Tight performance in Bayesian scheduling

Please check the document version of this publication:

• A submitted manuscript is the version of the article upon submission and before peer-review. There can be important differences between the submitted version and the official published version of record. People interested in the research are advised to contact the author for the final version of the publication, or visit the DOI to the publisher's website.
• The final author version and the galley proof are versions of the publication after peer review.
• The final published version features the final layout of the paper including the volume, issue and page numbers.

Link to publication

Sebastián Marbán, Cyriel Rutten,
Tjark Vredeveld

**Tight performance in Bayesian
Scheduling**

RM/10/052

METEOR

Tight performance in Bayesian Scheduling

Sebastián Marbán   Cyriel Rutten   Tjark Vredeveld

Department of Quantitative Economics, Maastricht University
P.O. Box 616, 6200 MD Maastricht, The Netherlands
Email: {s.marban, c.rutten, t.vredeveld}@maastrichtuniversity.nl

Abstract

We consider a stochastic scheduling problem which generalizes traditional stochastic scheduling by introducing parameter uncertainty. Two classes of independent jobs have to be processed by a single machine so as to minimize the sum of expected completion times. The processing times of the jobs are assumed to be exponentially distributed with parameters ϑ and μ, depending on the class of the job. We adopt a Bayesian framework in which μ is assumed to be known, whereas the value of ϑ is unknown. However, the scheduler has specific beliefs about the parameter ϑ. By processing jobs from the corresponding class, the scheduler can update his beliefs about this parameter yielding better future decision making.

For the traditional stochastic scheduling variant, in which the parameters are known, the policy that always processes a job with shortest expected processing time (SEPT) is an optimal policy. In this paper, we show that in the Bayesian framework the performance of SEPT is at most a factor 2 away from the performance of an optimal policy. Furthermore, we construct instances with non-degenerately distributed processing times for which this bound is tight. To our knowledge, this latter result is unique within stochastic scheduling. Finally, we remark that SEPT is asymptotically optimal when the number of jobs of one class tends to infinity, given a fixed number of jobs of the second class.

1 Introduction

Over the last few decades a vast amount of research has focused on stochastic scheduling problems, e. g., A full range of articles is concerned with criteria that guarantee the optimality of simple policies for special scheduling problems or the quality of non-optimal policies. All these papers have in common that the processing times of the jobs are random variables of which the parameters, like expected value, are known for certain. In this paper, we relax this assumption. That is, we study a stochastic scheduling problem in which also the parameters of the processing time distributions are uncertain. By adopting Bayesian methodology, we can learn about these parameters by processing jobs and observing their processing times. We study a policy, which is optimal for the classical stochastic scheduling problem, but not necessarily optimal in the Bayesian framework. We show that this policy is only a factor 2 away from optimality.

Problem definition. Given is a set of jobs, each of which needs to be scheduled on a single machine. This machine can process at most one job at a time and once a job has been initiated it must remain on the machine until completion, i. e., preemption of jobs is not allowed. Moreover, the machine and all jobs are available for processing from the beginning. The processing time of a job is a random variable. The goal is to minimize the total completion time, $\sum_j C_j$, in expectation. In traditional stochastic scheduling, it is assumed that the parameters of the processing time distributions are known. In this paper, we generalize by assuming that these parameters are uncertain. However, prior beliefs on their values do exist which will be updated over time when new jobs are completed.

We consider two job classes J_A and J_B, a setting that has been previously studied by Burnetas and Katehakis [3] and Hamada and Glazebrook [14]. Classes J_A and J_B consist of n_A and n_B independent jobs. The processing time of a job in class J_A is a random variable, which is exponentially distributed with parameter ϑ. Here, ϑ is assumed to be unknown but fixed. The processing time of a job in class J_B is a random variable which is exponentially distributed with known parameter μ. Let X and Y represent two generic random variables for the processing times of a job of class J_A and J_B, respectively. Note that $\mathbb{E}[Y] = 1/\mu$. Furthermore, let random variable X_j denote the processing time of the j^{th} job of class J_A.

Bayesian methodology offers a method to formally recognize the uncertainty regarding ϑ. For some $\theta > 0$, let $g(\theta) := \frac{\partial}{\partial \theta} \Pr[\vartheta \leq \theta]$ denote a probability density function. Intuitively, the probability expresses how strongly we believe that the value of ϑ is less or equal to θ, prior to seeing any realization of processing times of jobs of class J_A. We assume $g(\theta)$ to be a gamma distribution with parameters $\omega > 0$ and $\alpha > 1$. Once k jobs of class J_A have been completed with processing time realizations x_1 up to x_k, the beliefs with respect to the unknown value of ϑ will be updated and expressed by the probability density function

$$g(\theta|x_1, \ldots, x_k) := \frac{\partial}{\partial \theta} \Pr[\vartheta \leq \theta | X_1 = x_1, \ldots, X_k = x_k].$$

Since the gamma distribution provides a conjugate prior for the exponential distribution, $g(\theta|x_1, \ldots, x_k)$ is also a gamma distribution but with parameters $\omega_k := \omega + \sum_{i=1}^{k} x_i$ and $\alpha_k := \alpha + k$. This result was already mentioned in Section 9.4 of [7] and is also trivially derived from Bayes' theorem for probability density functions.

Updating beliefs towards ϑ results in updated beliefs regarding the processing times of uncompleted jobs in class J_A. The probability density function expressing these latter beliefs, after having completed k jobs of class J_A, is denoted by

$$f(x_{k+1}) := \frac{\partial}{\partial x_{k+1}} \Pr[X_{k+1} \leq x_{k+1} | X_1 = x_1, \ldots, X_k = x_k],$$

which is equal to

$$f(x_{k+1}) = \int_0^\infty f(x_{k+1}|\theta) g(\theta|x_1, \ldots, x_k) \partial\theta$$
$$= \int_0^\infty \theta e^{-\theta x_{k+1}} \frac{\omega_k^{\alpha_k}}{\Gamma(\alpha_k)} \theta^{\alpha_k - 1} e^{-\theta \omega_k} \partial\theta = \frac{\alpha_k \omega_k^{\alpha_k}}{(\omega_k + x_{k+1})^{\alpha_k + 1}},$$

where $f(x_{k+1}|\theta)$ is an exponential probability density function with parameter θ. Furthermore, straightforward integration yields the first moment of X_{k+1}:

$$\mathbb{E}[X_{k+1}|x_1, \ldots, x_k] = \int_0^\infty x_{k+1} f(x_{k+1}) \partial x_{k+1} = \frac{\omega_k}{\alpha_k - 1}.$$

The solution to a stochastic scheduling problem is not a simple schedule, but a so-called *scheduling policy*. We follow the notion of scheduling policies as proposed by Möhring, Radermacher, and Weiss [22, 23]. Roughly spoken, a scheduling policy makes scheduling decisions at certain *decision time points*, and these decisions are based on information on the observed past, as well as the a priori knowledge of the input data of the problem. The policy, however, must not anticipate information about the future, such as the actual realizations of the processing times of jobs that have not yet been completed.

