Truth and Regret in Online Scheduling

SHUCHI CHAWLA, University of Wisconsin-Madison

NIKHIL DEVANUR, Microsoft Research

JANARDHAN KULKARNI, Microsoft Research

RAD NIAZADEH, Cornell University

We consider a scheduling problem where a cloud service provider has multiple units of a resource available over time. Selfish clients submit jobs, each with an arrival time, deadline, length, and value. The service provider's goal is to implement a truthful online mechanism for scheduling jobs so as to maximize the social welfare of the schedule. Recent work shows that under a stochastic assumption on job arrivals, there is a single-parameter family of mechanisms that achieves near-optimal social welfare. We show that given any such family of near-optimal online mechanisms, there exists an online mechanism that in the worst case performs nearly as well as the best of the given mechanisms. Our mechanism is truthful whenever the mechanisms in the given family are truthful and prompt, and achieves optimal (within constant factors) regret.

We model the problem of competing against a family of online scheduling mechanisms as one of learning from expert advice. A primary challenge is that any scheduling decisions we make affect not only the payoff at the current step, but also the resource availability and payoffs in future steps. Furthermore, switching from one algorithm (a.k.a. expert) to another in an online fashion is challenging both because it requires synchronization with the state of the latter algorithm as well as because it affects the incentive structure of the algorithms.

We further show how to adapt our algorithm to a non-clairvoyant setting where job lengths are unknown until jobs are run to completion. Once again, in this setting, we obtain truthfulness along with asymptotically optimal regret (within polylogarithmic factors).

Additional Key Words and Phrases: Cloud Computing; Online Learning; Pricing

1 INTRODUCTION

We consider an online mechanism design problem inspired by the allocation and scheduling of cloud services. A scheduler allocates scarce resources to jobs arriving over time with the goal of maximizing economic efficiency or social welfare. The jobs are submitted by selfish users who can lie about the job's value, length, arrival, or deadline, so as to obtain a better allocation or pay a cheaper price. Our goal is to design an online mechanism that is truthful and obtains good welfare guarantees in the worst case.

The primary motivation for our work comes from resource allocation in the context of cloud computing. The scheduling model we study is motivated by spot markets for virtual machines in “the cloud”, such as the one run by Amazon. A similar model, called the economic resource

A full version of this paper is available at <https://arxiv.org/abs/1703.00484>.

Authors' email addresses: shuchi@cs.wisc.edu; nikdev@microsoft.com; jakul@microsoft.com; rad@cs.cornell.edu.

Shuchi Chawla was supported in part by NSF grants CCF-1320854 and CCF-1617505. Rad Niazadeh was supported by NSF grant CCF-1512964 and by a Google Ph.D. Fellowship. Portions of this work were completed while the first author was a visitor and the fourth author an intern at Microsoft Research Redmond.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored.

Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

EC'17, June 26–30, 2017, Cambridge, MA, USA. © 2017 Association for Computing Machinery.

ACM ISBN 978-1-4503-4527-9/17/06...\$15.00. DOI: <http://dx.doi.org/10.1145/3033274.3085119>

allocation or ERA model, has been proposed in the context of Microsoft's Azure as described by [Babaioff et al. \[2017\]](#).

There is a vast and rich literature on online mechanism design in settings where selfish agents participate in the mechanism over time. What makes the scheduling problem described above interesting is that even ignoring the users' incentive constraints, there are strong lower bounds in the worst case for the purely algorithmic problem of scheduling jobs with deadlines to maximize welfare. [Canetti and Irani \[1998\]](#) showed, in particular, that no online algorithm can achieve a less than polylogarithmic competitive ratio for this problem in comparison to the hindsight optimal schedule. On the other hand, [Lavi and Nisan \[2015\]](#) showed that no deterministic mechanism that is truthful with respect to all of the parameters can approximate social welfare better than a factor T in the worst case, where T is the time horizon, even for unit length jobs on a single machine. In the face of these strong negative results, a number of works have considered weakening various aspects of the model in order to obtain positive results, such as requiring a slackness condition on the jobs' deadlines [[Azar et al., 2015](#)], allowing the algorithm to make tardy decisions [[Hajiaghayi et al., 2005](#)], satisfying incentive compatibility with respect to only a few of the jobs' parameters [[Azar and Khaitsin, 2011](#), [Cole et al., 2008](#)], etc.

A Bayesian benchmark. In this paper we follow an alternate approach of competing against a benchmark inspired by a stochastic model for job arrivals. For many online problems, the hindsight optimum is too pessimistic and strong a benchmark to compete against. A classic example from algorithmic mechanism design is the digital goods auction for which no mechanism can compete against the hindsight optimum that obtains the entire social welfare. But picking the right benchmark to compete against, namely the optimal posted price, has led to the design of many beautiful mechanisms with robust and strong revenue guarantees. [Hartline and Roughgarden \[2008\]](#) advocate a general framework for generating an appropriate benchmark for such online problems—determine the class of all mechanisms that are optimal for the problem in an appropriate stochastic setting; compete against the best of these Bayesian optimal mechanisms for a worst case instance.

In this paper we apply the [Hartline and Roughgarden](#) approach to online scheduling.¹ We show that for any given finite class of truthful scheduling mechanisms, we can design an online mechanism that is competitive against the best of the given mechanisms. Our mechanism is truthful with respect to all of the parameters of the jobs' types, is computationally efficient, and requires no assumptions on the input instance. We achieve asymptotically optimal regret guarantees with respect to all of the involved problem parameters.

While the [Hartline and Roughgarden](#) agenda has been successfully applied in mechanism design settings, it has not yet seen much use in algorithmic settings where the worst case optimization problem is hard but positive results are known under stochastic assumptions.² Our hope is that this work will spur future work in this direction.

Our work is inspired by the recent work of [Chawla et al. \[2017\]](#) that shows that when jobs are drawn from an i.i.d. distribution in every time period, a family of simple mechanisms achieves near-optimal social welfare. [Chawla et al.](#)'s mechanism is a simple greedy best-effort mechanism based on posted prices. The mechanism announces a price per-unit of resource for each time period into the future. When a job arrives it gets scheduled in a best-effort FIFO manner in the cheapest slots that satisfy its requirements. [Chawla et al.](#) show that if the number of resources per time

¹Note, however, that while [Hartline and Roughgarden](#) advocate competing against the Bayesian optimal mechanism, we compete against a suite of mechanisms that contain an approximately optimal one in any Bayesian setting.

²The idea of combining several online algorithms into one that is nearly as good as the best of them has been explored previously, but in the limited context of metrical task systems and adaptive data structures. See, e.g., [Blum and Burch \[2000\]](#), [Blum et al. \[2002\]](#).

period is large enough, then for any underlying distribution over job types, there exists a set of prices such that this posted-pricing-FIFO mechanism achieves a $1 - \epsilon$ approximation to expected social welfare. Unfortunately, finding the best price to offer requires knowing the fine details of the distribution over job types and solving a large linear program.

Chawla et al.'s mechanisms are parameterized by a single price. Machine learning techniques have been successfully used to tune parameters of heuristics for a wide variety of problems [Hutter et al., 2009, Xu et al., 2008]. This is typically done in a batch setting, where past data is used to find a good setting of parameters for the heuristic. For an inherently online problem such as scheduling, we seek to do the parameter tuning itself in an online manner. Can we search for the right parameters as we are running the heuristic?

We model the online mechanism design problem as a sort of “learning from expert advice” problem where the experts are algorithms in the set of all Bayesian optimal algorithms. We present a general black-box reduction from the online scheduling problem to online learning algorithms for the experts setting. Previous work along these lines has looked at online settings where the mechanism gets a new instance of the problem in every time step, and the learning algorithm can adapt the parameters of the mechanism as more and more instances are seen. In the scheduling setting the algorithms run on a single instance of the problem, and we want to achieve low regret over all Bayesian optimal mechanisms for this specific instance.

Challenges. There are several novel challenges that the scheduling context imposes. First, in our setting, jobs can grab resources for multiple consecutive units of time. So the decision of scheduling a job in the current time step can affect the availability of resources in many future steps. Furthermore it is unclear at what time an algorithm should receive credit for scheduling a job – when the job starts, or when it ends, or throughout its execution? What if the algorithm uses preemption?

Second, while an online algorithm for the experts problem can cleanly switch from one expert to another during execution, in the scheduling setting switching can be tricky. On the one hand, we cannot abandon jobs that have already been started but not finished by the previous algorithm. On the other hand, we must try to match the state of the new algorithm that we are following, so as to obtain the same reward. Most importantly, switching impacts the *incentive properties* of the underlying mechanisms. Indeed there are several ways in which jobs may lie to affect their outcomes in the combined mechanism, even if they cannot in the underlying algorithms. For example, jobs may be able to influence the time at which a switch happens, or the next algorithm that is selected. Moreover, a job may be able to benefit indirectly by influencing the schedule of the algorithm that is used just prior to the algorithm that determines the job's allocation.

