arXiv:2110.13116v1 [cs.DS] 25 Oct 2021

Learning-Augmented Dynamic Power Management with Multiple States via New Ski Rental Bounds

Antonios Antoniadis

University of Twente

Enschede, The Netherlands

Christian Coester*

Tel Aviv University

Tel Aviv, Israel

Marek Eliáš[†]

Università Bocconi

Milan, Italy

Adam Polak[‡]

EPFL

Lausanne, Switzerland

Bertrand Simon

IN2P3 Computing Center, CNRS

Villeurbanne, France

Abstract

We study the online problem of minimizing power consumption in systems with multiple power-saving states. During idle periods of unknown lengths, an algorithm has to choose between power-saving states of different energy consumption and wake-up costs. We develop a learning-augmented online algorithm that makes decisions based on (potentially inaccurate) predicted lengths of the idle periods. The algorithm's performance is near-optimal when predictions are accurate and degrades gracefully with increasing prediction error, with a worst-case guarantee almost identical to the optimal classical online algorithm for the problem. A key ingredient in our approach is a new algorithm for the online ski rental problem in the learning augmented setting with tight dependence on the prediction error. We support our theoretical findings with experiments.

1   Introduction

Energy represents up to 70% of total operating costs of modern data centers [40] and is one of the major quality-of-service parameters in battery-operated devices. In order to ameliorate this, contemporary CPUs are equipped with sleep states to which the processor can transition during periods of inactivity. In particular, the ACPI-standard [24] specifies that each processor should possess, along with the active state $C0$ that is used for processing tasks, at least one sleep state $C1$. Modern processors generally possess more sleep states $C2, \ldots$; for example, current Intel CPUs implement at least 4 such C-states [18]. Apart from CPUs, such sleep states appear in many systems ranging from hard drives or mobile devices to the start-stop feature found in many cars, and are furthermore often employed when rightsizing data centers [2].

Intuitively, in a "deeper" sleep state, the set of switched-off components will be a superset of the corresponding set in a more shallow sleep state. This implies that the *running cost* for residing in that deeper state will be lower, but the *wake-up* cost to return to the active state $C0$ will be higher compared to a more shallow sleep state. In other words, there is a tradeoff between the running and the wake-up cost. During each idle period, a *dynamic power management (DPM)* strategy

---

*Christian Coester is supported by the Israel Academy of Sciences and Humanities & Council for Higher Education Excellence Fellowship Program for International Postdoctoral Researchers. Research was carried out while he was at CWI in Amsterdam and supported by NWO VICI grant 639.023.812.

†Research was carried out while Marek Eliáš was at CWI in Amsterdam and was supported by NWO GROOT project number OCENW.GROOT.2019.015 (OPTIMAL)

‡Supported by SNSF project *Lattice Algorithms and Integer Programming* (185030)

1

has to decide in which state the system resides at each point in time, without a-priori knowledge about the duration of the idle period. Optimally managing these sleep states is a challenging problem due to its online nature. On the one hand, transitioning the system to a too deep state could be highly suboptimal if the idle period ends shortly after. On the other hand, spending too much idle time in a shallow state would accumulate high running costs. The impact of DPM strategies in practice has been studied for instance in data centers, where each machine may be put to a sleep mode if no request is expected. See the study of Lim et al. [33] on multi-tier data centers.

The special case of 2-state DPM systems, i.e., when there is only a single sleep state (besides the active state), is essentially equivalent to the *ski rental* problem, one of the most classical problems and of central importance in the area of online optimization [38; 25]. This problem is defined as follows: A person goes skiing for an unknown number of days. On every day of skiing, the person must decide whether to continue renting skis for one more day or to buy skis. Once skis are bought there will be no more cost on the following days, but the cost of buying is much higher than the cost of renting for a day. It is easy to see that this captures a single idle period of DPM with a single sleep state whose running cost is 0: The rental cost corresponds to the running cost of the active state and the cost of buying skis corresponds to the wake-up cost; transitioning to the sleep state corresponds to buying skis. Given this equivalence, the known 2-competitive deterministic algorithm and $e/(e-1) \approx 1.58$-competitive randomized algorithm for ski rental carry over to 2-state DPM, and these competitive ratios are tight. In fact, it was shown by Irani et al. [27] and Lotker et al. [35] that the same competitive ratios carry over even to multi-state DPM. Ski rental, also known as *rent-or-buy* problem, is a fundamental problem appearing in many domains not restricted to computer hardware questions. For the AI community, this problem for example implicitly appears in expert learning with switching costs: paying the price to switch to a better expert allows to save expenses in the future.

Beyond these results for the classical online setting, [27] also gave a deterministic $e/(e-1)$-competitive algorithm for the case in which the length of the idle periods is repeatedly drawn from a fixed, and known, probability distribution. When the probability distribution is fixed but unknown they developed an algorithm that *learns* the distribution over time and showed that it performs well in practice. Although it is perhaps not always reasonable to assume a fixed underlying probability distribution for the length of idle periods, real-life systems do often follow periodical patterns so that these lengths can indeed be frequently predicted with adequate accuracy, see Chung et al. [17] for a specific example. Nevertheless, it is not hard to see that blindly following such predictions can lead to arbitrarily bad performance when predictions are faulty. The field of *learning-augmented* algorithms [37] is concerned with algorithms that incorporate predictions in a robust way.

In this work, we introduce multi-state DPM to the learning-augmented setting. Extending ideas of [27] and [35], we give a reduction from multi-state DPM to ski rental that is applicable to the learning-augmented setting. Although ski rental has been investigated through the learning-augmented algorithms lens before [39; 43], earlier work has focused on the optimal trade-off between *consistency* (i.e., the performance when predictions are accurate) and *robustness* (i.e., the worst-case performance). To apply our reduction from DPM to ski rental, we require more refined guarantees for learning-augmented ski rental. To this end we develop a new learning-augmented algorithm for ski rental that obtains the optimal trade-off between consistency and dependence on the prediction error. Our resulting algorithm for DPM achieves a competitive ratio arbitrarily close to 1 in case of perfect predictions and its performance degrades gracefully to a competitive ratio arbitrarily close to the optimal robustness of $e/(e-1) \approx 1.58$ as the prediction error increases.

1.1 Formal definitions

Problem definition. In the problem of *dynamic power management (DPM)*, we are given $k+1$ power states denoted by $0, 1, \ldots, k$, with power consumptions $\alpha_0 > \cdots > \alpha_k \geq 0$ and wake-up costs $\beta_0 < \cdots < \beta_k$. For state 0, we have $\beta_0 = 0$ and we call this the *active state*. The input is a series

2

of *idle periods* of lengths $\ell_1, \ldots, \ell_T$ received online, i.e., the algorithm does not know the length of the current period before it ends. During each period, the algorithm can transition to states with lower and lower power consumption, paying energy cost $x\alpha_i$ for residing in state i for time x. If j is the state at the end of the idle period, then it has to pay the wake-up cost β_j to transition back to the active state 0. The goal is to minimize the total cost.

In the *learning-augmented* setting, the algorithm receives at the beginning of the ith idle period a prediction $\tau_i \geq 0$ for the value of ℓ_i as additional input. We define $\eta_i := \alpha_0|\tau_i - \ell_i|$ to be the error of the ith prediction, and $\eta := \sum_i^T \eta_i$ to be the total prediction error.

(Continuous-time) *ski rental* is the special case of DPM with $k = 1$, $\alpha_1 = 0$ and a single idle period of some length ℓ. In this case, we call $\alpha := \alpha_0$ the rental cost, $\beta := \beta_1$ the buying cost, and ℓ the length of the ski season. In learning-augmented ski rental, we write the single prediction as $\tau := \tau_1$.

(ρ, μ)-competitiveness. Classical online algorithms are typically analyzed in terms of *competitive ratio*. A (randomized) algorithm $\mathcal{A}$ for an online minimization problem is said to be ρ-*competitive* (or alternatively, obtain a *competitive ratio* of ρ) if for any input instance,

$$cost(\mathcal{A}) \leq \rho \cdot \text{OPT} + c, \tag{1}$$

where $cost(\mathcal{A})$ and OPT denote the (expected) cost of $\mathcal{A}$ and the optimal cost of the instance and c is a constant independent of the *online* part of the input (i.e., the lengths ℓ_i in case of DPM). For the ski rental problem one requires $c = 0$, since the trivial algorithm that buys at time 0 has constant cost β.

In the learning-augmented setting, for $\rho \geq 1$ and $\mu \geq 0$, we say that $\mathcal{A}$ is (ρ, μ)-*competitive* if

$$cost(\mathcal{A}) \leq \rho \cdot \text{OPT} + \mu \cdot \eta \tag{2}$$

for any instance, where η is the prediction error. This corresponds to a competitive ratio of $\rho + \mu \frac{\eta}{\text{OPT}}$ (with $c = 0$). While this could be unbounded as $\eta/OPT \rightarrow \infty$, our DPM algorithm achieves a favorable competitive ratio even in this case (see Theorem 5, where we take the minimum over a range of pairs (ρ, μ), including $\mu = 0$).

For a (ρ, μ)-competitive algorithm, ρ is also called the *consistency* (i.e., competitive ratio in case of perfect predictions) while μ describes the dependence on the prediction error.

1.2 Our results

Our first result is a (ρ, μ)-*competitive algorithm* for ski rental that achieves the optimal μ corresponding to the given ρ. For $\rho \in [1, \frac{e}{e-1}]$, let

$$\mu(\rho) := \max\left\{\frac{1 - \rho\frac{e-1}{e}}{\ln 2}, \rho(1 - T)e^{-T}\right\}, \tag{3}$$

where $T \in [0, 1]$ is the solution to $T^2 e^{-T} = 1 - \frac{1}{\rho}$. Let $\tilde{\rho} \approx 1.16$ be the value of ρ for which both terms in the maximum yield the same value. The first term dominates for $\rho > \tilde{\rho}$ and the second term if $\rho < \tilde{\rho}$. Note that $\mu(1) = 1$ and $\mu\left(\frac{e}{e-1}\right) = 0$. See Figure 1 (left) for an illustration.

Theorem 1. *For any $\rho \in [1, \frac{e}{e-1}]$, there is a $(\rho, \mu(\rho))$-competitive randomized algorithm for learning-augmented ski rental, i.e., given a prediction with error η, its expected cost is at most $\rho \, \text{OPT} + \mu(\rho) \cdot \eta$.*

Note that $\rho < 1$ is impossible for any algorithm (due to the case $\eta = 0$) and $\rho > \frac{e}{e-1}$ is uninteresting since $\rho = \frac{e}{e-1}$ already achieves the best possible value of $\mu = 0$.

We also prove a lower bound showing that $\mu(\rho)$ defined in (3) is the best possible.

Theorem 2. *For any $\rho \in [1, \frac{e}{e-1}]$ and any (randomized) algorithm $\mathcal{A}$, there is a ski rental instance with some prediction error η such that the expected cost of $\mathcal{A}$ is at least $\rho \, \text{OPT} + \mu(\rho)\eta$.*

3
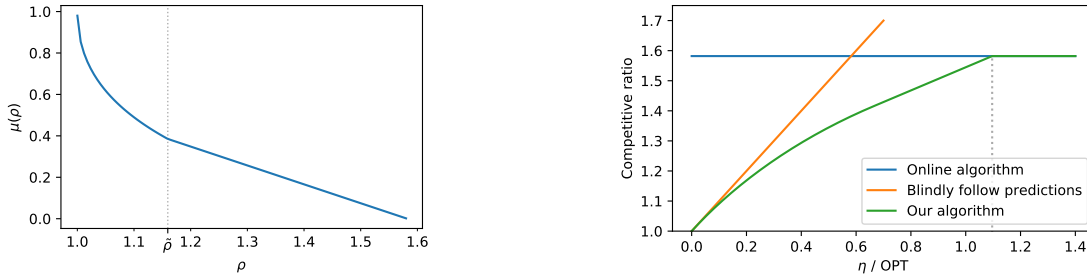
Figure 1: Illustration of $\mu(\rho)$ and of the resulting competitive ratio in function of $\eta/\textsc{Opt}$.

However, for most values of the prediction τ it is possible to achieve a better $\mu < \mu(\rho)$, and $\mu(\rho)$ only captures the worst case over all possible predictions τ. The proof of Theorem 1 is first sketched in Section 2 for clarity, after which we provide the complete analysis while describing how to achieve the best possible μ as a function of both ρ and τ. Theorem 2 is proved in Section 3.

In Section 4, we give a reduction from DPM to ski rental in the learning-augmented setting, provided that the ski rental algorithm satisfies a natural monotonicity property (defined formally in Section 4):

Lemma 3. *If there is a monotone (ρ, μ)-competitive ski rental algorithm, then there is a (ρ, μ)-competitive algorithm for DPM.*

Since our ski rental algorithm is monotone, this directly yields a $(\rho, \mu(\rho))$-competitive algorithm for DPM. From the special case $(\rho, \mu) = \left(\frac{e}{e-1}, 0\right)$, this theorem directly implies the following result for classical DPM (without predictions), which was first proved by Lotker et al. [35] for the equivalent multi-slope ski rental problem:

Corollary 4 ([35]). *There is a $\frac{e}{e-1}$-competitive randomized online algorithm for DPM (without predictions).*

Using techniques from online learning, in a way similar to [5], we show in Section 5 how to achieve "almost" $(\rho, \mu(\rho))$-competitiveness simultaneously for *all* ρ:

Theorem 5. *For any $\epsilon > 0$, there is a learning-augmented algorithm $\mathcal{A}$ for dynamic power management whose expected cost can be bounded as*

$$cost(\mathcal{A}) \leq (1+\epsilon) \min\left\{\rho \,\textsc{Opt} + \mu(\rho) \cdot \eta \,\big|\, \rho \in [1, \tfrac{e}{e-1}]\right\} + O\left(\tfrac{\beta_k}{\epsilon} \log \tfrac{1}{\epsilon}\right).$$

The above theorem gives a competitive ratio arbitrarily close to $\min\{\rho + \mu(\rho) \cdot \frac{\eta}{\textsc{Opt}}\}$, which is equal to 1 if $\eta = 0$ and never greater than $\frac{e}{e-1}$. In particular, we achieve a performance that degrades gracefully from near-optimal consistency to near-optimal robustness as the error increases.[1] See Figure 1 (right) for an illustration.

In Section 6, we illustrate the performance of these algorithms by simulations on synthetic datasets, where the dependence on the prediction error can be observed as expected from theoretical results.

1.3 Related work

Learning-Augmented Algorithms. Learning augmented algorithms have been a very active area of research since the seminal paper of Lykouris and Vassilvitskii [36]. We direct the interested reader to a survey [37] by Mitzenmacher and Vassilvitskii, as well as [7; 19; 36; 5; 41; 42; 34; 31] for recent results on secretary problems, paging, k-server as well as scheduling problems. In the following we survey some results in the area more closely related to our work.

---

[1] At first glance, our consistency and robustness might seem to contradict the lower bound of Wei and Zhang [43] for ski rental. However, [43] crucially uses $c = 0$ in the definition of competitiveness for ski rental.