In this paper, we analyze the performance of the Shortest Expected Processing Time policy (SEPT) with respect to an optimal policy (OPT). In traditional stochastic scheduling an optimal policy is one that minimizes the total completion time in expectation. In the Bayesian framework, an optimal policy also minimizes the total completion time in expectation, but in addition underlies the uncertainty about the parameter ϑ. The scheduling policy SEPT is defined as follows: when the machine is idle, start processing the job with shortest expected processing time. Within the

Bayesian framework, we allow SEPT to update the expected processing time of jobs in a class every time a job of this specific class has been completed.

In terms of decision making, SEPT and OPT could be interpreted as having a short-term and long-term view, respectively. SEPT only processes a job of class J_A in case the expected processing time of this job is minimal. OPT, however, might choose to process a job of class J_A for which the expected processing time is not necessarily minimal. As a trade-off, OPT benefits from the additional information which is acquired regarding the uncertain parameter ϑ. This information could then lead to better future decision making and a lower sum of completion times. From the above reasoning, one might expect SEPT to perform suboptimal. It is well-known that SEPT is optimal for stochastic single machine scheduling, see [25]. On the other hand, it follows from [3] and [14] that whenever the parameters of the processing times are uncertain, SEPT is not necessarily optimal. This raises the question how effective SEPT is within a Bayesian framework. A conclusive answer is presented in this paper.

Related work. Stochastic scheduling has been an active field of research for more than 40 years. In traditional stochastic scheduling, the processing times of jobs are random variables for which the parameters of the underlying distribution are known. Rothkopf [25] shows that WSEPT (Weighted Shortest Expected Processing Time) is an optimal policy for the stochastic single machine scheduling problem, where the objective is to minimize the sum of weighted expected completion times. Weiss [29, 30] analyzes the performance of WSEPT for the stochastic parallel machine scheduling problem. He shows asymptotic optimality of WSEPT for a certain class of processing time distributions. The first guarantee on the quality of an approximative policy was given by Möhring, Schulz, and Uetz [21]. Other approximative policies have been considered in [6, 19, 20, 27, 28].

The scheduling problem studied in this paper extends traditional stochastic scheduling by applying a Bayesian framework. So far there are only few papers within the field of scheduling that also apply this framework, mostly on a single machine. In the pioneering paper of Gittins and Glazebrook [11], the distributions of processing times of jobs depend all upon the same unknown parameter. The optimal schedule is obtained by calculating appropriate dynamic allocation indices, first proposed by [12]. Another Bayesian scheduling paper is by Hamada and Glazebrook [14]. They study the problem with multiple weighted job classes. The processing time distributions of these classes depend on either known or unknown parameters. Optimal policies are derived by formulating the problem as a dynamic program and using dynamic allocation indices similar to the ones in [9, 10].

Burnetas and Katehakis [3] derive dynamic programming optimality conditions for the same problem with two classes: one with known and one with unknown underlying parameter. For arbitrary processing time distributions, they show that an optimal policy does not depend on the number of jobs in the class with known parameter. As a result the optimality conditions simplify, allowing the problem to be modeled as a stopping problem. For exponential processing times, they thereafter obtain characterizations concerning the structure and properties of the optimal policy. Finally, a policy is given that approximates the decisions made by an optimal policy, in case of the number of jobs in the class with unknown parameter approaches infinity.

There are several other papers that combine stochastic scheduling with Bayesian methodology. In [13], Glazebrook and Owen employ a comparative study in which they seek to quantify the difference between using adaptive scheduling policies based on Bayesian methodology and non-adaptive classical stochastic scheduling policies. The difference is called the value of an adaptive solution. Upper bounds on this value are derived for several scheduling problems on a single machine. Further, Rieder and Weishaupt [24] study a problem in which jobs may change their class with some fixed but unknown probability. A Bayesian approach is taken to update beliefs regarding these probabilities. Finally, the paper of Hamada and Tamaki [15] constitutes the first extension to two machines, for two job classes. Again, a dynamic programming formulation is used to derive optimality conditions.

Bayesian methodology is also widely applied in research fields related to scheduling. In inventory management, for example, there is a large body of literature dealing with uncertain

demand distributions and Bayesian learning. Pioneered by [2, 26], some recent papers are given by [4, 16, 18]. The majority of these papers assumes that prices are exogenous and studies the problem of making optimal inventory decisions. Bayesian demand learning has also received a great deal of attention within the field of pricing, see [1, 5, 17]. All these papers are experimental in that they focus on developing heuristics and studying their computational aspects. The first, and so far only, paper to analyze the theoretical worst-case performance of a pricing heuristic is [8].

Motivation. In traditional stochastic problems, parameters are assumed to be fixed and known. Bayesian methodology however, allows a decision maker to be uncertain about their true value. Parameters are treated as random variables, which is not a description of their variability but a description of the uncertainty. All available information about the true value of the parameters is combined in one expression, a probability distribution. At first it includes all existing beliefs regarding the parameters which might be present prior to seeing any data. Over time these beliefs are updated in the light of new realizations.

In the field of scheduling, the majority of research papers applying Bayesian methodology focuses on analyzing optimal policies using a dynamic programming formulation and deriving dynamic allocation indices. However, applying such an optimal policy is often impractical since the computational complexity explodes with the size of the instance. For example, quantifying the dynamic allocation indices as derived in [14] is hard, even for moderate values of α and n_A, since it requires recursively solving a vast amount of non-linear equations. This calls for the need to develop policies of low computational effort which yield good qualitative performance. The policy SEPT, analyzed in this paper, is an example of such a policy.

So far there are only few papers that propose approximative policies and investigate their theoretical performance, for stochastic problems in a Bayesian framework. In particular, in Bayesian scheduling only the work of Glazebrook and Owen [13] considers and analyzes the quality of approximative policies. We distinct ourselves from their work by studying an adaptive approximative policy. Whereas they focus on the performance of policies that do not update their beliefs of jobs' processing times in light of new realizations, we study the policy SEPT which adapts its beliefs about the mean processing times of jobs, every time a job is completed.

Our contribution. In this paper, we provide a worst case performance guarantee on the expected performance of the Shortest Expected Processing Time policy with respect to the expected performance of an optimal policy. This performance guarantee depends on n_A and n_B and is tight. Furthermore, for arbitrary n_A and n_B, we show that this performance guarantee cannot exceed 2, but might be arbitrarily close. To our knowledge, this exemplifies the first tight performance guarantee in stochastic scheduling, where the tightness follows from non-degenerate processing time distributions. Finally, we remark that SEPT is asymptotically optimal whenever the number of jobs of one class remains fixed while the number of jobs of the second class tends to infinity.

2 Dynamic programming formulation

In this section, we formulate the problem at hand as a dynamic program, introduced by [14]. For notational convenience, we denote SEPT and OPT by, respectively, Π^S and Π^* in the remainder of the paper.