Our techniques. We design a switching procedure that overcomes all of these challenges. Each switch takes a few time steps to finish previously scheduled jobs, synchronize with the state of the new mechanism, as well as handle the scheduling of intermediate jobs in a manner so as to preserve their incentive structure. As a consequence of this “switching delay” our algorithm incurs a bounded cost for every switch. Using this switching procedure in tandem with a reduction to experts with switching costs allow us to obtain a sublinear regret guarantee in addition to truthfulness.

Our truthful switching technique is oblivious to the details of the underlying truthful mechanisms and works as long as the underlying mechanisms are prompt, that is, they announce the allocation and payment of a job right at the time of the job's arrival. To our knowledge this is the first result that combines truthful mechanisms online into an overall truthful mechanism.

Non-clairvoyance. Finally, one challenge faced by schedulers in a real-world setting is that jobs may not know how long they would take to run. Fortunately, the benchmark suggested by Chawla et al. [2017], namely the optimal posted-pricing-FIFO mechanism, continues to obtain near-optimal

welfare—the mechanism does not need to know jobs' lengths in order to make scheduling decisions, although the existence of a good posted price assumes that lengths along with other job parameters are drawn from some i.i.d. distribution.

Non-clairvoyance poses extra challenges in the online scheduling setting. We can no longer keep track of different algorithms' rewards or their states at the time of a switch. Synchronization with an algorithm's state at the time of a switch is crucial because for some algorithms, such as posted-price-FIFO in particular, out of sync execution can cost almost the entire social welfare of the algorithm. (See, e.g., Example 4.1.)³ We therefore consider a slightly modified benchmark where the performance of an algorithm is measured according to the welfare it accumulates if it is periodically (randomly) restarted with an empty state. These random restarts do not significantly affect the performance of algorithms such as posted-price-FIFO in the stochastic setting.

We give a reduction from the non-clairvoyant setting to a multi-armed bandit (MAB) problem. The main challenge in this reduction is to couple the random restarts of the expert algorithms with the times at which the MAB algorithm decides to switch experts, in a manner that ensures that the restarts are independent of the internal coin flips of the MAB algorithm. To do so, we partition the scheduling process into mini-batches synchronized with random restarts, and run the MAB algorithm over this batched instance. Because switching between algorithms happens exactly at the time of a random restart, it becomes possible for us to sync with the states of the expert algorithms.

Random restarts can once again break the incentive properties of the overall mechanism. We need to take care to ensure that jobs that are caught in the middle of a restart cannot benefit by misreporting their arrival or deadline. This necessitates a careful redesign of the switching protocol for non-clairvoyant settings. As in the clairvoyant setting, we obtain the optimal dependence within polylogarithmic factors of the regret on the time horizon.

Outline. In Section 3 we present a truthful online algorithm for the clairvoyant setting along with an upper bound on its regret. We extend both the truthfulness and regret guarantees to the non-clairvoyant setting in Section 4. Section 5 presents matching lower bounds on regret.

2 MODEL AND DEFINITIONS

2.1 The online job scheduling problem

An instance of the online job scheduling problem consists of a finite set of jobs $\mathcal{J}$, a time horizon T , and the number m of resources (machines) available per unit of time. Each job $j \in \mathcal{J}$ arrives at time $a_j \in [T]$, has a deadline $d_j \in [a_j, a_j + d_{\max}]$ and a processing length $l_j \in [0, \ell_{\max}]$. Assume all the deadlines are in the time horizon $[T]$. Completing each time unit of job $j \in \mathcal{J}$ generates a value-per-length $v_j \in [0, v_{\max}]$.

Online scheduling algorithms. An online scheduler is an algorithm that determines which jobs to schedule and when, and how much to charge each scheduled job. Each job j , at the time of its arrival, reports its arrival time, deadline, length, and value; this four-tuple is called the job's type. The scheduling algorithm determines whether or not to schedule the job (*admission control step*) and, if the job is scheduled, maps it to a set τ_j of time units (*scheduling step*) and charges it a payment p_j . If $|\tau_j \cap [a_j, d_j]| \geq l_j$, that is, the job is allocated at least l_j time units before its deadline, then the job obtains a utility of $(v_j - p_j) \cdot l_j$. The schedule produced by the algorithm is feasible if no more than m jobs are assigned to each time unit.

We now discuss various features of online scheduling algorithms:

³ In the full version of the paper we also show that in a continuous time setting no algorithm can get a sub-linear regret.

Preemption: We say that the algorithm is *non-preemptive* if the set τ_j consists of contiguous time units for every job j . In other words, when a job is started, the algorithm processes it without pausing until it is finished.

Truthfulness: A scheduling algorithm is *truthful* if for every job j , fixing the reported types of jobs in $\mathcal{J}_{-j}$, job j 's utility is maximized by reporting its true type. Jobs can misreport any of the four components of their type, however, following convention we assume that jobs cannot report an earlier arrival time.

Promptness: A scheduling algorithm is *prompt* if for every job j , the job's allocation and payment, (τ_j, p_j) , are determined at the time of the job's arrival. At times we will refer to a weaker property: an algorithm is *order respecting* if for every job j , the job's allocation and payment, (τ_j, p_j) , are functions of jobs in $\mathcal{J}$ that arrive prior to j and not of those jobs that arrive after j .

Clairvoyance: The *clairvoyant scheduling* problem is the setting where every job j reports its length l_j to the scheduling mechanism, together with other parts of its type, upon its arrival, whereas in the *non-clairvoyant scheduling problem* jobs do not report their lengths upon arrival. In fact, the scheduling mechanism observes the length of a job only after it completes the job. Since the length of job j is unknown prior to its completion in the non-clairvoyant scheduling, we have to slightly modify other aspects of the setting:

- we change the definition of deadline d_j to denote the latest time that j can be started.⁴
- We do not allow preemption in the non-clairvoyant setting.

With these two modifications, it is indeed guaranteed that if a job j is allocated at a time no later than its deadline, then it will be scheduled properly, i.e. it will be given enough time to be completed.

Let ONL denote an online scheduling algorithm, and let $J(\text{ONL}) = \mathcal{J} \cap \{j : |\tau_j \cap [a_j, d_j]| \geq l_j\}$ denote the set of jobs that receive service in ONL. We use $W_t(\text{ONL}, \mathcal{J})$ to denote the value generated by the algorithm at time unit t :

$$W_t(\text{ONL}, \mathcal{J}) = \sum_{j \in J(\text{ONL}): t \in \tau_j} v_j$$

The total value generated by the algorithm, a.k.a. its *social welfare*, is given by:

$$W(\text{ONL}, \mathcal{J}) = \sum_{t=1}^T W_t(\text{ONL}, \mathcal{J}) = \sum_{j \in J(\text{ONL})} v_j l_j$$

We drop the argument $\mathcal{J}$ when it is clear from the context.

Regret minimization in online scheduling. We consider an online learning problem, where we are given a finite set of scheduling algorithms and our goal is to compete with the best one in hindsight with respect to the social welfare objective. Let $\{\text{ALG}_1, \dots, \text{ALG}_n\}$ be the set of n online schedulers. Given an instance $\mathcal{J}$, let $\text{OPT}(\mathcal{J}) = \max_{i \in [n]} W(\text{ALG}_i, \mathcal{J})$ denote the social welfare obtained by the hindsight optimal algorithm on this instance. Let ONL denote our online scheduling algorithm. The regret of ONL is defined as:

$$\text{REG}(\text{ONL}) \triangleq \max_{\mathcal{J}} (\text{OPT}(\mathcal{J}) - W(\text{ONL}, \mathcal{J}))$$

⁴Note the difference with deadlines in the clairvoyant setting, where d_j was defined to be the latest time that j could be completed.

