in fact the pure strategy of sending job a to m_2 and job b to m_1. Since mechanism φ yields $[c, b]$, player 2 would be better off by sending his job to machine m_1 instead. However, this contradicts $\tilde{\pi} \in \mathcal{E}(\Gamma(\Lambda_{\{a,b\}}, \varphi))$. Hence, $\alpha \in \{\varphi^{\pi} : \pi \in \Pi \text{ and } Pr(\pi|\tilde{\pi}) > 0\}$. As α results with positive probability we conclude that φ has to yield $[a, b]$. So, with positive probability, say $q_\alpha > 0$, player 1's costs equal $(1 + (1 + 2)) = 4$ as φ yields $[a, b]$, and with probability $1 - q_\alpha$, player 1's costs equal $1 + 2 = 3$. But then, since φ yields $[a, c]$, player 1 would be better of by playing the pure strategy that consists of sending job a to m_2 and job

b to m_1: it would give (with probability 1) the lower costs $1+2=3$. This contradicts $\tilde{\pi} \in \mathcal{E}(\Gamma(\Lambda_{\{a,b\}}, \varphi))$. Similarly, assuming that $\{\varphi^\pi : \pi \in \Pi \text{ and } Pr(\pi|\tilde{\pi}) > 0\} \subseteq \{\beta, \gamma\}$ leads to a contradiction with $\tilde{\pi} \in \mathcal{E}(\Gamma(\Lambda_{\{a,b\}}, \varphi))$. Therefore, for each Nash equilibrium $\tilde{\pi} \in \mathcal{E}(\Gamma(\Lambda_{\{a,b\}}, \varphi))$,

$$\{\varphi^\pi : \pi \in \Pi \text{ and } Pr(\pi|\tilde{\pi}) > 0\} \not\subseteq \{\alpha, \delta\}, \{\beta, \gamma\}. \tag{22}$$

We now prove that for each Nash equilibrium $\tilde{\pi} \in \mathcal{E}(\Gamma(\Lambda_{\{a,c\}}, \varphi))$, and any subset $\Sigma \subseteq \Sigma^*(\Lambda_{\{a,c\}})$,

$$\{\varphi^\pi : \pi \in \Pi,\ Pr(\pi|\tilde{\pi}) > 0\} \not\subseteq \Sigma. \tag{23}$$

If $|\Sigma| \geq 3, 4$, then (23) follows from (20). If $|\Sigma| = 1$, then (23) follows from (22). If $|\Sigma| = 2$, then (23) follows from (20) together with (22). From (23) it immediately follows that for $\Lambda = \Lambda_{\{a,b\}}$ we have that each Nash equilibrium $\tilde{\pi} \in \mathcal{E}(\Gamma(\Lambda, \varphi))$ does not satisfy (13).

Case IV: φ yields $[c, b]$ and $[c, a]$.

Let $J_1 = \{b, c\}$ and $J_2 = \{a\}$. Let $\Lambda = \Lambda_{\{b,c\}}$. Assume there exists some Nash equilibrium $\tilde{\pi} \in \mathcal{E}(\Gamma(\Lambda, \varphi))$ that satisfies (13). Since φ yields $[c, a]$,

$$\{\varphi^\pi : \pi \in \Pi \text{ and } Pr(\pi|\tilde{\pi}) > 0\} \subseteq \{\alpha, \gamma\}. \tag{24}$$

As α or γ results we conclude that φ has to yield $[a, b]$. Suppose $\{\varphi^\pi : \pi \in \Pi \text{ and } Pr(\pi|\tilde{\pi}) > 0\} = \{\alpha, \gamma\}$. In α, job a is processed on machine m_1 and in γ it is processed on machine m_2. Similarly, in α, job c is processed on machine m_2 and in γ it is processed on machine m_1. From (24) and the fact that φ satisfies PC it follows that player 1 (player 2) sends his job c (his job a) to either machine with positive probability. Since jobs a and c are owned by different players, it follows that with positive probability the two jobs end up being processed on the same machine. Hence, (24) is violated. Therefore, either $\{\alpha\} = \{\varphi^\pi : \pi \in \Pi \text{ and } Pr(\pi|\tilde{\pi}) > 0\}$ or $\{\gamma\} = \{\varphi^\pi : \pi \in \Pi \text{ and } Pr(\pi|\tilde{\pi}) > 0\}$.