4

The ski rental problem has already been studied within the context of learning augmented algorithms. Here, the main objective was to optimize the tradeoff between consistency and robustness (performance on perfect predictions and worst-case performance). The first results are due to Purohit et al. [39] who propose a deterministic and a randomized algorithm. A hyperparameter allows to choose a prescribed consistency ρ and leads to a corresponding robustness. They also present a linear dependency on the error: their randomized algorithm is (ρ, ρ)-competitive for $\rho \geq 1$, with larger ρ allowing for better robustness. Note that such a guarantee of (ρ, ρ)-competitiveness is not valuable in our model where we do not focus on robustness as blindly following the predictions leads to a $(1, 1)$-competitive algorithm. Wei and Zhang [43] show that the consistency / robustness tradeoff achieved by the randomized algorithm of [39] is Pareto-optimal. Angelopoulos et al. [4] propose a deterministic algorithm achieving a Pareto-optimal consistency / robustness tradeoff, but with no additional guarantee when the error is small. Interestingly, these algorithms not focusing on (ρ, μ)-competitiveness are naturally monotone, so easily extend to DPM by Lemma 3, contrarily to the tight algorithm we present in this paper. Nevertheless, experimental data (Section 6) seem to indicate that our algorithm optimizing (ρ, μ)-competitiveness for ski rental leads to better algorithms for DPM. A variant with multiple predictions was also studied in [20].

As we will see in Section 5, DPM can be cast as a problem from the class of *Metrical Task Systems (MTS)*. Antoniadis et al. [5] gave a learning-augmented algorithm for MTS that can be interpreted as $(1, 4)$-competitive within their prediction setup.

A different problem related to energy conservation is the classical online speed scaling problem, which was recently studied in the learning-augmented setting by Bamas et al. [9].

DPM. The equivalence between 2-state DPM and ski rental is mentioned in [38]. Therefore the well-known 2-competitive deterministic and an $e/(e-1)$-competitive randomized algorithm [29] for the classical ski rental problem carry over to 2-state DPM, and these bounds are known to be tight.

Irani et al. [27] present an extension of the 2-competitive algorithm for two-state DPM to multi-state DPM that also achieves a competitive ratio of 2. Furthermore they give an $e/(e-1)$-competitive algorithm for the case that the lengths of the idle periods come from a fixed probability distribution.

Lotker et al. [35] consider what they call *multi-slope ski rental* which is equivalent to the DPM problem. Among other results, they show how to reduce a $(k+1)$-slope ski rental instance to k classical ski rental instances. The reduction from DPM to ski rental presented in this paper is similar, but more general in order to also be applicable in the presence of predictions with the introduced (ρ, μ)-competitiveness. They furthermore show how to compute the best possible randomized strategy for any instance of the problem.

There have been several previous approaches that try to predict the length of an idle interval (see, e.g., [17; 27], and the survey of Benini et al. [11]). However, the proposed approaches to use these predictions are not robust against a potentially high prediction error.

Augustine et al. [8] investigate a problem generalizing DPM where transition cost is paid for going to a deeper sleep state rather than waking up and these transition costs may be non-additive (i.e., it can be cheaper to skip states). Albers [2] studies the offline version of the problem with multiple, parallel devices and shows that it can be solved in polynomial time.

Irani et al. [28] introduced a 2-state problem where jobs that need to be processed have a release-time, a deadline and a required processing time. This gives further flexibility to the system to schedule the jobs and create periods of inactivity so as to maximize the energy-savings by transitioning to the sleep state. For the offline version, there is an exact polynomial-time algorithm due to Baptiste et al. [10]. Recently, a 3-approximation algorithm for the multiprocessor-case was developed [6].

Another related problem consists of deciding which components of a data-center should be powered on or off in order to process the current load on the set of active components (see, e.g., [3]). A similar problem, where jobs have individual processing times for each machine, was studied in [30; 32]. Helmbold et al. [22] considered the problem of spinning down the disk of a

5

mobile computer when idle times are expected, which is another instance of DPM.

Several surveys cover DPM, see for example [11; 1; 26].

2 New algorithm for ski rental

For $\rho \in [1, e/(e-1)]$, we present a $(\rho, \mu(\rho))$-competitive algorithm for (learning augmented) ski rental, proving Theorem 1. The next lemma shows that it suffices to give such an algorithm for $\alpha = \beta = 1$.

Lemma 6. *An algorithm $\mathcal{A}'$ that is (ρ, μ)-competitive for instances of the ski rental problem with $\alpha = \beta = 1$ implies a (ρ, μ)-competitive algorithm $\mathcal{A}$ for arbitrary $\alpha, \beta > 0$.*

Proof. Given (ρ, μ)-competitive algorithm $\mathcal{A}'$ for unit buying and rental costs, we define an algorithm $\mathcal{A}$ for an instance with renting cost α and buying cost β as follows: Given a prediction τ, simulate $\mathcal{A}'$ with prediction $\frac{\alpha}{\beta}\tau$. If $\mathcal{A}'$ buys at time t', then $\mathcal{A}$ buys at time $t = \frac{\beta}{\alpha}t'$.

Let $c(\ell)$ and $c'(\ell')$ denote the total expected costs of $\mathcal{A}$ and $\mathcal{A}'$, respectively, with corresponding lengths of the ski season ℓ and $\ell' = \frac{\alpha}{\beta}\ell$. First, note that $c(\ell) = \beta c'(\ell')$: If $\mathcal{A}'$ buys before ℓ', incurring cost 1, then $\mathcal{A}$ also buys, incurring cost β. Similarly, if $\mathcal{A}'$ rents for time $t' \leq \ell'$, paying cost t', then $\mathcal{A}$ rents for time $t = \frac{\beta}{\alpha}t'$, paying cost $\alpha t = \beta t'$. Since $\mathcal{A}'$ is (ρ, μ)-competitive, we have

$$c(\ell) = \beta c'(\ell') \leq \beta \left(\rho \min\{\tfrac{\alpha}{\beta}\ell, 1\} + \mu |\tfrac{\alpha}{\beta}\tau - \tfrac{\alpha}{\beta}\ell| \right) = \rho \, \text{OPT} + \mu\eta,$$

where $\text{OPT} = \min\{\alpha\ell, \beta\}$ is the cost of the offline optimum and $\eta = \alpha|\tau - \ell|$ is the prediction error of the original instance. $\square$

A key difference between proving ρ-competitiveness in the classical online setting and (ρ, μ)-competitiveness in the learning-augmented setting is the following. In the online setting without predictions, the optimal competitive ratio of $\rho = e/(e-1)$ is achieved by the greedy algorithm that at all times buys with the probability that keeps (1) (for $c = 0$) tight assuming the skiing season ends immediately after the current time. This crucially relies on the fact that the upper bound $\rho \cdot \text{OPT}$ in (1) is monotone (non-decreasing) and concave as a function of the length of the ski season. Neither monotonicity (if $\tau > 1$) nor concavity (regardless of τ) are satisfied for the upper bound $\rho \cdot \text{OPT} + \mu \cdot \eta$ in (2), which substantially complicates the description and especially the analysis of our algorithm. In particular, we will use the value of the prediction τ to determine the times when we will aim for (2) to be tight.

2.1 Description of the algorithm

We next describe our randomized algorithm for instances with $\alpha = \beta = 1$, which can then be used to solve arbitrary ski rental instances using Lemma 6. Our algorithm is fully specified by the cumulative distribution function (CDF) F_τ of the time when the algorithm buys skis. The algorithm then draws a $p \in [0,1]$ uniformly at random and buys at the earliest time $t \in [0,\infty)$ such that $F_\tau(t) \geq p$. The CDF F_τ will depend on the given prediction $\tau \geq 0$ as well as the fixed ρ and μ, which can be chosen as $\mu = \mu(\rho)$, see Equation (3). Depending on the precise value of τ, it might also be possible to choose a smaller μ as we show later in Section 2.3.

Definition of the CDF (see Figure 2) We denote by P_0 the probability of buying at time 0 and, for any $t > 0$, we denote by p_t the probability density of buying at time t, so that the probability that the algorithm buys by time x can be expressed as

$$F_\tau(x) = P_0 + \int_0^x p_t \, dt.$$

For convenience, we also specify the probability $P_\infty = 1 - (P_0 + \int_0^\infty p_t \, dt)$ of never buying.

To define P_0 and p_t, we distinguish three cases depending on the value of the prediction τ. Note that we always have $0 \leq \mu \leq 1 \leq \rho \leq \frac{e}{e-1}$.

6
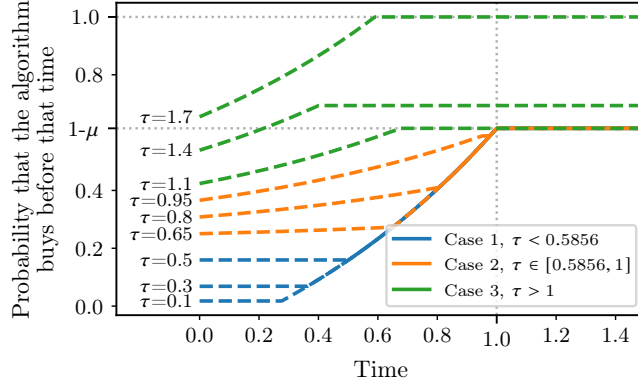
Figure 2: Our (ρ, μ)-competitive ski rental algorithm for $\rho = \tilde{\rho} \approx 1.1596$ and $\mu = \mu(\tilde{\rho}) \approx 0.3852$. The figure presents the cumulative distribution functions of the time of buying for several prediction values τ. Here $\alpha = \beta = 1$, i.e., at time $t = 1$ buying and renting has equal costs.

Case 1: $\mu\tau < \mu - \rho + 1$. We choose

$$P_0 = \frac{\tau(\rho - 1)}{1 - \tau}, \qquad p_t = \begin{cases} \rho e^{t-1} & \text{for } t \in (b, 1] \\ 0 & \text{otherwise} \end{cases}, \qquad P_\infty = \min\{\mu, 1 - P_0\},$$

where $b \in [\tau, 1]$ is chosen such that $P_0 + P_\infty + \int_b^1 \rho e^{t-1} dt = 1$, in order to have the sum of probabilities equal to 1. Note that if $P_0 \geq 1 - \mu$, we have $b = 1$ and $p_t = 0$ for all $t > 0$.

Case 2: $\mu - \rho + 1 \leq \mu\tau$ and $\tau \leq 1$. We choose

$$P_0 = \mu\tau, \qquad p_t = \begin{cases} (\mu\tau + \rho - \mu - 1)e^t & \text{for } t \leq a \\ \rho e^{t-1} & \text{for } t \in (b, 1] \\ 0 & \text{otherwise} \end{cases}, \qquad P_\infty = \min\{\mu, 1 - P_0\},$$

where $a \in [0, \tau]$ is chosen maximal such that $P_0 + P_\infty + \int_0^a (\mu\tau + \rho - \mu - 1)e^t dt \leq 1$, and $b \in [\tau, 1]$ is chosen so that $P_0 + P_\infty + \int_0^a (\mu\tau + \rho - \mu - 1)e^t dt + \int_b^1 \rho e^{t-1} dt = 1$ in order to have the sum of probabilities equal to 1. In case $\rho = \frac{e}{e-1}$, we have $\mu = 0$ and $(\mu\tau + \rho - \mu - 1)e^t = (\rho - 1)e^t = \rho e^{t-1}$, recovering the classical online algorithm of Karlin et al. [29].

Case 3: $\tau > 1$. If $\mu\tau \geq 1$, we buy at time 0. Otherwise, we choose

$$P_0 = \mu\tau, \qquad p_t = \begin{cases} (\mu\tau + \rho - \mu - 1)e^t & \text{if } t \leq T \\ 0 & \text{if } t > T \end{cases}, \qquad P_\infty = \rho - \mu - (\mu\tau + \rho - \mu - 1)e^T,$$

where T is the number closest to $\tau - 1$ that satisfies

$$e^T \leq \frac{\rho - \mu}{\mu\tau + \rho - \mu - 1} \qquad\qquad \text{(equivalently } P_\infty \geq 0) \qquad\qquad (4)$$

$$e^T \geq \frac{\rho - 2\mu}{\mu\tau + \rho - \mu - 1} \qquad\qquad \text{(equivalently } P_\infty \leq \mu). \qquad\qquad (5)$$

Thus, either $T = \tau - 1$ if this choice satisfies both bounds, or T is at an endpoint of the feasible interval prescribed by (4) and (5).

7

2.2 Sketch of the analysis

We expose the main ideas of the analysis in this section. Full analysis and the proof of Theorem 1 can be found in Section 2.3.

Our algorithm is (ρ, μ)-competitive if and only if for all $x \geq 0$ we have

$$cost(x) := P_0 + \int_0^x (1+t)\, p_t\, dt + \int_x^\infty x\, p_t\, dt + x P_\infty \leq \rho \min\{x, 1\} + \mu|\tau - x|, \qquad (6)$$

where $cost(x)$ denotes the expected cost of the algorithm in the case when $\ell = x$: If we intend to buy at some time t and $t < x$, we pay $1 + t$, otherwise we pay x. On the right hand side, $\min\{x, 1\}$ is the optimal cost and $|\tau - x|$ is the prediction error, assuming $\alpha = \beta = 1$.

We first sketch the analysis for Case 2, and then discuss the differences in Case 1. These cases are relatively simple. Case 3 is far more involved and we will only sketch the ideas.

Case 2: For the algorithm to be well defined, we need to choose μ such that a suitable $b \in [\tau, 1]$ exists. For $\mu = \mu(\rho)$, this is ensured by the inequality $\mu(\rho) \geq \frac{1 - \rho \frac{e-1}{e}}{\ln 2}$ from the definition of $\mu(\rho)$, which implies existence of such b for any value of τ. If $\tau = \ln 2$, then $\mu = \mu(\rho)$ is in fact the smallest possible, allowing only $b = \tau$. For other values of τ, suitable b exists also for smaller values of μ. We now show that (6) is satisfied.

Note that (6) is tight for $x = 0$, with both sides equal to $\mu\tau$. To obtain (6) for all $x > 0$, it suffices to show that the derivative of the left-hand side with respect to x is at most the derivative of the right-hand side (where derivatives exist). For $x \in (0, \infty) \setminus \{a, b, 1\}$, we have

$$\frac{d}{dx} cost(x) = (1+x)p_x + \int_x^\infty p_t\, dt - x p_x + P_\infty = p_x + \int_x^\infty p_t\, dt + P_\infty.$$

For $x \in (0, a)$ this yields

$$\frac{d}{dx} cost(x) = p_x + 1 - P_0 - (p_x - p_0) = 1 - \mu\tau + (\mu\tau + \rho - \mu - 1)e^0 = \rho - \mu,$$

which is equal to the derivative of the right-hand side of (6). For $x \in (a, b)$, $\frac{d}{dx} cost(x)$ is even smaller because p_x is 0, and the derivative of the right-hand side of (6) is $\rho - \mu$ or $\rho + \mu$. For $x \in (b, 1)$,

$$\frac{d}{dx} cost(x) = p_x + \int_x^\infty p_t\, dt + P_\infty = p_x + (p_1 - p_x) + P_\infty = \rho + P_\infty \leq \rho + \mu,$$

which is equal to the derivative of the right-hand side of (6). Finally, for $x > 1$ we have $\frac{d}{dx} cost(x) = P_\infty \leq \mu$ and the derivative of the right-hand side is also μ.

Case 1: The reason we cannot define p_t in the same way as in Case 2 is that p_t would be negative for $t \leq a$ (i.e., the algorithm would try to sell skis that it bought at time 0, which is not allowed). We therefore choose P_0 such that (6) is tight for $x = \tau$ if we do not buy in the interval $(0, \tau]$. The remainder of the proof of (6) is similar to Case 2. For $\mu = \mu(\rho)$, the existence of $b \in [\tau, 1]$ follows from the inequality $\mu \geq \rho(1 - T)e^{-T}$ in the definition of $\mu(\rho)$. Note that such μ is the smallest possible for $\tau = 1 - T$.