2.2 Learning from expert advice

We will reduce the regret minimization problem for online scheduling to the problem of learning from expert advice. In the latter, we are given n experts indexed by i . In each time step $t \in [T]$, the online algorithm must choose a (potentially random) expert, $i_t \in [n]$, to follow. An adversary then reveals a reward vector $\{r_t^{(i)}\}$. We assume that the adversary is oblivious, that is, it cannot observe the internal coin flips of the algorithm. The total payoff of expert i is given by $\sum_{t \in [T]} r_t^{(i)}$. The payoff of the algorithm is given by $\mathbb{E}[\sum_{t \in [T]} r_t^{(i_t)}]$, where the expectation is taken over the algorithm's internal coin flips. The regret of the algorithm is:

$$\left(\max_{i \in [n]} \sum_{t \in [T]} r_t^{(i)} \right) - \mathbb{E} \left[\sum_{t \in [T]} r_t^{(i_t)} \right]$$

Let R denote an upper bound on $r_t^{(i)}$ for any $i \in [n]$ and $t \in [T]$. Then, several different online algorithms are known to achieve a regret of $O(R\sqrt{T \log n})$, and this bound is tight [Cesa-Bianchi and Lugosi, 2006, Freund and Schapire, 1995, Kalai and Vempala, 2005].

Experts with switching costs. This is a variant of the problem of learning from expert advice in which the algorithm faces a *switching cost* of C units every time it switches from one expert to another in consecutive time steps. In particular, the payoff of the algorithm is given by $\mathbb{E}[\sum_{t \in [T]} r_t^{(i_t)}] - C |\{t \in [T] : i_t \neq i_{t-1}\}|$. The first term corresponds to the rewards and the second corresponds to the switching cost. Accordingly, the regret of the algorithm is:

$$\left(\max_{i \in [n]} \sum_{t \in [T]} r_t^{(i)} \right) - \mathbb{E} \left[\sum_{t \in [T]} r_t^{(i_t)} \right] + C \mathbb{E} [|\{t \in [T] : i_t \neq i_{t-1}\}|]$$

THEOREM 2.1 (KALAI AND VEMPALA [2005]). *There is an algorithm Expert-ALG(C) for the experts problem with a switching cost of C such that*

$$\text{REG}(\text{Expert-ALG}(C)) \leq O \left(\sqrt{R(R+C)T \log n} \right).$$

Multi-armed bandit setting. In the multi-armed bandit (MAB) setting, the online algorithm may only observe the reward $r_t^{(i_t)}$ of the expert that it selects at time t , and cannot observe the remaining rewards. Algorithms for MAB typically mix some exploration alongside following the recommendation of an online learning algorithm for the full-information setting.

THEOREM 2.2 (AUER ET AL. [1995]). *There is an algorithm Bandit-ALG for the MAB problem with*

$$\text{REG}(\text{Bandit-ALG}) \leq O \left(R \sqrt{T n \log n} \right).$$

We can further consider an extension of the MAB setting to the setting with a switching cost. For this problem, we mainly use that there is no algorithm with a regret of $o(T^{2/3})$ [Dekel et al., 2014], to get a similar lower bound for our problem. The precise statement of their result is in Section 5.

3 THE CLAIRVOYANT SETTING

In this section we consider the online scheduling problem in the clairvoyant setting, namely where every job reports its length (in addition to the rest of its type) at the time of its arrival. We are given n online scheduling algorithms, $\text{ALG}_1, \dots, \text{ALG}_n$, and our goal is to design an online algorithm that minimizes regret relative to the best of the n algorithms in hindsight. We begin by showing how to switch between algorithms in a way that preserves truthfulness in Section 3.1. In Section 3.2 we

present a reduction from this problem to the problem of learning from expert advice with switching costs. In Section 5 we prove that the regret guarantee we obtain from the reduction is optimal.

3.1 Truthful switching

In this section, we show how to switch between truthful mechanisms while preserving truthfulness and making sure the loss in welfare is bounded. We consider the following setting. Let A and B be two order respecting truthful scheduling mechanisms. Our goal is to switch from mechanism A to mechanism B at time 0. (This is just a normalization of the time index for ease of notation.) We consider show how to perform this switch in the clairvoyant setting, and extend our algorithm to the non-clairvoyant setting in Section 4.2. The loss in welfare from our switching algorithm is captured in the following lemma.

LEMMA 3.1. *Given order respecting truthful mechanisms A and B , there exists an order respecting truthful mechanism C that obtains welfare at least*

$$\sum_{t \leq 0} W_t(A) + \sum_{t \geq 1} W_t(B) - 2v_{\max}d_{\max}m.$$

In particular, all jobs that arrive by time 0 and are completed by mechanism A are also completed by C . We can compose any number of “switching” steps, losing an additive $2v_{\max}d_{\max}m$ amount in welfare each time.

THEOREM 3.2. *Suppose we wish to switch among many order respecting truthful mechanisms as follows: start with A_0 at time 1, switch to A_1 at time t_1 , then to A_2 at time t_2 and so on till you switch to A_L at time t_L for some $L \in \mathbb{Z}_+$. Let $t_0 = 0$ and $t_{L+1} = T$ for notational convenience. Then there is an order respecting truthful mechanism whose welfare is at least*

$$\sum_{i=0}^L \sum_{t \in (t_i, t_{i+1}]} W_t(A_i) - 2Lv_{\max}md_{\max}.$$

PROOF. Let B_1 be the mechanism obtained by applying Lemma 3.1 to switch from A_0 to A_1 at time t_1 . Apply the lemma again to switch from B_1 to A_2 at time t_2 ; let the resulting mechanism be B_2 . Continuing this way, we apply the lemma to switch from B_i to A_{i+1} at time t_{i+1} , to get mechanism B_{i+1} , for all i up to $L - 1$. The resulting mechanism at the end, B_L , is the one we want. $\square$

In the rest of this section we prove Lemma 3.1.

3.1.1 The online switching algorithm.

Definition 3.3. The mechanism C (Lemma 3.1) is as follows.

- (1) For jobs that arrive by time 0, mimic mechanism A and return the same allocation, schedule and prices. Observe that jobs that are scheduled in this step are terminated by time $d_{\max}$.
- (2) Mark the remaining time slots in $[1, d_{\max}]$ as *unavailable*. This means that for all jobs j that were not considered in the previous step (because $a_j > 0$) and have deadline $d_j \leq d_{\max}$, we decline service and charge a price of 0.
- (3) For all remaining jobs, i.e., jobs j with $a_j > 0$ and $d_j > d_{\max}$, consider the jobs in the order of arrival, and do the following:
 - (a) If B rejects j , then reject j .
 - (b) If there are not enough slots available to cover j 's length prior to its deadline, reject j .
 - (c) Otherwise, accept and schedule j in a “best effort” manner. Specifically, assign to the job all of the slots that it gets in B and that are still available in C 's schedule. If any of these

slots is unavailable, replace it with the earliest available slot in C 's schedule. We call these newly assigned slots the “replacement” slots for job j . Charge j the same payment as in mechanism B .

Design choices. We explain the design choices made in the above mechanism. In step (1) we continue to process jobs that arrive by time 0 according to mechanism A . If we abruptly stop mechanism A , then there may be an incentive for some jobs to lie so that they get scheduled by time 0. In step (2) we make the remaining slots unavailable. Why not directly go to step (3) and schedule jobs that B has accepted in a best effort manner? One of the properties we need for truthfulness to hold is that jobs that arrive after time 0 finish at a time in mechanism C that is no earlier than their finish time in mechanism B . Otherwise there may be an incentive for a job to lie so that it gets accepted in B but is scheduled to finish after its true deadline, whereas mechanism C ends up scheduling it within its true deadline. Lemma 3.4 below shows that the algorithm C satisfies this property. Another option is to simply reject all jobs that arrive during the time interval $[1, d_{\max}]$. This would not be truthful since it incentivizes a job that arrives in this interval to report a later arrival time. In step (3) (a) if we start considering jobs that B rejected because we have some more available slots than B , we might break the truthfulness of B . Finally, in step (3) (c) we first assign the same slots to the job as in B in order to ensure the no early completion property. Assigning the remaining available slots in the chronological order is crucial for the welfare analysis.

3.1.2 Truthfulness. We begin by proving the no early completion property.

LEMMA 3.4. *Any replacement slot assigned in step (3) (c) is always later in time relative to the unavailable slot it replaces.*

PROOF. Suppose one of the slots assigned to a job j in mechanism B , say at time t , is unavailable in mechanism C . If $t \leq d_{\max}$ then by construction the replacement slot is later. Otherwise, t itself is a replacement slot for some other job j' . The arrival time of j' , $a_{j'}$ is no larger than a_j because jobs are processed in FIFO order. Since replacement slots are assigned in chronological order, all slots in $[a_{j'}, t - 1]$ must have been unavailable when t was assigned to j' . Now all slots in $[a_j, t]$ are unavailable when we consider job j , so its replacement for slot t can only be later. $\square$

LEMMA 3.5. *Mechanism C is truthful.*

PROOF. We consider three cases depending on which step of the mechanism handles the job. Recall that we assume that jobs cannot report an earlier arrival time.

