



Performance of the smallest-variance-first rule in appointment sequencing

LSE Research Online URL for this paper: <http://eprints.lse.ac.uk/104585/>

Version: Accepted Version

Article:

de Kemp, Madelon A, Mandjes, Michel and Olver, Neil (2020) Performance of the smallest-variance-first rule in appointment sequencing. *Operations Research*. ISSN 0030-364X (In Press)

Reuse

Items deposited in LSE Research Online are protected by copyright, with all rights reserved unless indicated otherwise. They may be downloaded and/or printed for private study, or other acts as permitted by national copyright laws. The publisher or other rights holders may allow further reproduction and re-use of the full text version. This is indicated by the licence information on the LSE Research Online record for the item.

PERFORMANCE OF THE SMALLEST-VARIANCE-FIRST RULE IN APPOINTMENT SEQUENCING

MADELON A. DE KEMP, MICHEL MANDJES, NEIL OLVER

ABSTRACT. A classical problem in appointment scheduling, with applications in health care, concerns the determination of the patients' arrival times that minimize a cost function that is a weighted sum of mean waiting times and mean idle times. One aspect of this problem is the *sequencing problem*, which focuses on ordering the patients. We assess the performance of the *smallest-variance-first* (SVF) rule, which sequences patients in order of increasing variance of their service durations. While it is known that SVF is not always optimal, it has been widely observed that it performs well in practice and simulation. We provide a theoretical justification for this observation by proving, in various settings, quantitative *worst-case* bounds on the ratio between the cost incurred by the SVF rule and the minimum attainable cost. We also show that, in great generality, SVF is asymptotically optimal, i.e., the ratio approaches 1 as the number of patients grows large. While evaluating policies by considering an approximation ratio is a standard approach in many algorithmic settings, our results appear to be the first of this type in the appointment scheduling literature.

SUBJECT CLASSIFICATION. Health care: appointment scheduling. Scheduling: stochastic appointment sequencing. Optimization: approximation algorithm.

AREA OF REVIEW. Stochastic Models.

AFFILIATIONS. M. de Kemp and M. Mandjes are with Korteweg-de Vries Institute, University of Amsterdam. N. Olver is with the Department of Econometrics & Operations Research, Vrije Universiteit Amsterdam, and is also affiliated with CWI, Amsterdam. The research for this paper is partly funded by the NWO Gravitation Programme NETWORKS, Grant Number 024.002.003 (de Kemp, Mandjes) and NWO Vidi grant 016.Vidi.189.087 (Olver).

1. INTRODUCTION

Setting up appointment schedules plays an important role in health care and various other domains. The main challenge lies in efficiently running the system, but at the same time providing the customers an acceptable level of service. The service level can be expressed in terms of the waiting times the customers are facing, and the system efficiency in terms of the service provider’s idle time. The problem of generating an optimal schedule is generally formulated as minimizing a cost function (or simply “cost”) that is a weighted average of the expected idle time and the expected waiting times. As most literature on this topic focuses on applications in health care, we refer throughout this paper to customers as patients, and to the server as the doctor.

The problem of scheduling appointments can be split into two parts: one needs to determine the amount of time scheduled for each appointment, and one needs to determine in which order the patients should arrive. These problems are usually referred to as the *scheduling* problem and *sequencing* problem, respectively. This paper will primarily focus on the sequencing problem (but also includes results on the combined sequencing and scheduling problem), in a context with a single doctor seeing a sequence of patients. We impose the common assumptions that the service times of the patients form a sequence of independent random variables, while they arrive punctually at the scheduled times (which we will refer to as *epochs*). In this setting, a variety of techniques is available that determines for a given order the optimal arrival epochs; see, e.g., [2, 25] and references therein. However, much less is known about the efficient computation of “good” sequences. Already for a relatively modest number of patients, the number of possible sequences is huge, thus seriously complicating the search for an optimal order. An appointment scheduling review paper from 2017 [1] states that the optimal sequencing problem is one of the main open problems in the area:

“[...] one of the biggest challenges for future research is to find optimal (or near-optimal) solutions to more realistic appointment sequencing problems.”

A number of papers consider the sequencing (or combined sequencing and scheduling) problem and develop various stochastic programming models for it [3, 10, 27, 29]. However, the resulting optimization problems are very difficult to solve. Variants of the problem have been shown to be NP-hard [24, 29], indicating that this difficulty is inherent.

In a popular alternative approach one considers simple heuristics for the sequencing problem. The most frequently used heuristic is to order the patients by the variance of their service times, from smallest to largest. Throughout this paper we refer to this sequence as the SVF (smallest-variance-first) sequence. The intuition for using the SVF sequence is

that an unusually long service time in the beginning could cause many later patients to have to wait, and the SVF sequence aims to reduce the risk of this occurring. An additional appeal lies in the fact that SVF is simple to implement, as it only requires knowledge of the variances of the service times. Simulation experiments have revealed that the SVF rule typically performs remarkably well. In some cases it can be formally shown to be optimal; for instance (imposing some distributional assumptions) in the case of two patients [14, 39]. Recently, however, Kong et al. [24] provided instances showing SVF need not be optimal, even for cases involving relatively simple service times (e.g., uniform or lognormal). Despite the SVF sequence appearing promising in simulations, little is known about its theoretical performance, or of any other simple heuristic for that matter.

In this paper, we propose a new direction of research for the sequencing problem: *finding sequences that provably perform well*. Instead of finding an optimal sequence, such research aims at finding performance bounds on easily-computed sequences. Considering previous research, the SVF sequence is the obvious candidate for such an easily-computed and well-performing sequence, and will therefore be our focus. The precise quantity of interest to us will be the *ratio* of the cost of the schedule coming from the SVF sequence, and the cost of the schedule coming from the optimal sequence.

Our main goal in this paper is to prove upper bounds on this ratio – known as the *approximation ratio* – in various settings. This direction of study is very standard in the algorithmic community when considering intractable (NP-hard) problems, for example in machine scheduling (see [16, 17, 35] and references therein). However, it has not been studied in the appointment sequencing context. Note that for *typical* problem instances the SVF sequence could perform significantly better than suggested by an upper bound on the approximation ratio, as the bound must also hold for *worst-case* instances.

1.1. Main contributions. In the first part of the paper we concentrate exclusively on the effect of the sequence, using the simplest choice of schedule: each patient is assigned a slot of length equal to its mean service time. In other words, the arrival time of any patient is set equal to the sum of the mean service times of all preceding patients. This is certainly not the optimal solution to the scheduling problem, in the sense of minimizing the cost function introduced above, but it has the advantage of being very simple and easily applicable, and also completely independent of the choice of tradeoff in the cost function between doctor idle time and patient waiting time. Owing to these attractive properties, this “mean-based” type of schedule is a commonly used approach in practice, as was stated in, e.g., the survey paper [1] and in [12].

We start, in Section 2, by arguing that without any restrictions on the service-time distributions, no bound on the approximation ratio of SVF is possible, both under mean-based schedules and under optimally-spaced schedules (i.e., schedules in which the arrival epochs are chosen so as to minimize the cost function). We do so by constructing an example involving only two patients, but in which the service-time distributions are rather artificial. Then we present the ordering assumption on the service-time distributions that will be required for most of our results. Importantly, a large class of distributions meets this assumption, including e.g. the exponential distribution, but also the lognormal distributions that are frequently used in the health care context (such as the ones identified in [6]).

For the mean-based scheduling rule, we have performed an extensive set of numerical experiments for exponential and lognormal service time distributions, to gain insight into the performance of the SVF rule for typical problem instances. These experiments indicate that the SVF sequence performs very well for these distributions, with approximation ratios being below 1.01. However, an example at the start of Section 3 indicates that there are instances that are covered by the ordering assumption where the approximation ratio is 1.52, implying that it is not possible to prove a better bound than 1.52 if this assumption is in place.

Section 3 focuses on the mean-based scheduling rule. Under the ordering assumption we prove that the approximation ratio of SVF is at most 2 for symmetric service-time distributions, and at most 4 in general. In other words, we show that *for all instances* (i.e., for all numbers of patients and all service-time distributions satisfying the assumption imposed) the SVF cost is at most four times the optimal cost. We also consider two special cases:

- Service times are evidently nonnegative, but one could consider the situation that normal distributions are used as an approximation of the actual distributions of service times. In Section 3.2, we prove that then the approximation ratio is at most $4(\sqrt{2} - 1) \approx 1.6569$. While we do not believe that our result here is sharp, it indicates that the performance of SVF for well-behaved service-time distributions is most likely substantially better than suggested by the bounds 2 and 4 mentioned above.
- In Section 3.3 we bridge the gap between the upper bound of 2 for symmetric distributions and the general upper bound of 4, by developing a method that isolates the effect of asymmetry. For the lognormal distributions fitted to real data in [6], this method results in an approximation ratio of at most 3.43.

The problem of finding the optimal sequence becomes increasingly difficult as the number of patients grows larger, and thus the need for considering heuristics such as the SVF rule is more important when a substantial number of patients is involved. In Section 4, we show for mean-based schedules that as the number of patients grows large, the approximation ratio

tends to 1. This result requires only a very weak assumption on the service-time distributions (the ordering assumption is not needed here). The important practical implication of this result is that SVF is close to optimal in settings where the number of patients is substantial.

In Section 5 we shift our attention to the performance of SVF when optimally-spaced schedules (i.e., with arrival epochs minimizing the cost function), rather than mean-based schedules, are being used. In this setting, our numerical experiments indicate that the SVF sequence tends to be optimal for exponential and lognormal service time distributions; we do however present an example satisfying the ordering assumption in which SVF is no longer optimal.

We proceed by proving bounds on the approximation ratio of the SVF rule for the combined sequencing and scheduling problem. Here, we wish to compare a heuristic for this combined problem to the overall optimal schedule, over all possible sequences and schedules. Observe that the simple mean-based scheduling rule may in general lead to high cost, because waiting times could easily propagate. We therefore consider a simple alternative scheduling rule, suggested by Charnetski [8]: the slot assigned to a patient is equal to its mean service time, plus some multiple α of the standard deviation of its service time (where this α is optimized). Again under some assumptions, we show that this scheduling rule, combined with the SVF sequencing rule, yields a cost that is (relative to the optimal cost) off by at most a constant factor. Because of its frequent use in health care contexts [6, 22], we pay special attention to the case of lognormally distributed service times. Using a slightly different scheduling heuristic (the interarrival time being a multiple of the mean service time), we find an upper bound on the approximation ratio. Applying this result to the data in Çayırılı et al. [6], we find an upper bound of 2.90 in the case that in the cost function the waiting and idle times are equally important.

We finish Section 5 by giving an example to demonstrate that for optimally-spaced schedules the SVF sequence is no longer asymptotically optimal as the number of patients tends to infinity, in contrast to the result for mean-based schedules that was presented in Section 4.

1.2. Further related work. We proceed by providing an account of the related literature, without attempting to give a full overview; for more extensive reviews on the appointment scheduling and sequencing literature, we refer the reader to, e.g., Ahmadi-Javid et al. [1], Çayırılı and Veral [5], and Gupta and Denton [15].