Let $(n_A, n_B, \omega, \alpha) \in \mathbb{Z}_+ \times \mathbb{Z}_+ \times \mathbb{R}_{>0} \times \mathbb{R}_{>1}$ be a state vector encompassing all relevant information of the state the system is in. It consists of the number of jobs in class J_A and J_B as well as the parameters of the current belief for ϑ. Note that we do not include the known parameter μ in this state vector as it does not change during the process. If in state $(n_A, n_B, \omega, \alpha)$, a job of class J_A is processed and completed having realization x, then the state changes to $(n_A - 1, n_B, \omega + x, \alpha + 1)$. On the other hand, if in state $(n_A, n_B, \omega, \alpha)$ a job of class J_B is processed and completed, then the state changes to $(n_A, n_B - 1, \omega, \alpha)$. Let $\mathbb{E}\left[\Pi^S(n_A, n_B, \omega, \alpha)\right]$ and $\mathbb{E}\left[\Pi^*(n_A, n_B, \omega, \alpha)\right]$ denote the expected sum of completion times when, respectively, Π^S and Π^* are adopted from state

$(n_A, n_B, \omega, \alpha)$ onwards. Further, let $\mathbb{E}\left[\Pi_A^*(n_A, n_B, \omega, \alpha)\right]$ denote the sum of the expected completion times of a policy which first processes a job of class J_A (assuming $n_A \geq 1$) and follows an optimal policy afterwards. Similarly for $\mathbb{E}\left[\Pi_B^*(n_A, n_B, \omega, \alpha)\right]$. An optimal policy can be modeled by the following dynamic program:

$$\mathbb{E}\left[\Pi^*(n_A, n_B, \omega, \alpha)\right] = \min\left\{\mathbb{E}\left[\Pi_A^*(n_A, n_B, \omega, \alpha)\right]; \mathbb{E}\left[\Pi_B^*(n_A, n_B, \omega, \alpha)\right]\right\} \quad \forall\, n_A, n_B \geq 1$$

and

$$\mathbb{E}\left[\Pi^*(n_A, 0, \omega, \alpha)\right] = \frac{\omega}{\alpha - 1} \sum_{i=1}^{n_A} i \quad \forall\, n_A \geq 0,$$

$$\mathbb{E}\left[\Pi^*(0, n_B, \omega, \alpha)\right] = \frac{1}{\mu} \sum_{i=1}^{n_B} i \qquad \forall\, n_B \geq 0.$$

As the length of the first job to be processed by a policy influences the completion time of all jobs, straightforward calculations show that for all $n_A \geq 1$ and $n_B \geq 1$

$$\mathbb{E}\left[\Pi_A^*(n_A, n_B, \omega, \alpha)\right] = (n_A + n_B)\frac{\omega}{\alpha - 1} + \int_x \mathbb{E}\left[\Pi^*(n_A - 1, n_B, \omega + x, \alpha + 1)\right] f(x)\partial x,$$

$$\mathbb{E}\left[\Pi_B^*(n_A, n_B, \omega, \alpha)\right] = (n_A + n_B)\frac{1}{\mu} + \mathbb{E}\left[\Pi^*(n_A, n_B - 1, \omega, \alpha)\right],$$

where $f(x)$ is the probability density function of the processing time of a job of class J_A.

3 Performance bounds on scheduling policies

In this section, we present several useful bounds on the performance of policies for the Bayesian scheduling problem. First, we state a trivial lower bound on the performance of an arbitrary policy. This bound is based on the fact that in any policy jobs of a class have to wait for other jobs of the same class. Hence, in constructing the lower bound we neglect waiting times caused by jobs having to wait for jobs of a different class.

Lemma 1. *Let Π be an arbitrary scheduling policy. Then, for any n_A, n_B, ω, μ and $\alpha > 1$,*

$$\mathbb{E}\left[\Pi(n_A, n_B, \omega, \alpha)\right] \geq \mathbb{E}\left[\Pi(n_A, 0, \omega, \alpha)\right] + \mathbb{E}\left[\Pi(0, n_B, \omega, \alpha)\right]$$

$$= \frac{(n_A + 1)n_A}{2}\frac{\omega}{\alpha - 1} + \frac{(n_B + 1)n_B}{2\mu}.$$

Next, we give an upper bound on the performance of policies OPT and SEPT. Due to space limitations. In order to obtain the result, we make use of the following fact: whenever SEPT processes a job of class J_B the expected processing times of neither class are updated, causing SEPT to process one more job of class J_B, if available. This leads to the observation below.

Observation 1. *SEPT processes all jobs of class J_B consecutively.*

Lemma 2. *For any n_A, n_B, ω, μ and $\alpha > 1$,*

$$\mathbb{E}\left[\Pi^*(n_A, n_B, \omega, \alpha)\right] \leq \mathbb{E}\left[\Pi^S(n_A, n_B, \omega, \alpha)\right]$$

$$\leq \frac{(n_A + 1)n_A}{2}\frac{\omega}{\alpha - 1} + \frac{(n_B + 1)n_B}{2\mu} + n_A n_B \min\left\{\frac{1}{\mu}, \frac{\omega}{\alpha - 1}\right\}.$$

Proof. The first inequality trivially follows from the optimality of OPT. Therefore, we continue by proving the second inequality. In case SEPT starts by processing a job of class J_B, i.e., $\frac{1}{\mu} < \frac{\omega}{\alpha - 1}$, it follows from Observation 1 that SEPT will process all jobs of class J_B consecutively. Hence,

$$\mathbb{E}\left[\Pi^S(n_A, n_B, \omega, \alpha)\right] = \frac{(n_A + 1)n_A}{2}\frac{\omega}{\alpha - 1} + \frac{(n_B + 1)n_B}{2\mu} + n_A n_B \frac{1}{\mu}$$

$$= \frac{(n_A + 1)n_A}{2}\frac{\omega}{\alpha - 1} + \frac{(n_B + 1)n_B}{2\mu} + n_A n_B \min\left\{\frac{1}{\mu}, \frac{\omega}{\alpha - 1}\right\}.$$

On the other hand, if SEPT starts by processing a job of class J_A, i.e., $\frac{\omega}{\alpha-1} \leq \frac{1}{\mu}$, the upper bound can be proved by induction on n_A. For $n_A = 1$ and arbitrary n_B, it is easily verified that

$$\mathbb{E}\left[\Pi^{\mathrm{S}}(1, n_B, \omega, \alpha)\right] = (n_B + 1)\frac{\omega}{\alpha - 1} + \frac{(n_B + 1)n_B}{2\mu}$$

$$= \frac{\omega}{\alpha - 1} + \frac{(n_B + 1)n_B}{2\mu} + n_B \min\left\{\frac{1}{\mu}, \frac{\omega}{\alpha - 1}\right\}.$$

Now assume that

$$\mathbb{E}\left[\Pi^{\mathrm{S}}(n, n_B, \omega, \alpha)\right] \leq \frac{(n+1)n}{2}\frac{\omega}{\alpha - 1} + \frac{(n_B + 1)n_B}{2\mu} + n \cdot n_B \frac{\omega}{\alpha - 1}$$

holds for all $n \leq n_A - 1$ and any n_B. We will verify that this statement is also valid for $n = n_A$.

$$\mathbb{E}\left[\Pi^{\mathrm{S}}(n_A, n_B, \omega, \alpha)\right] = (n_A + n_B)\frac{\omega}{\alpha - 1} + \int_0^\infty \mathbb{E}\left[\Pi^{\mathrm{S}}(n_A - 1, n_B, \omega + x, \alpha + 1)\right] f(x)\partial x$$

$$\leq (n_A + n_B)\frac{\omega}{\alpha - 1} + \int_0^\infty \left(\frac{(n_A - 1)(n_A + 2n_B)}{2}\frac{\omega + x}{\alpha} + \frac{(n_B + 1)n_B}{2\mu}\right) f(x)\partial x$$

$$= \frac{(n_A + 1)n_A}{2}\frac{\omega}{\alpha - 1} + \frac{(n_B + 1)n_B}{2\mu} + n_A n_B \frac{\omega}{\alpha - 1}$$

$$= \frac{(n_A + 1)n_A}{2}\frac{\omega}{\alpha - 1} + \frac{(n_B + 1)n_B}{2\mu} + n_A n_B \min\left\{\frac{1}{\mu}, \frac{\omega}{\alpha - 1}\right\}.$$

where the inequality follows from the induction step. $\qquad\square$

We finally present a recursive upper bound on the performance of an optimal policy.