Assume, without loss of generality, that $\{\alpha\} = \{\varphi^\pi : \pi \in \Pi \text{ and } Pr(\pi|\tilde{\pi}) > 0\}$. Then, both players use a pure strategy and player 1's costs equal $(1+2)+3 = 6$ as φ yields $[a, b]$. However, since φ yields $[c, a]$, player 1 would be better off by playing the strategy that consists of sending job b to m_2 and job c to m_1 as it would yield the lower costs of $2+3=5$. This contradicts $\tilde{\pi} \in \mathcal{E}(\Gamma(\Lambda_{\{b,c\}}, \varphi))$. Therefore, for $\Lambda = \Lambda_{\{b,c\}}$ we have that each Nash equilibrium $\tilde{\pi} \in \mathcal{E}(\Gamma(\Lambda, \varphi))$ does not satisfy (13). $\square$

Remark 5. Inspection of the proof of Theorem 5 shows that we only need a weaker version of Machine Anonymity. More precisely, it suffices to impose (11) and (12) for permutations ρ that swap the pre–scheduled batches of two machines (and keep the pre–schedule for all other jobs intact). $\diamond$

5 Simulations

In this section we investigate the behavior of PoA in relation with the tight bound of PoA as described in Theorems 3 and 4. For this purpose we simulate four different classes of scheduling problems that are classified by the number of players and jobs.

First, we report the settings of the simulations. The inputs for each simulation are the number of players, the number of jobs per players, the processing time of each job and the number of machines. We distinguish four different classes. The first class selects randomly the number of players, i.e. $|N|$, from the set $\{2,3,4,5\}$. The number of jobs for player i, i.e. J_i, is randomly drawn from the set $\{1,2,3,4,5\}$ and the processing times of each job is randomly drawn from the interval $(0,1)$ of real numbers. Finally, a number is randomly selected from the set $\{1,2,...,10\}$ that represents the number of machines. The second class differs from the first one only by the set of jobs which is replaced by $\{10,11,12,13,14,15\}$. The third class differs from the first one only by the set of players which is replaced by $\{5,6,7,8,9,10\}$. Finally, the fourth class differs from the first one by replacing the set of players by $\{5,6,7,8,9,10\}$ and the set of machines by $\{2,3,...,50\}$.

For each class 10000 simulations are executed which results in the PoA for each simulated scheduling problem. The resulting numerical data is depicted by means of box plots in Figure 1. In each box plot the central line is the median, the central circle is the average, the edges of the box are the 25th and 75th percentiles, and the whiskers extend to the 2.5th (or lower) and 97.5th (or upper) percentiles. We have also included the graph of the function $x \mapsto \frac{3x-1}{2x}$, which by Theorem 4 gives the tight bound of the price of anarchy for any integer $x \geq 2$ when there are x machines. Finally, in each box plot, the medians are connected by the graph of a piecewise linear function.

Figure 1 a (b,c,d) represents the first (second, third, fourth) class of simulated scheduling problems and allows us to make the following observations.

First, we consider the percentage loss of performance with respect to optimal schedules. We observe that the average loss in performance is at most 23 %. This value is attained in the fourth class with $|M| = 16$. Here, the average PoA ≈ 1.23. This is considerably lower than the tight upperbound of PoA in this situation: 1.468, which reflects a loss of performance of 46%. Moreover, the relative distance from the PoA averages to the corresponding tight upper bounds is at least 15 %. This value is attained in the third case with $|M| = 10$. If we consider the upper percentile of the PoA, then we observe that this distance is at least 4%. However, in most situations it is more than 10%.

Second, taking the first class as starting point, we see that augmenting the number of jobs per players (i.e. shifting from the first class to the second class), the average PoA is reduced drastically. Moreover, the distance between the lower and upper percentiles is very small in comparison with this distance in the first class. This seems counter intuitive, but in these two cases we have only a few players. The difference between the two cases are the number of jobs. In the second case we have more jobs. So, the probability to obtain a bad schedule is smaller than in the case we have only a few jobs. Therefore, the payoff of the worst Nash equilibrium in the second case is considerably lower than the payoff of the worst Nash equilibrium in the first case. If the number of players is