As we mentioned, Kong et al. [24] gave examples showing that SVF is not in general optimal. In some very specific cases, optimality of SVF has been demonstrated. For only

two patients, the SVF sequence is optimal when the service times are both exponentially distributed or both uniformly distributed [39], or more generally, when the two service times are comparable according to a certain convex ordering [14]. For three patients, Kong et al. [24] find sufficient conditions for the SVF sequence to be optimal, when the time scheduled for each appointment is equal to the mean service time. (We have verified that this result can be extended to four patients using the same methods.)

Kemper et al. [19] analyze a sequential optimization approach, meaning that the arrival time of a patient is optimized without taking into account its impact on later patients. They show that under this rather different notion of optimality, and if the service times come from the same scale-family, then SVF does provide the best ordering.

One line of research focuses on comparing various sequencing heuristics (including SVF) through simulation. Denton et al. [11] consider a model similar to ours, and discuss the effectiveness of a number of simple sequencing heuristics using simulation, based on real surgery data. The SVF heuristic performed best of all the heuristics they considered. Mak et al. [27] consider a model where waiting time costs may be different for different patients; by relying on tractable approximations, they also find that SVF performs well. Klassen and Rohleder [22] and Rohleder and Klassen [33] consider an appointment scheduling model where not all patient information is known in advance; rather, patients must be scheduled as they call in to make an appointment (and so without information about patients who call later). Once again, it was empirically found that it worked best to put patients with low-variance service times early in the schedule.

A number of papers model variants of the combined sequencing and scheduling problem as stochastic integer or linear programming problems. Solving these programs is very challenging however, and generally exact results were only obtained for small instances. Works along these lines include Denton and Gupta [10], Mancilla and Storer [29] and Berg et al. [3]. For larger instances, it was necessary to resort to heuristics such as SVF for the sequencing problem. We mention Vanden Bosch and Dietz [37] who propose instead a local search heuristic to iteratively improve the sequence by finding pairs of patients who can be swapped to reduce the cost.

There are also a number of papers which take a robust optimization approach [23, 28, 30]. Here, instead of working with explicitly given service-time distributions, the goal is to find a schedule minimizing the worst-case expected cost given only that the distributions meet certain constraints (such as certain given moments). The main advantage of robust optimization is that only the constraints are needed rather than full distributional information. Note the

difference with our approach: robust optimization minimizes the *cost* of solution for the worst-case distributions satisfying the constraints, while our approach bounds the *approximation ratio*. The bound on the approximation ratio indicates the performance compared to optimal for any problem instance. In contrast, a robust optimization solution might not be anywhere near optimal for a typical problem instance, but will be good for the worst-case instances. Robust optimization results and bounds on approximation ratios could thus be seen as complementary results.

Among the works on robust optimization, that by Mak et al. [28] is most relevant to us, as they prove (under mild assumptions) that in their framework SVF is optimal. In their setup, the joint distribution of the service times could be any distribution matching known moments for individual service times (e.g., means and variances). However, the worst-case distributions corresponding to the optimal interarrival times are typically highly correlated; these results do not carry over to a setup in which independence is assumed. Kong et al. [23] consider a model in which not only the means and variances but also the covariances are specified. However, only the scheduling problem is discussed, and sequencing is not taken into account. Mittal et al. [30] discuss another robust model, in which each service time can take any value in a given interval. They observe that a $(1 + \epsilon)$ -approximation algorithm for the combined scheduling and sequencing problem can be obtained, for any $\epsilon > 0$.

Finally, we would like to point out the relation with the area of machine scheduling (see the book by Pinedo [31] for more background). The main difference between machine scheduling and appointment scheduling is that in the former the arrival times of the jobs (being the patients in our setting) are given, while in the latter these are decision variables. The machine scheduling problem that is most closely related to the setup that we consider, can be found in Guda et al. [13]. In the problem considered there, the due dates and sequence of jobs are to be decided, in order to minimize a weighted average of expected earliness and tardiness around the due dates. It is shown that the SVF rule is optimal, under a specific assumption on the service times of the jobs. It should be noted, though, that in the model considered in Guda et al. [13] all jobs are present from the start, implying that there is no idle time. Compared to our model, this greatly simplifies the evaluation of the cost function, thus facilitating finding an optimal solution.

2. MODEL AND PRELIMINARIES

Throughout this paper we consider a problem instance with n patients, numbered 1 up to n . We denote the service time of patient i in this problem instance by B_i , which has mean μ_i and variance σ_i^2 ; we assume that B_1, \dots, B_n are independent random variables.

As pointed out in the introduction, one should distinguish between the scheduling problem and the sequencing problem. The sequencing problem, on which we primarily focus, is to decide which patient is assigned which appointment slot (given a certain scheduling rule). The sequence is denoted by a permutation $\tau \in \mathcal{S}_n$ (where \mathcal{S}_n denotes the set of all permutations on $\{1, \dots, n\}$). The value $\tau(i)$ will denote the index of the patient that is assigned to appointment slot i . The scheduling problem is to decide the interarrival times between patients, given the sequence in which they arrive. We use x_j to denote the interarrival time between patient j and the next patient, i.e., the length of the appointment slot reserved for patient j . The vector $\mathbf{x} = (x_1, x_2, \dots, x_n)$ will be referred to as the schedule.

2.1. Performance metrics. We proceed by introducing the cost function we work with in this paper, which is based on the patients' waiting times and the doctor's idle times. Let W_i denote the waiting time of the patient in appointment slot i . Let I_i be the idle time before the start of appointment slot i after the previous patient has been served. Given a sequence τ and interarrival times x_i , the waiting times and idle times satisfy the Lindley recursions [26]:

$$W_{i+1} = (W_i + B_{\tau(i)} - x_{\tau(i)})^+, \quad I_{i+1} = (W_i + B_{\tau(i)} - x_{\tau(i)})^-, \quad (1)$$

where we use the compact notation $x^+ := \max\{0, x\}$ and $x^- := \max\{0, -x\}$.

We use a parameter $\omega \in (0, 1)$ to indicate the relative importance of idle time and waiting time. As a cost function, we seek to minimize

$$C_\omega(\tau, \mathbf{x}) := \omega \sum_{i=1}^n \mathbb{E}I_i + (1 - \omega) \sum_{i=1}^n \mathbb{E}W_i, \quad (2)$$

a weighted average of the expected total idle time and expected total waiting time. Observe that this cost function depends on the sequence τ , on the schedule \mathbf{x} , and on the patient service-time distributions $\mathbf{B} = (B_1, \dots, B_n)$. We generally suppress the dependence on \mathbf{B} , but we may write $C_\omega(\mathbf{B}, \tau, \mathbf{x})$ if we wish to be explicit. As an aside, we mention that an approach to estimate ω in a practical context can be found in [32].

Throughout this paper, without loss of generality we let the patients be indexed such that $\sigma_1^2 \leq \sigma_2^2 \leq \dots \leq \sigma_n^2$. The SVF sequence is then the sequence given by the identity permutation id given by $\text{id}(i) = i$, which we compare with the optimal sequence, i.e., the sequence that minimizes (2). To compare these sequences, we study the ratio between the cost functions under the SVF sequence and the optimal sequence. If this ratio is close to 1, then this is evidence that the SVF sequence performs well.

In this paper we specifically focus on two settings. In the first place we consider the performance of SVF for *mean-based schedules*, which are given by $\mathbf{x} = \boldsymbol{\mu}$. We then consider

the approximation ratio

$$\varrho_\omega(\mathbf{B}) := \frac{C_\omega(\mathbf{B}, \text{id}, \boldsymbol{\mu})}{\min\{C_\omega(\mathbf{B}, \tau, \boldsymbol{\mu}) : \tau \in \mathcal{S}_n\}}.$$

We will write just ϱ_ω when the service-time distributions under consideration are unambiguous. In the second place, we consider the performance of SVF for *optimally-spaced schedules*. In this context we compare the SVF sequence id along with a given schedule \mathbf{x} with the optimal combination of sequence and schedule (i.e., a sequencing τ and schedule \mathbf{y} that minimize the cost function). This means that here we consider the approximation ratio

$$r_\omega(\mathbf{B}, \mathbf{x}) := \frac{C_\omega(\mathbf{B}, \text{id}, \mathbf{x})}{\min\{C_\omega(\mathbf{B}, \tau, \mathbf{y}) : \tau \in \mathcal{S}_n, \mathbf{y} \in \mathbb{R}_+^n\}}.$$

Once again, we will omit \mathbf{x} and \mathbf{B} when their choice is unambiguous.

The main objective of this paper is to prove, under specific assumptions, upper bounds on ϱ_ω and r_ω . Such an upper bound then guarantees that the SVF sequence always has a cost function of at most such an upper bound times the optimal cost function. We also show, under a mild condition, that $\varrho_\omega(\mathbf{B})$ converges to 1 as the number of patients tends to infinity, thus proving that the SVF sequence is asymptotically optimal when mean-based schedules are used.

Remark 2.1. Service times are inherently nonnegative, but our framework (based on the Lindley recursions (1)) carries over to situations where the B_i are allowed to take negative values. This might be useful if the true distributions of service times can be approximated using distributions that can take negative values (with some small probability), for example normal distributions. If such distributions that can take negative values form a good fit to the data in some application, the theoretical performance of the SVF rule for these distributions gives a meaningful indication for the performance of the SVF rule in this application.

2.2. Preliminaries. We need the following well-known results concerning the waiting and idle times. It follows by iterating the Lindley recursion (1) that W_{k+1} is the maximum of a random walk with steps $B_{\tau(k)} - x_{\tau(k)}, B_{\tau(k-1)} - x_{\tau(k-1)}, \dots, B_{\tau(1)} - x_{\tau(1)}$, that is,

$$W_{k+1} = \max \left\{ 0, \max \left\{ \sum_{i=j}^k B_{\tau(i)} - x_{\tau(i)} : j \in \{1, \dots, k\} \right\} \right\}. \quad (3)$$

In the setting of mean-based schedules $\mathbf{x} = \boldsymbol{\mu}$, we introduce the notation $X_i := B_i - \mu_i$, and the random walk

$$S_j := \sum_{i=1}^j X_{\tau(k-i+1)}.$$

Note that S_j also depends on k and τ , but we keep this dependence implicit to ease the notation. We then find for the mean-based schedule that

$$W_{k+1} = \max\{0, S_1, \dots, S_k\}. \quad (4)$$

Note that (1) implies $W_{i+1} - I_{i+1} = W_i + B_{\tau(i)} - x_{\tau(i)}$. Summing over i we find the identity

$$\sum_{i=1}^n I_i + \sum_{i=1}^n B_i = \sum_{i=1}^{n-1} x_{\tau(i)} + W_n + B_{\tau(n)}, \quad (5)$$

which corresponds to the total time until all n patients have been served (the ‘‘session end time’’). For a given schedule, this relation can be used to express the expected total idle time in terms of the expected waiting time of the last patient. Therefore, we can focus on the waiting times, and derive results for the idle time from (5).

The following example shows that, if one does not impose any conditions on the service-time distributions, one can construct examples in which the approximation ratio is unbounded.