Lemma 3. *For any n_A, n_B, ω, μ and $\alpha > 1$,*

$$\mathbb{E}\left[\Pi^*(n_A, n_B, \omega, \alpha)\right] - \mathbb{E}\left[\Pi^*(n_A, n_B - 1, \omega, \alpha)\right]$$
$$\leq \mathbb{E}\left[\Pi^*(n_A, n_B - 1, \omega, \alpha)\right] - \mathbb{E}\left[\Pi^*(n_A, n_B - 2, \omega, \alpha)\right] + 1/\mu$$

Proof. We provide a proof by induction on n_A. In case $n_A = 0$, it follows immediately that

$$\mathbb{E}\left[\Pi^*(0, n_B - 1, \omega, \alpha)\right] - \mathbb{E}\left[\Pi^*(0, n_B - 2, \omega, \alpha)\right] + 1/\mu = \frac{(n_B + 1)n_B}{2\mu} - \frac{(n_B - 1)n_B}{2\mu}$$

$$= \mathbb{E}\left[\Pi^*(0, n_B, \omega, \alpha)\right] - \mathbb{E}\left[\Pi^*(0, n_B - 1, \omega, \alpha)\right]$$

for any n_B, ω, μ and $\alpha > 1$. Now assume that

$$\mathbb{E}\left[\Pi^*(n, n_B, \omega, \alpha)\right] - \mathbb{E}\left[\Pi^*(n, n_B - 1, \omega, \alpha)\right]$$
$$\leq \mathbb{E}\left[\Pi^*(n, n_B - 1, \omega, \alpha)\right] - \mathbb{E}\left[\Pi^*(n, n_B - 2, \omega, \alpha)\right] + 1/\mu$$

for all $n \leq n_A - 1$ and any n_B, ω, μ and $\alpha > 1$. We will verify that this statement is also valid for $n = n_A$. In order to do so, we first assume that $\mathbb{E}\left[\Pi_{\mathrm{B}}^*(n_A, n_B - 1, w, a)\right] < \mathbb{E}\left[\Pi_{\mathrm{A}}^*(n_A, n_B - 1, w, a)\right]$. Then,

$$\mathbb{E}\left[\Pi^*(n_A, n_B, \omega, \alpha)\right] - \mathbb{E}\left[\Pi^*(n_A, n_B - 1, \omega, \alpha)\right]$$
$$\leq \mathbb{E}\left[\Pi_{\mathrm{B}}^*(n_A, n_B, \omega, \alpha)\right] - \mathbb{E}\left[\Pi_{\mathrm{B}}^*(n_A, n_B - 1, \omega, \alpha)\right]$$
$$= \mathbb{E}\left[\Pi^*(n_A, n_B - 1, \omega, \alpha)\right] - \mathbb{E}\left[\Pi^*(n_A, n_B - 2, \omega, \alpha)\right] + 1/\mu,$$

where the inequality follows from the assumption and the optimality of Π^*. Secondly, assume that $\mathbb{E}\left[\Pi_{\mathrm{B}}^*(n_A, n_B - 1, \omega, \alpha)\right] \geq \mathbb{E}\left[\Pi_{\mathrm{A}}^*(n_A, n_B - 1, \omega, \alpha)\right]$. This gives

$$
\begin{aligned}
\mathbb{E}\left[\Pi^*(n_A, n_B, \omega, \alpha)\right] &- \mathbb{E}\left[\Pi^*(n_A, n_B - 1, \omega, \alpha)\right] \\
&\leq \mathbb{E}\left[\Pi_{\mathrm{A}}^*(n_A, n_B, \omega, \alpha)\right] - \mathbb{E}\left[\Pi_{\mathrm{A}}^*(n_A, n_B - 1, \omega, \alpha)\right] \\
&= \int_0^\infty \mathbb{E}\left[\Pi^*(n_A - 1, n_B, \omega + x, \alpha + 1)\right] f(x)\partial x \\
&\quad - \int_0^\infty \left(\mathbb{E}\left[\Pi^*(n_A - 1, n_B - 1, \omega + x, \alpha + 1)\right] - \frac{\omega + x}{\alpha}\right) f(x)\partial x \\
&\leq \int_0^\infty \mathbb{E}\left[\Pi^*(n_A - 1, n_B - 1, \omega + x, \alpha + 1)\right] f(x)\partial x \\
&\quad - \int_0^\infty \left(\mathbb{E}\left[\Pi^*(n_A - 1, n_B - 2, \omega + x, \alpha + 1)\right] - \frac{\omega + x}{\alpha}\right) f(x)\partial x + \frac{1}{\mu} \\
&= \mathbb{E}\left[\Pi_{\mathrm{A}}^*(n_A, n_B - 1, \omega, \alpha)\right] - \mathbb{E}\left[\Pi_{\mathrm{A}}^*(n_A, n_B - 2, \omega, \alpha)\right] + 1/\mu \\
&\leq \mathbb{E}\left[\Pi^*(n_A, n_B - 1, \omega, \alpha)\right] - \mathbb{E}\left[\Pi^*(n_A, n_B - 2, \omega, \alpha)\right] + 1/\mu,
\end{aligned}
$$

where the first and last inequality follow from the assumption and the optimality of Π^*, and the second inequality from the induction step. $\qquad\square$

4 Structural analysis of OPT

In this section, we show that OPT and SEPT produce the same schedule in case OPT starts by processing a job of class J_B. This result builds upon the property of OPT processing all jobs of class J_B consecutively. To establish this property, we first show that in case it is optimal to start processing a job of class J_A, it remains optimal to start processing a job of class J_A when the number of jobs in class J_B is increased by an arbitrary number.