- (1) Suppose that $a_j \leq 0$, which means that the job gets processed in step (1) and gets an allocation and payment as per mechanism A . If the job reports an arrival time > 0 , then it is not processed in step (1) and gets none of the slots in $[a_j, d_j]$, because $d_j \leq d_{\max}$, and all those slots are marked unavailable at the beginning of step (2). Any other misreport means that the job still gets processed in step (1). Now we can appeal to the fact that algorithm A is truthful to assert that the job does not benefit from misreporting its type.
- (2) Suppose that $a_j > 0$ and $d_j \leq d_{\max}$, which means it is processed in step (2). In this case, regardless of its actual report the job gets no slots in its time window.
- (3) Suppose that $a_j > 0$ and $d_j > d_{\max}$, which means it is processed in step (3). Since the job cannot report an earlier arrival time, it cannot be processed in step (1), and reporting a deadline $\leq d_{\max}$ means it gets no slots. Hence the only misreports we need to consider are such that the job is still processed in step (3).

The truthfulness of B should now imply that no misreport can be beneficial in C as well. This is almost true since, for instance, the price paid is the same in both (on acceptance). However,

there is a possibility that misreporting a later deadline in B (possibly combined with a misreport of other parameters) results in a lower price, but that in B 's schedule the job finishes after its true deadline d_j . This would be a non-beneficial misreport in B but could be beneficial in C if it actually finishes earlier than d_j in C , while enjoying the lower price. Lemma 3.4 ensures this does not happen. $\square$

3.1.3 Welfare. Define a time slot $t > d_{\max}$ to be “free” if mechanism C schedules fewer jobs in time t than mechanism B . The *number* of free slots at time t is the difference, given that it is non-negative, and zero otherwise. We first argue that there are few free slots in C 's schedule.

LEMMA 3.6. *All replacement slots occur before the first free slot.*

PROOF. Let t be the first free slot. Consider a job that arrives before t . This job is not assigned any replacement slots after t since t is free and hence available, and replacement slots are assigned in chronological order. We will argue that jobs arriving after t have no replacement slots, i.e., they get the same slots as in B . This is by induction on the arrival order of these jobs. Consider the very first such job. All earlier jobs arrive before t by definition, and have no replacement slots after t as already argued, therefore all of the slots assigned to this job in B 's schedule are available. This is the base case. The argument for the inductive case is almost exactly the same. $\square$

LEMMA 3.7. *The total number of free slots is at most $md_{\max}$. In particular, if t is the earliest time of a free slot, then all the free slots are in the interval $[t, t + d_{\max}]$. In other words, mechanisms B and C get synchronized after time $t + d_{\max}$.*

PROOF. From Lemma 3.6, there are no replacement slots after t . Any job that arrives at t or later and is scheduled in B gets the same slots in C as in B , and hence there are no free slots corresponding to such a job. All the free slots must correspond to jobs that arrive before t , and are therefore in the interval $[t, t + d_{\max}]$. $\square$

We are now ready to prove the main lemma of this section.

PROOF OF LEMMA 3.1. It is easy to see that mechanism C is also order respecting. Since we already showed that the mechanism is truthful in Lemma 3.5, we only need to argue about the welfare.

Any job that arrives before time 0 and is accepted by A is also accepted by C and completed, therefore it gets the same welfare as A upto time 0. Now we argue about the total loss in welfare during the time $t \geq 1$. Let $\ell_B(t)$ (resp. $\ell_C(t)$) be the number of jobs scheduled at time $t \geq 1$ by mechanism B (resp. mechanism C), and let $\ell_F(t)$ be the number of free slots at time t . By the definition of a free slot, we have that

$$\sum_{t \geq 1} \ell_B(t) \leq \sum_{t \geq d_{\max}+1} (\ell_C(t) + \ell_F(t)) + md_{\max}.$$

The set of jobs accepted by C is a subset of the set of jobs accepted by B , due to steps (2) and (3a). The total length of all jobs that are accepted by B but not by C is equal to $\sum_{t \geq 1} (\ell_B(t) - \ell_C(t)) \leq \sum_{t \geq d_{\max}+1} \ell_F(t) + md_{\max} \leq 2md_{\max}$, where the last inequality is from Lemma 3.7. Thus the total loss is at most $2mv_{\max}d_{\max}$. $\square$

3.2 Reduction to experts with switching costs

Let $\text{Expert-ALG}(C)$ denote an online algorithm for the problem of learning from expert advice with switching cost C that achieves the regret guarantee of Theorem 2.1. Expert-ALG is given an instance with n experts, indexed by $i \in [n]$. It specifies for every time step $t \in [T]$ a random

expert i_t , and then receives a reward vector $\{r_t^{(i)}\}$. Our online scheduling algorithm, that we call *Follow-The-Switcher* or FTS, simulates Expert-ALG in a black-box fashion and follows its advice on which expert, a.k.a. algorithm, to run at every time step.

Definition 3.8. Given the n online scheduling algorithms, $\text{ALG}_1, \dots, \text{ALG}_n$, the Follow-The-Switcher, a.k.a. FTS, algorithm simulates the online algorithm Expert-ALG(C) with C set to $2v_{\max}d_{\max}m$. It then proceeds as follows.

At each time $t \in [T]$:

- (1) Simulate algorithms $\text{ALG}_1, \dots, \text{ALG}_n$ on the freshly arrived set of jobs.
 - (2) Query Expert-ALG to obtain the index $i_t \in [n]$.
 - (3) If $i_t \neq i_{t-1}$, then switch from $\text{ALG}_{i_{t-1}}$ to ALG_{i_t} as described in Section 3.1. Otherwise continue running the same algorithm $\text{ALG}_{i_{t-1}} = \text{ALG}_{i_t}$.
 - (4) Set $r_t^{(i)} \leftarrow W_t(\text{ALG}_{i_t})$ for all $i \in [n]$. Send the reward vector $\{r_t^{(i)}\}$ to Expert-ALG.
-

The following theorem now immediately follows from Theorem 3.2.

THEOREM 3.9. *Let $C = 2v_{\max}d_{\max}m$. Then the Follow-The-Switcher (FTS) algorithm, described in Definition 3.8, admits the following regret-bound:*

$$\text{REG}(\text{FTS}) \leq \text{REG}(\text{Expert-ALG}(C)) \leq O\left(mv_{\max}\sqrt{Td_{\max}\log n}\right).$$

4 NON-CLAIRVOYANT SETTING

In this section, we look at regret minimization in the non-clairvoyant setting; we recall the main differences here. Every job reports all parts of its type except its length, and the algorithm only observes the length of a job when (and if) it is completed. Hence a non-clairvoyant algorithm cannot plan for a complete schedule ahead of time. The algorithm maintains a queue of unfinished jobs and at every time t decides which job to schedule from this queue at that time. The deadline for a job is now the number of time slots that the job is willing to wait out. If a job passes its deadline, which means that the number of time slots that a job j waits since its arrival exceeds a threshold d_j , then the job is deleted from the queue. Due to this reason, the notion of promptness is not quite applicable to the non-clairvoyant setting. In its place, we use the order respecting property, which states that only jobs arriving earlier can influence the allocation and payments for a given job. We assume that there is a total order on the arrival time of the jobs, by breaking ties arbitrarily in case multiple jobs arrive at the same time.

Similar to the clairvoyant setting, we have a set of n online scheduling algorithms, $\text{ALG}_1, \dots, \text{ALG}_n$, and we aim to design an online algorithm that minimizes the regret relative to the best of these algorithms in hindsight. We begin our discussion in Section 4.1, where we show that this benchmark by itself is impossible to compete with, which motivates a reasonable modification. In Section 4.2 we show how to switch between two mechanisms truthfully. Finally in Section 4.3, we show how to use multi-armed bandit algorithms to get tight regret bounds.

4.1 Scheduling algorithms with random restarts

Robustness of the benchmark. In the non-clairvoyant setting, it is easy to come up with examples showing that the welfare obtained by an online scheduling mechanism is very sensitive to *timing* in the adversarial model of jobs, i.e., by slightly changing the starting time of the mechanism the obtained welfare can be drastically different. This has been demonstrated in Example 4.1.