Case 3: The first step in the analysis of Case 3 is to derive an inequality involving ρ, μ, τ and T that is equivalent to the algorithm being (ρ, μ)-competitive. Denoting by $\mu_\tau(\rho)$ the minimal μ satisfying this inequality, it suffices to show that $\mu_\tau(\rho) \leq \mu(\rho)$ for all $\tau > 1$. However, the difficulty is that no closed-form expression for $\mu_\tau(\rho)$ exists. However, we are still able to show that $\tau \mapsto \mu_\tau(\rho)$ can have a local maximum only if $T = \tau - 1$, and therefore $\sup_{\tau > 1} \mu_\tau(\rho)$ is achieved either for $\tau \to 1$ or when $T = \tau - 1$. This allows us to eliminate τ from the aforementioned inequality, and we can then show that $\mu = \mu(\rho)$ satisfies the remaining inequality (with tightness occurring for $\rho \leq \tilde{\rho}$ and $\tau = T + 1$).

8
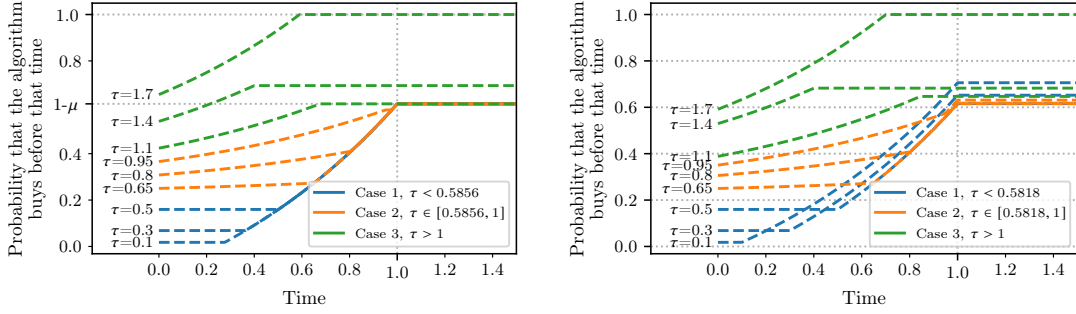
Figure 3: Our (ρ, μ)-competitive ski rental algorithm for $\rho = \tilde{\rho} \approx 1.1596$ and a) $\mu = \mu(\tilde{\rho}) \approx 0.3852$ on the left and b) $\mu = \mu_\tau(\tilde{\rho})$ on the right. The figures present the cumulative distribution functions of the time of buying for several prediction values τ.

2.3 Generalization and complete analysis

We will show a stronger result than that of Theorem 1 by considering a generalization of our algorithm to additional values of μ. For fixed $\rho \in \left[1, \frac{e}{e-1}\right]$ and a prediction $\tau \geq 0$, let $\mu_\tau(\rho)$ denote the best value of μ such that a (ρ, μ)-competitive algorithm exists for instances with this particular prediction τ. We will allow any $\mu \in [\mu_\tau(\rho), 1]$ and show that the resulting algorithm is (ρ, μ)-competitive algorithm for instances with prediction τ. As this will include the value $\mu = \mu(\rho)$, Theorem 1 will follow.

The benefit of this more general algorithm is that it achieves two goals at once: On the one hand, choosing $\mu = \mu_\tau(\rho)$ yields the best possible guarantee for any given prediction τ. On the other hand, choosing $\mu = \mu(\rho)$ independently of τ, we achieve an algorithm that satisfies the monotonicity property required by our reduction from DPM to ski rental (Section 4). Although a more specialized algorithm description only for the case $\mu = \mu_\tau(\rho)$ would suffice to prove Theorem 1 (since $\mu_\tau(\rho) \leq \mu(\rho)$), we remark that this algorithm would fail to be monotone, see Figure 3.

The definition of the more general algorithm is verbatim as in Section 2.1, except we no longer fix $\mu = \mu(\rho)$ but instead consider arbitrary $\mu \in [\mu_\tau(\rho), 1]$ for $\mu_\tau(\rho)$ as defined in the next subsection.

2.3.1 Definition of $\mu_\tau(\rho)$

For $\tau \leq 1$, define

$$\mu_\tau(\rho) := \begin{cases} -\frac{\rho-1}{1-\tau} + \rho e^{\tau-1} & \text{if } (1-\tau)e^{\tau-1} > 2 - \frac{2}{\rho} \\ \frac{e^\tau\left(1-\rho\frac{e-1}{e}\right)}{2-(1-\tau)e^\tau} & \text{if } (1-\tau)e^{\tau-1} \leq 2 - \frac{2}{\rho}. \end{cases} \tag{7}$$

Note that the denominators in both expressions are strictly positive, since $\tau < 1$ in the first case and $e^\tau(1-\tau) \leq 1$ for all $\tau \in \mathbb{R}$.

For $\tau > 1$, we define $\mu_\tau(\rho) = \min\{\mu \mid (\mu, T) \in M_\tau(\rho) \text{ for some } T\}$, where $M_\tau(\rho)$ is the set of all pairs $(\mu, T) \in \mathbb{R}^2$ satisfying the following constraints:

$$\rho\tau + (T-\tau)(\mu\tau + \rho - \mu - 1)e^T \leq \rho \tag{8}$$

$$\rho - 2\mu \leq (\mu\tau + \rho - \mu - 1)e^T \leq \rho - \mu \tag{9}$$

Note that $(1/\tau, 0) \in M_\tau(\rho)$, and therefore $\mu_\tau(\rho) \leq 1/\tau$.

Lemma 7. *For each $\rho \in \left[1, \frac{e}{e-1}\right]$ and $\tau \geq 0$, we have $\mu_\tau(\rho) \leq \mu(\rho)$.*

Proof sketch. The two cases for $\tau \leq 1$ are simple: If $(1-\tau)e^{\tau-1} > 2 - \frac{2}{\rho}$, the upper bound can be found by substituting $T = 1-\tau$ and computing the derivative with respect to T. If $(1-\tau)e^{\tau-1} \leq 2-\frac{2}{\rho}$,

9

the upper bound can be found by taking derivatives with respect to τ. This case is tight if $\tau = \ln 2$. The proof for the case $\tau > 1$ is more involved and given in Section 2.3.6. □

For $\rho \in \left[1, \frac{e}{e-1}\right]$, $\tau \geq 0$ and $\mu \in [\mu_\tau(\rho), 1]$, we denote by $\mathcal{A}_{\rho,\mu,\tau}$ the algorithm described in Section 2.1 for ski rental instances with prediction τ, except we no longer require $\mu = \mu(\rho)$. Our goal is to prove the following theorem, which combined with Lemma 7 yields Theorem 1.

Theorem 8. *For any $\rho \in \left[1, \frac{e}{e-1}\right]$, $\tau \geq 0$ and $\mu \in [\mu_\tau(\rho), 1]$, algorithm $\mathcal{A}_{\rho,\mu,\tau}$ is well-defined and (ρ, μ)-competitive for ski rental instances with prediction τ.*

The following lemma shows that for $\tau \leq 1$, we can essentially replace the two conditions in the definition of $\mu_\tau(\rho)$ by the conditions of Case 1 and Case 2 from the definition of our algorithm.

Lemma 9. *If $\tau \leq 1$ and $\mu \geq \mu_\tau(\rho)$, then*

$$\mu \geq \begin{cases} -\frac{\rho-1}{1-\tau} + \rho e^{\tau-1} & \text{if } \mu\tau < \mu - \rho + 1 \text{ (Case 1)} \\ \frac{e^\tau\left(1-\rho\frac{e-1}{e}\right)}{2-e^\tau(1-\tau)} & \text{if } \mu - \rho + 1 \leq \mu\tau \text{ (Case 2)}. \end{cases} \quad (10)$$

Proof. *Case 1* ($\mu\tau < \mu - \rho + 1$): Suppose that (10) fails. Then

$$\rho - 1 < \mu(1-\tau) \leq -\rho + 1 + \rho(1-\tau)e^{\tau-1},$$

where the first inequality uses the condition of Case 1 and the second inequality uses $\tau \leq 1$ and the assumption that (10) fails. Rearranging, we get

$$2 - \frac{2}{\rho} < (1-\tau)e^{\tau-1},$$

but this is precisely the first condition in (7). Thus, the premise $\mu \geq \mu_\tau(\rho)$ implies inequality (10).

Case 2 ($\mu - \rho + 1 \leq \mu\tau$): If $(1-\tau)e^{\tau-1} \leq 2 - \frac{2}{\rho}$, then (10) follows immediately from $\mu \geq \mu_\tau(\rho)$. So suppose $(1-\tau)e^{\tau-1} > 2 - \frac{2}{\rho}$. Our goal is to derive a contradiction. We have

$$\rho - 1 \geq \mu(1-\tau) \geq \mu_\tau(\rho)(1-\tau) = -\rho + 1 + \rho(1-\tau)e^{\tau-1},$$

where the first inequality uses the condition of Case 2, the second inequality uses $\tau \leq 1$ and $\mu \geq \mu_\tau(\rho)$, and the equation uses our assumption $(1-\tau)e^{\tau-1} > 2 - \frac{2}{\rho}$ and the definition of $\mu_\tau(\rho)$. Rearranging, we get

$$2 - \frac{2}{\rho} \geq (1-\tau)e^{\tau-1},$$

yielding the desired contradiction. □

2.3.2 Condition for competitiveness

For given τ, the algorithm achieves the desired (ρ, μ)-competitiveness if and only if for all $x \geq 0$ we have

$$cost(x) := P_0 + \int_0^x (1+t)\, p_t\, dt + \int_x^\infty x\, p_t\, dt + xP_\infty \leq \rho \min\{x, 1\} + \mu|\tau - x|, \quad (11)$$

where $cost(x)$ denotes the expected cost of the algorithm in the case when $\ell = x$: If we intend to buy at some time t and $t < x$, we pay $1 + t$, otherwise we pay x. On the right hand side, $\min\{x, 1\}$ is the optimal cost and $|\tau - x|$ is the prediction error, assuming $\alpha = \beta = 1$. We denote by $rhs(x) := \rho \min\{x, 1\} + \mu|\tau - x|$ the right hand side of (11).

10

The following identity will be useful. It holds for almost all $x > 0$ (excluding only those x where p_x is discontinuous):

$$\frac{d}{dx} cost(x) = (1 + x)p_x + \int_x^\infty p_t dt + P_\infty - xp_x$$

$$= p_x + \int_x^\infty p_t dt + P_\infty \tag{12}$$

$$= p_x + 1 - P_0 - \int_0^x p_t dt, \tag{13}$$

where the last equation uses the fact that probabilities add up to 1.

2.3.3   Case 1 ($\mu\tau < \mu - \rho + 1$)

Note that $\tau < 1$ in this case.

Well-definedness.   We need to argue that $P_0 \leq 1$ and that $b \in [\tau, 1]$ as stated in the algorithm description exists. By the condition of Case 1, we have $\rho - 1 \leq \mu(1 - \tau)$ and therefore $P_0 = \frac{\tau(\rho - 1)}{1 - \tau} \leq \mu\tau \leq 1$, where the last inequality uses $\mu \leq 1$ and $\tau < 1$.

Regarding existence of b, if $P_\infty = 1 - P_0$ then clearly $b = 1$. Otherwise, we choose

$$b = 1 + \ln \frac{\mu + \frac{\rho - 1}{1 - \tau}}{\rho}.$$

Then

$$P_0 + P_\infty + \int_b^1 \rho e^{t-1} dt = \frac{\tau(\rho - 1)}{1 - \tau} + \mu + \rho - \rho e^{b-1}$$

$$= \frac{\tau(\rho - 1)}{1 - \tau} + \rho - \frac{\rho - 1}{1 - \tau}$$

$$= 1,$$

as required by the definition of b. Note that $b \leq 1$ since otherwise the integral would be negative, and hence the left hand side would be less than 1 since $P_\infty \leq 1 - P_0$. Moreover, by Lemma 9 we have $b \geq \tau$. So indeed $b \in [\tau, 1]$, i.e., the algorithm is well-defined in Case 1.

Competitiveness.   We need to show that (11) is satisfied. First note that

$$cost(\tau) = P_0 + \int_\tau^\infty \tau p_t dt + \tau P_\infty = P_0 + \tau(1 - P_0) = (1 - \tau)\frac{\tau(\rho - 1)}{1 - \tau} + \tau = \rho\tau,$$

making (11) tight for $x = \tau$. To conclude (11) for all x, it suffices to show the following equivalence for all $x \in (0, \tau) \cup (\tau, b) \cup (b, 1) \cup (1, \infty)$:

$$\frac{d}{dx} cost(x) \leq \frac{d}{dx} rhs(x) \iff x \geq \tau. \tag{14}$$

For $x < b$, from (13) we get

$$\frac{d}{dx} cost(x) = 1 - P_0 = 1 - \frac{\tau(\rho - 1)}{1 - \tau} = \frac{1 - \tau\rho}{1 - \tau}.$$

If $x < \tau$, using the condition of Case 1 we conclude $\frac{d}{dx} cost(x) > \rho - \mu = \frac{d}{dx} rhs(x)$. For $x \in (\tau, b)$, (14) follows from

$$\frac{d}{dx} cost(x) = \frac{1 - \tau\rho}{1 - \tau} \leq \rho + \mu = \frac{d}{dx} rhs(x).$$

11

For $x \in (b, 1)$, (12) shows

$$\frac{d}{dx} cost(x) = p_x + \int_x^\infty p_t dt + P_\infty = p_x + (p_1 - p_x) + P_\infty = \rho + P_\infty \le \rho + \mu \le \frac{d}{dx} rhs(x).$$

Finally, for $x > 1$ we have $\frac{d}{dx} cost(x) = P_\infty \le \mu = \frac{d}{dx} rhs(x)$.

We note the following observation, which will be useful later when we prove the lower bound.

Observation 10. *In Case 1, if $b = \tau$ and $P_\infty = \mu$, then (11) is tight for all $x \ge \tau$.*

2.3.4 Case 2 ($\mu - \rho + 1 \le \mu\tau$ and $\tau \le 1$)

Well-definedness. The only non-trivial part is to argue about the existence of b. If $a < \tau$ or $P_\infty = 1 - P_0$, then $b = 1$. Otherwise, $a = \tau$ and $P_\infty = \mu$ and we choose

$$b = 1 + \ln \frac{\mu(2 - (1-\tau)e^\tau) + e^\tau(\rho - 1)}{\rho}.$$

Then

$$
\begin{aligned}
P_0 + P_\infty &+ \int_0^a (\mu\tau + \rho - \mu - 1)e^t dt + \int_b^1 \rho e^{t-1} dt \\
&= \mu\tau + \mu + (\mu\tau + \rho - \mu - 1)(e^\tau - 1) + \rho - \rho e^{b-1} \\
&= 2\mu + 1 + (\mu\tau + \rho - \mu - 1)e^\tau - \mu(2 - (1-\tau)e^\tau) - e^\tau(\rho - 1) \\
&= 1
\end{aligned}
$$

as required by the definition of b. We have $b \le 1$ since otherwise the left hand side would be less than 1 (by definition of a). Moreover, by Lemma 9,

$$b \ge 1 + \ln \frac{e^\tau \left(1 - \rho \frac{e-1}{e}\right) + e^\tau(\rho - 1)}{\rho} = \tau.$$

So $b \in [\tau, 1]$, as required.