Example 2.2. Suppose we have two patients, and the service time of patient i is given by

$$B_i = \begin{cases} \mu_i + \frac{1}{a_i} & \text{with probability } a_i^2 \\ \mu_i - \frac{1}{a_i} & \text{with probability } a_i^2 \\ \mu_i & \text{with probability } 1 - 2a_i^2, \end{cases} \quad (6)$$

for some values $\mu_i > 0$ and $a_i \leq 1/\sqrt{2}$. Then $\mathbb{E}B_i = \mu_i$, and $\text{Var}B_i = 2$, and so either of the two possible sequences could be considered the SVF sequence. We take the SVF sequence to be given by $\tau(i) = i$; one can of course slightly perturb the distributions so that this is the *unique* SVF ordering.

Suppose we use the mean-based schedule given by $\mathbf{x} = \boldsymbol{\mu}$. The cost function for the SVF sequence $\tau(i) = i$ is then

$$\omega \mathbb{E}I_2 + (1 - \omega) \mathbb{E}W_2 = \omega \mathbb{E}(B_1 - \mu_1)^- + (1 - \omega) \mathbb{E}(B_1 - \mu_1)^+ = \omega a_1 + (1 - \omega) a_1 = a_1.$$

In the same way, the cost function for the other sequence is equal to a_2 . If we take a_1 to be bigger than a_2 , we conclude $\rho_\omega = a_1/a_2$, which can be arbitrarily large. As a consequence, it is necessary to impose assumptions on the service-time distributions in order to bound ρ_ω . The construction can easily be extended to one corresponding to any larger number of patients, by introducing additional patients with deterministic service times. Therefore, this example also shows that some assumption is necessary for the asymptotic result in Section 4.

The example applies also when an optimally-spaced, rather than mean-based, scheduling rule is used. Fixing $\omega = \frac{1}{2}$, for two patients with service times as in (6), the cost function is

$$\frac{1}{2}\mathbb{E}(B_{\tau(1)} - x_{\tau(1)})^- + \frac{1}{2}\mathbb{E}(B_{\tau(1)} - x_{\tau(1)})^+.$$

By Lemma C.4, the $x_{\tau(1)}$ that minimizes this cost function is given by $x_{\tau(1)} = \mu_{\tau(1)}$, which is the mean of $B_{\tau(1)}$. So the situation is unchanged, and also for the optimal scheduling rule no bound on the approximation ratio can be found without imposing further assumptions.

Example 2.2 shows that it will be necessary to impose assumptions on the service-time distributions in order to be able to bound the approximation ratio. Notice that the three-point distributions used in the example can be considered as artificial, in the sense that they substantially differ from distributions that are frequently used in the context of health care. One would like to use an assumption under which the approximation ratio can be bounded, but which does not rule out relevant, commonly used service-time distributions.

Not all service time distributions might be sufficiently comparable through their variance alone. We will need the concepts of a *convex ordering* and *dilation ordering* on random variables. More background on the convex ordering, dilation ordering and related concepts can be found in [34].

Definition 2.3. The random variable A is said to be smaller in the convex order than the random variable B if $\mathbb{E}\phi(A) \leq \mathbb{E}\phi(B)$ for all convex functions $\phi : \mathbb{R} \rightarrow \mathbb{R}$ for which the expectations exist. This will be denoted by $A \leq_{\text{cx}} B$. If $A - \mathbb{E}A \leq_{\text{cx}} B - \mathbb{E}B$, then A is said to be smaller than B in the dilation order, denoted as $A \leq_{\text{dil}} B$.

Note that $A \leq_{\text{cx}} B$ implies $A \leq_{\text{dil}} B$, and $A \leq_{\text{dil}} B$ implies $\text{Var}A \leq \text{Var}B$. Throughout this paper, unless stated otherwise, we will make the following assumption on the service time distributions.

Assumption 2.4 (dilation ordering). We have $B_1 \leq_{\text{dil}} B_2 \leq_{\text{dil}} \dots \leq_{\text{dil}} B_n$.

We remark that this is the condition under which [14] proves optimality of SVF for two patients. Note also that this assumption implies $\text{Var}(B_1) \leq \dots \leq \text{Var}(B_n)$. Examples of instances satisfying this assumption include all B_i having exponential distributions (by Theorem 3.A.18 in [34]), and all B_i having lognormal distributions such that both $\mathbb{E}[\ln B_1] \leq \dots \leq \mathbb{E}[\ln B_n]$ and $\text{Var}(\ln B_1) \leq \dots \leq \text{Var}(\ln B_n)$, as proved in Appendix D. The following lemma, taken from [34], is useful when checking whether given random variables satisfy a convex order.

Lemma 2.5. *The random variables A and B satisfy $A \leq_{\text{cx}} B$ if and only if there exists a coupling $\widehat{A} \stackrel{\text{d}}{=} A$ and $\widehat{B} \stackrel{\text{d}}{=} B$ such that $\mathbb{E}[\widehat{B}|\widehat{A}] = \widehat{A}$.*

3. BOUNDS ON PERFORMANCE UNDER MEAN-BASED SCHEDULES

In this section we study ϱ_ω , the approximation ratio under the mean-based schedule. To gain some insight as to what approximation ratio we can expect for typical problem instances, we start by discussing the output of a series of numerical experiments, described in more detail in Appendix E. It should be kept in mind that, as finding the optimal sequences in principle requires comparing all $n!$ candidates and is therefore computationally very expensive, such experiments can only be done for relatively small values of n . We recall that this is also the main reason why we study approximation ratios: computing the optimal sequence for larger instances is prohibitively slow, which explains why it is useful to gain insight into the performance of heuristics.

We considered exponential and lognormal service-time distributions. The exponential distributions have the advantage of an efficiently computable cost function using the machinery developed in [38]. For exponentially distributed service times the main conclusion is that the SVF sequence typically performs within 1% of optimal ($\varrho_\omega \leq 1.01$). Lognormal distributions are more realistic from a practical perspective [6, 22]. No exact computational scheme being available, we resorted to a discrete-time approximation to evaluate the cost; see Appendix E for more background on the implementation. In the experiments we performed the approximation ratios were uniformly below 1.01.

The numerics discussed in the previous paragraph suggest that SVF typically performs very well. However, the service-time distributions featuring in the next example lead to a much higher approximation ratio.

Example 3.1. Suppose the service time of patient i is given by

$$B_i = \begin{cases} m+1 & \text{with probability } m/(m+1) \\ 0 & \text{with probability } 1/(m+1), \end{cases}$$

for $i < n-1$, and $B_{n-1} \stackrel{\text{d}}{=} B_n \stackrel{\text{d}}{=} KB_1$, for some parameters m and K . We set $\omega = 1$. We compare the SVF sequence with the sequence τ given by $\tau(1) = n-1$, $\tau(n-1) = 1$ and $\tau(i) = i$ for $i \neq 1, n-1$. For $n = 10$ patients, $m = 10$ and $K = 10$ this results on a lower bound of 1.29 on the approximation ratio ϱ_1 . We can, however, obtain considerably higher approximation ratios: for example for $n = 50\,000$ patients, $m = 500$ and $K = 5\,000$ we find that the approximation ratio is larger than 1.52. Further experiments revealed that similar examples for even bigger values of n only marginally increase this lower bound. Note

that this problem instance satisfies Assumption 2.4, the dilation ordering assumption. This means that it is impossible to prove a better bound on ϱ_ω than 1.52 without imposing further assumptions.




Figure 1. The approximation ratio ϱ_1 as a function of the parameter p of Bernoulli random variables. Each instance consists of 200 patients, with 198 of them having a Bernoulli distributed service time with the probability to take the larger value being equal to p , whereas the other two patients have a service time distributed as 50 times the same Bernoulli distribution.

The distributions in Example 3.1 could be considered as quite extreme: they are two point distribution that almost always take a value slightly larger than the mean, and very rarely take a value much less than the mean. Our numerics indicate that this behavior seems to be necessary to obtain a large gap between the performance of SVF and the optimal sequence. To further study the effect of such “skewness”, we performed the following experiment. We evaluate the approximation ratio ϱ_1 for instances with $n = 200$ patients, with $n - 2$ of them having a Bernoulli distributed service time with the probability to take the larger value being equal to p , whereas the other two patients have a service time distributed as 50 times the same Bernoulli distribution. Figure 1 shows the approximation ratio ϱ_1 as a function of p . In this case we are able to compute the optimal sequence, despite the relatively large value of n , due to the fact that many patients have the same service-time distribution, so that the number of possible distinct sequences is drastically reduced. It can be seen that SVF is optimal for $p \leq 1/2$, and then grows as p increases, until it decreases again for p very close to one.

The above numerics suggest that SVF performs very well with both exponential and lognormal distributions. These distributions have tails only to the right, and *not* to the left. If one “flips” the lognormal distributions, i.e., replace each service time B_i with $-B_i$, we can obtain substantially larger approximation ratios; for one particular instance, we find 1.0345

Theorem	Assumptions on service time distributions	Bound on ϱ_ω
3.3	dilation ordering and symmetry	2
3.4	dilation ordering	4
3.8	normal distributions	$4(\sqrt{2} - 1) \approx 1.66$
3.11	dilation ordering	between 2 and 4

Table 1. Summary of main results of Section 3

as opposed to 1.0086 (see Appendix E for more details). These examples suggest that some measure of skewness is expected to be a key driver in whether SVF performs near-optimally. As this experiment intends to investigate the effect of skewness on the approximation ratio we have not pursued using maximally realistic service times. Indeed, $-B_i$ is not a realistic service time as it is always negative. However, we can also use $c - B_i$ instead of $-B_i$ without affecting the cost per sequence, and choose c so large that $c - B_i$ will be positive with arbitrarily large probability. Neglecting the probability mass $c - B_i$ has on the negative numbers will not have any substantial effect on the cost.

Now that we have developed some intuition on the performance of SVF, we concentrate in the rest of this section on obtaining upper bounds for ϱ_ω . This amounts to giving an upper bound on the cost when using the SVF sequence, and a lower bound on the cost that is valid for any sequence, hence also for the optimal sequence. The waiting times and idle times under the SVF sequence will be denoted by W_i^{SVF} and I_i^{SVF} respectively, and the waiting times and idle times under the optimal sequence will be denoted by W_i^{OPT} and I_i^{OPT} .

The remainder of this section is structured as follows. In Section 3.1 we prove results under rather general assumptions: Theorem 3.3 and Theorem 3.4. More specifically, these theorems give bounds on the approximation ratio ϱ_ω , when we assume that the service times are symmetrically distributed and follow a dilation order (Theorem 3.3), and when we only assume that they follow a dilation order (Theorem 3.4). In Section 3.2 we consider the special case of normally distributed service times. Theorem 3.8 gives an improved bound on ϱ_ω in this case. In Section 3.3 we discuss a method for improving numerically upon the bound of Theorem 3.4; informally, the more symmetric the service-time distributions, the closer the resulting bound is to the value stated in Theorem 3.3. Table 1 provides a summary of the results covered by this section.

3.1. General results. Recall that we impose Assumption 2.4 (the dilation ordering assumption). For the bound on ϱ_ω that is presented in Theorem 3.3 we also make the following assumption.

Assumption 3.2 (symmetry). *The B_i have symmetric distributions around their mean.*

Examples of instances satisfying both the ordering and symmetry assumption include all B_i having normal distributions, all B_i having uniform distributions and all B_i having Laplace distributions. For all three examples, the ordering assumption follows from [34], Theorem 3.A.18. Note that Example 2.2 consists of symmetric service time distributions, so even under the symmetry assumption it is necessary to impose the dilation ordering assumption in order to obtain any bound on ρ_ω .