Example 4.1. Consider running FIFO scheduling with pricing admission control at $p = 1$. Suppose we have two sets of jobs J_1 and J_2 , with $v_j = 1$ for all the jobs. Suppose $(a_j, l_j) = (4i + 1, 4)$ for all $j \in J_1$ and $(a_{j'}, l_{j'}) = (4i + 4, 2)$ for all $j' \in J_2$, where $i = 0, 1, \dots, (\frac{T}{4} - 1)$. All of them have immediate deadlines, which means they need to be scheduled when they arrive. If we schedule jobs in J_1 , we generate a welfare equal to T . Now consider starting at a time t which does not corresponds to the arrival times of jobs in J_1 . Then we only end up scheduling jobs in J_2 (and in the non-clairvoyant setting we will not even notice how valuable jobs in J_1 were!) and get only welfare equal to $T/2$.

The example is robust to variations such as picking a random job when you start (by having many identical copies of jobs in J_2), or picking a random starting time.

As it is clear from this example, the welfare obtained from such a mechanism cannot be a reasonable benchmark for our regret minimization, as one has to think very carefully about when to start running such a mechanism to calculate the benchmark. Otherwise, the benchmark mechanism could easily get tricked into a false start. In other words, we need to define the benchmark in a way that is *robust* to this sort of timing issues, independent of the choice of scheduling mechanism defining the benchmark.

Syncing issues. In the non-clairvoyant setting, our scheduling mechanism is not able to *simulate* an arbitrary candidate scheduling mechanism ALG_i starting from an arbitrary time, since there is no way of knowing its state. Accordingly, following the decisions of a bandit algorithm, similar to what we did in Section 3.2, is generally not possible in the non-clairvoyant setting.

However, if both our mechanism and the new switched scheduling mechanism restart from a fresh state at exactly the same time after switching (e.g. slightly after the switching time when our mechanism is done with its current jobs) then our mechanism can sync with the new scheduling mechanism.

To address above concerns and to be able to design a truthful online scheduling mechanism that achieves a meaningful regret bound, we introduce a couple of new ingredients in our model and redefine our benchmark. We start by defining the notion of a *random restart* formally as following. An important property of the way we restart is that it preserves truthfulness: a mechanism that was truthful to begin with is still truthful with a restart. We use ideas from the clairvoyant truthful switching: think of the restart as a switch from A to A .

Definition 4.2. Given an online scheduling mechanism A , we define the *restart* at time t as follows.

- During $[t : t + (\ell_{\max} + d_{\max})]$, mechanism $\mathcal{M}$ continues working on the jobs that have arrived before t .
- If a job j arrives during $[t : t + (\ell_{\max} + d_{\max})]$, modify it as follows.
 - Shift its arrival time to the end of this interval, i.e. $a_j \leftarrow t + (\ell_{\max} + d_{\max}) + 1$.
 - Adjust the deadline of the job so that it reflects the time lost during the interval $[t, t + d_{\max} + \ell_{\max}]$. This might mean some jobs are past their deadline. These jobs are rejected.
 - Preserve the arrival order. Use the tie breaking rule to make sure the arrival order of jobs whose starting time was set to $t + d_{\max} + \ell_{\max}$ is the same as in the original instance.

Having the formal definition of a restart, we ask the following question: how can one define a robust benchmark in the non-clairvoyant scheduling problem, given a set of candidate scheduling mechanisms? Here is an adaptation of our previous benchmark, i.e. welfare of the best-in-hindsight scheduling mechanism, for the non-clairvoyant setting.

Definition 4.3. Given candidate scheduling mechanisms $\text{ALG}_1, \dots, \text{ALG}_n$, and a parameter $\gamma \in [0, 1]$, the *random-restarting* mechanisms $\overline{\text{ALG}}_1, \dots, \overline{\text{ALG}}_n$ are defined to be the original candidate

mechanisms accompanied by independent random restarts of probability γ at every time $t \in [T]$. We define the *random-restarting benchmark* for an instance $\mathcal{J}$ to be $\overline{\text{OPT}}(\mathcal{J}) = \max_{i \in [n]} \mathbb{E}[W(\overline{\text{ALG}}_i, \mathcal{J})]$.

The benefits of using benchmark $\overline{\text{OPT}}(\mathcal{J})$ are twofold. First, this benchmark is robust to timing issues, because the scheduling mechanism generating the benchmark restarts independently at random at every time t with probability γ . Therefore, the benchmark loses no more than the generated welfare between two consecutive random-restarts due to timing issues. Second, while this property does not hold in general, for a large family of scheduling mechanisms (such as posted-pricing-FIFO) and under the *stochastic model* of jobs (e.g. see [Chawla et al., 2017]), the welfare loss due to independent (but infrequent) random restarts will easily be bounded. This property of the pair (stochastic model, online scheduling mechanisms), which we call *robustness-to-welfare-loss*, is formalized as following.

Definition 4.4. Given a distribution over jobs $\mathcal{D}$, the random-restarting benchmark $\overline{\text{OPT}}(\mathcal{J})$ with parameter γ is robust-to-welfare-loss in expectation over stochastic jobs $\mathcal{D}$ if

$$\mathbb{E}_{\mathcal{J} \sim \mathcal{D}}[\overline{\text{OPT}}(\mathcal{J})] \geq \mathbb{E}_{\mathcal{J} \sim \mathcal{D}}[\text{OPT}(\mathcal{J})] - \gamma \cdot T \cdot v_{\max}(\ell_{\max} + d_{\max})$$

Moreover, a random-restarting online scheduling mechanism $\overline{\text{ALG}}_i$ with parameter γ is robust-to-welfare-loss in expectation under stochastic jobs $\mathcal{D}$ if

$$\mathbb{E}_{\mathcal{J} \sim \mathcal{D}} \left[\sum_{t \in [T]} W_t(\overline{\text{ALG}}_i, \mathcal{J}) \right] \geq \mathbb{E}_{\mathcal{J} \sim \mathcal{D}} \left[\sum_{t \in [T]} W_t(\text{ALG}_i, \mathcal{J}) \right] - \gamma \cdot T \cdot v_{\max}(\ell_{\max} + d_{\max})$$

Clearly, if all of the random-restarting mechanisms $\overline{\text{ALG}}_1, \dots, \overline{\text{ALG}}_n$ are robust-to-welfare-loss under stochastic job model $\mathcal{D}$, then $\overline{\text{OPT}}(\mathcal{J})$ will also be robust-to-welfare-loss under $\mathcal{D}$.

4.2 Truthful switching in the non-clairvoyant setting

We now show how to switch between truthful non-clairvoyant mechanisms. We reiterate a subtle aspect of truthfulness in the non-clairvoyant setting: a job that tries to influence the mechanism by switching its position in the arrival order can be treated the same way as a job reporting a later arrival time: it cannot be beneficial to do this given that the mechanism is truthful. Posted-pricing-FIFO is an example of a mechanism that is both truthful and order respecting: all jobs whose values are less than a threshold price are rejected, and the rest of the jobs are scheduled in arrival order.

We claim that the random restart algorithm of Definition 4.2 preserves the truthfulness of the underlying scheduling mechanism. The following lemma is proved in the full version of the paper.

LEMMA 4.5. *If A is an order respecting truthful mechanism, then A with (arbitrary) restarts is also order respecting and truthful.*

Combining ideas from the truthful switching algorithm in the clairvoyant setting and the truthful random restart algorithm, we develop a truthful switching algorithm for non-clairvoyant settings that switches from a mechanism A to a mechanism B at time 0.

Definition 4.6. The mechanism C is as follows.

- (1) For jobs that arrive by time 0, mimic mechanism A and return the same allocation, schedule and prices as A . All these jobs are completed by time $d_{\max} + \ell_{\max}$.
- (2) For all remaining jobs, i.e., jobs j with $a_j \geq 1$ run mechanism B on these with the following modifications.
 - (a) If the arrival time of a job is $< d_{\max} + \ell_{\max}$, set its arrival time to $d_{\max} + \ell_{\max} + 1$.

- (b) Adjust the deadline of the job so that it reflects the time lost during the interval $[1, d_{\max} + \ell_{\max}]$. This might mean some jobs are past their deadline. These jobs are rejected.
- (c) Preserve the arrival order. Use the tie breaking rule to make sure the arrival order of jobs whose starting time was set to $d_{\max} + \ell_{\max}$ is the same as in the original instance.

LEMMA 4.7. *The state of the algorithm C at time $d_{\max} + \ell_{\max}$ is the same as the state of the algorithm B at time $d_{\max} + \ell_{\max}$, given that B is restarted during the interval $[1, d_{\max} + \ell_{\max}]$.*

PROOF. This follows by observing that Step (2) of mechanism C is identical to the modifications made during a restart. Any job that arrives by time 0 does not influence the state of mechanism B at time $d_{\max} + \ell_{\max}$ in either case. $\square$

We obtain the following theorem (see the full version for a proof).