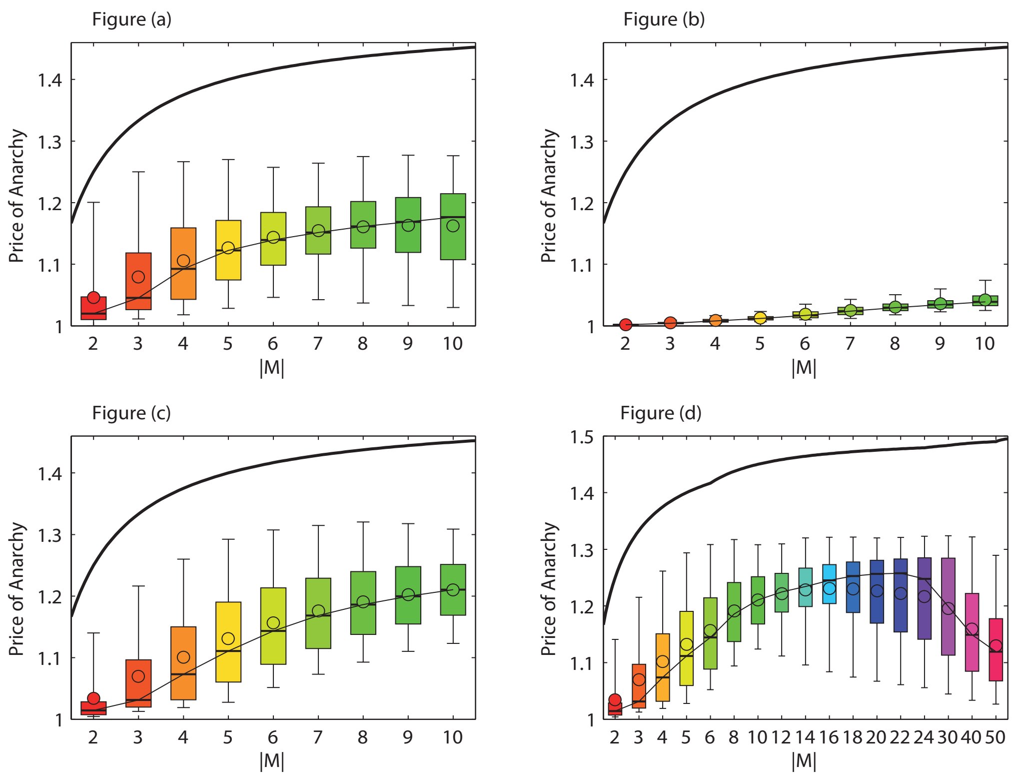

Figure 1: Box plots of price of anarchy for uniformly distributed processing times

augmented (i.e. shifting from the first class to the third class) we observe that both the average PoA and the distance between the lower and upper percentile slightly increase.

Third, if we augment the number of machines, the average PoA increases relatively quickly. This implies that the optimal costs decrease more rapidly than do the costs associated with the worst Nash equilibrium. After the PoA reaches a maximum it decreases relatively slowly. In fact, we can argue that when the number of machines tends to infinite (and all other parameters do not change), the price of anarchy tends to 1. This is an immediate consequence of the fact that when there is a very large number of machines any optimal schedule processes at most one job on each machine and the probability that a worse Nash equilibrium assigns two jobs to the same machine tends to 0.

Fourth, we see that for small and large numbers of machines, the average price of anarchy is located above the median price of anarchy. For intermediate values, the average is below the median.

6 Concluding Remarks

This paper focused on the costs of outsourcing decisions being made individually rather than cooperatively. We identified tight bounds for the price of anarchy and the impossibility for a natural mechanism to enforce the existence of equilibria that result in first-best schedules. Note that this impossibility implies that under natural assumptions on the mechanism the gap between first-best and the worst Nash equilibrium cannot be avoided, but also that the gap between first-best and the best Nash equilibrium cannot be avoided. Hence, though our choice between a focus on the best or the worst Nash equilibrium (focus on decentralization costs or price of anarchy) may have seen arbitrary in favor of an analysis using the PoA, the analysis of coordinating mechanisms is not affected by this choice. Additionally, one might wonder what the impact of this choice would be on the bounds, which is what we will analyze next.

Formally, for a game Γ, the **decentralization cost** is defined as

$$DC(\Gamma) = \frac{\min_{\tilde{\pi} \in \mathcal{E}(\Gamma)} \sum_{j \in J} \tilde{C}_j(\sigma^{\tilde{\pi}})}{\sum_{j \in J} C_j(\sigma^*)},$$

where σ^* is an optimal schedule. Obviously, for any game Γ, $1 \leq DC(\Gamma) \leq PoA(\Gamma)$.