In this section, we prove the following theorems. Recall that Example 3.1 discusses an instance satisfying the dilation ordering assumption where $\rho_\omega \geq 1.52$.

Theorem 3.3. *Under the dilation ordering and symmetry assumptions, we have $\rho_\omega \leq 2$.*

Theorem 3.4. *Under the dilation ordering assumption, we have $\rho_\omega \leq 4$.*

A first key insight is that it suffices to prove bounds on $\mathbb{E}W_{k+1}$ for given k when the first k slots are constrained to contain patients $1, \dots, k$. This is made explicit in the next lemma.

Lemma 3.5. *Let $\mathbb{E}W_{k+1}^{\text{OPT}'}$ denote the expected waiting time of the patient in appointment slot $k+1$, under the sequence that minimizes this expected waiting time, subject to the constraint that $\tau(i) \leq k$ for all $i = 1, \dots, k$, i.e. the first k patients are assigned to the first k slots. Suppose $\mathbb{E}W_{k+1}^{\text{SVF}}/\mathbb{E}W_{k+1}^{\text{OPT}'}$ $\leq \rho'$ for all k . Then, under the dilation ordering assumption, $\rho_\omega \leq \rho'$.*

Proof. Taking expectations in (5) and using that $x_i = \mu_i$, we find that $\sum_{i=1}^n \mathbb{E}I_i = \mathbb{E}W_n$. Hence, the cost function (2) equals $\omega \mathbb{E}W_n + (1 - \omega) \sum_{i=1}^n \mathbb{E}W_i$. Our goal is thus to bound the ratio

$$\rho_\omega = \frac{\omega \mathbb{E}W_n^{\text{SVF}} + (1 - \omega) \sum_{i=1}^n \mathbb{E}W_i^{\text{SVF}}}{\omega \mathbb{E}W_n^{\text{OPT}} + (1 - \omega) \sum_{i=1}^n \mathbb{E}W_i^{\text{OPT}}}. \quad (7)$$

Now note that $W_{k+1} = \max\{0, S_1, \dots, S_k\}$ is a convex function in each of the X_i , as it is the maximum of linear functions in X_i . Under Assumption 2.4 this implies that $\mathbb{E}W_{k+1}^{\text{OPT}} \geq \mathbb{E}W_{k+1}^{\text{OPT}'}$, as each step X_i with $i > k$ can be replaced by some X_j with $j \leq k$ and as $X_j \leq_{\text{cx}} X_i$ this lowers the expected waiting time. Now

$$\mathbb{E}W_{k+1}^{\text{SVF}}/\mathbb{E}W_{k+1}^{\text{OPT}} \leq \mathbb{E}W_{k+1}^{\text{SVF}}/\mathbb{E}W_{k+1}^{\text{OPT}'} \leq \rho',$$

and so also the ratio in (7) is bounded by ρ' , which completes the proof. \square

The following lemma is a second key insight to be used in the proofs of both Theorem 3.3 and Theorem 3.4.

Lemma 3.6. *Under the symmetry assumption, the random variable W_{k+1} is stochastically dominated by $|S_k|$, and thus*

$$\mathbb{E}W_{k+1} \leq \mathbb{E}|S_k| = \mathbb{E}|X_{\tau(1)} + X_{\tau(2)} + \cdots + X_{\tau(k)}|.$$

Proof. Recall that we have $W_{k+1} = \max\{0, S_1, \dots, S_k\}$ from (4). Under the symmetry assumption, the steps X_i of the random walk S have a symmetric distribution around zero, and hence the same is true for the S_i .

Let $T(a) = \inf\{i : S_i \geq a\}$, and note that $\mathbb{P}(W_{k+1} \geq a) = \mathbb{P}(T(a) \leq k)$. To bound this probability, we look at the random walk reflected in a after $T(a)$. This reflected process \widehat{S}_i is defined by

$$\widehat{S}_i = \begin{cases} S_i & \text{if } i < T(a) \\ 2a - S_i & \text{if } i \geq T(a). \end{cases} \quad (8)$$

We have $S_{T(a)} \geq a$, so $\widehat{S}_{T(a)} = 2a - S_{T(a)} \leq a \leq S_{T(a)}$. As the X_i have symmetric distributions, the increments of S_i and \widehat{S}_i for $i \geq T(a)$ have the same distribution. Therefore, we see that \widehat{S}_i is stochastically dominated by S_i , for every i . We conclude that $\mathbb{P}(\widehat{S}_k > a) \leq \mathbb{P}(S_k > a)$ for all a .

Now note that $W_{k+1} \geq a$ implies that either $S_k \geq a$ or $\widehat{S}_k = 2a - S_k > a$. As these are disjoint events we now have

$$\mathbb{P}(W_{k+1} \geq a) = \mathbb{P}(S_k \geq a) + \mathbb{P}(\widehat{S}_k > a) \leq \mathbb{P}(S_k \geq a) + \mathbb{P}(S_k > a) \leq \mathbb{P}(|S_k| \geq a).$$

This holds for any $a \geq 0$, so W_{k+1} is stochastically dominated by $|S_k|$, as was claimed. \square

Proof of Theorem 3.3. As $W_{k+1} = \max\{0, S_1, \dots, S_k\}$, we have

$$W_{k+1} \geq S_k^+ = (X_{\tau(1)} + \cdots + X_{\tau(k)})^+.$$

Note that $\tau(i) \leq k$ for all $i \leq k$ when we consider $\mathbb{E}W_{k+1}^{\text{OPT}'}$, so now

$$\mathbb{E}W_{k+1}^{\text{OPT}'} \geq \mathbb{E}(X_1 + \cdots + X_k)^+. \quad (9)$$

On the other hand, by Lemma 3.6,

$$\mathbb{E}W_{k+1}^{\text{SVF}} \leq \mathbb{E}|X_1 + \cdots + X_k| = 2\mathbb{E}(X_1 + \cdots + X_k)^+ \leq 2\mathbb{E}W_{k+1}^{\text{OPT}'}$$

As $\mathbb{E}W_{k+1}^{\text{SVF}}/\mathbb{E}W_{k+1}^{\text{OPT}'}$ is now bounded by 2, Theorem 3.3 follows from Lemma 3.5. \square

Proof of Theorem 3.4. Note that Lemma 3.5 and the lower bound (9) are valid without the symmetry assumption being needed. We therefore only need an upper bound on $\mathbb{E}W_{k+1}^{\text{SVF}}$.

Let X'_1, X'_2, \dots, X'_n have the same distributions as respectively X_1, X_2, \dots, X_n such that all these random variables are independent. Let W'_{k+1} be the maximum of the random walk with steps $X_k - X'_k, X_{k-1} - X'_{k-1}, \dots, X_1 - X'_1$. As

$$\mathbb{E}[X_i - X'_i | X_i] = X_i,$$

we see using Lemma 2.5 that $X_i \leq_{\text{cx}} X_i - X'_i$. Note that $W_{k+1} = \max\{0, S_1, \dots, S_k\}$ is a convex function in X_i , as it is the maximum of functions linear in X_i . Therefore, each time we replace a step X_i with a step $X_i - X'_i$ the expected maximum of the random walk will increase, so $\mathbb{E}W_{k+1}^{\text{SVF}} \leq \mathbb{E}W'_{k+1}$.

Now note that the steps $X_i - X'_i$ all have a symmetric distribution, so we can apply Lemma 3.6 to find

$$\begin{aligned} \mathbb{E}W_{k+1}^{\text{SVF}} &\leq \mathbb{E}W'_{k+1} \leq \mathbb{E}|X_1 + \dots + X_k - (X'_1 + \dots + X'_k)| \\ &\leq 2\mathbb{E}|X_1 + \dots + X_k| = 4\mathbb{E}(X_1 + \dots + X_k)^+ \leq 4\mathbb{E}W_{k+1}^{\text{OPT}'}. \end{aligned}$$

As $\mathbb{E}W_{k+1}^{\text{SVF}}/\mathbb{E}W_{k+1}^{\text{OPT}'}$ is now bounded by 4, the result follows from Lemma 3.5. \square

Remark 3.7. In case the scheduled session end time equals the expected total service time, the overtime reads

$$W_{n+1} = (W_n + X_{\tau(n)})^+,$$

which can also be included in the cost function. As such, overtime is handled similarly to waiting time, and consequently the results of Theorems 3.3 and 3.4 remain valid when some extra term $c\mathbb{E}W_{n+1}$ with $c > 0$ is added to the cost function.

3.2. Normally distributed service times. The results of Theorems 3.3 and 3.4 can be strengthened for specific service-time distributions. One such result is the following.

Theorem 3.8. *When the B_i are all normally distributed we have $\rho_\omega \leq 4(\sqrt{2} - 1)$.*

In order to prove Theorem 3.8, we need the following two lemmas, giving stronger bounds on $\mathbb{E}W_{k+1}^{\text{SVF}}$ and $\mathbb{E}W_{k+1}^{\text{OPT}'}$. The proofs of these lemmas, that hold for any symmetrically distributed service times, can be found in Appendix A.

Lemma 3.9. *Under the symmetry assumption,*

$$\mathbb{E}W_{k+1} \leq \mathbb{E}(X_{\tau(1)} + \dots + X_{\tau(k)})^+ + \mathbb{E}(X_{\tau(1)} + \dots + X_{\tau(k-1)})^+.$$

Lemma 3.10. *Under the symmetry assumption, for any ℓ ,*

$$\mathbb{E}W_{k+1} \geq \frac{1}{2} \left(\mathbb{E} (X_{\tau(1)} + \cdots + X_{\tau(k)})^+ + \mathbb{E} (X_{\tau(1)} + \cdots + X_{\tau(\ell)})^+ + \mathbb{E} (X_{\tau(\ell+1)} + \cdots + X_{\tau(k)})^+ \right).$$

Proof of Theorem 3.8. Note that normal distributions satisfy both the dilation ordering and symmetry assumption. Now the sum $X_1 + \cdots + X_i$ again has a normal distribution, with mean zero and variance $\Sigma_i^2 := \sigma_1^2 + \cdots + \sigma_i^2$. For the SVF sequence we now have, using Lemma 3.9, that

$$\mathbb{E}W_{k+1}^{\text{SVF}} \leq \frac{1}{\sqrt{2\pi}} (\Sigma_k + \Sigma_{k-1}). \quad (10)$$

Now we still need an expression for a lower bound on $\mathbb{E}W_{k+1}^{\text{OPT}'}$. Let $\tilde{\Sigma}_i^2 := \sigma_{\tau(1)}^2 + \cdots + \sigma_{\tau(i)}^2$ be the variance of $X_{\tau(1)} + \cdots + X_{\tau(i)}$. From Lemma 3.10 it then follows that

$$\mathbb{E}W_{k+1} \geq \frac{1}{2} \left(\tilde{\Sigma}_k + \tilde{\Sigma}_\ell + \sqrt{\tilde{\Sigma}_k^2 - \tilde{\Sigma}_\ell^2} \right).$$