Lemma 4. *For any* $n_A, n_B > 1, \omega, \mu$ *and* $\alpha > 1$, *if*

$$\mathbb{E}\left[\Pi_{\mathrm{B}}^*(n_A, n_B - 1, \omega, \alpha)\right] \geq \mathbb{E}\left[\Pi_{\mathrm{A}}^*(n_A, n_B - 1, \omega, \alpha)\right],$$

then

$$\mathbb{E}\left[\Pi_{\mathrm{B}}^*(n_A, n_B, \omega, \alpha)\right] \geq \mathbb{E}\left[\Pi_{\mathrm{A}}^*(n_A, n_B, \omega, \alpha)\right].$$

Proof. Assume that

$$\mathbb{E}\left[\Pi_{\mathrm{B}}^*(n_A, n_B - 1, \omega, \alpha)\right] \geq \mathbb{E}\left[\Pi_{\mathrm{A}}^*(n_A, n_B - 1, \omega, \alpha)\right],$$

for any $n_A, n_B > 1, \omega, \mu$ and $\alpha > 1$. Since,

$$
\begin{aligned}
\mathbb{E}\left[\Pi_{\mathrm{B}}^*(n_A, n_B - 1, \omega, \alpha)\right] &= (n_A + n_B - 1)(1/\mu) + \mathbb{E}\left[\Pi^*(n_A, n_B - 2, \omega, \alpha)\right] \\
&\leq (n_A + n_B - 1)(1/\mu) + \mathbb{E}\left[\Pi_{\mathrm{A}}^*(n_A, n_B - 2, \omega, \alpha)\right] \\
&= (n_A + n_B - 1)\frac{1}{\mu} + (n_A + n_B - 2)\frac{\omega}{\alpha - 1} \\
&\quad + \int_0^\infty \mathbb{E}\left[\Pi^*(n_A - 1, n_B - 2, \omega + x, \alpha + 1)\right] f(x)\partial x,
\end{aligned}
$$

and

$$
\begin{aligned}
\mathbb{E}\left[\Pi_{\mathrm{A}}^*(n_A, n_B - 1, \omega, \alpha)\right] &= (n_A + n_B - 1)\frac{\omega}{\alpha - 1} \\
&\quad + \int_0^\infty \mathbb{E}\left[\Pi^*(n_A - 1, n_B - 1, \omega + x, \alpha + 1)\right] f(x)\partial x,
\end{aligned}
$$

7

rewriting yields

$$
\begin{aligned}
\frac{n_A + n_B}{\mu} - \frac{\omega}{\alpha - 1} &\geq \frac{1}{\mu} + \int_0^\infty \left(\mathbb{E}\left[\Pi^*(n_A - 1, n_B - 1, \omega + x, \alpha + 1)\right] \right. \\
&\qquad\qquad \left. - \mathbb{E}\left[\Pi^*(n_A - 1, n_B - 2, \omega + x, \alpha + 1)\right] \right) f(x)\partial x \\
&\geq \int_0^\infty \left(\mathbb{E}\left[\Pi^*(n_A - 1, n_B, \omega + x, \alpha + 1)\right] \right. \\
&\qquad\qquad \left. - \mathbb{E}\left[\Pi^*(n_A - 1, n_B - 1, \omega + x, \alpha + 1)\right] \right) f(x)\partial x, \qquad (1)
\end{aligned}
$$

where the last inequality follows from Lemma 4. Then,

$$
\begin{aligned}
\mathbb{E}\left[\Pi_{\mathrm{B}}^*(n_A, n_B, \omega, \alpha)\right] &= \frac{n_A + n_B}{\mu} + \mathbb{E}\left[\Pi_{\mathrm{A}}^*(n_A, n_B - 1, \omega, \alpha)\right] \\
&= \frac{n_A + n_B}{\mu} + (n_A + n_B - 1)\frac{\omega}{\alpha - 1} \\
&\quad + \int_0^\infty \mathbb{E}\left[\Pi^*(n_A - 1, n_B - 1, \omega + x, \alpha + 1)\right] f(x)\partial x \\
&\geq (n_A + n_B)\frac{\omega}{\alpha - 1} + \int_0^\infty \mathbb{E}\left[\Pi^*(n_A - 1, n_B, \omega + x, \alpha + 1)\right] f(x)\partial x \\
&= \mathbb{E}\left[\Pi_{\mathrm{A}}^*(n_A, n_B, \omega, \alpha)\right],
\end{aligned}
$$

where the inequality follows from (1). $\qquad\square$

It follows immediately from the lemma above that if it is optimal to start processing a job of class J_B, then all jobs of class will be processed consecutively until none remains.

Theorem 1. *OPT processes all jobs of class J_B consecutively.*

Proof. Assume that $\mathbb{E}\left[\Pi_{\mathrm{B}}^*(n_A, n_B, \omega, \alpha)\right] < \mathbb{E}\left[\Pi_{\mathrm{A}}^*(n_A, n_B, \omega, \alpha)\right]$ for any $n_A, n_B > 1, \omega, \mu$ and $\alpha > 1$. We want to show that $\mathbb{E}\left[\Pi_{\mathrm{B}}^*(n_A, n_B - 1, \omega, \alpha)\right] < \mathbb{E}\left[\Pi_{\mathrm{A}}^*(n_A, n_B - 1, \omega, \alpha)\right]$, in which case the theorem trivially follows. By contradiction, assume that $\mathbb{E}\left[\Pi_{\mathrm{B}}^*(n_A, n_B - 1, \omega, \alpha)\right] \geq \mathbb{E}\left[\Pi_{\mathrm{A}}^*(n_A, n_B - 1, \omega, \alpha)\right]$. Then, $\mathbb{E}\left[\Pi_{\mathrm{B}}^*(n_A, n_B, \omega, \alpha)\right] \geq \mathbb{E}\left[\Pi_{\mathrm{A}}^*(n_A, n_B, \omega, \alpha)\right]$ by Lemma 4, yielding a contradiction to our original assumption. $\qquad\square$ We remark that the result given in the theorem above also follows from Theorem 3.4 and Remark 3.5 of the work of Burnetas and Katehakis [3]. However, our line of reasoning is build upon different arguments and bounds, which provide new insight into the problem at hand. Burnetas and Katehakis introduce a regret function which is the difference between the expected sum of completion times of an optimal policy and a lower bound on this value. They show that instead of minimizing the expected sum of completion times of an optimal policy, one might instead minimize this regret function. Finally, they establish that deciding whether it is optimal to process either a job of class J_A or a job of class J_B does not depend on the number of jobs of class J_B which still have to be processed. Hence, all jobs of class J_B will be processed consecutively.

Below we show that if OPT starts by processing a job of class J_B, then the expected processing time of a job of class J_A exceeds that of a job of class J_B.

Theorem 2. *If OPT starts by processing a job of class J_B, then $\frac{1}{\mu} \leq \frac{\omega}{\alpha - 1}$.*

Proof. Let $\mathbb{E}\left[\Pi_{\mathrm{B}}^*(n_A, n_B, \omega, \alpha)\right] < \mathbb{E}\left[\Pi_{\mathrm{A}}^*(n_A, n_B, \omega, \alpha)\right]$. From Theorem 1, we have that in any optimal policy all jobs of class J_B are processed consecutively. Hence,

$$
\begin{aligned}
\mathbb{E}\left[\Pi^*(n_A, n_B, \omega, \alpha)\right] &= \frac{(n_A + 1)n_A}{2}\frac{\omega}{\alpha - 1} + \frac{(n_B + 1)n_B}{2\mu} + \frac{n_A n_B}{\mu} \\
&\leq \frac{(n_A + 1)n_A}{2}\frac{\omega}{\alpha - 1} + \frac{(n_B + 1)n_B}{2\mu} + n_A n_B \frac{\omega}{\alpha - 1},
\end{aligned}
$$

8

where the inequality follows from Lemma 2. Rewriting gives $\frac{1}{\mu} \leq \frac{\omega}{\alpha-1}$, and the result follows. □

From Observation 1 and Theorem 1 we have that both SEPT and OPT process all jobs of class J_B consecutively. In the theorem above we established that whenever OPT starts processing a job of class J_B, then so does SEPT. Combining these insights yields the following corollary.