Competitiveness. Note that (11) is tight for $x = 0$, with both sides equal to $\mu\tau$. To obtain (11) for all $x > 0$, it suffices to show that $\frac{d}{dx} cost(x) \le \frac{d}{dx} rhs(x)$ whenever both derivatives exist.

For $x \in (0, a)$, from (13) we get

$$\frac{d}{dx} cost(x) = p_x + 1 - P_0 - \int_0^x p_t dt = 1 - \mu\tau + (\mu\tau + \rho - \mu - 1)e^0 = \rho - \mu = \frac{d}{dx} rhs(x).$$

For $x \in (a, b)$, $\frac{d}{dx} cost(x)$ is even smaller because p_x is 0, and $\frac{d}{dx} rhs(x) \in \{\rho - \mu, \rho + \mu\}$. For $x \in (b, 1) \cup (1, \infty)$, the proof of $\frac{d}{dx} cost(x) \le \frac{d}{dx} rhs(x)$ is identical to Case 1.

Observation 11. *In Case 2, if $a = b = \tau$ and $P_\infty = \mu$, then (11) is tight for all $x \ge 0$.*

2.3.5 Case 3 ($\tau > 1$)

The case $\mu\tau \ge 1$ is trivial, so we assume throughout this section that $\mu < 1/\tau$.

Well-definedness. The existence of T follows from $\mu < 1/\tau < 1$. We only need to check that probabilities add up to 1:

$$P_0 + P_\infty + \int_0^\infty p_t d_t = \mu\tau + \rho - \mu - (\mu\tau + \rho - \mu - 1)e^T + (\mu\tau + \rho - \mu - 1)(e^T - 1) = 1,$$

where the first equation uses positivity of T:

Fact 12. $T > 0$.

Proof. We either have $T \ge \tau - 1 > 0$ or the upper bound (4) is tight. In the latter case, $T = \ln \frac{\rho - \mu}{\mu\tau + \rho - \mu - 1} > \ln 1 = 0$ since $\mu\tau < 1$. $\qquad\square$

12

Competitiveness. For $\tau > 1$ and $\mu \in [\mu_\tau(\rho), 1/\tau)$, denote by $T(\tau, \mu)$ the corresponding value of T chosen by the algorithm. We may drop τ in the notation when it is clear from the context.

Lemma 13. $T(\tau, \mu)$ *is a non-increasing function of* μ.

Proof. It suffices to show that $e^{T(\mu)}$ is non-increasing in μ. By definition, $e^{T(\mu)}$ is the projection of $e^{\tau-1}$ onto the interval

$$\left[\frac{\rho - 2\mu}{\mu\tau + \rho - \mu - 1}, \frac{\rho - \mu}{\mu\tau + \rho - \mu - 1} \right].$$

Since both endpoints of the interval are decreasing functions of μ (since $\tau > 1$), it follows that $T(\mu)$ is non-increasing in μ. $\quad\square$

Lemma 14. $T(\tau, \mu) < \tau$ *for every* $\mu \in [\mu_\tau(\rho), 1/\tau)$. *Moreover, if* $\mu = \mu_\tau(\rho)$*, then* $(\mu, T(\tau, \mu)) \in M_\tau(\rho)$.

Proof. If the second statement holds, then since $\tau > 1$, constraint (8) implies that $T(\mu_\tau(\rho)) < \tau$. By Lemma 13, this implies the first statement. It remains to show the second statement.

Let $\mu = \mu_\tau(\rho)$ and observe that there exists T such that $(\mu, T) \in M_\tau(\rho)$.

By taking derivatives, one can see that the left-hand side of (8) is concave in T for $T < \tau - 2$ and convex in T for $T > \tau - 2$, with a local minimum at $T = \tau - 1$. Since constraint (8) is violated for $T \to -\infty$, but there exists T such that $(\mu, T) \in M_\tau(\rho)$, we conclude that $T = \tau - 1$ is a global minimum of the left-hand side of (8) and $(\mu, T(\mu)) \in M_\tau(\rho)$ by definition of $T(\mu)$. $\quad\square$

Lemma 15. *For every fixed* $\tau > 1$ *and* $\mu \in [\mu_\tau(\rho), 1/\tau)$*, the following statements are equivalent:*

1. *The algorithm is* (ρ, μ)*-competitive for instances with prediction* τ.

2. $cost(\tau) \leq \rho$.

3. *Constraint* (8) *is satisfied for* $T = T(\tau, \mu)$.

In the positive case, $T \leq 1$.

Proof. The implication from the first to the second statement is immediate.

Using (13) and (12), we observe that the cost is a piecewise linear function of x:

$$\frac{d}{dx} cost(x) = \begin{cases} 1 - \mu\tau + (\mu\tau + \rho - \mu - 1)\left[e^x - \int_0^x e^t\, dt \right] & = \rho - \mu & \text{if } x < T \\ P_\infty & \text{if } x > T \end{cases} \tag{15}$$

Since $T < \tau$ by Lemma 14, this yields

$$cost(\tau) = cost(0) + (\rho - \mu)T + (\tau - T)P_\infty$$
$$= \mu\tau + (\rho - \mu)T + (\tau - T)\left[\rho - \mu - (\mu\tau + \rho - \mu - 1)e^T \right]$$
$$= \rho\tau + (\tau - T)(\mu\tau + \rho - \mu - 1)e^T$$

Thus, $cost(\tau)$ is precisely the left-hand side of (8), which shows the equivalence of the second and third statement.

It remains to show that the second and third statement imply $T \leq 1$ and that (11) holds for *all* x.

We must have $T \leq 1$ since otherwise we get a contradiction to the second statement via

$$cost(\tau) > cost(1) = \mu\tau + \rho - \mu \geq \rho.$$

For $x = 0$, (11) is satisfied with equality. For $x \in (0, T)$, by (15) and since $T \leq 1 < \tau$ we have $\frac{d}{dx} cost(x) = \frac{d}{dx} rhs(x)$, so (11) is tightly satisfied for all $x \in [0, T]$. By assumption we also satisfy it for $x = \tau$. Since $cost(x)$ is linear in x for $x > T$ and $rhs(x)$ is concave for $x \in [T, \tau]$ (linearly increasing for $x \in [T, 1]$ and linearly decreasing for $x \in [1, \tau]$), this shows that (11) also holds for $x \in [T, \tau]$. For $x > \tau$, we have $\frac{d}{dx} cost(x) = P_\infty \leq \mu = \frac{d}{dx} rhs(x)$ (where the inequality uses (5)), so (11) holds there as well. See Figure 4 for an illustration of these functions. $\quad\square$

13
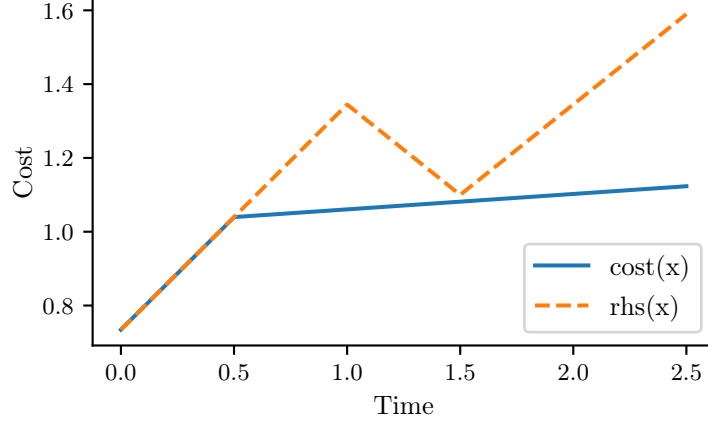
Figure 4: Evolution of *cost* and *rhs* in function of the time, for $\rho = 1.1$, $\mu = \mu(\rho)$, $\tau = 1.5$, so $T = 0.5$.

The following observation is immediate from the proof of Lemma 15.

Observation 16. *Let $\tau > 1$ and $\mu \in [\mu_\tau(\rho), 1/\tau)$. Then (11) is tight for all $x \in [0, T]$. If $P_\infty = \mu$ and (11) is tight for $x = \tau$, then it is tight for all $x \geq \tau$.*

For $\tau > 1$ and $\mu \in [\mu_\tau(\rho), 1/\tau)$, we may write $cost(\tau, \mu)$ instead of $cost(\tau)$ to emphasize the additional dependence on μ. By (the proof of) Lemma 15, $cost(\tau, \mu)$ is equal to the left-hand side of (8), i.e.,

$$cost(\tau, \mu) = \rho\tau + (T(\tau, \mu) - \tau)(\mu\tau + \rho - \mu - 1)e^{T(\tau, \mu)} \tag{16}$$

It remains to argue that the equivalent statements of Lemma 15 are indeed true for all $\mu \in [\mu_\tau(\rho), 1/\tau)$. For $\mu = \mu_\tau(\rho)$, this follows from $(\mu, T(\mu) \in M_\tau(\rho)$ as shown in Lemma 14. The following lemma implies that this remains true for larger μ.

Lemma 17. *For every $\tau > 1$, the function $\mu \mapsto cost(\tau, \mu)$ is decreasing.*

Proof. Consider first the effect of increasing μ within some interval I such that one of the bounds (4) or (5) is tight for all $\mu \in I$. This means that the total probability of buying before time τ is the same for all $\mu \in I$, and thus the buying cost contribution to $cost(\tau, \mu)$ is the same for all $\mu \in I$. However, increasing μ means that the algorithm buys more aggressively (i.e., earlier), which leads to a decrease in the expected rent cost. Thus, increasing μ in I leads to a decrease of $cost(\tau, \mu)$.

Now, consider the effect of increasing μ within an interval I where neither (4) nor (5) is tight. Then $T(\tau, \mu) = \tau - 1$ is constant for all $\mu \in I$. In this case, due to (16) and since $\tau > 1$, $cost(\tau, \mu)$ is again decreasing in μ. $\qquad\square$

2.3.6 Bounding $\mu_\tau(\rho)$ in Case 3

The goal of this section is to prove the following Lemma:

Lemma 18. *For any $\rho \in \left[1, \frac{e}{e-1}\right]$ and $\tau > 1$,*

$$\mu_\tau(\rho) \leq \max\left\{\frac{\rho + e - \rho e}{2}, \rho(1 - T)e^{-T}\right\},$$

where $T \in [0, 1]$ is the solution to $T^2 e^{-T} = 1 - \frac{1}{\rho}$.

14
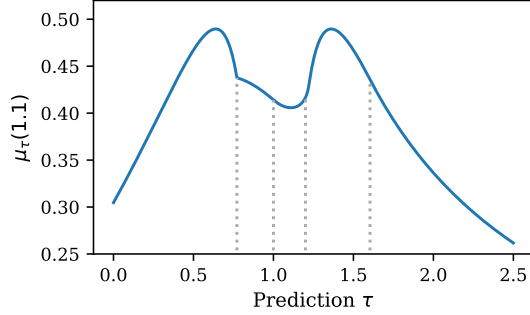
Figure 5: $\mu_\tau(\rho)$ as a function of τ for $\rho = 1.1$. The vertical lines separate the five cases of interest. From left to right, we distinguish Case 1, Case 2, and three subcases of Case 3 ($\tau > 1$) where: $P_\infty = \mu$; $0 < P_\infty < \mu$; $P_\infty = 0$.

Since $2 > e\ln(2)$, this will complete the proof of Lemma 7.

The difficulty in this case, compared to upper bounding $\mu_\tau(\rho)$ in Cases 1 and 2, arises from the fact that there exists no closed-form expression for $\mu_\tau(\rho)$.

The lemma is trivial for $\rho = 1$ since $\mu_\tau(1) \leq 1/\tau < 1$ and due to the solution $T = 0$. Therefore, assume throughout the remainder of this section that $\rho > 1$. Since we view ρ as fixed, we may write μ_τ instead of $\mu_\tau(\rho)$.

By the next Lemma, the interval $[\mu_\tau, 1/\tau)$ is indeed non-empty and constraint (8) is tight for $\mu = \mu_T$ and $T = T(\tau, \mu)$.

Lemma 19. *For all $\tau > 1$, we have $\mu_\tau < 1/\tau$ and $cost(\tau, \mu_\tau) = \rho$.*

Proof. The inequality $\mu_\tau < 1/\tau$ can be seen by noting that for $\mu < 1/\tau$, we have $(\mu, 0) \in M_\tau(\rho)$ provided that μ is close enough to $1/\tau$. Recall that we assumed $\rho > 1$.

For the other statement, the direction "$\leq$" is immediate from the fact that $cost(\tau, \mu)$ is the left-hand side of (8). For the other direction, consider first the case $\mu_\tau = 0$. Then $(\mu\tau + \rho - \mu - 1)e^{T(\tau,\mu)} = \rho$ by definition of $T(\tau, \rho)$, so $cost(\tau, \mu) = \rho T(\tau, \rho) = \rho \ln \frac{\rho}{\rho - 1} \geq \rho$, where the last inequality uses $\rho \leq \frac{e}{e-1}$.

Consider now $\mu_\tau > 0$. If constraint (8) were strict for $T = T(\tau, \mu_\tau)$, then by continuity of the constraints (8) and (9) one could find a pair $(\mu, T) \in M_\tau(\rho)$ with $\mu < \mu_\tau$, contradicting the minimality of μ_τ. $\qquad\square$

As we will see in the next lemma, $\tau \mapsto \mu_\tau$ can have a maximum for $\tau > 1$ only if $T(\tau, \mu_\tau) = \tau - 1$. In fact, there can be at most one local maximum for $\tau > 1$ (cf. Figure 5).

Lemma 20. *Suppose $\tau \mapsto \mu_\tau$ has a maximum at $\tau^* > 1$, and write $\mu^* := \mu_{\tau^*} = \max_\tau \mu_\tau$. Then $T(\tau^*, \mu^*) = \tau^* - 1$.*

Proof. Suppose towards a contradiction that $T(\tau^*, \mu^*) \neq \tau^* - 1$. We will show that there exists $\tau \neq \tau^*$ such that $cost(\tau, \mu^*) > cost(\tau^*, \mu^*)$. Combined with the fact that $cost(\tau^*, \mu^*) = \rho = cost(\tau, \mu_\tau)$ (Lemma 19), we get $cost(\tau, \mu^*) > cost(\tau, \mu_\tau)$. But then Lemma 17 implies that $\mu^* < \mu_\tau$, contradicting the choice of μ^*.

Since $T(\tau^*, \mu^*) \neq \tau^* - 1$ this means that one of the bounds (4) or (5) would be violated by the choice $T = \tau^* - 1$. Since the bound is continuous in τ, there is some open interval I around τ^* such that for all $\tau \in I$, the same bound (4) or (5) would still be violated by the choice $T = \tau - 1$. So for all $\tau \in I$, the according bound is tight, i.e.,

$$T(\tau, \mu^*) = \ln \frac{\rho - i\mu^*}{\mu^*\tau + \rho - \mu^* - 1} \qquad \text{for all } \tau \in I,$$

15

with $i = 1$ or $i = 2$ depending on which of the two bounds would have been violated. Using (16) we get

$$
\begin{aligned}
cost(\tau, \mu^*) &= \rho\tau + (T(\tau, \mu^*) - \tau)(\rho - i\mu^*) \\
&= i\mu^*\tau + T(\tau, \mu^*) \cdot (\rho - i\mu^*) \qquad \text{for all } \tau \in I.
\end{aligned}
$$

We next compute the derivatives with respect to τ:

$$
\frac{d}{d\tau} T(\tau, \mu^*) = \frac{-\mu^*}{\mu^*\tau + \rho - \mu^* - 1} \qquad \text{for all } \tau \in I
$$

and therefore

$$
\begin{aligned}
\frac{d}{d\tau} cost(\tau, \mu^*) &= i\mu^* - \frac{\mu^*}{\mu^*\tau + \rho - \mu^* - 1} \cdot (\rho - i\mu^*) \\
&= \frac{i\mu^*}{\mu^*\tau + \rho - \mu^* - 1} \left(\mu^*\tau + \frac{i-1}{i}\rho - 1 \right) \qquad \text{for all } \tau \in I.
\end{aligned}
$$

Thus, inside I, we see that $cost(\tau, \mu^*)$ is decreasing in τ for $\tau < \left(1 - \frac{i-1}{i}\rho\right)/\mu^*$ and increasing in τ for $\tau > \left(1 - \frac{i-1}{i}\rho\right)/\mu^*$. In particular, $cost(\tau, \mu^*)$ is maximized when τ approaches one of the two boundaries of I. So we can find $\tau \in I$ with $cost(\tau, \mu^*) > cost(\tau^*, \mu^*)$, as desired. $\qquad\square$

To obtain an upper bound on μ_τ, the previous lemma now allows us to eliminate τ from our equations by either replacing it by $T - 1$ or letting it tend towards 1.