Recall that $\mathbb{E}W_{k+1}^{\text{OPT}'}$ was the optimal expected waiting time when $\tau(i) \leq k$ whenever $i \leq k$. Therefore, we have $\tilde{\Sigma}_k = \Sigma_k$ and $\sigma_k^2 = \max\{\sigma_{\tau(1)}^2, \dots, \sigma_{\tau(k)}^2\}$. Now note that

$$\Sigma_k + \tilde{\Sigma}_\ell + \sqrt{\Sigma_k^2 - \tilde{\Sigma}_\ell^2}$$

is largest when $\tilde{\Sigma}_\ell^2$ is as close to $\frac{1}{2}\Sigma_k^2$ as possible. As σ_k^2 is largest of the $\sigma_{\tau(i)}^2$ with $i \leq k$, we can always choose ℓ such that

$$\frac{1}{2}\Sigma_{k-1}^2 \leq \tilde{\Sigma}_\ell^2 \leq \frac{1}{2}\Sigma_{k-1}^2 + \sigma_k^2.$$

This choice of ℓ provides us with the lower bound

$$\mathbb{E}W_{k+1}^{\text{OPT}'} \geq \frac{1}{2} \frac{1}{\sqrt{2\pi}} \left(\Sigma_k + \sqrt{\frac{1}{2}\Sigma_{k-1}^2} + \sqrt{\frac{1}{2}\Sigma_{k-1}^2 + \sigma_k^2} \right),$$

valid for any sequence. Comparing with (10), we obtain

$$\frac{\mathbb{E}W_{k+1}^{\text{SVF}}}{\mathbb{E}W_{k+1}^{\text{OPT}'}} \leq 2(\Sigma_k + \Sigma_{k-1}) / \left(\Sigma_k + \sqrt{\frac{1}{2}\Sigma_{k-1}^2} + \sqrt{\frac{1}{2}\Sigma_{k-1}^2 + \sigma_k^2} \right).$$

As $\Sigma_k^2 = \Sigma_{k-1}^2 + \sigma_k^2$, this fraction only depends on the relative size of Σ_{k-1}^2 compared to σ_k^2 . Suppose that $\Sigma_{k-1}^2 = c\sigma_k^2$, for some $c \geq 0$. Then $\Sigma_k^2 = (c+1)\sigma_k^2$, and the fraction becomes

$$\frac{2(\sqrt{c+1} + \sqrt{c})}{\sqrt{c+1} + \sqrt{\frac{1}{2}c} + \sqrt{\frac{1}{2}c + 1}} =: f(c).$$

It can easily be seen that f is increasing, and that $f(c) \rightarrow 4(\sqrt{2} - 1)$ as $c \rightarrow \infty$.

We now know that $\mathbb{E}W_{k+1}^{\text{SVF}}/\mathbb{E}W_{k+1}^{\text{OPT}'} \leq 4(\sqrt{2} - 1) \approx 1.6569$. By Lemma 3.5 the same is then also true for the cost function. This proves Theorem 3.8. \square

3.3. Numerically improving the bound of Theorem 3.4. Under the dilation ordering assumption, we have proved that $\rho_\omega \leq 4$, and we also proved that $\rho_\omega \leq 2$ when the service times have symmetric distributions. This suggests that one can find an upper bound on ρ_ω between 2 and 4 for service-time distributions that have some degree of symmetry, but are not fully symmetric.

Here we introduce a method to split the service time distributions into a symmetric and a nonsymmetric part, thus isolating the effect of the asymmetry on the upper bound. This can be used to numerically compute an upper bound on ρ_ω (lower than 4, that is) for given problem instances. In this section we do so for lognormal service time distributions that fit real data in [6]. We still impose the dilation ordering assumption.

We introduce the method for continuously distributed service times to simplify the exposition, noting that extending the method to non-continuous distributions is straightforward. Suppose X_i has density $f_i(x)$. We set $g_i(x) := \min\{f_i(x), f_i(-x)\}$, $h_i(x) := f_i(x) - g_i(x)$ and $p_i := \int_{\mathbb{R}} h_i(x) dx$. Then we let U_i be a random variable with density $g_i(x)/(1 - p_i)$. We let A_i be a random variable, independent of U_i , with density $h_i(x)/p_i$. Let J_i be a Bernoulli variable taking the value one with probability p_i , independent of A_i and U_i . We thus have

$$X_i \stackrel{d}{=} U_i(1 - J_i) + A_i J_i.$$

Note that this construction is such that U_i has a symmetric distribution around zero. One thus has that U_i corresponds to the symmetric part of X_i , whereas A_i to the nonsymmetric part. Note that $\mathbb{E}X_i = 0$ and $\mathbb{E}U_i = 0$, so we must have $\mathbb{E}A_i = 0$. Let A'_i have the same distribution as A_i , so that A'_i is independent of all the other random variables. Since $\mathbb{E}[A'_i J_i | J_i] = 0$, we have

$$\mathbb{E} [U_i(1 - J_i) + (A_i - A'_i)J_i | U_i(1 - J_i) + A_i J_i] = U_i(1 - J_i) + A_i J_i.$$

By Lemma 2.5 we conclude

$$X_i \leq_{\text{cx}} U_i(1 - J_i) + (A_i - A'_i)J_i.$$

As W_{k+1} is a convex function in each of the X_i , we can then replace each X_i by this upper bound in convex order to get an upper bound on $\mathbb{E}W_{k+1}^{\text{SVF}}$. Using Lemma 3.6, we then find the following upper bound:

$$\mathbb{E}W_{k+1}^{\text{SVF}} \leq \mathbb{E}|X_1 + \dots + X_k| + \mathbb{E}|A_1 J_1 + \dots + A_k J_k|. \quad (11)$$

Group	Mean	Standard deviation
Return	15.50	5.038
New	19.09	6.85

Table 2. Parameters of the lognormal distributions fitted by [6].

Note that the second term in the right-hand side of (11) can be numerically evaluated to any desired level of precision, e.g. by simulation. The resulting upper bound can thus be compared numerically to the lower bound $\mathbb{E}W_{k+1}^{\text{OPT}'} \geq \mathbb{E}(X_1 + \dots + X_k)^+$, for each k , valid under the dilation ordering assumption. Combining the above with Lemma 3.5, this leads to the bound given in the next theorem.

Theorem 3.11. *Under the dilation ordering assumption, we have*

$$\varrho_\omega \leq 2 + 2 \max\{\mathbb{E}|A_1 J_1 + \dots + A_k J_k| / \mathbb{E}|X_1 + \dots + X_k| : k = 1, \dots, n\}. \quad (12)$$

The more symmetric the service times and thus the random variables X_i , the smaller the p_i and hence also the upper bound in (12). When the service times are completely symmetric, the asymmetric parts A_i will be zero, and we recover the upper bound of 2 of Theorem 3.3.

Note that the upper bound in Theorem 3.11 is much easier to numerically compute or simulate than ϱ_ω itself, as for the latter one needs to go over all $n!$ possible sequences to find the optimal one. Also, this method can be used to find an upper bound on ϱ_ω for any problem instance where the service times come from a finite set of distributions and an upper bound on n is given, as illustrated in the next example.

Example 3.12. We base this example on the distributions fitted to health care data in [6]. There, patients were divided in two groups: new and return patients. For both groups, lognormal distributions were found as a good fit to the data used in the paper, with parameters as shown in Table 2. We checked that problem instances coming from these two distributions satisfy both $\mathbb{E}[\ln B_1] \leq \dots \leq \mathbb{E}[\ln B_n]$ and $\text{Var}(\ln B_1) \leq \dots \leq \text{Var}(\ln B_n)$, and therefore satisfy the dilation ordering assumption. It was also mentioned that the doctor that provided the data sees 10 patients per session.

We now consider 11 problem instances, each of them corresponding to $n = 10$ patients. In the k -th instance $k - 1$ patients have the first lognormal service-time distribution, whereas the remaining $11 - k$ patients have the other lognormal distribution, with $k = 1, \dots, 11$. We compute the upper bound in (12) for each of these instances through simulation. Note that we can simulate $A_i J_i$ by first drawing a random instance of X_i , and then setting A_i equal to X_i with probability $h_i(X_i)/f_i(X_i)$ and equal to zero otherwise. Simulating

1 000 000 random instances of X_1, \dots, X_n then allows us to compute $\mathbb{E}|X_1 + \dots + X_k|$ and $\mathbb{E}|A_1 J_1 + \dots + A_k J_k|$ for $k = 1, \dots, n$, and thus the upper bound in (12) for each of the 11 problem instances. Note that smaller problem instances are automatically included because we compute $\mathbb{E}|X_1 + \dots + X_k|/\mathbb{E}|A_1 J_1 + \dots + A_k J_k|$ for all k . We find that $\rho_\omega \leq 3.43$ for any problem instance consisting of at most 10 patients, with service times that follow one of the two lognormal distributions. We thus conclude that in this example $\rho_\omega \leq 3.43$ for any problem instance.

4. ASYMPTOTIC OPTIMALITY OF SVF UNDER MEAN-BASED SCHEDULES

The problem of finding the optimal sequence gets increasingly difficult as the number of patients grows large. In the numerical experiments described in Appendix E we could only find the optimal sequence for up to 10 or 11 patients, but in some applications more than 20 patients need to be scheduled in a session (see e.g. [22]). Therefore, we would like to know how the approximation ratio behaves as n grows larger. In this section we assess the performance of the SVF sequence as the number of patients tends to infinity. Throughout this section we assume that the schedule is mean-based: the time planned for each appointment is equal to the corresponding mean service time. The goal in this section is to prove that the SVF sequence is asymptotically optimal as the number of patients tends to infinity, under a weak assumption.

We consider the setting in which we are given, for each value of n , a vector $\mathbf{B}_n = (B_{n,1}, B_{n,2}, \dots, B_{n,n})$ of service-time distributions. For $i \leq n$, let $\mu_{n,i}$ and $\sigma_{n,i}^2$ denote the mean and variance of $B_{n,i}$, and let $X_{n,i} := B_{n,i} - \mu_{n,i}$ for all $i \leq n$. Similarly, $W_{n,i}$, and $I_{n,i}$ are all with respect to the service-time distributions \mathbf{B}_n , and an implicit fixed permutation $\tau \in \mathcal{S}_n$. Note that Example 2.2 shows that ρ_ω cannot be bounded without imposing some assumption, even when n tends to infinity. We do not need Assumption 2.4 (dilation ordering assumption) in this section. Instead, we impose the following assumption, similar to the Lyapunov condition of the Lyapunov-version of the central limit theorem (CLT), concerning the order of magnitude of the $(2 + \delta)$ -th moments of the service times. The difference between our assumption and the conventional Lyapunov condition is the supremum over all $n \geq k$ and all sequences τ .

Assumption 4.1. *We assume that there exists a $\delta > 0$ such that, as $k \rightarrow \infty$,*

$$q_k := \sup_{n \geq k, \tau \in \mathcal{S}_n} \frac{1}{\sqrt{\sum_{i=1}^k \sigma_{n,\tau(i)}^2}} \sum_{i=1}^k \mathbb{E}|X_{n,\tau(i)}|^{2+\delta} \rightarrow 0.$$

The main result of this section is the following.