Corollary 1. *If OPT starts by processing a job of class J_B, then OPT and SEPT yield the same schedule.*

5 Performance of SEPT in a Bayesian setting

In the previous section we observed that SEPT and OPT yield the same schedule and hence identical objective values for the case that OPT starts by processing a job of class J_B. In this section we analyze the worst case performance of SEPT with respect the performance of OPT in its full generality. We provide a tight upper bound on the performance guarantee depending on the parameters n_A and n_B. Moreover, we show that this upper bound cannot exceed 2.

5.1 Performance guarantee for SEPT

Theorem 3. *For any n_A, n_B, ω, μ and $\alpha > 1$,*

$$\frac{\mathbb{E}\left[\Pi^{\mathrm{S}}(n_A, n_B, \omega, \alpha)\right]}{\mathbb{E}\left[\Pi^{*}(n_A, n_B, \omega, \alpha)\right]} < \frac{n_A^2 + n_B^2 + 2n_A n_B + n_A + n_B}{n_A^2 + n_B^2 + n_A + 3n_B}.$$

Proof. From Corollary 1, it follows that in case OPT starts by processing a job of class J_B, both policies produce the same schedule and performance. Therefore, we assume that OPT starts by processing a job of class J_A. By Lemma 1, we then have

$$\mathbb{E}\left[\Pi_{\mathrm{A}}^{*}(n_A, n_B, \omega, \alpha)\right] = (n_A + n_B)\frac{\omega}{\alpha-1} + \int_0^\infty \mathbb{E}\left[\Pi^{*}(n_A - 1, n_B, \omega + x, \alpha + 1)\right] f(x)\partial x$$

$$\geq \frac{(n_B+1)n_B}{2\mu} + \frac{(n_A+1)n_A}{2}\frac{\omega}{\alpha-1} + n_B\frac{\omega}{\alpha-1}. \tag{2}$$

Combining (2) and Lemma 2, we obtain

$$\frac{\mathbb{E}\left[\Pi^{\mathrm{S}}(n_A, n_B, \omega, \alpha)\right]}{\mathbb{E}\left[\Pi^{*}(n_A, n_B, \omega, \alpha)\right]} \leq \frac{\frac{(n_A+1)n_A}{2}\frac{\omega}{\alpha-1} + \frac{(n_B+1)n_B}{2\mu} + n_A n_B \min\left\{\frac{1}{\mu}, \frac{\omega}{\alpha-1}\right\}}{\frac{(n_B+1)n_B}{2\mu} + \frac{(n_A+1)n_A}{2}\frac{\omega}{\alpha-1} + n_B\frac{\omega}{\alpha-1}}$$

$$\leq \frac{\frac{(n_A+1)n_A}{2}\frac{\omega}{\alpha-1} + \frac{(n_B+1)n_B}{2}\min\left\{\frac{1}{\mu}, \frac{\omega}{\alpha-1}\right\} + n_A n_B \min\left\{\frac{1}{\mu}, \frac{\omega}{\alpha-1}\right\}}{\frac{(n_B+1)n_B}{2}\min\left\{\frac{1}{\mu}, \frac{\omega}{\alpha-1}\right\} + \frac{(n_A+1)n_A}{2}\frac{\omega}{\alpha-1} + n_B\frac{\omega}{\alpha-1}} \tag{3}$$

$$\leq \frac{\frac{(n_A+1)n_A}{2}\frac{\omega}{\alpha-1} + \frac{(n_B+1)n_B}{2}\frac{\omega}{\alpha-1} + n_A n_B\frac{\omega}{\alpha-1}}{\frac{(n_B+1)n_B}{2}\frac{\omega}{\alpha-1} + \frac{(n_A+1)n_A}{2}\frac{\omega}{\alpha-1} + n_B\frac{\omega}{\alpha-1}} \tag{4}$$

$$= \frac{n_A^2 + n_B^2 + 2n_A n_B + n_A + n_B}{n_A^2 + n_B^2 + n_A + 3n_B}.$$

Here, the second inequality follows from the fact that for any constants a, b, c and d, $\frac{a+b}{a+c} \leq \frac{d+b}{d+c}$ in case $b \geq c > 0$ and $a \geq d > 0$. The third inequality can be seen from the fact that the ratio is increasing in $\frac{1}{\mu}$ for $\frac{1}{\mu} \leq \frac{\omega}{\alpha-1}$. Note that (3) and (4) cannot both hold with equality simultaneously, hence the strict inequality in the theorem. □

9

5.2 Tightness of the performance guarantee

We show that a series of instances exist for which the ratio of the performance of SEPT over the performance of OPT approaches the performance guarantee of the previous subsection arbitrarily close. In order to obtain this result, we need make use of the following proposition.

Proposition 1.
$$\lim_{\beta \downarrow 1} \frac{1}{\beta} \left(\frac{\beta - 1}{\beta} \right)^{\beta - 1} = 1.$$

Proof. Let $v(\beta) = \beta \ln \beta = \ln \beta / (1/\beta)$. It follows trivially from l'Hôpital's rule that

$$\lim_{\beta \downarrow 0} v(\beta) = 0. \tag{5}$$

Further, let $w(\beta) = \frac{1}{\beta} \left(\frac{\beta-1}{\beta} \right)^{\beta-1}$. Then,

$$\ln w(\beta) = \ln \frac{1}{\beta} \left(\frac{\beta - 1}{\beta} \right)^{\beta - 1} = \ln \frac{(\beta - 1)^{\beta - 1}}{\beta^{\beta}}$$
$$= (\beta - 1) \ln(\beta - 1) - \beta \ln \beta$$

and
$$\lim_{\beta \downarrow 1} \ln w(\beta) = \lim_{\beta \downarrow 1} \left\{ (\beta - 1) \ln(\beta - 1) - \beta \ln \beta \right\} = 0,$$

where the last equality follows from (5). Since $u = \exp(\ln u)$, this brings us to

$$\lim_{\beta \downarrow 1} w(\beta) = \lim_{\beta \downarrow 1} \exp(\ln w(\beta)) = \exp(0) = 1,$$

which concludes the proof. $\square$

Theorem 4. *For any n_A and n_B, there exists parameter settings ω, μ and $\alpha > 1$, such that*

$$\frac{\mathbb{E} \left[\Pi^{\mathrm{S}}(n_A, n_B, \omega, \alpha) \right]}{\mathbb{E} \left[\Pi^{*}(n_A, n_B, \omega, \alpha) \right]} > \frac{n_A^2 + n_B^2 + 2n_A n_B + n_A + n_B}{n_A^2 + n_B^2 + n_A + 3n_B} - \epsilon,$$

for any $\epsilon > 0$.