Proof of Lemma 18. For $\mu = \mu_\tau$ and $T = T(\tau, \mu_\tau)$, we get from Lemma 19 and equation (16) that

$$
\rho = \rho\tau - (\tau - T)(\mu\tau + \rho - \mu - 1)e^T.
$$

Rearranging to isolate μ, we get

$$
\begin{aligned}
\mu &= \frac{\rho(\tau - 1) - (\tau - T)(\rho - 1)e^T}{(\tau - T)(\tau - 1)e^T} \\
&= \frac{\rho}{(\tau - T)e^T} - \frac{\rho - 1}{\tau - 1}.
\end{aligned} \tag{17}
$$

Since $\mu \geq 0$, it must be the case that $T(\tau, \mu_\tau) \to 1$ as $\tau \to 1$, so the lower bound (5) is tight for τ close to 1. Writing (5) with equality and plugging in the limit values $\tau = T = 1$, this shows that

$$
\frac{\rho - 2\mu_\tau}{\rho - 1} \to e \qquad \text{as } \tau \to 1,
$$

so

$$
\mu_\tau \to \frac{\rho + e - \rho e}{2} \qquad \text{as } \tau \to 1.
$$

If μ_τ is *not* maximized when $\tau \to 1$, then let $\tau > 1$ be such that $\mu_\tau = \max_{\tau > 1} \mu_\tau$. By Lemma 20, we then have $\tau = T + 1$. Plugging this into (17), we get

$$
\mu_\tau = \frac{\rho}{e^T} - \frac{\rho - 1}{T}.
$$

Taking derivatives shows that the right hand side is maximized when $T \in (0, 1]$ is the solution to $T^2 e^{-T} = \frac{\rho - 1}{\rho}$. Then the $\frac{\rho - 1}{T}$ term is equal to $\rho T e^{-T}$, so that overall $\mu \leq \rho(1 - T)e^{-T}$. $\qquad\square$

16

3 Proof of Theorem 2: Lower bound for ski rental

We will show for any fixed τ that one cannot be better than $(\rho, \mu_\tau(\rho))$-competitive (with μ_τ as defined in Section 2.3.1). This will imply Theorem 2 due to the following lemma:

Lemma 21. *For any $\rho \in \left[1, \frac{e}{e-1}\right]$ there exists $\tau \geq 0$ such that $\mu_\tau(\rho) = \mu(\rho)$.*

Proof. Note that the two expressions in the definition (3) of $\mu(\rho)$ yield the same value for $\tilde{\rho} \approx 1.16$.

Whenever $\rho \geq \tilde{\rho}$, the first term in the definition (3) dominates. In this case, we choose $\tau = \ln 2$. Such τ belongs to the second case of the definition (7) of $\mu_\tau(\rho)$ whenever $(1-\tau)e^{\tau-1} \leq 2 - \frac{2}{\rho}$. This is equivalent to $\rho \geq \frac{e}{e-1+\ln 2} \approx 1.127$, which is in fact smaller than $\tilde{\rho}$. Therefore, for any $\rho \geq \tilde{\rho} > 1.127$, we have

$$\mu_\tau(\rho) = \frac{e^\tau(1 - \rho\frac{e-1}{e})}{2 - (1-\tau)e^\tau} = \frac{1 - \rho\frac{e-1}{e}}{\ln 2} = \mu(\rho).$$

Now, consider $\rho < \tilde{\rho}$. Here, the second term in the definition (3) of $\mu(\rho)$ dominates. For such ρ, there is a (unique) $T \in \left[0, \frac{1}{2}\right]$ such that $T^2 e^{-T} = 1 - \frac{1}{\rho}$. We choose $\tau = 1 - T$. This belongs to the first case of the definition (7) of $\mu_\tau(\rho)$, since $(1-\tau)e^{\tau-1} > 2 - \frac{2}{\rho}$ is equivalent to $1 - \frac{1}{\rho} > T(2 - \frac{2}{\rho})$, which is satisfied since $T < \frac{1}{2}$. Thus,

$$\begin{aligned}
\mu_\tau(\rho) &= -\frac{\rho-1}{1-\tau} + \rho e^{\tau-1} \\
&= \rho\left(-\frac{T^2 e^{-T}}{T} + e^{-T}\right) \\
&= \rho(1-T)e^{-T} \\
&= \mu(\rho),
\end{aligned}$$

where the second equation uses the definition of T and the choice of τ. $\qquad\square$

The goal in the remainder of this section is to prove the following theorem, which combined with Lemma 21 yields Theorem 2.

Theorem 22. *For all $\rho \in \left[1, \frac{e}{e-1}\right]$, $\tau \geq 0$ and $\widehat{\mu} < \mu_\tau(\rho)$, there exists no $(\rho, \widehat{\mu})$-competitive algorithm for ski rental instances with prediction τ.*

Let $\rho \in \left(1, \frac{e}{e-1}\right]$ and a prediction $\tau \geq 0$ be fixed for the remainder of this section and write $\mu_\tau = \mu_\tau(\rho)$. Note that by continuity, it suffices to show the theorem for $\rho > 1$.

The idea to prove Theorem 22 is to show that any algorithm must have at least the same expected cost as our algorithm on some instance for which the analysis of our algorithm is tight. Denote by $cost(x)$ the expected cost of our algorithm for $\mu = \mu_\tau$ when the skiing season ends at time x, and denote by $\mathrm{OPT}(x) = \min\{x, 1\}$ the optimal cost and by $\eta(x) = |\tau - x|$ the prediction error. The following lemma summarizes the cases for which the analysis of our algorithm is tight.

Lemma 23. *For $\mu = \mu_\tau$, our algorithm satisfies $cost(x) = \rho\,\mathrm{OPT}(x) + \mu_t\eta(x)$ in the following cases:*

- *In Case 1 for all $x \geq \tau$.*

- *In Case 2 for all $x \geq 0$.*

- *In Case 3 for all $x \in [0, T] \cup \{\tau\}$, and if $P_\infty = \mu$, then also for all $x > \tau$.*

Proof. **Case 1:** By Observation 10, it suffices to show that $b = \tau$ and $P_\infty = \mu$. By definition of $\mu = \mu_\tau$, we have

$$\begin{aligned}
P_0 + \mu + \int_\tau^1 \rho e^{t-1} dt &= \frac{\tau(\rho-1)}{1-\tau} - \frac{\rho-1}{1-\tau} + \rho e^{\tau-1} + \rho - \rho e^{\tau-1} \\
&= 1.
\end{aligned}$$

17

This implies $P_\infty = \mu$ since otherwise we would have $1 = P_0 + P_\infty < P_0 + \mu \le 1$, a contradiction. By definition of b, this also yields $b = \tau$.

Case 2: By Observation 11, it suffices to show that $a = b = \tau$ and $P_\infty = \mu$. We have

$$P_0 + \mu + \int_0^\tau (\mu\tau + \rho - \mu - 1)e^t dt + \int_\tau^1 \rho e^{t-1} dt = \mu\tau + \mu + (\mu\tau + \rho - \mu - 1)(e^\tau - 1) + \rho - \rho e^{\tau-1} dt$$

$$= \mu\left(2 + (\tau - 1)e^\tau\right) + 1 + e^\tau(\rho - 1 - \rho/e)$$

$$= e^\tau\left(1 - \rho\frac{e-1}{e}\right) + 1 + e^\tau(\rho - 1 - \rho/e)$$

$$= 1,$$

implying $P_\infty = \mu$ by the same argument as before, and therefore also $a = b = \tau$ by definition of a and b.

Case 3: Follows from Observation 16 and Lemma 19. $\qquad\square$

Let $\mathcal{A}$ denote our algorithm (for prediction τ and $\mu = \mu_\tau$) and F its CDF for the time of buying. Note that F is continuous.

Suppose that a $(\rho, \widehat{\mu})$-competitive algorithm $\widehat{\mathcal{A}}$ for instances with prediction τ exists for some $\widehat{\mu} < \mu_\tau(\rho)$, and denote by $\widehat{F}$ its CDF. We may assume without loss of generality that $\widehat{F}$ is continuous (otherwise, $\widehat{F}$ can be approximated arbitrarily well by a continuous CDF whose corresponding algorithm is $(\rho, \widehat{\mu}')$-competitive for $\widehat{\mu}'$ arbitrarily close to $\widehat{\mu}$). Denote by $\widehat{cost}(x)$ the expected cost of $\widehat{\mathcal{A}}$ when the skiing season has length x. To prove Theorem 22, it suffices to show that there exists x for which $cost(x)$ is tight (as per Lemma 23) and $\widehat{cost}(x) > cost(x)$.

Let $y = \inf\{t \ge 0 : \widehat{F}(t) \ne F(t)\}$. Note that $y < \infty$ must exist, since otherwise $\widehat{\mathcal{A}}$ is the same as $\mathcal{A}$, which cannot achieve $\widehat{\mu} < \mu_\tau(\rho)$ by Lemma 23. By continuity of $\widehat{F}$ and F, we either have $y = 0$ or $\widehat{F}(t) = F(t)$ for all $t \le y$. Also by continuity, there exists $y' > y$ such that either $\widehat{F}(t) > F(t)$ for all $t \in (y, y']$ or $\widehat{F}(t) < F(t)$ for all $t \in (y, y']$. We may assume $y' < y + 1$.

Lemma 24. *If $\widehat{F}(t) > F(t)$ for all $t \in (y, y']$, then there exists $z > y$ such that $\widehat{cost}(z) > cost(z)$.*

Proof. If $y = 0$ and $\widehat{F}(0) > F(0)$, then $\widehat{cost}(0) > cost(0)$ and the desired $z > 0$ must exist by continuity of $cost$ and $\widehat{cost}$.

Otherwise, $\widehat{F}(t) = F(t)$ for all $t \le y$ and therefore the expected cost incurred until time y is the same for both algorithms. Let $z = \arg\max_{z \in (y,y']}\{\widehat{F}(z) - F(z)\}$. In the time interval $(y, z]$, the expected buying cost incurred by $\widehat{\mathcal{A}}$ is precisely $\widehat{F}(z) - F(z)$ greater than that of $\mathcal{A}$, and the rental cost saved by $\widehat{\mathcal{A}}$ compared to $\mathcal{A}$ is at most $(z - y)(\widehat{F}(z) - F(z))$. Since $z - y \le y' - y < 1$, this saving in rental cost cannot make up for the greater buying cost, and therefore $\widehat{cost}(z) > cost(z)$. $\qquad\square$

Lemma 25. *If $\widehat{F}(t) < F(t)$ for all $t \in (y, y']$, then there is $z > y$ such that $\widehat{cost}(z) > cost(z)$. In particular, we can choose $z = \min\{z > y \mid \widehat{F}(z) = F(z)\}$ if it exists. Otherwise, z can be chosen very large so that, e.g., $z > \tau$.*

Proof. If $\widehat{F}(t) < F(t)$ for all $t > y$, then for z large enough it holds that $\widehat{cost}(z) > cost(z)$, since $\widehat{\mathcal{A}}$'s greater rental cost will eventually outweigh its smaller buying cost. Clearly, if we choose z sufficiently large, we also get the additional property that $z > \tau$.

Otherwise, by continuity we can choose $z = \min\{z > y \mid \widehat{F}(z) = F(z)\}$. Since $\widehat{F}(z) = F(z)$, the expected buying cost is the same for both algorithms. On the other hand, since $\widehat{F}(t) < F(t)$ for all $t \in (y, z)$, i.e., $\widehat{\mathcal{A}}$ is delaying buying, the expected rental cost of $\widehat{\mathcal{A}}$ is higher than that of $\mathcal{A}$. $\qquad\square$

To complete the proof of Theorem 22, we will show that in each of the three cases from the definition of our algorithm, there exists an x satisfying the condition of Lemma 23 and such that $\widehat{cost}(x) > cost(x)$ or, when $x \ne \tau$ (so that $\eta > 0$), $\widehat{cost}(x) \ge cost(x)$.

18

Case 1: By Lemma 23, $cost(x)$ is tight for all $x \geq \tau$.

Recall from the definition of our algorithm that $F(0) = F(\tau)$.

If $\widehat{F}(\tau) > F(0) = F(\tau)$, then $\widehat{\mathcal{A}}$'s expected buying cost up until time τ exceeds that of $\mathcal{A}$ by exactly $\widehat{F}(\tau) - F(0)$, and its reduction in renting cost is at most $\tau(\widehat{F}(\tau) - F(0))$. Since $\tau < 1$ in Case 1, the reduction cannot make up for the excess and we get $\widehat{cost}(\tau) > cost(\tau)$.

Otherwise, $\widehat{F}(\tau) \leq F(\tau)$. Lemmas 24 and 25 give us z such that $\widehat{cost}(z) > cost(z)$. If $z \geq \tau$, we are done. Otherwise, we have

$$\widehat{cost}(\tau) \geq \widehat{cost}(z) + (\tau - z)(1 - \widehat{F}(\tau)) > cost(z) + (\tau - z)(1 - F(\tau)) = cost(\tau),$$

since $\widehat{cost}(z) > cost(z)$ and $\widehat{F}(\tau) \leq F(\tau)$.

Case 2: By Lemma 23, $cost(x)$ is tight for all $x \geq 0$. Thanks to Lemmas 24 and 25, we are done.

Case 3: By Lemma 23, $cost(x)$ is tight for all $x \in [0, T] \cup \{\tau\}$, and if $P_\infty = \mu$, then also for all $x > \tau$.

If $\widehat{F}(0) \geq F(0)$, then we choose $x = 0 \neq \tau$, noting that $\widehat{cost}(0) = \widehat{F}(0) \geq F(0) = cost(0)$.

Otherwise $\widehat{F}(0) < F(0)$. Lemma 25 gives us z such that $\widehat{cost}(z) > cost(z)$. Recall that T is chosen within the interval defined by (4) and (5) as close as possible to $\tau - 1$. Thus, we observe that $T < \tau - 1$ only if (4) is tight and $P_\infty = 0$. Similarly, $T > \tau - 1$ only if (5) is tight and $P_\infty = \mu$. We consider two cases.