Theorem 4.2. *Under Assumption 4.1, $\varrho_\omega(\mathbf{B}_n) \rightarrow 1$ as $n \rightarrow \infty$.*

To obtain some intuition about why such a result is reasonable, consider normally distributed service times. Fix some n and $k \leq n$ (both large), and consider the waiting time $W_{n,k+1}$ corresponding to some fixed permutation τ . Let $(S_j)_{j \leq k}$ be the random walk for which $W_{n,k+1}$ is the maximum; so $S_j = \sum_{i=1}^j X_{n,\tau(k+1-i)}$. Because the steps are normally distributed, we can embed this random walk within a standard Brownian motion $V(t)$. More precisely, we can define the coupling $S_j = V(\theta_j)$, where $\theta_j := \sum_{i=1}^j \sigma_{n,\tau(k+1-i)}^2$. Assumption 4.1 then tells us simply that no individual service time has a non-negligible fraction of the total variance, in the large n limit. Thus the gaps $\theta_{j+1} - \theta_j$ between sample times become negligibly small in comparison to the interval of interest, namely $[0, \theta_k]$. So the maximum $W_{n,k+1}$ of the random walk is very well approximated by the supremum of the Brownian motion on the interval $[0, \theta_k]$.

This supremum is of course completely understood. By the reflection principle, we know that $\sup_{0 \leq t \leq \theta_k} V(t)$ has the distribution of the absolute value of a normal with variance θ_k and mean zero. Since the SVF ordering minimizes θ_k (simultaneously for all choices of k), it is thus optimal in the limit.

Extending this intuition to general distributions and making it rigorous requires care. The plan is to apply the reflection principle to the random walk, so that hopefully the magnitude of the final position of the walk is a good estimate of the maximum of the unreflected walk. More precisely, fix some value $a > 0$, and let $T(a)$ denote the first step at which the random walk $(S_j)_{j \leq k}$ reaches a (or $T(a) = \infty$ if this event does not occur). Now define

$$\widehat{S}_j(a) = \begin{cases} S_j & \text{if } j < T(a) \\ 2a - S_j & \text{if } j \geq T(a) \end{cases} \quad \text{and} \quad \widetilde{S}_j(a) = \begin{cases} S_j & \text{if } j < T(a) \\ 2S_{T(a)} - S_j & \text{if } j \geq T(a); \end{cases}$$

cf. (8). So $\widehat{S}_j(a)$ is the walk reflected immediately as it ‘‘crosses’’ a for the first time, and \widetilde{S}_j the walk reflected from the first step after crossing a (see Figure 2).

The event that $W_{n,k+1}$ is at least a occurs precisely when $T(a)$ is finite, which implies that either $S_k \geq a$ or $\widehat{S}_k(a) = 2a - S_k > a$. As these are disjoint events, we have

$$\mathbb{P}(W_{n,k+1} \geq a) = \mathbb{P}(S_k \geq a) + \mathbb{P}(\widehat{S}_k(a) > a);$$

this is along similar lines to the proof of Lemma 3.6. Note that the processes $\widehat{S}_j(a)$ and $\widetilde{S}_j(a)$ have the same increments, except for step $T(a)$. In this step the increments differ by $2(S_{T(a)} - a)$, twice the amount by which the random walk ‘‘overshoots’’ level a . As this

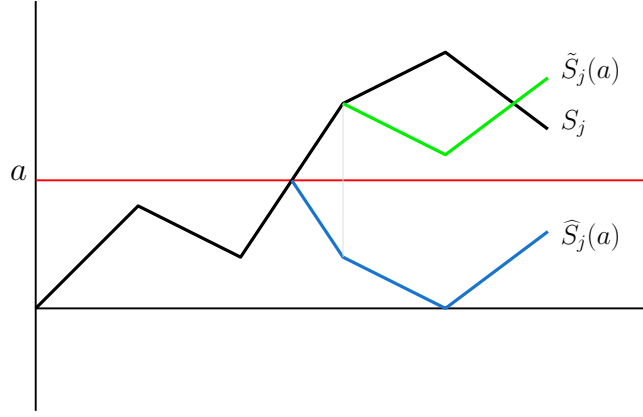


Figure 2. The processes S_j , $\widehat{S}_j(a)$ and $\widetilde{S}_j(a)$. The upper horizontal line indicates level a .

overshoot is nonnegative and bounded by $\max_{1 \leq i \leq k} X_{n,\tau(i)}$, we find that

$$\widetilde{S}_k(a) \geq \widehat{S}_k(a) \geq \widetilde{S}_k(a) - 2 \max_{1 \leq i \leq k} X_{n,\tau(i)}. \quad (13)$$

This leads to the estimates

$$\begin{aligned} \mathbb{P}(S_k \geq a) + \mathbb{P}(\widetilde{S}_k(a) > a) &\geq \mathbb{P}(W_{n,k+1} \geq a) \\ &\geq \mathbb{P}(S_k > a) + \mathbb{P}\left(\widetilde{S}_k(a) > a + 2 \max_{1 \leq i \leq k} X_{n,\tau(i)}\right). \end{aligned} \quad (14)$$

We now come to the first difficulty: we would like to estimate the probabilities $\mathbb{P}(S_k > a)$ and $\mathbb{P}(\widetilde{S}_k(a) > a)$ by applying a CLT result. However, the reflected process (\widetilde{S}_j) does not have independent steps (unless the service time distributions are symmetric around their mean, which we do not assume). Fortunately, (\widetilde{S}_j) is a martingale, and so we can apply the following CLT-type result for martingales. Here and in the remainder of this section, Z denotes a standard normal random variable.

Theorem 4.3 (Heyde and Brown [18]). *Let (ξ_i, \mathcal{F}_i) be a sequence of martingale differences, and let $Y_j = \xi_1 + \dots + \xi_j$ be the corresponding martingale. Suppose that the conditional variance, given by*

$$\sum_{i=1}^k \mathbb{E}[\xi_i^2 | \mathcal{F}_{i-1}],$$

is equal to one for some k . Then for any $\delta > 0$ there exists a constant C_δ that depends on δ only, such that

$$\sup_{x \in \mathbb{R}} |\mathbb{P}(Y_k > x) - \mathbb{P}(Z > x)| \leq C_\delta \left(\sum_{i=1}^k \mathbb{E}|\xi_i|^{2+\delta} \right)^{1/(3+\delta)}.$$

This is exactly what we need to obtain the following proposition (the full derivation can be found in Appendix B).

Proposition 4.4. *For any $k, n \geq k$ and permutation τ , we have*

$$\frac{\mathbb{E}W_{n,k+1}}{\sqrt{\sum_{j=1}^k \sigma_{n,\tau(j)}^2}} \leq \mathbb{E}|Z| + 2(C_\delta + 1)\sqrt{q_k}^{-1/(3+\delta)},$$

and C_δ is as in Theorem 4.3.

Next, we come to the lower bound. Considering (14), the second difficulty becomes apparent: we do not make any boundedness assumptions on the service time distributions, and so $\max_{1 \leq i \leq k} X_{n,\tau(i)}$ could be very large. This is remedied with a truncation argument.

We will consider the random walk $(S'_j)_{j \leq k}$ with steps $X_{n,\tau(i)} \mathbb{1}_{X_{n,\tau(i)} \leq c_{n,k}}$ instead of $X_{n,\tau(i)}$ (where $\mathbb{1}_E$ denotes the indicator of an event E). Here, $c_{n,k}$ is a bound depending on n and k , but not i . Let $W'_{n,k+1}$ be the maximum of the new random walk; clearly $\mathbb{E}W_{n,k+1} \geq \mathbb{E}W'_{n,k+1}$. Further, defining $\widehat{S}'_j(a)$ and $\widetilde{S}'_j(a)$ analogously to $\widehat{S}_j(a)$ and $\widetilde{S}_j(a)$ but with respect to the walk (S'_j) , (13) now yields

$$\widehat{S}_k(a) \geq \widetilde{S}_k(a) - 2c_{n,k}.$$

This allows us to apply Theorem 4.3 to obtain a lower bound on $\mathbb{E}W'_{n,k+1}$ and hence $\mathbb{E}W_{n,k+1}$. (A technical complication is that S'_j and $\widetilde{S}'_j(a)$ are no longer quite martingales; however, this can be overcome.) The value $c_{n,k}$ must be chosen carefully in a way that balances the need to sufficiently bound the overshoot (thus making (14) effective) and the need to not affect the steps too much (since this would decrease the maximum significantly). The result is the following proposition; further details can be found in Appendix B.

Proposition 4.5. *Under Assumption 4.1, for each $\varepsilon > 0$ there exists a K depending on ε only, such that for all $k \geq K, n \geq k$ and permutations τ ,*

$$\frac{\mathbb{E}W_{n,k+1}}{\sqrt{\sum_{i=1}^k \sigma_{n,\tau(i)}^2}} \geq (1 - \varepsilon)\mathbb{E}|Z|.$$

Combining Proposition 4.4 and Proposition 4.5 yields Theorem 4.2 in a straightforward manner. For completeness, the details can be found in Appendix B.

Remark 4.6. As in Remark 3.7, when the scheduled session end time is equal to the expected total service time, the expected overtime $\mathbb{E}W_{n,n+1}$ can be handled similarly to waiting time, and the result of Theorem 4.2 is also valid when some extra term $c\mathbb{E}W_{n,n+1}$ (with $c > 0$) is added to the cost function.

Remark 4.7. In this remark we assess the rate at which $\varrho_\omega(\mathbf{B}_n)$ converges to 1. Suppose that $\inf_{n,i}\{\sigma_{n,i}^2\} > 0$ and $\sup_{n,i}\{\mathbb{E}|X_{n,i}|^{2+\delta}\} < \infty$. Then $q_k = O(k^{-\delta/2})$ indeed converges to zero. Following the steps of the proof of Theorem 4.2 we then find

$$\varrho_\omega(\mathbf{B}_n) = 1 + O\left(\frac{K}{n}\right) + O\left(\sqrt{q_K}^{1/(3+\delta)}\right) = 1 + O\left(\frac{K}{n}\right) + O\left(K^{-\delta/(12+4\delta)}\right).$$

To obtain some insight into the convergence rate, observe that by choosing K in a way that these terms are balanced, it follows that

$$\varrho_\omega(\mathbf{B}_n) = 1 + O\left(n^{-\delta/(12+5\delta)}\right).$$

Note that for many practical distributions, including lognormal distributions, all moments exist. In such a case, Theorem 4.2 can be applied with any choice of δ .

5. BOUNDS ON PERFORMANCE FOR OPTIMALLY-SPACED SCHEDULES

The previous sections focused on mean-based schedules, i.e. schedules where the inter-arrival times are equal to the mean service times. Rather than mean-based schedules, one would preferably use optimally-spaced schedules. The problem of finding optimally-spaced schedules is well understood (see e.g. [2, 25]). In this section we consider the performance of the SVF sequence compared to the optimal combination of sequence and schedule. For this case we have also performed extensive numerical experiments to gain insight into what performance can be expected for SVF for practical distributions, described in more detail in Appendix F. As in the mean-based case, we focused on exponential and lognormal service times. Our main finding was that for all experiments we performed, the SVF sequence was the optimal sequence for all experiments with these distributions. It is noted, though, that such experiments can only be done for relatively small numbers of n , due to the complexity involved in computing the optimal sequence and schedule.