Proof. First, set $\mu = 1$. Then,

$$\mathbb{E} \left[\Pi^{*}(n_A, n_B, \omega, \alpha) \right] \leq \mathbb{E} \left[\Pi_A^{*}(n_A, n_B, \omega, \alpha) \right]$$
$$= (n_A + n_B) \frac{\omega}{\alpha - 1} + \int_0^{\infty} \mathbb{E} \left[\Pi^{*}(n_A - 1, n_B, \omega + x, \alpha + 1) \right] f(x) \partial x$$
$$\leq (n_A + n_B) \frac{\omega}{\alpha - 1} + \frac{(n_A - 1)n_A}{2} \frac{\omega}{\alpha - 1} + \frac{(n_B + 1)n_B}{2}$$
$$+ (n_A - 1)n_B \int_0^{\infty} \min \left\{ 1, \frac{\omega + x}{\alpha} \right\} \frac{\alpha \omega^{\alpha}}{(\omega + x)^{\alpha + 1}} \partial x$$
$$= n_B \frac{\omega}{\alpha - 1} + \frac{(n_A + 1)n_A}{2} \frac{\omega}{\alpha - 1} + \frac{(n_B + 1)n_B}{2}$$
$$+ (n_A - 1)n_B \frac{\omega}{\alpha - 1} \left[1 - \frac{1}{\alpha} \left(\frac{\omega}{\alpha} \right)^{\alpha - 1} \right]. \tag{6}$$

Here, the second inequality follows from Lemma 2, and the last equality follows from the derivation below.

$$\int_0^\infty \min\left\{1, \frac{\omega + x}{\alpha}\right\} \frac{\alpha \omega^\alpha}{(\omega + x)^{\alpha+1}} \partial x = \int_0^{\alpha-\omega} \frac{\omega^\alpha}{(\omega + x)^\alpha} \partial x + \int_{\alpha-\omega}^\infty \frac{\alpha \omega^\alpha}{(\omega + x)^{\alpha+1}} \partial x$$

$$= \omega^\alpha \left[\frac{1}{\alpha-1} \frac{1}{\omega^{\alpha-1}} - \frac{1}{\alpha-1} \frac{1}{\alpha^{\alpha-1}}\right] + \left(\frac{\omega}{\alpha}\right)^\alpha$$

$$= \frac{\omega}{\alpha-1}\left[1 - \frac{1}{\alpha}\left(\frac{\omega}{\alpha}\right)^{\alpha-1}\right].$$

Next, set $\omega = \alpha - 1$. This gives

$$\mathbb{E}\left[\Pi^{\mathrm{S}}(n_A, n_B, \alpha - 1, \alpha)\right] = \frac{(n_A + 1)n_A}{2} + \frac{(n_B + 1)n_B}{2} + n_A n_B. \tag{7}$$

Combining (6) and (7), we define

$$z(n_A, n_B, \alpha) := \frac{\frac{(n_A+1)n_A}{2} + \frac{(n_B+1)n_B}{2} + n_A n_B}{n_B + \frac{(n_A+1)n_A}{2} + \frac{(n_B+1)n_B}{2} + (n_A - 1)n_B\left[1 - \frac{1}{\alpha}\left(\frac{\alpha-1}{\alpha}\right)^{\alpha-1}\right]}.$$

Furthermore, define

$$\rho(n_A, n_B) := \frac{n_A^2 + n_B^2 + 2n_A n_B + n_A + n_B}{n_A^2 + n_B^2 + n_A + 3n_B}$$

Clearly, $z(n_A, n_B, \alpha) < \rho(n_A, n_B)$, and it follows from Proposition 1 that

$$\lim_{\alpha \downarrow 1} z(n_A, n_B, \alpha) = \rho(n_A, n_B).$$

Therefore, for any $\epsilon > 0$ there exists an $\tilde{\alpha} > 1$ such that

$$z(n_A, n_B, \tilde{\alpha}) > \rho(n_A, n_B) - \epsilon. \tag{8}$$

From equations (6) and (7), we also have that

$$\frac{\mathbb{E}\left[\Pi^{\mathrm{S}}(n_A, n_B, \alpha - 1, \alpha)\right]}{\mathbb{E}\left[\Pi^*(n_A, n_B, \alpha - 1, \alpha)\right]} \geq z(n_A, n_B, \alpha), \tag{9}$$

for all n_A, n_B and $\alpha > 1$ whenever $\mu = 1$ and $\omega = \alpha - 1$. Hence, combining Theorem 3 and equations (8) and (9), we find that for any epsilon $\epsilon > 0$ there exists an $\tilde{\alpha} > 1$, for which

$$\rho(n_A, n_B) > \frac{\mathbb{E}\left[\Pi^{\mathrm{S}}(n_A, n_B, \tilde{\alpha} - 1, \tilde{\alpha})\right]}{\mathbb{E}\left[\Pi^*(n_A, n_B, \tilde{\alpha} - 1, \tilde{\alpha})\right]} > \rho(n_A, n_B) - \epsilon.$$

$\square$

5.3 Constant factor performance guarantee

It follows from Theorem 3 that the performance guarantee of SEPT will not exceed 2. Furthermore, extending the proof of Theorem 4, it shows that a series of instances can be constructed for which the ratio of the performance of SEPT over the performance of OPT approaches the upper bound of 2 arbitrary close.

Corollary 2. *For any n_A, n_B, ω, μ and $\alpha > 1$,*

$$\frac{\mathbb{E}\left[\Pi^{\mathrm{S}}(n_A, n_B, \omega, \alpha)\right]}{\mathbb{E}\left[\Pi^*(n_A, n_B, \omega, \alpha)\right]} < 2,$$

and there exist parameters settings n_A, n_B, ω, μ and $\alpha > 1$ such that

$$\frac{\mathbb{E}\left[\Pi^{\mathrm{S}}(n_A, n_B, \omega, \alpha)\right]}{\mathbb{E}\left[\Pi^{*}(n_A, n_B, \omega, \alpha)\right]} > 2 - \epsilon$$

for any $\epsilon > 0$.

Proof. From Theorem 3, it follows that

$$\frac{\mathbb{E}\left[\Pi^{\mathrm{S}}(n_A, n_B, \omega, \alpha)\right]}{\mathbb{E}\left[\Pi^{*}(n_A, n_B, \omega, \alpha)\right]} < \frac{n_A^2 + n_B^2 + 2n_A n_B + n_A + n_B}{n_A^2 + n_B^2 + n_A + 3n_B} \leq \frac{4n_{\max}^2 + 2n_{\max}}{2n_{\max}^2 + 4n_{\max}} < 2,$$

where $n_{\max} = \max\{n_A, n_B\}$. We show that this bound is tight by extending the instance provided in Theorem 4. We additionally set $n_A = n_B = n$ and define

$$z(n, \alpha) := \frac{4n^2 + 2n}{2n^2 + 4n + (n-1)n\left[1 - \frac{1}{\alpha}\left(\frac{\alpha-1}{\alpha}\right)^{\alpha-1}\right]} < 2.$$