Assuming that there are $|M| = 2$ machines, in Example 4 we exhibit a series of scheduling games Γ^ϵ such that for any $\rho < \frac{5}{4} = \frac{3|M|-1}{2|M|}$ there is a game $\Gamma^{\epsilon(\rho)}$ with $DC(\Gamma^{\epsilon(\rho)}) = PoA(\Gamma^{\epsilon(\rho)}) \geq \rho$. Therefore, since the bound in Theorem 3 is tight for the price of anarchy (Theorem 4), the same bound is tight for the decentralized cost as well.

Example 4. (Same tight bound for price of anarchy and decentralization cost.) Let $M = \{m_1, m_2\}$ and $N = \{1, 2\}$. Let $J_1 = \{a, c\}$ and $J_2 = \{b, d\}$. Let $\epsilon \in (0, \frac{1}{4})$. Suppose $(p_a, p_b, p_c, p_d) = (\epsilon, 2\epsilon, 1 - 2\epsilon, 1 - \epsilon)$. Note $p_a < p_b < p_c < p_d$. Consider the associated game Γ^ϵ. Similarly to the scheduling game discussed in Examples 1 and 2, each strategy can fully be described by indicating which jobs are sent to m_1 (the complement is sent to m_2). And, as before, there is a unique Nash equilibrium in mixed strategies. For player 1 it is always strictly better to play $\{a\}$ than $\emptyset$, and it is always strictly better to play $\{c\}$ than $\{a, c\}$. Similarly, for player 2 it is always strictly better to play $\{b\}$ than $\emptyset$, and it is always strictly better to play $\{d\}$ than $\{b, d\}$. So, in any Nash equilibrium, strategies $\emptyset$ (for both players), $\{a, c\}$ (for player 1), and $\{b, d\}$ (player 2) receive probability 0. Hence, it suffices to restrict attention to the reduced game described in Table 4.

$1\backslash 2$	$\{b\}$	$\{d\}$
$\{a\}$	$\boldsymbol{1-\epsilon}$, 2	$1+\epsilon$, $\boldsymbol{1+2\epsilon}$
$\{c\}$	$1+\epsilon$, $\boldsymbol{1+2\epsilon}$	$\boldsymbol{1-\epsilon}$, 2

Table 4: Table of Example 4

Applying standard game–theoretic tools one can readily show that the strategy–profile in which each of the strategies $\{a\}$, $\{c\}$, $\{b\}$, and $\{d\}$ receives probability 1/2 constitutes the unique Nash equilibrium $\tilde{\pi}^-$ in mixed strategies. Hence, the costs induced by the unique Nash equilibrium are

$$\tilde{c}_1(\tilde{\pi}^-) + \tilde{c}_2(\tilde{\pi}^-) = [\frac{1}{2}(1-\epsilon) + \frac{1}{2}(1+\epsilon)] + [\frac{1}{2}(2) + \frac{1}{2}(1+2\epsilon)] = \frac{5}{2} + \epsilon.$$

For any optimal schedule σ^* the associated costs equal

$$c_1(\sigma^*) + c_2(\sigma^*) = \epsilon + 2\epsilon + (\epsilon + 1 - 2\epsilon) + (2\epsilon + 1 - \epsilon) = 2 + 3\epsilon.$$

Hence, from the unicity of the Nash equilibrium $\tilde{\pi}^-$ it follows that

$$DC(\Gamma^\epsilon) = PoA(\Gamma^\epsilon) = \frac{\frac{5}{2} + \epsilon}{2 + 3\epsilon}.$$

Note that $\lim_{\epsilon\to 0} DC(\Gamma^\epsilon) = \lim_{\epsilon\to 0} PoA(\Gamma^\epsilon) = \frac{5}{4}$, which for $|M| = 2$ coincides with the tight bound established for the price of anarchy in Theorems 3 and 4. $\diamond$

For more than two machines the tight bound for the PoA of the current paper is, of course, a bound for the decentralization costs as well. An interesting open problem would be the identification of a tight bound for the decentralization costs in case of more than two machines.