Kong et al. [24] give an example demonstrating that SVF is not always optimal for optimally-spaced schedules, but their example does not satisfy the dilation ordering assumption (Assumption 2.4). The following example shows that even when Assumption 2.4 applies, SVF may not be optimal.

Example 5.1. Suppose we have $n = 7$ patients, and we set $\omega = \frac{1}{2}$. The first three patients have service times that take value 0 or 2, each with probability $\frac{1}{2}$. The other four patients have service times that take values 0 or 4 each with probability $\frac{1}{4}$ and that take value 2 with probability $\frac{1}{2}$. As $B_i - 2$, for $i = 4, 5, 6, 7$, has the same distribution as $B_1 - 1 + A$, where A takes values -1 or 1 each with probability $\frac{1}{2}$ independent of B_1 , we can see through Lemma 2.5 that $B_1 \leq_{\text{dil}} B_i$. So this example satisfies the dilation ordering assumption.

We computed the optimal interarrival times and corresponding cost for each possible sequences (per the method described in Appendix F). The ratio between the optimal cost for the SVF sequence and the optimal cost overall was found to be 1.0787, so the SVF sequence is not optimal.

The goal of the remainder of this section is to prove upper bounds on the approximation ratio r_ω . In Section 5.1, we will do so when the service-time distributions are from the same location-scale family, leading to Theorem 5.3. In Section 5.2, we discuss the implications of this theorem for specific location-scale families of interest. Then in Section 5.3, we move beyond location-scale families and consider the lognormally distributed service times, which are of particular interest in practice [6, 22]. Finally, in Section 5.4 we show that for optimally-spaced schedules SVF is *not* asymptotically optimal, in stark contrast to the situation for mean-based schedules.

5.1. Location-scale family of service times. We impose the following assumption.

Assumption 5.2. *The distributions $(B_i)_{1 \leq i \leq n}$ form a location-scale family. In other words, there exists a random variable B having mean zero and variance one such that $B_i \stackrel{d}{=} \mu_i + \sigma_i B$.*

Note that, by [34], Theorem 3.A.18, this assumption implies Assumption 2.4. Due to Example 2.2, no bound on the approximation ratio r_ω can be found without imposing *some* assumption on the service-time distributions.

As well as using SVF as a sequencing rule, we must now specify an (ideally simple) scheduling rule as well. We use a schedule of the form $\mathbf{x} = \boldsymbol{\mu} + \alpha \boldsymbol{\sigma}$ for some $\alpha > 0$, in line with a suggestion by [8]. So we include additional slack in the schedule after each patient proportional to its standard deviation, in order to prevent the propagation of delays. We will choose

$$\alpha = \sqrt{\frac{1-\omega}{2\omega} + \frac{\sigma_{n-1}}{2 \sum_{i=1}^{n-1} \sigma_i}}; \quad (15)$$

the rationale behind this choice will become clear from the proof, where it will turn out to optimize the bound we obtain. Let Q_B denote the quantile function of B , i.e. $Q_B(y) = \inf\{x : y \leq \mathbb{P}(B \leq x)\}$. Define $B(\omega) = B - Q_B(1 - \omega)$, and

$$K(B, \omega) = \sqrt{2\omega} / [\omega \mathbb{E}B(\omega)^- + (1 - \omega) \mathbb{E}B(\omega)^+].$$

The main result of this section is the following.

Theorem 5.3. *Suppose that, for the SVF sequence, we use the schedule $\mathbf{x} = \boldsymbol{\mu} + \alpha \boldsymbol{\sigma}$, with α given by (15). Under Assumption 5.2, we have $r_\omega \leq K(B, \omega)$.*

This result follows immediately from the bounds on the cost function $C_\omega(\tau, \mathbf{x})$ given in the following two propositions, that are proved in Appendix C.

Proposition 5.4. *Suppose α is given by (15). Under Assumption 5.2,*

$$C_\omega(\text{id}, \boldsymbol{\mu} + \alpha\boldsymbol{\sigma}) \leq \sqrt{2\omega} \sum_{i=1}^{n-1} \sigma_i.$$

Proposition 5.5. *Under Assumption 5.2, for any sequence and schedule,*

$$C_\omega(\tau, \mathbf{x}) \geq [\omega \mathbb{E}B(\omega)^- + (1 - \omega) \mathbb{E}B(\omega)^+] \sum_{i=1}^{n-1} \sigma_i.$$

The idea behind the proof of Proposition 5.4 is as follows. We use that the waiting time can be expressed as the maximum of a random walk, as per (3). An upper bound for this maximum can now be found by using a comparison with another random walk that has i.i.d. steps, each distributed as the step of the original random walk with the largest variance. This upper bound is found by noting that if one (i) splits the steps in two parts, and (ii) multiplies the last part by some constant larger than one (leaving the first part unchanged), then the maximum increases. For the maximum of the new i.i.d. random walk, the classical *Kingman's bound* can be applied. After thus finding an upper bound on the expected waiting time, the expected idle time can then also be bounded using (5).

The idea behind the proof of Proposition 5.5 is to write

$$\omega \mathbb{E}I_{k+1} + (1 - \omega) \mathbb{E}W_{k+1} = \omega \mathbb{E}(W_k + B_{\tau(k)} - x_{\tau(k)})^- + (1 - \omega) \mathbb{E}(W_k + B_{\tau(k)} - x_{\tau(k)})^+,$$

and minimize this over $W_k - x_{\tau(k)}$. This minimization problem is the classical *newsvendor problem*, which has a known solution. This results in a lower bound on the cost function that is independent of the schedule. This lower bound can also be easily minimized over the sequences, resulting in Proposition 5.5.

When $\omega = \frac{1}{2}$, i.e. when waiting time and idle time are equally important, we know that $Q_B(\frac{1}{2})$ is equal to the median of B . In this case, we have $K(B, \frac{1}{2}) = 2/\mathbb{E}|B - m|$, where m is the median of B .

5.2. Examples. In this subsection we present examples of location-scale families for which we can compute $K(B, \omega)$ or $K(B, \frac{1}{2})$ from Theorem 5.3, so as to obtain insight into the magnitude of the constant $K(B, \omega)$. For the location-scale families of normal, uniform, shifted exponential and Laplace distributions, the results are shown in Table 3. For normal distributions $K(B, \omega)$ is not shown, as the expression does not simplify (with respect to the one presented in Theorem 5.3).

Location-scale family	$K(B, \omega)$	$K(B, \frac{1}{2})$
Normal	See Theorem 5.3	$\sqrt{2\pi} \approx 2.51$
Uniform	$\frac{1}{1-\omega} \sqrt{\frac{2}{3\omega}}$	$\frac{4}{3}\sqrt{3} \approx 2.31$
Shifted exponential	$-\frac{\sqrt{2}}{\sqrt{\omega} \ln(\omega)}$	$2/\ln(2) \approx 2.89$
Laplace	$\frac{2\sqrt{\omega}}{\min\{\omega, 1-\omega\}(1 - \ln(2 \min\{\omega, 1-\omega\}))}$	$2\sqrt{2} \approx 2.83$

Table 3. The values of $K(B, \omega)$ and $K(B, \frac{1}{2})$ for some location-scale families.

Now consider the case of Pareto (of type II, that is) distributions. A random variable X has such a distribution if

$$\mathbb{P}(X > x) = \left(1 + \frac{x - \mu}{\sigma}\right)^{-\beta} \text{ for } x \geq \mu,$$

for certain parameters $\mu, \sigma > 0, \beta$. The Pareto distributions with fixed parameter β form a location-scale family. Suppose that the B_i have Pareto distributions with fixed parameter $\beta > 2$. Then

$$K(B, \omega) = \sqrt{\frac{2\omega\beta}{\beta-2}} \left/ \left[-2\omega^{-\beta(\beta-1)+1} + \omega^{-\beta(\beta-1)} - (\beta-1)\omega^{\beta+1} + \beta\omega \right] \right.$$

In addition,

$$K(B, \frac{1}{2}) = 2\sqrt{\frac{\beta}{\beta-2}} \left/ \left[\beta - \left(\frac{1}{2}\right)^\beta (\beta-1) \right] \right.$$

For most typical location-scale families, the value of $K(B, \frac{1}{2})$ is between 2 and 3. However, in the Pareto case the value becomes much larger when β approaches two. Also, for ω close to either one of the extremes 0 or 1, the constant $K(B, \omega)$ blows up.

5.3. Lognormally distributed service times. In this subsection we focus on the case of lognormal service times. The scheduling rule that we use here is such that the interarrival times are a multiple (larger than 1) of the mean service times, rather than the mean service times increased by a multiple of the corresponding standard deviations. As will become clear below, this type of scheduling rule turns out to be convenient for the case of lognormal distributions. Note that, with $m_i := \mathbb{E}[\ln B_i]$ and $s_i^2 := \text{Var}(\ln B_i)$, we have that $\mathbb{E}B_i = \exp(m_i + s_i^2/2)$, so that the rule involves both parameters of the lognormal distribution. We have the following result.

Theorem 5.6. *Suppose the B_i are lognormally distributed with $m_1 \leq \dots \leq m_n$ and $s_1^2 \leq \dots \leq s_n^2$. When we use the schedule $\mathbf{x} = (1 + \alpha)\boldsymbol{\mu}$ for the SVF sequence, with*

$$\alpha = \frac{1}{\sqrt{2\omega}} \sqrt{(\exp(s_{n-1}^2) - 1)},$$

then

$$r_\omega \leq 2\omega\alpha \cdot [(1 - \omega)\mathbb{P}(Z \geq Q_Z(1 - \omega) - s_1) - \omega\mathbb{P}(Z \leq Q_Z(1 - \omega) - s_1)]^{-1}.$$

The proof can be found in Appendix C. It is based on a similar idea as the one used in the proof of Theorem 5.3. The main difference is in the upper bound, where it needs to be proved that the i.i.d. random walk used for comparison indeed has a larger expected maximum. For lognormal distributions, we use a convex ordering among the stepsize distributions to prove this, noting that the maximum of a random walk is a convex function in the stepsizes.

As an example, we apply Theorem 5.6 to the data found in [6]. Recall from Section 3.3 that the patients were divided into “new” and “return” patients, with service times fitted by lognormal distributions with parameters given in Table 2. It can be checked that any problem instance containing a mix of these patient groups satisfies the assumptions of Theorem 5.6. For any such problem instance, the largest possible s_{n-1} corresponds to a new patient, and the smallest possible s_1 corresponds to a return patient. When setting $\omega = \frac{1}{2}$, calculating the upper bound in Theorem 5.6, we find for this example that $r_{\frac{1}{2}} \leq 2.90$.

5.4. Example where SVF is not asymptotically optimal. In Section 4 we showed that the SVF sequence is asymptotically optimal as $n \rightarrow \infty$ when mean-based schedules are used. Perhaps surprisingly, this is no longer true when optimally-spaced schedules are used. The example described in this section satisfies the condition used in the mean-based case, Assumption 4.1. Note that the example does not satisfy Assumption 2.4 (the dilation ordering assumption).