Since by Proposition 1,

$$\lim_{n \to \infty} \lim_{\alpha \downarrow 1} z(n, \alpha) = \lim_{n \to \infty} \frac{4n^2 + 2n}{2n^2 + 4n} = 2,$$

it follows that for any $\epsilon > 0$ there exists values for $\tilde{n}$ and $\tilde{\alpha}$ such that

$$z(\tilde{n}, \tilde{\alpha}) > 2 - \epsilon. \tag{10}$$

Further, from (6), (7) and (10), there exist values $\tilde{n}$ and $\tilde{\alpha}$ such that

$$2 > \frac{\mathbb{E}\left[\Pi^{\mathrm{S}}(\tilde{n}, \tilde{n}, \tilde{\alpha} - 1, \tilde{\alpha})\right]}{\mathbb{E}\left[\Pi^{*}(\tilde{n}, \tilde{n}, \tilde{\alpha} - 1, \tilde{\alpha})\right]} \geq z(\tilde{n}, \tilde{n}, \tilde{\alpha}) = z(\tilde{n}, \tilde{\alpha}) > 2 - \epsilon,$$

for any $\epsilon > 0$, and the result follows. Note that $z(\tilde{n}, \tilde{n}, \tilde{\alpha})$ is defined in the proof of Theorem 4. $\square$

6 Concluding remarks

In this paper, we showed that the performance of SEPT can be a factor of 2 away from the performance of an optimal policy when both n_A and n_B tend to infinity. For given n_A and n_B, Theorem 3 provides a performance guarantee only depending on these two parameters. It immediately follows from this result that the performance guarantee will be close to one in case the number of jobs in one class is of a different order than the number of jobs in the second class. To be more explicit, when the number of jobs in one class is fixed while the number of jobs in the second class tends to infinity, then the performance guarantee will go to one, yielding asymptotic optimality of SEPT.

Next, we remark that the performance guarantees provided in Theorem 3 and Corollary 2 are highly sensitive to the values of the parameters and in particular to the value of α. To illustrate this sensitivity, take $\alpha \geq 2$. Then, it can be shown that the performance guarantee for arbitrary n_A and n_B reduces from 2 to a tight bound of $2/(2 - 1/e) \approx 1.225$.

References

[1] V.F. Araman and R. Caldentey. Dynamic pricing for nonperishable products with demand learning. *Operations Research*, 57(5):1169–1188, 2009.

[2] K.S. Azoury. Bayes solution to dynamic inventory models under unknown demand distribution. *Management Science*, 31(9):1150–1160, 1985.

[3] A.N. Burnetas and M.N. Katehakis. On sequencing two types of tasks on a single processor under incomplete information. *Probability in the Engineering and Informational Sciences*, 7(1):85–119, 1993.

[4] L. Chen and E.L. Plambeck. Dynamic inventory management with learning about the demand distribution and substitution probability. *Manufacturing & Service Operations Management*, 10(2):236–256, 2008.

[5] E. Cope. Bayesian strategies for dynamic pricing in e-commerce. *Naval Research Logistics*, 54(3):265–281, 2006.

[6] B. C. Dean. *Approximation Algorithms for Stochastic Scheduling Problems*. PhD thesis, Massachusetts Institute of Technology, 2005.

[7] M.H. DeGroot. *Optimal Statistical Decisions*. McGraw-Hill, New York, 1970.

[8] F.F. Farias and B. Van Roy. Dynamic pricing with a prior on market response. *Operations Research*, 58(1):16–29, 2010.

[9] J.C. Gittins. Bandit processes and dynamic allocation indices (with discussion). *Journal of the Royal Statistical Society, Series B (Methodological)*, 41(2):148–177, 1979.

[10] J.C. Gittins. *Multi-armed bandit allocation indices*. Wiley, N.Y., 1989.

[11] J.C. Gittins and K.D. Glazebrook. On Bayesian models in stochastic scheduling. *Journal of Applied Probability*, 14(3):556–565, 1977.

[12] J.C. Gittins and D.M. Jones. A dynamic allocation index for the sequential design of experiments. In *Progress in Statistics*, pages 241–266. 1974.

[13] K.D. Glazebrook and R.W. Owen. On the value of adaptive solutions to stochastic scheduling problems. *Mathematics of Operations Research*, 20(1):65–89, 1995.

[14] T. Hamada and K.D. Glazebrook. A Bayesian sequential single machine scheduling problem to minimize the expected weighted sum of flowtimes of jobs with exponential processing times. *Operations Research*, 41(5):924–934, 1993.

[15] T. Hamada and M. Tamaki. Some results on a Bayesian sequential scheduling on two identical parallel processors. *Journal of the Operations Research Society of Japan*, 42(14):316–329, 1999.

[16] M.A. Lariviere and E.L. Porteus. Stalking information: Bayesian inventory management with unobserved lost sales. *Management Science*, 45(3):346–363, 1999.

[17] K.Y. Lin. Dynamic pricing with real-time demand learning. *Operations Research*, 174(1):522–538, 2003.

[18] X. Lu, J.-S. Song, and K. Zhu. Inventory control with unobservable lost sales and bayesian updates. Working paper, 2006.

[19] N. Megow, M. Uetz, and T. Vredeveld. Models and algorithms for stochastic online scheduling. *Mathematics of Operations Research*, 31(3):513–525, 2006.

[20] N. Megow and T. Vredeveld. Approximation results for preemptive stochastic online scheduling. In *Y. Azar and T. Erlebach (eds): Proceedings of the 14th European Symposium on Algorithms (ESA 2006)*, pages 516–527, 2006. Lecture Notes in Computer Science; Vol. 4168.

[21] R.H. Möhring, A.S. Schulz, and M. Uetz. Approximation in stochastic scheduling:the power of lp-based priority policies. *Journal of ACM*, 46(6):924–942, 1999.

[22] Rolf H. Möhring, Franz Josef Radermacher, and Gideon Weiss. Stochastic scheduling problems I: General strategies. *ZOR – Zeitschrift für Operations Research*, 28:193–260, 1984.

[23] Rolf H. Möhring, Franz Josef Radermacher, and Gideon Weiss. Stochastic scheduling problems II: Set strategies. *ZOR – Zeitschrift für Operations Research*, 29(3):A65–A104, 1985.

[24] U. Rieder and J. Weishaupt. Customer scheduling with incomplete information. *Probability in the Engineering and Informational Sciences*, 9(2):269–284, 1995.

[25] M.H. Rothkopf. Scheduling with random service times. *Management Science*, 12(9):703–713, 1966.

[26] H. Scarf. Bayes solutions of the statistical inventory problem. *The Annals of Mathematical Statistics*, 30(2):490–508, 1959.

[27] A.S. Schulz. New old algorithms for stochastic scheduling. In *Algorithms for Optimization with Incomplete Information*, 2005. Dagstuhl Seminar Proceedings, Vol. 05031.

[28] M. Skutella and M. Uetz. Stochastic machine scheduling with precedence constraints. *SIAM Journal on Computing*, 34(4):788–802, 2005.

[29] G. Weiss. Approximation results in parallel machines stochastic scheduling. *Annals of Operations Research*, 26(1):195–242, 1990.

[30] G. Weiss. Turnpike optimality of Smith's rule in parallel machines stochastic scheduling. *Mathematics of Operations Research*, 17(2):255–270, 1992.