THEOREM 4.8. *Given order respecting truthful mechanisms A and B in the non-clairvoyant setting, switching mechanism C in Definition 4.6 is order respecting, truthful, and obtains welfare at least*

$$\sum_{t \leq d_{\max} + \ell_{\max}} W_t(A) + \sum_{t \geq 1 + d_{\max} + \ell_{\max}} W_t(B),$$

given that A and B are restarted at time 1.

4.3 Reduction to multi-armed bandits

In this section, we show how to design a truthful online learning algorithm that minimizes the regret relative to the random-restarting benchmark, i.e. best-in-hindsight of random-restarting truthful mechanisms $\text{ALG}_1, \dots, \text{ALG}_n$. Similar to Section 3.2, we consider a relevant adversarial Multi-Armed Bandits problem or MAB, where we have an arm for each of the n online scheduling mechanisms and we have bandit feedback, meaning that any algorithm only observes the reward of the arm it plays and not the other arms.

In such a setting, we assume having query access to a MAB algorithm Bandit-ALG that admits the optimal regret bound $O(R\sqrt{Tn \log n})$ in Theorem 2.2, where n is the number of arms, T is the time horizon and R is an upper-bound on the reward of an arm. Moreover, for the sake of simplicity, we assume Bandit-ALG does not need to know the time horizon T or rewards range R in advance, and it only needs to know these quantities are bounded. This assumption can be removed by using a *doubling trick*: given black-box access to a bandit algorithm $\text{Bandit-ALG}_1(R, T)$ that requires knowing R and T , one can come up with another black-box algorithm Bandit-ALG_2 with the same asymptotic regret bound that does not need this information. This reduction is explained in the full version of the paper.

Our proposed algorithm, which we call *Follow-The-Bandit-Switcher* or *FTBS*, uses Bandit-ALG in a black-box fashion when looped in with the right rewards. It then follows Bandit-ALG 's advice by truthful switching between arms, as suggested by Theorem 4.8. This helps the FTBS to aggregate truth and welfare guarantees of mechanisms $\text{ALG}_1, \dots, \text{ALG}_n$ even in the non-clairvoyant setting. We formally define the FTBS as following.

Definition 4.9. Given the n online scheduling mechanisms, $\text{ALG}_1, \dots, \text{ALG}_n$, query access to online bandit algorithm Bandit-ALG , and parameter $\gamma \in [0, 1]$, the Follow-The-Bandit-Switcher mechanism proceeds as follows.

Initialize $\hat{t} = 1$. (This index is counter for number of Heads)

Query Bandit-ALG for the initial arm i_1 , and run algorithm ALG_{i_1} .

At each time $t \in [T]$:

- (1) Flip an independent coin $\kappa(t)$ with $\Pr(\kappa(t) = \text{Heads}) = \gamma$.

- (2) If coin $\kappa(t)$ shows a Heads,
- (a) Let $t' = \max\{t'' < t : \kappa(t'') = \text{Heads}\}$. If no such an integer exists, let $t' = 1$.
 - (b) Update the sequence of bandit rewards between two consecutive Heads:
$$r_x^{(i_t)} = \begin{cases} 0, & \text{if } x \in [t', \min(t' + \ell_{\max} + d_{\max}, t - 1)] \\ W_x(\text{FTBS}), & \text{if } x \in [\min(t' + \ell_{\max} + d_{\max}, t - 1) + 1, t - 1] \end{cases}$$
 - (c) Set $R_{\hat{t}}^{(i_{t'})} \leftarrow \sum_{x \in [t', t-1]} r_x^{(i_{t'})}$ and send this bandit feedback to Bandit-ALG.
 - (d) Set $\hat{t} \leftarrow \hat{t} + 1$.
 - (e) Let the new arm chosen by Bandit-ALG be $i_{\hat{t}} \in [n]$.
 - (f) If $i_{\hat{t}} \neq i_{\hat{t}-1}$, switch from $\text{ALG}_{i_{\hat{t}-1}}$ to $\text{ALG}_{i_{\hat{t}}}$ using the mechanism in Section 4.2.
 - (g) Otherwise, restart $\text{ALG}_{i_{\hat{t}}}$ at time t .
- (3) If coin $\kappa(t)$ shows a Tails, continue running $\text{ALG}_{i_{\hat{t}}}$.

We now state and prove a tight regret bound (up to logarithmic factor) for FTBS.

THEOREM 4.10. *The Follow-The-Bandit-Switcher (FTBS) scheduling mechanism, described in Definition 4.9, admits the following regret-bound if $\gamma = (\ell_{\max} + d_{\max})^{-2/3} T^{-1/3} (n \log(n))^{1/3}$:*

$$\text{REG}(\text{FTBS}) \leq O(m \cdot v_{\max}(\ell_{\max} + d_{\max})^{1/3} (n \log(n))^{1/3} T^{2/3} \log T) = \tilde{O}(T^{2/3})$$

where $\text{REG}(\text{FTBS})$ is the regret relative to $\overline{\text{OPT}}$, i.e. the random-restarting benchmark as in Definition 4.3.

We finally consider stochastic jobs, and we focus on benchmarks that are robust-to-welfare-loss in expectation under this stochastic model, as described in Definition 4.4. The following corollary is immediate.

COROLLARY 4.11. *Given a distribution $\mathcal{D}$ over jobs and a robust-to-welfare-loss benchmark $\text{OPT}(\mathcal{J})$ as in Definition 4.4, there exists a scheduling mechanism whose expected regret relative to $\mathbf{E}_{\mathcal{J} \sim \mathcal{D}}[\text{OPT}(\mathcal{J})]$ is bounded by $\tilde{O}(T^{2/3})$.*

PROOF OF THEOREM 4.10. Let $C_1 \triangleq v_{\max}(\ell_{\max} + d_{\max})m$ and $C_2 \triangleq v_{\max}m$ be constants. Fix a sequence of coins $\mathbf{K}[T] \triangleq [\kappa(1), \dots, \kappa(T)]$ and let $t_1, t_2, \dots, t_{\hat{T}}$ be the times $t \in [T]$ that coin $\kappa(t)$ shows a Heads. As a convention, let $t_0 \triangleq 1$ and $t_{\hat{T}+1} \triangleq \tilde{T} \geq T$ be the next time that coin $\kappa(t)$ flips a Heads if we continue flipping after T . By abuse of notation, we will use $\overline{\text{ALG}}_i$ to denote mechanism ALG_i restarted at exactly these times. For each $i \in [n]$ and $x \in [1, \hat{T}]$, let $R_x^{(i)}$ be:

$$R_x^{(i)} \triangleq \sum_{t=\min(t_{x-1}+\ell_{\max}+d_{\max}, t_x-1)+1}^{t_x-1} W_t(\overline{\text{ALG}}_i) \geq \sum_{t=t_{x-1}}^{t_x-1} W_t(\overline{\text{ALG}}_i) - C_1$$

Also, for the last interval $[t_{\hat{T}}, T]$ and for each $i \in [n]$, let $R_{\hat{T}+1}^{(i)}$ be

$$R_{\hat{T}+1}^{(i)} \triangleq \sum_{t=\min(t_{\hat{T}}+\ell_{\max}+d_{\max}, t_x-1)+1}^T W_t(\overline{\text{ALG}}_i) \geq \sum_{t=t_{\hat{T}}}^T W_t(\overline{\text{ALG}}_i) - C_1$$

Therefore, by summing over all intervals $\{[t_{x-1}, t_x]\}_{x \in [\hat{T}]} \cup [t_{\hat{T}}, T]$, we have:

$$\forall i \in [n] : \sum_{x \in [\hat{T}+1]} R_x^{(i)} \geq W(\overline{\text{ALG}}_i) - C_1(\hat{T} + 1) \quad (1)$$