- $\widehat{F}(\tau) \geq F(\tau)$. By continuity of $\widehat{F}$ and Lemma 25, we have $z \leq \tau$ and $\widehat{F}(z) = F(z)$. If $z \in [0, T] \cup \{\tau\}$, we are done, so we may assume that $T < z < \tau$. By definition of T, either $P_\infty = 0$ or $\tau - T \leq 1$. If $P_\infty = 0$, then $\widehat{F}(t) = F(t) = 1$ for any $t \geq z$ and we have $\widehat{cost}(\tau) = \widehat{cost}(z) > cost(z) = cost(\tau)$. Otherwise, by considering the additional buying and rental cost after time z, we have

$$
\begin{aligned}
\widehat{cost}(\tau) &\geq \widehat{cost}(z) + \big(\widehat{F}(\tau) - \widehat{F}(z)\big) + (\tau - z)\big(1 - \widehat{F}(\tau)\big) \\
&> cost(z) + (\tau - z)\big(\widehat{F}(\tau) - \widehat{F}(z)\big) + (\tau - z)\big(1 - \widehat{F}(\tau)\big) \\
&= cost(z) + (\tau - z)(1 - F(\tau)) \\
&= cost(\tau),
\end{aligned}
$$

  where the strict inequality follows from $\widehat{cost}(z) > cost(z)$ and $\tau - z < \tau - T \leq 1$, and the equations use $\widehat{F}(z) = F(z) = F(\tau)$, recalling that our algorithm does not buy after time T.

- $\widehat{F}(\tau) < F(\tau)$. If there exists $t \in [0, T]$ such that $\widehat{F}(t) = F(t)$, then $z \in [0, T]$ by Lemma 25 and we are done. Otherwise, using $\widehat{F}(\tau) < F(\tau)$, monotonicity and continuity of the CDFs, and the fact that $F(t)$ is constant for $t \geq T$, we get that $\widehat{F}(t) < F(t)$ for all $t \in [0, \tau]$. Then $z > \tau$ by Lemma 25. If $P_\infty = \mu$, then $cost(z)$ is tight and therefore it is enough to consider the case $P_\infty < \mu$, where $\tau - T \geq 1$ by definition of T.

  In this case, we choose $x = \tau$. To compare $\widehat{cost}(\tau)$ and $cost(\tau)$, we will separately consider the contribution coming from buying cost, rental cost in the interval $[0, T]$ and rental cost in the interval $[T, \tau]$. Since $\widehat{F}(t) \leq \widehat{F}(\tau) < F(\tau) = F(t)$ for all $t \in [T, \tau]$, the rental cost of $\widehat{\mathcal{A}}$ in the time interval $[T, \tau]$ exceeds that of $\mathcal{A}$ by at least $(\tau - T)\big(F(\tau) - \widehat{F}(\tau)\big) \geq F(\tau) - \widehat{F}(\tau)$, where the inequality is due to $\tau - T \geq 1$. The benefit in buying cost of $\widehat{\mathcal{A}}$ compared to $\mathcal{A}$, which amounts to $F(\tau) - \widehat{F}(\tau)$, can at best make up for this. Since $\widehat{F}(t) < F(t)$ for all $t \in [0, \tau]$, the rental cost of $\widehat{\mathcal{A}}$ in the remaining time interval $[0, T)$ is strictly greater than that of $\mathcal{A}$, so that overall $\widehat{cost}(\tau) > cost(\tau)$. Here, we used that $T > 0$ by Fact 12.[2]

---

[2]This implicitly assumes that $\rho > 1$, since in Section 2.3.5 we assumed $\mu < 1/\tau$, which in Section 2.3.6 was proved to be true for $\mu = \mu_\tau$ assuming $\rho > 1$. However, by continuity Theorem 22 also extends to the case $\rho = 1$. Alternatively, we can get strict inequality even for $T = 0$ by noting that then $\tau - T = \tau > 1$.

19

4 Reduction from DPM to ski rental

We now give a reduction from DPM to ski rental in the learning-augmented setting (Lemma 3), provided that the ski rental algorithm satisfies the following monotonicity property: We say that a ski rental algorithm for rental cost $\alpha = 1$ and buying cost $\beta = 1$ is *monotone* if its CDF F_τ for the buying time when given prediction τ satisfies

$$F_\tau(t) \leq F_{\tau'}(t) \qquad \text{for all } t \geq 0 \text{ and } \tau < \tau'.$$

Intuitively, this property is very natural: The longer the predicted duration of skiing, the greater should be our probability of buying. Indeed, our algorithm satisfies this property:

Lemma 26. *For $\mu = \mu(\rho)$, the (μ, ρ)-competitive ski rental algorithm from Section 2 is monotone, i.e., its CDF F_τ when given prediction τ satisfies $F_\tau(t) \leq F_{\tau'}(t)$ for all $t \geq 0$ and $\tau < \tau'$.*

Proof. One can verify that at the boundary between adjacent cases of the algorithm, both cases define the same probability distribution (i.e., $\tau = 1 - \frac{\rho-1}{\mu}$ as the boundary between Case 1 and Case 2, and $\tau = 1$ as the boundary between Case 2 and Case 3). In particular, for $\tau = 1$ this can be seen by separately considering $\mu \geq 1/2$ (where Case 3 would give $T = 0$) and $\mu < 1/2$ (which makes (5) tight, so that Case 3 would give $P_\infty = \mu$). It remains to show that $F_\tau(t)$ is non-decreasing as τ increases *within* one of the three cases.

Cases 1 and 2 ($\tau \leq 1$). Here, P_0 is increasing in τ and so is $(\tau\mu + \rho - \mu - 1)e^t$. Therefore, $F_\tau(t) = P_0 + \int_0^t p_t dt$ is increasing in τ as long as $t \leq b$. If $F_\tau(\tau) \geq 1 - \mu$, then $p_t = 0$ for each $t > \tau$ and we are done. Otherwise, note that $F_\tau(t) = 1 - \mu$ for any $t \geq 1$. And, going backwards in time, $F_\tau(t)$ has the same value for any τ until it reaches $F_\tau(b)$ at $t = b$.

Case 3 ($\tau > 1$). Both P_0 and $(\tau\mu + \rho - \mu - 1)e^t$ are increasing in τ, therefore $F_\tau(t)$ is increasing in τ until it reaches $1 - P_\infty$ at $t = T$. Therefore, it is enough to show that P_∞ is non-increasing in τ. Either one of (4) and (5) is tight, or we have $T = \tau - 1$. In the first case, P_∞ is equal to 0 or μ, i.e., independent on τ. Otherwise, $T = \tau - 1$ implies that e^T is increasing in τ, and therefore P_∞ is decreasing in τ. $\qquad\square$

As mentioned earlier, for many τ one could actually achieve a better $\mu_\tau(\rho) < \mu(\rho)$. However, somewhat surprisingly the optimal such algorithm would *not* be monotone. The monotonicity of our algorithm therefore crucially relies on our specific description (in particular the choice of a and b), which only aims for $(\rho, \mu(\rho))$-competitiveness with $\mu(\rho) = \sup_\tau \mu_\tau(\rho)$.

Combining Theorem 1, Lemma 3 and Lemma 26, we get:

Corollary 27. *For every $\rho \in [1, \frac{e}{e-1}]$, there is a $(\rho, \mu(\rho))$-competitive algorithm for DPM.*

To prove Lemma 3, it suffices to describe a (ρ, μ)-competitive algorithm for the special case of DPM with a single idle period: Running such an algorithm for each individual period yields a (ρ, μ)-competitive algorithm for DPM with any number of idle periods, since we can simply sum inequality (2) over all periods to obtain the corresponding inequality for the entire instance.

Consider now a single idle period of length ℓ for DPM. We first recall some observations of Irani et al. [27] about the optimal *offline* algorithm: It is easy to see that the optimal offline algorithm would transition to some state j only once at the beginning of the period and remain there throughout the period, paying cost $\alpha_j \ell + \beta_j$. Thus, state j is preferred over state $j - 1$ if and only if $\alpha_{j-1}\ell + \beta_{j-1} > \alpha_j\ell + \beta_j$, or equivalently $\ell > t_j := \frac{\beta_j - \beta_{j-1}}{\alpha_{j-1} - \alpha_j}$. We may assume without loss of generality that $t_1 < \cdots < t_k$: Indeed, suppose $t_{j+1} \leq t_j$, then state j is redundant because whenever j is preferred over $j - 1$, then $j + 1$ is preferred over j. Defining $t_0 := 0$ and $t_{k+1} := +\infty$, we get a partition $[0, +\infty) = \bigcup_{j=0}^k I_j$, where $I_j = [t_j, t_{j+1})$. We can then express the cost of the offline optimum as

$$\text{OPT} = \alpha_{j^*}\ell + \beta_{j^*}, \text{ with } j^* \text{ such that } \ell \in I_{j^*}. \qquad (18)$$

In the online setting, we of course do not know ℓ. The idea of our algorithm (similar to [35]) is to simulate k ski rental algorithms $\mathcal{A}_1, \ldots, \mathcal{A}_k$ in parallel, where the task of $\mathcal{A}_j$ is to decide

20

whether it is time to transition from the state $j-1$ to j. For this, we choose $\mathcal{A}_j$ to be an algorithm for ski rental with rental cost $\alpha_{j-1} - \alpha_j$ and buying cost $\beta_j - \beta_{j-1}$. Let F_τ be the CDF of the buying time of a *monotone* ski rental algorithm (for $\alpha = \beta = 1$) when given prediction τ. Recalling our reduction from arbitrary α and β to the case $\alpha = \beta = 1$ in Lemma 6, the CDF of $\mathcal{A}_j$ is given by

$$F^j(t) := F_{\tau/t_j}\left(t/t_j\right). \tag{19}$$

An outline of our algorithm is given in Algorithm 1.

Algorithm 1: DPM with a single idle period

---
for $j=1,\ldots,k$ **do**
  Let F^j be as defined by (19), induced by a monotone (ρ, μ)-competitive ski rental
  algorithm;
Choose $p \in [0, 1]$ uniformly at random;
At any time t: choose state $j = \max\{j : F^j(t) \ge p\}$;

Proof of Lemma 3. We now show that Algorithm 1 is (ρ, μ)-competitive for DPM instances with a single idle period, completing the proof of Lemma 3.

Denote by $A_j := \inf\{t \mid F^j(t) \ge p\}$ the random variable for the time when $\mathcal{A}_j$ triggers the transition to state j. As shown in Lemma 26, the monotonicity of the ski rental algorithm yields $F^{j-1}(t) \ge F^j(t)$ for all t. Therefore, $A_1 \le \cdots \le A_k$.

For a prediction τ and an idle period of length $x = \ell$, let

$$cost_j = (\beta_j - \beta_{j-1}) \cdot \mathbb{1}_{A_j \le x} + (\alpha_{j-1} - \alpha_j) \cdot \min\{A_j, x\}$$

be the random variable for the cost of $\mathcal{A}_j$. Since $\mathcal{A}_j$ is (ρ, μ)-competitive, its expected cost is at most

$$E(cost_j) \le \rho \cdot \min\{(\alpha_{j-1} - \alpha_j)x, \beta_j - \beta_{j-1}\} + \mu \cdot (\alpha_{j-1} - \alpha_j) \cdot |\tau - x|.$$

Let $\bar{j}$ be the maximal j such that $A_j \le x$. Writing $A_0 := 0$ and $A_{k+1} := \infty$, the cost of Algorithm 1 is

$$cost = \beta_{\bar{j}} + \sum_{j=1}^{\bar{j}} (A_j - A_{j-1})\alpha_{j-1} + (x - A_{\bar{j}})\alpha_{\bar{j}}$$

$$= \sum_{j=1}^{k} (\beta_j - \beta_{j-1}) \cdot \mathbb{1}_{A_j \le x} + \sum_{j=1}^{k+1} (\min\{A_j, x\} - \min\{A_{j-1}, x\})\alpha_{j-1}$$

$$= \sum_{j=1}^{k} (\beta_j - \beta_{j-1}) \cdot \mathbb{1}_{A_j \le x} + \sum_{j=1}^{k} (\alpha_{j-1} - \alpha_j) \min\{A_j, x\} + \alpha_k x,$$

which is precisely $\alpha_k x + \sum_{j=1}^{k} cost_j$. Taking expectations, we get

$$E(cost) \le \alpha_k x + \rho \sum_{j=1}^{k} \min\left\{(\alpha_{j-1} - \alpha_j)x, \beta_j - \beta_{j-1}\right\} + \mu \cdot |\tau - x| \sum_{j=1}^{k} (\alpha_{j-1} - \alpha_j)$$

$$= \alpha_k x + \rho\left((\alpha_{j^*} - \alpha_k)x + \beta_{j^*} - \beta_0\right) + \mu \cdot |\tau - x| \cdot (\alpha_0 - \alpha_k),$$

where j^* denotes the optimal state such that $\text{OPT} = \alpha_{j^*}x + \beta_{j^*}$. The equality holds since

$$\beta_j - \beta_{j-1} \le x \cdot (\alpha_{j-1} - \alpha_j) \iff t_j \le x \iff j \le j^*.$$

The theorem now follows since $\beta_0 = 0$, $\alpha_k \ge 0$ and $1 \le \rho$. □

21

4.1 Conversion to a prudent algorithm

It was shown by Lotker et al. [35] that any DPM algorithm can be converted (online) into one that assigns a non-zero probability to at most two adjacent power states, and the resulting algorithm can only have smaller expected cost than the original algorithm. More precisely, for a given time, denote by p_i the probability that the algorithm is in power state i. We call the probability vector $(p_0, \ldots, p_k)$ *prudent* if it is of the form $(0, \ldots, 0, p_{i-1}, p_i, 0, \ldots, 0)$ for some $i = 1, \ldots k$. A DPM algorithm is prudent if its probability vector is prudent at all times. Note that our DPM algorithm is *not* prudent: Indeed, our ski rental algorithm typically buys with non-zero probability already at time 0, and thus the DPM algorithm obtained via the above reduction assigns non-zero probability to *all* power states at time 0. Since conversion to a prudent algorithm using the method of [35] reduces an algorithm's cost, any implementation should apply this conversion, which we now describe.

To make the DPM algorithm prudent, we need to convert each non-prudent probability vector $p = (p_0, \ldots, p_k)$ into a prudent one $\tilde{p}$. Let $B_p = \sum_i p_i \beta_i$ be the expected wake-up cost of vector p, and let $m = \max\{i \mid \beta_i \le B_p\}$ be the deepest state whose wake-up cost is at most B_p. Note that if $m = k$, then $p = (0, \ldots, 0, 1)$ is already prudent. Otherwise, $m < k$ and define

$$\tilde{p}_i = \begin{cases} \frac{\beta_{m+1} - B_p}{\beta_{m+1} - \beta_m} & i = m \\ 1 - \tilde{p}_m & i = m+1 \\ 0 & otherwise. \end{cases}$$

It was shown in [35, Theorem 4.2] that when replacing any non-prudent vector p by a prudent vector $\tilde{p}$ in this way, the resulting algorithm pays the same expected wake-up cost but less running cost than the original non-prudent algorithm.

5 Finding the best trade-off online

Our goal is to design an algorithm whose performance almost matches that of Corollary 27 simultaneously for *all* ρ, proving Theorem 5. It will be useful to view DPM as a Metrical Task System.

Metrical Task Systems (MTS).   Metrical Task Systems (MTS), introduced by Borodin et al. [13], is a broad class of online problems containing many other problems as special cases. In MTS, we are given a metric space M of *states*. We start at a predefined initial state x_0. At each time $t = 1, 2, \ldots, T$, we are presented with a *cost function* $c_t \colon M \to \mathbb{R}_+$. Then, we have to choose our new state x_t and pay $\text{dist}(x_{t-1}, x_t) + c_t(x_t)$, where $\text{dist}(x_{t-1}, x_t)$ is the distance between x_{t-1} and x_t in M. The objective is to minimize the overall cost incurred over time.