References

Aydinliyim, T. and G. Vairaktarakis (2010). Coordination of outsourced operations to minimize weighted flow time and capacity booking costs. *Manufacturing & Service Operations Management*, 12, 236–255.

Aydinliyim, T. and L. Varaktarakis (2011). Sequencing strategies and coordination issues in outsourcing and subcontracting operations. In *Planning Production and Inventories in the Extended Enterprise*, volume 151, 269–319.

Bukchin, Y. and E. Hanany (2007). Decentralization cost in scheduling: a game-theoretic approach. *Manufacturing & Service Operations Management*, 9 (3), 263–275.

Bukchin, Y. and E. Hanany (2011). Decentralization cost in supply chain jobshop scheduling with minimum flowtime objective. *Research memo*, Tel Aviv University, Israel.

Cai, X. and G. Vairaktarakis (2012). Coordination of outsourced operations at a third-party facility subject booking, overtime, and tardiness costs. *Operations Research*, 60, 1436–1450.

Cole, R., J. Correa, V. Gkatzelis, and V. Morrokni (2013). Decentralized utilitarian mechanisms for scheduling games. *Games and Economic Behavior*, (on line).

Correa, J. and M. Queyranne (2012). Efficiency of equilibria in restricted uniform machine scheduling with total weighted completion time as social cost. *Naval Research Logistics*, 95, 384–395.

Curiel, I., G. Pederzoli, and S. Tijs (1989). Sequencing games. *European Journal of Operational Research*, 40, 344–351.

Curiel, I., J. Potters, V. Rajendra Prasad, S. Tijs, and B. Veltman (1994). Sequencing and cooperation. *Operations Research*, 42, 566–568.

Fernandez, A. E., P. Borm, P. Calleja, and H. Hamers (2008). Sequencing games with repeated players. *Annals of Operations Research*, 158, 189–203.

Hain, R. and M. Mitra (2004). Simple sequencing problems with interdependent costs. *Games and economic behavior*, 48, 271–291.

Hall, N. and Z. Liu (2010). Capacity allocation and scheduling in supply chains. *Operations Research*, 58, 1711–1725.

Hamers, H., P. Borm, and S. Tijs (1995). On games corresponding to sequencing situations with ready times. *Mathematical Programming*, 70, 1–13.

Heydenreich, B., R. Müller, and M. Uetz (2006). Decentralization and mechanism design for online machine scheduling. *Algorithm Theory*, 4059, 136–147.

Heydenreich, B., R. Müller, and M. Uetz (2007). Games and mechanism design in machine scheduling - an introduction. *Production and Operations Management*, 16, 437–454.

Hoeksma, R. and M. Uetz (2012). The price of anarchy for minsum related machine scheduling. In *Approximation and Online Algorithms (WAOA 2011), R. Solis-Oba and G. Persiano (eds.)*, 404–413. Lecture Notes in Computer Science 7164.

Immorlica, N., L. Li, V. Mirrokni, and A. Schulz (2009). Coordination mechanisms for selfish scheduling. *Theoretical computer science*, 410, 1589–1598.

Koutsoupias, E. and C. Papadimitriou (1999). Worst-case equilibria. In *Lecture notes in Computer Science, 16th annual symposium on theoretical aspects of computer science.*, 404–413. Springer, Germany.

Lawler, E., J. Lenstra, A. Rinnooy Kan, and D. Shmoys (1993). Sequencing and scheduling: algorithms and complexity. In *Logistics of production and inventory*, edited by S. Graves, A. Rinnooy Kan, and P. Zipkin, 445–522. North Holland, The Netherlands.

Moulin, H. (2008). The price of anarchy of serial, average and ecremental cost sharing. *Economic Theory*, 36, 379–405.

van den Nouweland, A., M. Krabbenborg, and J. Potters (1992). Flowshops with a dominant machine. *European Journal of Operational Research*, 62, 38–46.

PINEDO, M. (2002). *Scheduling: theory, algorithms and systems*. Prentice Hall, Englewood Cliffs, NJ.

ROUGHGARDEN, T. (2006). *Selfish routing and the price of anarchy*. MIT Press, Cambridge.

SLIKKER, M. (2006). Relaxed sequencing games have a nonempty core. *Naval Research Logistics*, 53 (4), 235–242.

VAIRAKTARAKIS, G. (2013). Noncooperative games for subcontracting operations. *Manufacturing & Service Operations Management*, 15, 148–158.