Note that after the coin shows a Heads at time t_x , our mechanism either stays with the same algorithm and does a restart or the truthful switching from $\text{ALG}_{i_{x-1}}$ to ALG_{i_x} as in Section 4.2. Also, Theorem 4.8 and the definition of truthful restart in Definition 4.2 guarantee that our scheduling

mechanism will be synced with the restarting algorithms $\overline{\text{ALG}}_{i_x}$ after $\ell_{\max} + d_{\max}$ units of time, and therefore the mechanism generates a welfare that is at least the welfare generated by the new arm in the interval $[\min(t_{x-1} + \ell_{\max} + d_{\max}, t_x - 1) + 1, t_x - 1]$. Formally speaking,

$$\forall x \in [\hat{T}] : \sum_{t=t_{x-1}}^{t_x} W_t(\text{FTBS}) \geq R_x^{(i_x)}, \quad \sum_{t=t_{\hat{T}}}^T W_t(\text{FTBS}) \geq R_{\hat{T}}^{(i_{\hat{T}})}$$

and therefore, by summing over all intervals $\{[t_{x-1}, t_x]\}_{x \in [\hat{T}]} \cup [\hat{T}, T]$, we have:

$$W(\text{FTBS}) \geq \sum_{x \in [\hat{T}+1]} R_x^{(i_x)} \quad (2)$$

Conditioned on the sequence of coin flips $\mathbf{K}[T]$ (and therefore times $t_1, t_2, \dots, t_{\hat{T}}$), the rewards $\{R_x^{(i)}\}_{x \in [\hat{T}+1], i \in [n]}$ define an (oblivious) adversarial instance of a MAB problem. For this adversarial instance, time horizon is indeed the random variable $(\hat{T} + 1)$, i.e. number of Heads in the sequence $\mathbf{K}[T]$ plus one. Moreover, if we let $\hat{U} = \max_{j \in [\hat{T}+1]} U_j$ where

$$j \in [1, \hat{T} + 1] : U_j \triangleq t_j - t_{j-1}$$

then the rewards of this instance would also be upper-bounded by the random variable $C_2 \hat{U}$. Then, due to the optimal regret bound of Bandit-ALG (Theorem 2.2), we have:

$$\forall i \in [n] : \mathbf{E} \left[\sum_{x \in [\hat{T}+1]} R_x^{(i)} | \mathbf{K}[T] \right] - \mathbf{E} \left[\sum_{x \in [\hat{T}+1]} R_x^{(i_x)} | \mathbf{K}[T] \right] \leq O(\hat{U} \hat{T}^{1/2} C_2 (n \log n)^{1/2}) \quad (3)$$

Combining Inequalities (1), (2) and (3) and taking expectations:

$$\text{REG}(\text{FTBS}) \leq O(\mathbf{E}[\hat{U} \hat{T}^{1/2} C_2 (n \log n)^{1/2}]) + \mathbf{E}[C_1(\hat{T} + 1)] \quad (4)$$

Now, $\mathbf{E}[C_1(\hat{T} + 1)] = O(\gamma \cdot C_1 T)$. To bound the other term, we use the following fact, proved in [Eisenberg, 2008], about independent and identically distributed geometric random variables.

LEMMA 4.12 ([EISENBERG, 2008]). *If $g_1, \dots, g_k$ are i.i.d. and $g_i \sim \text{Geometric}(\gamma)$, then*

$$\mathbf{E} \left[\max_{j \in [k]} g_j \right] \leq H_k \cdot \gamma^{-1} = O(\gamma^{-1} \log k)$$

Note that conditioned on $\hat{T}$, random variables $\{U_j\}_{j \in [\hat{T}+1]}$ are $(\hat{T} + 1)$ i.i.d. $\text{Geometric}(\gamma)$ random variables. Using Lemma 4.12 we have:

$$\mathbf{E}[\hat{U} \hat{T}^{1/2} | \hat{T}] = \hat{T}^{1/2} \mathbf{E} \left[\max_{j \in [\hat{T}]} U_j | \hat{T} \right] \leq \hat{T}^{1/2} \cdot H_{\hat{T}} \cdot \gamma^{-1} = \gamma^{-1} \cdot O(\hat{T}^{1/2} \log(\hat{T}))$$

Now, function $f(x) = x^{1/2} \log(x)$ is concave. By taking expectation and using Jensen's inequality, we further upper-bound this term. Hence:

$$\mathbf{E}[\hat{U} \hat{T}^{1/2}] \leq \gamma^{-1} \cdot O(f(\mathbf{E}[\hat{T}])) = \gamma^{-1/2} \cdot T^{1/2} \log(\gamma T)$$

So, we can upper-bound the RHS of (4) by

$$O(\gamma^{-1/2} T^{1/2} \log(\gamma T) C_2 (n \log n)^{1/2}) + O(\gamma \cdot C_1 T)$$

By setting $\gamma = (\frac{C_2}{C_1})^{2/3} T^{-1/3} (n \log n)^{1/3}$, we get the desired regret bound. $\square$

5 LOWER BOUNDS

We first state a lower bound of $\Omega(\sqrt{T})$ on the regret for the clairvoyant scheduling problem. The proof of this theorem can be found in the full version of the paper. We note that this lower bound can be extended to $\Omega(m\sqrt{T})$ when there are m machines by simply having m copies of the same set of jobs every time. Similarly the lower bound can be made to scale linearly with $v_{\max}$ and $\sqrt{d_{\max}}$, matching the upper bound we give in Theorem 3.9.

THEOREM 5.1. *There exists an instance of the clairvoyant scheduling problem where the regret of any online algorithm relative to the hindsight optimal posted-pricing-FIFO algorithm is $\Omega(\sqrt{T})$.*

Next we show a lower bound of $\tilde{\Omega}(T^{2/3})$ on the regret for the non-clairvoyant setting that matches our upper bound within polylogarithmic factors. Our lower bound follows by a reduction from the lower bound given in [Dekel et al., 2014] for the multi-armed bandit problem with switching costs.

For the bandit problem with switching costs with n actions, [Dekel et al., 2014] show that there exists a sequence of loss functions $\ell_1, \ell_2, \dots, \ell_T$, where $\ell_i : [n] \rightarrow [0, 1]$, such that any online algorithm incurs a regret of at least $\tilde{\Omega}(n^{1/3}T^{2/3})$. We use this loss sequence to define an instance of the non-clairvoyant scheduling problem as follows. First we give a reduction to regret against the welfare benchmark without random restarts. Later we show how to extend the lower bound to apply against the random restart benchmark. In our lower bound instance, we fix $n = 2$, and

let $\ell_i(1)$ and $\ell_i(2)$ denote the losses of actions 1 and 2 in round i as defined in [Dekel et al., 2014]. We map each round of the game to 8 time steps; that is, round i corresponds to the time interval $[8i, 8(i+1) - 1]$. Our instance has 4 sets of jobs J_1, J_2, J_3 , and J_4 , as shown in the above figure. In each round, one job from each set arrives. Jobs in the set J_1 arrive at the beginning of each round; that is, at time steps $8i$ for $i = 0, 1, 2, \dots, T$. The processing length of a job $j \in J_1$ that arrives in the round i is 6 with probability $p_i(1)$ and 8 with probability $(1 - p_i(1))$, where $p_i(1) = 1/2 + \ell_i(1)/2$. Observe that processing lengths of the jobs in J_1 depend on losses defined by [Dekel et al., 2014]. Further, the jobs in J_1 have a value of 1 per unit length. In round i , a job from J_2 arrives at time $8(i+1) - 2$ for $i = 0, 1, 2, 3, \dots, T$, and has a processing length of 2. The value per unit length of jobs in J_2 is 3.

The set J_3 consists of jobs that arrive at time instants $8i - 3$ for $i = 1, 2, \dots, T$, and have value per unit length of 2. The processing length of job $j \in J_3$ released in the round i is 4 with probability $p_i(2)$ and 2 with probability $1 - p_i(2)$, where $p_i(2) = \ell_i(2)$. Similar to the jobs in J_1 , the processing lengths of jobs in J_3 depend on the losses defined by the result of [Dekel et al., 2014]. Finally, the jobs in set J_4 are released at time steps $8i - 1$ for $i = 1, 2, \dots, T$, and have a processing length of 2 and value per unit length of 3.

The deadlines of all jobs are equal to their arrival times. (This condition is not necessary but simplifies the construction.) Hence, if a job is not scheduled upon its arrival, the algorithm loses the job.

Let ALG_1 and ALG_2 denote the two posted price scheduling mechanisms with prices 1 and 2 and FIFO scheduling policy. (Note that other posted price mechanisms, for example one that posts a price of 3, lose a constant factor in each round, hence we do not consider them.) The algorithms ALG_1 and ALG_2 correspond to the action set in the bandit problem with switching costs. The following two lemmas follow from the construction of lower bound instance. See the full version of the paper for a proof.

LEMMA 5.2. *For all rounds $i = 0, 1, \dots, T$, the expected value of ALG_1 in round i is $10 - 2\ell_i(1)$ and expected value of ALG_2 in round i is $10 - 2\ell_i(2)$.*

LEMMA 5.3. *If an online algorithm switches from ALG_2 to ALG_1 in any round i , it incurs a loss of at least 6.*

From Lemmas 5.2 and 5.3, we get the following theorem.