To formulate DPM as a Metrical Task System, we choose states $0, 1, \ldots, k$ corresponding to the power states, with distances $\text{dist}(i, j) = \frac{1}{2}|\beta_i - \beta_j|$, so that the cost of switching from the state 0 to j and back is β_j. We choose 0 as the initial state. We discretize time in the DPM instance using time steps of some small length $\delta > 0$. At each time step belonging to some idle period, we issue a cost function c such that $c(j) = \delta \alpha_j$ for each $j = 0, \ldots, k$. At the end of each idle period, we issue a cost function where $c(0) = 0$ and $c(j) = +\infty$ for $j = 1, \ldots, k$, which forces any algorithm to move back to the active state.

We use the result of Blum and Burch [12] to combine multiple instances of our algorithm with different parameters ρ.

Theorem 28 (Blum and Burch [12])**.** *There is an algorithm which, given N online algorithms $A_1, \ldots A_N$ for an MTS with diameter D and $\epsilon_1 < 1/2$, achieves expected cost at most*

$$(1 + \epsilon_1) \cdot \min_i \{cost(A_i)\} + O(D/\epsilon_1) \ln N.$$

22

Proof of Theorem 5 For $\epsilon_2 > 0$, we choose a set $P \subset [1, \frac{e}{e-1}]$ of $O(1/\epsilon_2)$ values of ρ as follows:

$$P := \left\{1, \tilde{\rho}, \frac{e}{e-1}\right\} \cup \left\{\rho_i \,\middle|\, \mu(\rho_i) = (1 + i\epsilon_2)\mu(\tilde{\rho}), \text{ where } i = 1, \ldots, \left\lfloor\frac{1 - \mu(\tilde{\rho})}{\mu(\tilde{\rho})\epsilon_2}\right\rfloor\right\},$$

where $\tilde{\rho} \approx 1.16$ is as defined below equation (3). We claim that

$$\min_{\rho \in P}\{\rho \operatorname{OPT} + \mu(\rho)\eta\} \le (1 + \epsilon_2)\min_{\rho \in [1, e/(e-1)]}\{\rho \operatorname{OPT} + \mu(\rho)\eta\}.$$

If the minimizer ρ of the right-hand side lies in $\left[\tilde{\rho}, \frac{e}{e-1}\right]$, then it is one of the endpoints $\rho \in \left\{\tilde{\rho}, \frac{e}{e-1}\right\}$ since $\mu(\rho)$ is linear in this interval. If the minimizer is $\rho \in [1, \tilde{\rho})$, then there exists i such that $\rho_- \le \rho \le \rho_i$, where $\rho_- = \rho_{i+1}$ if it exists and $\rho_- = 1$ otherwise. Here, we have $\mu(\rho_-) \le (1+\epsilon_2)\mu(\rho_i) \le (1 + \epsilon_2)\mu(\rho)$, since μ is a decreasing function. Thus, $\rho_- \operatorname{OPT} + \mu(\rho_-)\eta \le \rho \operatorname{OPT} + (1 + \epsilon_2)\mu(\rho)\eta$.

Combining our $(\rho, \mu(\rho))$-competitive algorithms for DPM for $\rho \in P$ using Theorem 28, the resulting algorithm has cost at most

$$(1 + \epsilon_1)(1 + \epsilon_2)\min_{\rho \in [1, e/(e-1)]}\{\rho \operatorname{OPT} + \mu(\rho)\eta\} + O\left(\frac{\beta_k}{\epsilon_1} \cdot \ln\frac{1}{\epsilon_2}\right),$$

since $D = \beta_k/2$ in our case. This concludes the proof. □

5.1 Remarks on shifting/dynamic regret

Instead of the connection to MTS and the result of Blum and Burch [12], we could also use other online learning techniques. In particular, we could change the parameter ρ of a variant of our algorithm with bounded cost per iteration (see Section 5.1.1 below) in the beginning of each idle period based on any algorithm for the *experts problem* with *vanishing regret*, see [14] for reference. Under a mild assumption that the average length of the idle periods is a positive number independent of the number of idle periods, this would also lose only a factor $(1 + \epsilon)$ compared to the best ρ.

Of special interest are the results on *shifting/dynamic* regret, see work of Chen et al. [16], Cesa-Bianchi et al. [15], and Hall and Willett [21]. They allow us to reach a cost comparable not only to the algorithm with the best fixed ρ, but also to the best strategy of switching between multiple values of ρ a bounded number of times. For example, the additive regret of Fixed Share Forecaster [23] with respect to a strategy with m switches, is at most

$$O\left(\sqrt{mT \ln\frac{1}{\epsilon_2} \ln\frac{m}{T}}\right),$$

when T is the number of idle periods.

This kind of result can be very useful in scenarios where well-predictable parts of the input are interleaved with unpredictable or adversarial sequences which make the total prediction error high, forcing any static strategy to use large ρ. The shifting regret approach allows us to make use of the good predictions outside of the adversarial parts of the sequence.

5.1.1 Bounded cost per iteration

The cost of Algorithm 1 in one iteration may be arbitrary high (in case of an infinite prediction error), making it hard to use with some of the online learning techniques mentioned above. However, we can make sure that the cost of Algorithm 1 is bounded in each iteration at a small cost. In fact, it is enough to ensure bounded cost for the ski rental algorithm from Section 2 used in the construction of Algorithm 1. Note that if $P_\infty > 0$, then our algorithm might never buy skis, thus leading to unbounded cost. We can achieve bounded cost on the expense of worsening $\mu(\rho)$ by a factor of $(1 + \epsilon)$ in the following way. For $t < u = \frac{\beta}{\alpha}(3 + 1/\epsilon)$, we define $F_\tau(t)$ as for our ski rental algorithm from Section 2 (for general α and β), and $F_\tau(t) = 1$ for any $t \ge u$, making sure no more costs are incurred in times $t > u$.

23

Claim 29. *The algorithm with modified F_τ is $\big(\rho,\,(1+\epsilon)\mu(\rho)\big)$-competitive and its cost is never larger than $\beta(4+1/\epsilon)$.*

Proof. The bound on the cost is easy to see: In the worst case, we buy at time u, paying $\alpha\cdot\frac{\beta}{\alpha}(3+1/\epsilon)$ for renting until u and additional β for buying.

We will show that the algorithm is $\big(\rho,\,(1+\epsilon)\mu(\rho)\big)$-competitive under the assumption $\alpha=\beta=1$. The result for general α and β then follows from Lemma 6.

First case: $\tau \le 3$. Here, $cost(x) \le \rho\,\mathrm{OPT}+\mu(\rho)\cdot\eta$ for any $x \le u$ since the modified algorithm has the same behaviour as the one in Section 2. If $x \ge u$, we pay an additional cost of P_∞ for buying at time u:

$$
\begin{aligned}
cost(x) &\le \rho\,\mathrm{OPT}+\mu(\rho)\cdot|\tau - x| + P_\infty\\
&\le \rho\,\mathrm{OPT}+\mu(\rho)\cdot\big(1+\frac{1}{|\tau - x|}\big)|\tau - x|\\
&\le \rho\,\mathrm{OPT}+\mu(\rho)(1+\epsilon)\cdot|\tau - x|,
\end{aligned}
$$

since $P_\infty \le \mu(\rho)$, $\tau \le 3$ and $x \ge 3 + 1/\epsilon$.

Second case: $\tau > 3$. Here, we claim that $P_\infty = 0$ and therefore $F_{\rho,\tau}(1) = 1$, i.e., no modification is needed. Note that $P_\infty = 0$ whenever (4) is tight, i.e.,

$$
e^{\tau-1} \ge \frac{\rho - \mu}{\rho - \mu + \tau\mu - 1}. \tag{20}
$$

If we have $\mu > \mu(\tilde\rho) \approx 0.36$, then $\tau\mu - 1 > 0$ for any $\tau > 3$, making the right-hand side of (20) smaller than 1. Otherwise, $\mu(\rho) = \frac{1-\rho\frac{e-1}{e}}{\ln 2} > 1 - \rho\frac{e-1}{e}$, and we have

$$
\frac{\rho - \mu}{\rho - \mu + \tau\mu - 1} \le \frac{\rho}{\rho + (\tau-1)\mu - 1} < \frac{\rho}{\rho - \rho\frac{e-1}{e}} = e.
$$

making the right-hand side of (20) smaller than e, while $e^{\tau-1} > e$. □

6 Experiments

In this section, we present an experimental evaluation of our algorithms compared to existing learning-augmented ski rental algorithms for the ski rental and DPM problems[3]. We use a synthetic dataset, which was introduced by Purohit et al. [39], and a real-world dataset, which is based on smartphone traces from [44]. The results of our experiments[4] suggest that the performance of learning-augmented algorithms indeed degrades smoothly when the error increases, providing solutions which are better, for medium errors, than naive algorithms trusting the predictions and online (predictionless) algorithms. In the experiments, the performance of our algorithms degrades more smoothly when the prediction error increases than previous learning-augmented online algorithms do. This is expected, since consistency-robustness trade-offs of previous algorithms optimize the two extreme scenarios of perfect predictions and adversarially bad predictions, whereas the notion of (ρ,μ)-competitiveness also captures the case of useful but imperfect predictions.

For the ski rental problem, in addition to the classical $e/(e-1)$-competitive online algorithm, we consider the following algorithms: FTP, which blindly follows the prediction (i.e., it either buys at time 0 or never); PSK, the randomized algorithm from [39]; and ADJKR, the deterministic algorithm from [4]. As three algorithms – PSK, ADJKR, and ours – each depend on a hyperparameter, we set them in such a way that we obtain the same consistency (ρ in the notation of this paper). For example, the consistency of $\rho = 1.216$ corresponds to the parameter $\lambda \approx \ln(3/2)$ for PSK, as selected in [39].

---

[3]We run prior ski rental algorithms on DPM using Lemma 3 and Theorem 28, as discussed in Sections 4 and 5.

[4]The code and datasets we used to run our experiments are available at `https://github.com/adampolak/dpm`.

24

(a) Algorithms with fixed consistency ρ. (b) Combined algorithms utilising Theorem 28.
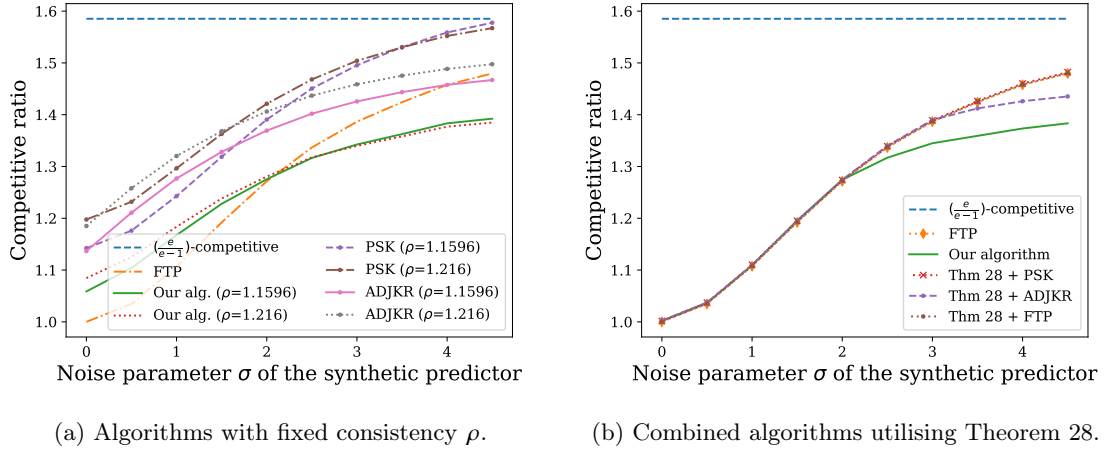
Figure 6: Performance achieved for the two-state DPM (i.e., ski rental) problem on the synthetic dataset PSK$_4$. Each synthetic prediction is equal to the exact request plus a random noise drawn from a normal distribution of mean 0 and standard deviation σ.

We consider either two power states (of respective power consumption $\{0, 1\}$ and wake-up costs $\{1, 0\}$, which correspond to ski rental) or four power states, whose respective power consumption are $\{1, 0.47, 0.105, 0\}$ and wake-up costs are $\{0, 0.12, 0.33, 1\}$, values corresponding to the *active*, *idle*, *stand-by* and *sleep* states of an IBM mobile hard-drive [27]. For four power states, we convert algorithms initially designed for two power states (i.e., multi-round ski rental) using Lemma 3 and convert the resulting algorithms to their prudent variants as discussed in Section 4.1.

We either run algorithms with fixed values of ρ, handling each idle period independently from the past ones, or we use Theorem 28 to let the algorithms adjust the value of ρ across multiple idle periods. For randomized algorithms (i.e., PSK and ours) we use $\rho \in \{1, 1.1, 1.16, 1.3, 1.4, 1.5, \frac{e}{e-1}\}$; values 1 and $\frac{e}{e-1}$ correspond to FTP and the classical randomized $e/(e-1)$-competitive online algorithm. For the deterministic ADJKR algorithm, we do not combine it with the randomized $e/(e-1)$-competitive algorithm; instead, we additionally use $\rho \in \{1.6, 1.7, 1.8, 1.9, 2\}$, where $\rho = 2$ corresponds to the classical deterministic 2-competitive online algorithm. Finally, for FTP, we combine it only with the classical randomized online algorithm. The value of the parameter ϵ_1 is set to 0.1.

We repeated each experiment 10 times. As the maximum standard deviation was smaller than 0.025, we do not print error bars on the charts and only display the average result.

Synthetic scenario. For the synthetic dataset, we generate both the input data and predictions following the approach proposed by Purohit et al. [39]. We use two datasets, each composed of 10 000 independently generated durations (of the idle periods/ski seasons). For the first dataset PSK$_4$, durations are drawn uniformly from $[0, 4]$ as originally proposed in [39]. For the second dataset PSK$_8$, durations are drawn uniformly from $[0, 8]$, in order to offer instances in which also the deepest power state might be chosen by the optimal offline algorithm. We feed the learning-augmented algorithms with synthetic predictions generated as follows: each prediction is equal to the exact request plus a random noise drawn from a normal distribution of mean 0 and standard deviation σ (rounding any negative predictions to 0). The performance is shown as the competitive ratio observed in function of σ.

In Figure 6a, we plot the results for two values of the consistency ρ for the two-state DPM problem (i.e., for multi-round ski rental; however the algorithms in that figure have fixed values of ρ and do not carry any information between consecutive rounds). Figure 6b shows the performance of the algorithms with Theorem 28 applied to automatically adjust the value of ρ. We can observe that for low error the algorithms perform equally well (because they quickly determine that $\rho = 1$

25
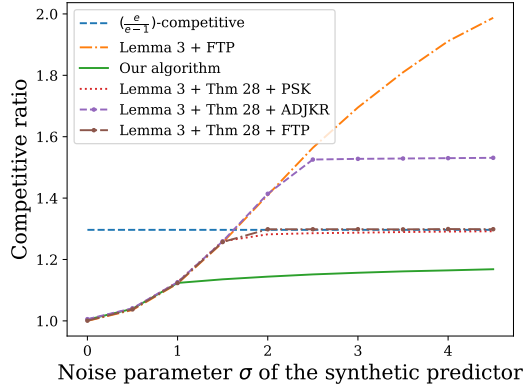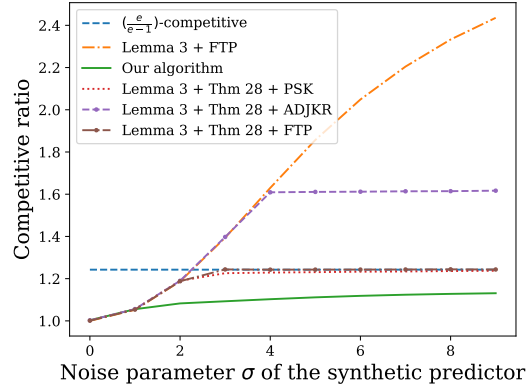
(a) Dataset PSK$_4$.
(b) Dataset PSK$_8$.