THEOREM 5.4. *Minimax regret of non-clairvoyant scheduling problem is at least $\tilde{\Omega}(T^{2/3})$.*

PROOF. Consider an online scheduling algorithm for the non-clairvoyant scheduling problem. At the beginning of each round i , it can either follow ALG_1 and obtain a value of $10 - 2\ell_i(1)$ or follow ALG_2 and get a value $10 - 2\ell_i(2)$. Furthermore, since we are in the non-clairvoyant setting, the online algorithm won't know the value it can obtain from the algorithm it is not following. Since, the value obtained by ALG_1 and ALG_2 are exactly the same in each round except for the terms $-\ell_i(1)$ and $-\ell_i(2)$, the regret of the non-clairvoyant scheduling algorithm is equal to the regret it suffers on the losses at each round i . Moreover, switching from ALG_2 to ALG_1 in any round incurs a cost of 6. Hence, our scheduling instance corresponds to bandit with switching cost problem, where switching from action 2 to action 1 incurs a cost. The result of [Dekel et al., 2014] shows that the problem has a minimax regret of least $\tilde{\Omega}(T^{2/3})$, when there is a switching cost between any pair of actions. However, it is easy to modify the proof in [Dekel et al., 2014], where there is a switching cost only between action 2 to action 1, losing a factor of 2 in the regret bound [Dekel et al., 2014]. This completes our reduction. $\square$

Now we argue that Theorem (5.4) can be extended to random restarting benchmarks. Before that we make the following simple observation regarding the lower bound instance in Theorem (5.4). Consider the posted price mechanisms ALG_1 and ALG_2 as defined in the proof of Theorem (5.4). $\overline{\text{ALG}}_1(\gamma), \overline{\text{ALG}}_2(\gamma)$ denote the random restarting versions of them with restart parameter γ .

LEMMA 5.5. *For some constant $c \geq 11$, we have:*

$$j \in \{1, 2\} : \mathbf{E}[W(\overline{\text{ALG}}_j(\gamma), J)] \geq W(\text{ALG}_j, J) - c \cdot \gamma \cdot T$$

PROOF. Proof of the lemma follows immediately from the observation that if ALG_j for $j = 1, 2$ restarts at time t , and t is in round i , then it loses values of jobs in that round. The expected number of restarts by the algorithms is at most $\gamma \cdot T$. As the total value of jobs in each round is at most 11, the statement of the lemma holds for $c \geq 11$. $\square$

To extend the lower bound to random restarting benchmarks, we need the following theorem from [Dekel et al., 2014] for the bandit with switching costs problem.

THEOREM 5.6 ([DEKEL ET AL., 2014]). *Let $\mathcal{A}$ be a multi-armed bandit algorithm that guarantees an expected regret (without switching costs) of $O(T^\alpha)$ then there exists a sequence of loss functions that forces $\mathcal{A}$ to make $\tilde{\Omega}(T^{2(1-\alpha)})$ switches.*

Combining Lemma 5.5 and above theorem, we prove the following theorem.

THEOREM 5.7. *No online algorithm can achieve a regret $O(T^{2/3-\epsilon})$ for any $\epsilon > 0$ against a random restarting benchmark with restart parameter $\gamma \in (T^{-1}, T^{-1/3}]$.*

PROOF. Proof is by contradiction. Suppose there is an online non-clairvoyant algorithm $\mathcal{A}$ that achieves a regret of $O(T^{2/3-\epsilon})$ for some $\epsilon > 0$. From Lemma 5.5, this implies that it achieves a regret of at most $O(T^{2/3-\epsilon})$ against the non-restarting benchmark of Theorem 5.4. Our proof of Theorem 5.4 gives a reduction from the bandit with switching cost problem to the non-clairvoyant scheduling problem. Therefore, we can invoke Theorem 5.6, which implies that $\mathcal{A}$ does at least $\tilde{\Omega}(T^{2(1-(2/3-\epsilon))})$ switches. Since there are only two actions in our lower bound instance, $\mathcal{A}$ still pays a switching cost of at least $\tilde{\Omega}(T^{2(1-(2/3-\epsilon))}) = \tilde{\Omega}(T^{2/3+2\epsilon})$ in switching from ALG_2 to ALG_1 . This gives a contradiction to the regret of $\mathcal{A}$ being $O(T^{2/3-\epsilon})$ against the non-restarting benchmark, and completes the proof. $\square$

REFERENCES

- Peter Auer, Nicolo Cesa-Bianchi, Yoav Freund, and Robert E Schapire. 1995. Gambling in a rigged casino: The adversarial multi-armed bandit problem. In *Foundations of Computer Science, 1995. Proceedings., 36th Annual Symposium on*. IEEE, 322–331.
- Yossi Azar, Inna Kalp-Shaltiel, Brendan Lucier, Ishai Menache, Joseph Seffi Naor, and Jonathan Yaniv. 2015. Truthful Online Scheduling with Commitments. In *Proceedings of the Sixteenth ACM Conference on Economics and Computation*. ACM, 715–732.
- Yossi Azar and Ety Khaitsin. 2011. Prompt mechanism for ad placement over time. In *International Symposium on Algorithmic Game Theory*. Springer, 19–30.
- Moshe Babaioff, Yishay Mansour, Noam Nisan, Gali Noti, Carlo Curino, Nar Ganapathy, Ishai Menache, Omer Reingold, Moshe Tennenholtz, and Erez Timnat. 2017. ERA: A Framework for Economic Resource Allocation for the Cloud. In *Proceedings of the 26th International Conference on World Wide Web Companion*. 635–642.
- Avrim Blum and Carl Burch. 2000. On-line Learning and the Metrical Task System Problem. *Machine Learning* 39, 1 (2000), 35–58. DOI: <http://dx.doi.org/10.1023/A:1007621832648>
- Avrim Blum, Shuchi Chawla, and Adam Kalai. 2002. Static Optimality and Dynamic Search-optimality in Lists and Trees. In *Proceedings of the Thirteenth Annual ACM-SIAM Symposium on Discrete Algorithms (SODA '02)*. Society for Industrial and Applied Mathematics, Philadelphia, PA, USA, 1–8. <http://dl.acm.org/citation.cfm?id=545381.545382>
- Ran Canetti and Sandy Irani. 1998. Bounding the power of preemption in randomized scheduling. *SIAM J. Comput.* 27, 4 (1998), 993–1015.
- Nicolo Cesa-Bianchi and Gábor Lugosi. 2006. *Prediction, learning, and games*. Cambridge university press.
- Shuchi Chawla, Nikhil Devanur, Alexander Holroyd, Anna Karlin, James Martin, and Balasubramanian Sivan. 2017. Stability of service under time-of-use pricing. In *STOC*. To appear.
- Richard Cole, Shahar Dobzinski, and Lisa Fleischer. 2008. Prompt mechanisms for online auctions. In *International Symposium on Algorithmic Game Theory*. Springer, 170–181.
- Ofer Dekel, Jian Ding, Tomer Koren, and Yuval Peres. 2014. Bandits with Switching Costs: T2/3 Regret. In *Proceedings of the Forty-sixth Annual ACM Symposium on Theory of Computing (STOC '14)*. ACM, New York, NY, USA, 459–467. DOI: <http://dx.doi.org/10.1145/2591796.2591868>
- Bennett Eisenberg. 2008. On the expectation of the maximum of IID geometric random variables. *Statistics & Probability Letters* 78, 2 (2008), 135–143.
- Yoav Freund and Robert E Schapire. 1995. A decision-theoretic generalization of on-line learning and an application to boosting. In *European conference on computational learning theory*. Springer, 23–37.
- Mohammad T. Hajiaghayi, Robert D. Kleinberg, Mohammad Mahdian, and David C. Parkes. 2005. Online auctions with re-usable goods. In *Proceedings of the 6th ACM conference on Electronic commerce*. ACM, 165–174.
- J. Hartline and T. Roughgarden. 2008. Optimal Mechanism Design and Money Burning. In *Proc. 39th ACM Symp. on Theory of Computing*.
- Frank Hutter, Holger H Hoos, Kevin Leyton-Brown, and Thomas Stützle. 2009. ParamILS: an automatic algorithm configuration framework. *Journal of Artificial Intelligence Research* 36, 1 (2009), 267–306.
- Adam Kalai and Santosh Vempala. 2005. Efficient algorithms for online decision problems. *J. Comput. System Sci.* 71, 3 (2005), 291–307.
- Ron Lavi and Noam Nisan. 2015. Online ascending auctions for gradually expiring items. *Journal of Economic Theory* 156 (2015), 45–76.
- Lin Xu, Frank Hutter, Holger H Hoos, and Kevin Leyton-Brown. 2008. SATzilla: portfolio-based algorithm selection for SAT. *Journal of artificial intelligence research* 32 (2008), 565–606.