Figure 7: Performance achieved for the four-state DPM problem on the synthetic datasets by algorithms utilising Theorem 28.
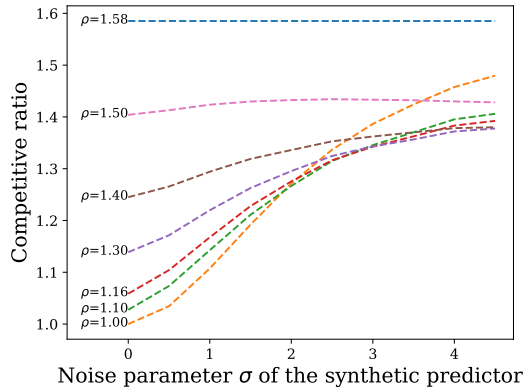
Figure 8: Performance of our ski rental algorithm for different values of ρ on PSK$_4$ dataset. Recall that for $\rho = e/(e-1) \approx 1.58$, our algorithm emulates exactly the classical $e/(e-1)$-competitive online algorithm, while for $\rho = 1$ it emulates the FTP algorithm.
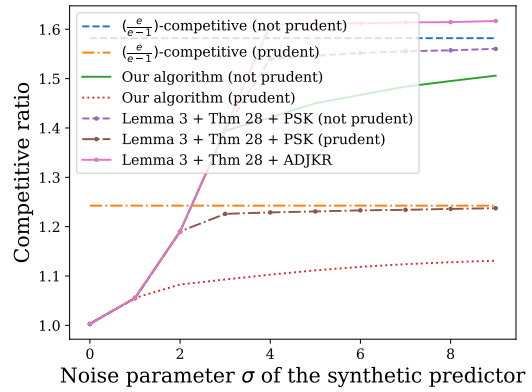
Figure 9: Impact of *prudence* (see Section 4.1) on performance achieved for the four-state DPM problem on PSK$_8$ dataset by algorithms utilising Theorem 28.
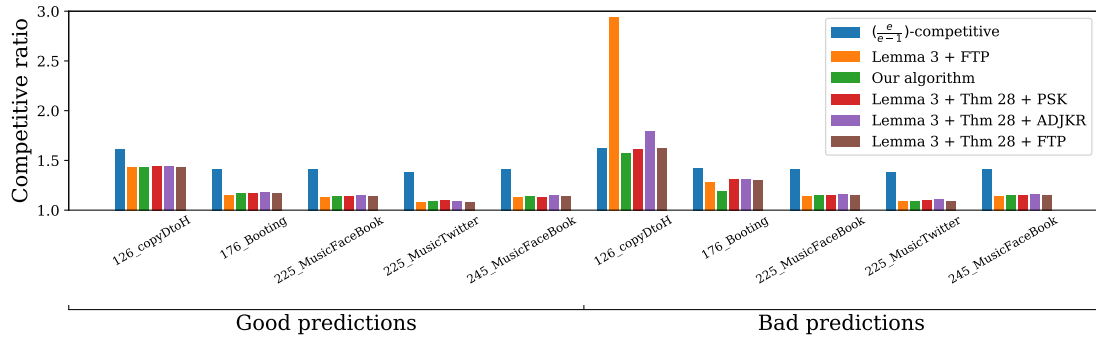
Figure 10: Performance achieved by algorithms utilising Theorem 28 for the four-state DPM problem on the real-world dataset based on I/O traces from Nexus 5 smartphone [44] with predictions generated by *Share* algorithm of Herbster and Warmuth [23] as proposed by Helmbold et al. [22].

26

is best, which corresponds to executing FTP). As the prediction error increases, our algorithm performs better than the other algorithms on this dataset.

Figure 7 presents analogous results to Figure 6b for four-state DPM on both datasets, PSK_4 and PSK_8. The figure shows that the algorithms PSK and ADJKR (combined with Lemma 3 and Theorem 28) essentially perform as well as the better of the two algorithms FTP and the classical online algorithm without predictions. This reflects the fact that these algorithms were designed to optimize the trade-off between consistency and robustness in every single idle period, thus aiming primarily at the two extremes of perfect predictions and adversarially bad predictions. In contrast, the way in which our algorithm's performance degrades as the error increases is (near-)optimal in terms of (ρ, μ)-competitiveness, and the algorithm's consistency and robustness are still optimal over many idle periods (but not necessarily for each idle period individually). Accordingly, our algorithm achieves a significant improvement over previous algorithms in the regime of medium-sized errors and even when predictions are only very weakly correlated with the truth.

The impact of the choice of the consistency ρ on our algorithm is depicted in Figure 8. Lower values of ρ imply a better consistency, so a better performance when σ is small, but may lead to worse results for larger prediction errors.

Finally, Figure 9 shows that converting randomized algorithms to prudent ones (see Section 4.1) is crucial for achieving good experimental performance. Four-state DPM algorithms with that modification ablated perform significantly worse than their prudent variants on the PSK_8 dataset. Recall that since ADJKR is deterministic, there exists no non-prudent version of this algorithm.

Real-world scenario. Even though some of the previous works on DPM [22; 27] include real-world experiments, these papers are rather aged and the datasets used in them seem no longer available. For that reason, in order to test the algorithms in a real-world-inspired scenario, we created a dataset based on I/O traces[5] from a Nexus 5 smartphone [44]. We took the five largest traces and extracted from them durations of idle periods between requests. Since we could not find a power states specification for that device, we used the same power states as for the synthetic experiments, i.e., an IBM mobile hard-drive's power states reported in [27]. Because of that, we had to scale up the idle period durations so that their order of magnitude becomes similar to that in the synthetic experiments.

For this experiment, instead of resorting to synthetic predictions, we implemented a simple actual predictor, proposed by Helmbold et al. [22] in the context of spinning down disks of mobile computers. The predictor adapts the *Share* learning algorithm of Herbster and Warmuth [23], which is based on the multiplicative weights update method. Since it is interesting to evaluate learning-augmented algorithms both in the presence of good and bad predictions, we consider two variants of that predictor. The good variant uses hyperparameters proposed in [22]; the bad variant has the rate parameter of weight updates negated.

Figure 10 presents the results of the experiment. In particular, on each dataset, either our algorithm performs better than all the others, or the nonrobust FTP is the best one and all robust learning-augmented algorithms are almost equally good.

References

[1] S. Albers. Energy-efficient algorithms. *Commun. ACM*, 53(5):86–96, 2010.

[2] S. Albers. On energy conservation in data centers. *ACM Trans. Parallel Comput.*, 6(3): 13:1–13:26, 2019.

[3] S. Albers and J. Quedenfeld. Optimal algorithms for right-sizing data centers. In C. Scheideler and J. T. Fineman, editors, *Proceedings of SPAA'18,*, pages 363–372. ACM, 2018.

[4] S. Angelopoulos, C. Dürr, S. Jin, S. Kamali, and M. Renault. Online Computation with Untrusted Advice. In *Proceedings of ITCS'20*, volume 151, pages 52:1–52:15, 2020.

---

[5] We obtained the traces from `http://iotta.snia.org/traces/block-io`.

27

[5] A. Antoniadis, C. Coester, M. Eliás, A. Polak, and B. Simon. Online metric algorithms with untrusted predictions. In *Proceedings of ICML'20*, pages 345–355. PMLR, 2020.

[6] A. Antoniadis, N. Garg, G. Kumar, and N. Kumar. Parallel machine scheduling to minimize energy consumption. In *Proceedings of SODA'20*, pages 2758–2769. SIAM, 2020.

[7] A. Antoniadis, T. Gouleakis, P. Kleer, and P. Kolev. Secretary and online matching problems with machine learned advice. In *Proceedings of NeurIPS'20*, 2020.

[8] J. Augustine, S. Irani, and C. Swamy. Optimal power-down strategies. *SIAM J. Comput.*, 37 (5):1499–1516, 2008.

[9] É. Bamas, A. Maggiori, L. Rohwedder, and O. Svensson. Learning augmented energy minimization via speed scaling. In *Proceedings of NeurIPS'20*, 2020.

[10] P. Baptiste, M. Chrobak, and C. Dürr. Polynomial-time algorithms for minimum energy scheduling. *ACM Trans. Algorithms*, 8(3):26:1–26:29, 2012.

[11] L. Benini, A. Bogliolo, and G. D. Micheli. A survey of design techniques for system-level dynamic power management. *IEEE Trans. Very Large Scale Integr. Syst.*, 8(3):299–316, 2000.

[12] A. Blum and C. Burch. On-line learning and the metrical task system problem. *Machine Learning*, 39(1):35–58, 2000.

[13] A. Borodin, N. Linial, and M. E. Saks. An optimal on-line algorithm for metrical task system. *J. ACM*, 39(4):745–763, 1992.

[14] N. Cesa-Bianchi and G. Lugosi. *Prediction, learning, and games*. Cambridge University Press, 2006. ISBN 978-0-521-84108-5.

[15] N. Cesa-Bianchi, P. Gaillard, G. Lugosi, and G. Stoltz. A new look at shifting regret. *CoRR*, abs/1202.3323, 2012. URL http://arxiv.org/abs/1202.3323.

[16] N. Chen, G. Goel, and A. Wierman. Smoothed online convex optimization in high dimensions via online balanced descent. In *Proceedings of COLT'18*, volume 75, pages 1574–1594. PMLR, 2018.

[17] E. Chung, L. Benini, A. Bogliolo, Y. Lu, and G. D. Micheli. Dynamic power management for nonstationary service requests. *IEEE Trans. Computers*, 51(11):1345–1361, 2002.

[18] J. R. J. Dirk Blevins, Michael Loewe. White Paper, Inter Corporation: Designing Systems without a Suspend Supply. https://www.intel.com/content/dam/www/public/us/en/documents/white-papers/systems-without-suspend-supply-paper.pdf, 2008. [Online; accessed 14-May-2021].

[19] P. Dütting, S. Lattanzi, R. P. Leme, and S. Vassilvitskii. Secretaries with advice. *CoRR*, abs/2011.06726, 2020.

[20] S. Gollapudi and D. Panigrahi. Online algorithms for rent-or-buy with expert advice. In *Proceedings of ICML'19*, pages 2319–2327, 2019.

[21] E. Hall and R. Willett. Dynamical models and tracking regret in online convex programming. In *Proceedings of ICML'13*, volume 28:1, pages 579–587. PMLR, 2013.

[22] D. P. Helmbold, D. D. Long, and B. Sherrod. A dynamic disk spin-down technique for mobile computing. In *Proceedings of the 2nd annual international conference on Mobile computing and networking*, pages 130–142, 1996.

28

[23] M. Herbster and M. K. Warmuth. Tracking the best expert. In *Proceedings of ICML'95*, pages 286–294. Morgan Kaufmann, 1995.

[24] Hewlett-Packard Corporation, Intel Corporation, Microsoft Corporation, Phoenix Technologies Ltd., and Toshiba Corporation. Advanced Configuration and Power Interface Specification (ACPI). `https://uefi.org/acpi/specs`, 2013. [Online; accessed 14-May-2021].

[25] S. Irani and A. R. Karlin. Online computation. In D. S. Hochbaum, editor, *Approximation Algorithms for NP-Hard Problems*, page 521–564. PWS Publishing Co., USA, 1996. ISBN 0534949681.

[26] S. Irani and K. Pruhs. Algorithmic problems in power management. *SIGACT News*, 36(2): 63–76, 2005.

[27] S. Irani, S. Shukla, and R. Gupta. Online strategies for dynamic power management in systems with multiple power-saving states. *ACM Trans. Embed. Comput. Syst.*, 2(3):325–346, 2003.

[28] S. Irani, S. K. Shukla, and R. Gupta. Algorithms for power savings. *ACM Trans. Algorithms*, 3(4):41, 2007.

[29] A. R. Karlin, M. S. Manasse, L. A. McGeoch, and S. S. Owicki. Competitive randomized algorithms for nonuniform problems. *Algorithmica*, 11(6):542–571, 1994.

[30] S. Khuller, J. Li, and B. Saha. Energy efficient scheduling via partial shutdown. In *SODA*, pages 1360–1372. SIAM, 2010.

[31] S. Lattanzi, T. Lavastida, B. Moseley, and S. Vassilvitskii. Online scheduling via learned weights. In *Proceedings of SODA'20*, pages 1859–1877, 2020.

[32] J. Li and S. Khuller. Generalized machine activation problems. In *Proceedings of SODA'11*, pages 80–94. SIAM, 2011.

[33] S.-H. Lim, B. Sharma, B. C. Tak, and C. R. Das. A dynamic energy management scheme for multi-tier data centers. In *(IEEE ISPASS) IEEE International Symposium on Performance Analysis of Systems and Software*, pages 257–266. IEEE, 2011.

[34] A. Lindermayr, N. Megow, and B. Simon. Double coverage with machine-learned advice. *CoRR*, abs/2103.01640, 2021. URL `https://arxiv.org/abs/2103.01640`.

[35] Z. Lotker, B. Patt-Shamir, and D. Rawitz. Rent, lease, or buy: Randomized algorithms for multislope ski rental. *SIAM J. Discret. Math.*, 26(2):718–736, 2012.

[36] T. Lykouris and S. Vassilvitskii. Competitive caching with machine learned advice. In *Proceedings of ICML'18*, pages 3302–3311, 2018.

[37] M. Mitzenmacher and S. Vassilvitskii. Algorithms with predictions. In T. Roughgarden, editor, *Beyond the Worst-Case Analysis of Algorithms*, pages 646–662. Cambridge University Press, 2020.

[38] S. Phillips and J. Westbrook. On-line algorithms: Competitive analysis and beyond. In M. J. Atallah, editor, *Algorithms and Theory of Computation Handbook*, chapter 10. CRC Press, 1999.

[39] M. Purohit, Z. Svitkina, and R. Kumar. Improving online algorithms via ML predictions. In *Proceedings of NeurIPS'18*, pages 9684–9693, 2018.

[40] M. Rareshide. Power in the data center and its cost across the U.S. `https://info.siteselectiongroup.com/blog/power-in-the-data-center-and-its-costs-across-the-united-states`, 2017. [Online, acessed 23-May-2021].

29

[41] D. Rohatgi. Near-optimal bounds for online caching with machine learned advice. In *Proceedings of SODA'20*, pages 1834–1845, 2020.

[42] A. Wei. Better and simpler learning-augmented online caching. In *Proceedings of AP-PROX/RANDOM'20*, volume 176 of *LIPIcs*, pages 60:1–60:17, 2020.

[43] A. Wei and F. Zhang. Optimal robustness-consistency trade-offs for learning-augmented online algorithms. In *Proceedings of NeurIPS'20*, 2020.

[44] D. Zhou, W. Pan, W. Wang, and T. Xie. I/O characteristics of smartphone applications and their implications for eMMC design. In *2015 IEEE International Symposium on Workload Characterization*, pages 12–21. IEEE Computer Society, 2015. doi: 10.1109/IISWC.2015.8.

30