The example is as follows. We set $\omega = \frac{1}{2}$, and we have two groups of patients, each consisting of $n/2$ patients (where n will be large). Let c be some positive integer that we will eventually allow to grow very large, and let a be a fixed constant larger than 1 that we will choose later. The service time distributions of the two groups are given by

$$B^{(1)} = \begin{cases} c & \text{with prob. } p := \frac{1}{1+c^2}, \\ -\frac{1}{c} & \text{with prob. } 1 - p, \end{cases} \quad B^{(2)} = \begin{cases} a & \text{with prob. } \frac{1}{2}, \\ -a & \text{with prob. } \frac{1}{2}. \end{cases}$$

These are chosen to have mean zero for convenience; by shifting them (which would shift the optimal schedule by the same amount), these can be adjusted to have nonnegative support.

The variance of the service times of the group 1 (group 2) patients is 1 ($a^2 > 1$, respectively). This means that the SVF sequence first serves all patients of group 1, followed by all patients from group 2. We will compare the SVF sequence with the *mixed sequence*, where patients of group 1 and 2 are served alternately starting with group 1. As we will see, for large enough c the mixed sequence will be cheaper under optimal scheduling, thus refuting the asymptotic optimality of SVF. Since we will consider a large- n limit, we will consider the average per-patient cost, consisting of the average expected waiting time and idle time per patient.

We begin by analyzing the SVF sequence. As we need a lower bound on its cost, we need to determine the optimal schedule. Fortunately, in the limit of large n , the situation simplifies. Each group can be treated as a stationary queue, with an associated expected waiting and idle time for each group at stationarity (patients at the beginning of the group, before stationarity is achieved, make a negligible contribution to the total cost). We use $W^{(i)}$ and $I^{(i)}$ to denote the distributions of waiting time and idle time for patients in group i . We need to determine the optimal stationary interarrival times $x^{(1)}$ and $x^{(2)}$ for patients in groups 1 and 2, respectively. As a first remark, the average expected idle time of a patient in group i is simply $x^{(i)}$; this is a consequence of the Lindley recursions (1) combined with stationarity, which yields

$$\mathbb{E} \left[W^{(i)} - I^{(i)} \right] = \mathbb{E} \left[W^{(i)} + B^{(i)} - x^{(i)} \right] = \mathbb{E} W^{(i)} - x^{(i)}.$$

The average expected waiting time is more challenging, but we can compute it exactly in the limit as $c \rightarrow \infty$; we defer the argument to Appendix C.2.

Lemma 5.7. *For the SVF sequence in the limit $c \rightarrow \infty$, the optimal interarrival times for groups 1 and 2 are $x^{(1)} = \frac{1}{2}\sqrt{2}$ and $x^{(2)} = a$ respectively. Patients in group 1 incur an average expected waiting time of $\mathbb{E}W^{(1)} = \frac{1}{2}\sqrt{2}$, whereas patients in group 2 incur no waiting time. This yields an average total cost of $\frac{1}{2}(a + \sqrt{2})$.*

We now move to the mixed sequence. Here, we need an upper bound on the cost, and so we can use any scheduling rule we like. We will simply use the same interarrival times as were used in the SVF sequence: a patient in group i has an interarrival time of $x^{(i)}$. Just as with the SVF sequence, in the limit as $n \rightarrow \infty$ the average expected idle time is $\frac{1}{2}(x^{(1)} + x^{(2)}) = \frac{1}{2}(a + \frac{1}{2}\sqrt{2})$. The main challenge is to bound the average expected waiting time.

Let $V^{(i)}$ be the waiting time experienced by a group i patient at stationarity. Then $V^{(i)}$ is the maximum of the random walk with alternating steps distributed as $B^{(1)} - x^{(1)}$ and

$B^{(2)} - x^{(2)}$. The random walk for $V^{(1)}$ starts with a step $B^{(2)} - x^{(2)}$ and the random walk for $V^{(2)}$ starts with a step $B^{(1)} - x^{(1)}$. This difference vanishes in the limit as c grows large (see Appendix C.2 for the proof).

Lemma 5.8. $\lim_{c \rightarrow \infty} \mathbb{E}V^{(1)} = \lim_{c \rightarrow \infty} \mathbb{E}V^{(2)}$.

So we focus on $\lim_{c \rightarrow \infty} \mathbb{E}V^{(1)}$. Note that the group 2 steps, distributed as $B^{(2)} - a$, are never positive, so the maximum of the alternating walk can only be attained after a group 1 step. Therefore, we can combine the two alternating steps into one step, and $V^{(1)}$ is the maximum of a random walk $(S_j^{(1)})_{j \geq 0}$ with steps $B_j^{(1)} + B_j^{(2)} - (a + \frac{1}{2}\sqrt{2})$. (Here, $B_j^{(i)}$ is an independent copy of $B^{(i)}$, for each $i \in \{1, 2\}$ and $j \in \mathbb{N}$.) Observe that $\mathbb{E}V^{(1)}$ is never bigger than $\mathbb{E}W^{(1)}$, the expected maximum of a random walk with steps distributed as $B^{(1)} - \frac{1}{2}\sqrt{2}$, as $B^{(2)} - a \leq 0$, so as a preliminary bound we have $\mathbb{E}V^{(1)} \leq \mathbb{E}W^{(1)} = \frac{1}{2}\sqrt{2}$. Recalling that in the SVF sequence the group 2 patients experience no waiting time, we see that our goal is to show that in fact $\mathbb{E}V^{(1)} < \frac{1}{2}\mathbb{E}W^{(1)} = \frac{1}{4}\sqrt{2}$.

The simplest way to bound $\mathbb{E}V^{(1)}$ is via a bound by [21], which says that if $(S_i)_{i \geq 0}$ is a random walk with i.i.d. steps distributed according to Z , where $\mathbb{E}Z < 0$, then

$$\mathbb{E} \max_{i \geq 0} S_i \leq -\frac{\text{Var } Z}{2\mathbb{E}Z}.$$

Applied to the random walk $S^{(1)}$,

$$\mathbb{E}V^{(1)} \leq \frac{\text{Var}[B^{(1)} + B^{(2)}]}{2(a + \frac{1}{2}\sqrt{2})} = \frac{1 + a^2}{2a + \sqrt{2}}.$$

This is always larger than $\frac{1}{2}\sqrt{2}$, for any $a \geq 1$, so this bound does not suffice. Instead, let us consider the random walk $(\tilde{S}_j)_{j \geq 0}$ with steps $B_i^{(1)} - (a + \frac{1}{2}\sqrt{2})$; let \tilde{V} be its maximum. Kingman's bound for this random walk yields

$$\mathbb{E}\tilde{V} \leq \frac{\text{Var } B^{(1)}}{2(a + \frac{1}{2}\sqrt{2})} = \frac{1}{2a + \sqrt{2}}.$$

As long as we choose $a > \frac{3}{2}\sqrt{2}$, this is below $\frac{1}{2}\mathbb{E}W^{(1)}$. All that remains is to show the following lemma.

Lemma 5.9. $\lim_{c \rightarrow \infty} \mathbb{E}V^{(1)} \leq \lim_{c \rightarrow \infty} \mathbb{E}\tilde{V}$.

We sketch the proof here, and relegate the details to Appendix C.2. Consider a typical realization of $(\tilde{S}_j)_{j \geq 0}$. It consists primarily of downward steps of size $a + \frac{1}{2}\sqrt{2} + 1/c \approx a + \frac{1}{2}\sqrt{2}$, with occasional large upwards jumps of size c . These upwards jumps, occurring with probability $p = 1/(1 + c^2) \approx 1/c^2$ are very rare. One can prove that with overwhelming probability, the maximum of $S^{(1)}$ occurs within the first $2c$ steps; condition on this for the remainder. If

there are no upwards jumps in the first $2c$ steps, then \tilde{V} and $V^{(1)}$ are both zero; this happens with probability $(1-p)^{2c} \approx 1 - 2pc \approx 1 - 2/c$. Independently of whether there are any such upwards jumps, the difference between the random walks for \tilde{V} and for $V^{(1)}$ after j steps is $\sum_{i=1}^j B_i^{(2)}$. By an appropriate concentration bound, we can deduce that $S_j^{(1)} \leq \tilde{S}_j + c^{2/3}$ for all $j \leq 2c$, with extremely high probability. So with probability about $1 - 2/c$, $V^{(1)} = \tilde{V} = 0$, and with probability about $2/c$, $V^{(1)} \leq \tilde{V} + c^{2/3}$. This implies that $\mathbb{E}V^{(1)} \leq \mathbb{E}\tilde{V} + c^{-1/3}$, yielding the lemma.

With the optimized choice $a = \frac{1}{2}\sqrt{2} + \sqrt{3}$, a quantitative lower bound of $r_{1/2} \geq \frac{1}{2} + \frac{1}{4}\sqrt{6} > 1.11$ can be obtained.

6. DISCUSSION AND DIRECTIONS FOR FURTHER RESEARCH

We have shown that under quite general conditions, the SVF sequence yields a constant-factor approximation. Furthermore, we have seen that additional information about the instance, such as knowing that the service-time distributions fall within a certain class, or that the number of patients is large, can lead to substantial improvements of our worst-case bounds.

For mean-based schedules, Example 3.1 and Theorem 3.4 show that the worst-case approximation ratio lies between 1.52 and 4; for symmetric service-time distributions we found the upper bound 2. It would be interesting to reduce the gap between the lower and upper bound; we suspect neither bound is tight. In particular, the upper bound on the cost of the SVF sequence appears to be a strong bound only in the regime of many patients, with service times of similar variances; the lower bound on the cost of arbitrary sequences, on the other hand, appear strong in situations where only a few service times with large variance have a significant impact on the cost function. This suggests that more refined arguments, possibly considering multiple regimes, could lead to an improved upper bound. Improving our bounds for special cases (such as normal and lognormal distributions), or considering other practically relevant service-time distributions, would also be of interest.

When optimizing over both the sequence and the schedule, we obtained bounds for location-scale families and for lognormally-distributed service times. These bounds obtained are not uniform: in the former case, the bounds depend on ω and the location-scale family, and in the latter case, on the parameters of the lognormal distributions. A constant-factor approximation that does not depend on these quantities, or that holds in greater generality (e.g., to all distributions satisfying the dilation ordering assumption), remains an open question. The SVF sequencing rule remains a promising candidate, but a more sophisticated choice of scheduling rule will certainly be needed.

We found that the SVF sequence is asymptotically optimal as the number of patients tends to infinity under a mild Lyapunov-type assumption, but that this is no longer the case when optimally-spaced schedules are used, as shown in Section 5.4. Such an asymptotic optimality for optimally-spaced schedules might still hold under more strict assumptions on the service times. Finding what these assumptions should be and proving asymptotic optimality under these assumptions is an interesting open problem.

To the best of our knowledge, this is the first paper that assesses whether an easily computed sequence performs provably well, rather than trying to find the optimal sequence for a special (typically low-dimensional) instance, or comparing heuristics through simulation. Finding the optimal sequence is an important (but inherently difficult) problem, and we hope our approach triggers more research in this direction.

APPENDIX A. PROOFS CORRESPONDING TO SECTION 3

Here we prove Lemma 3.9 and Lemma 3.10, that are used in Section 3.

Lemma 3.9. *Under Assumption 3.2,*

$$\mathbb{E}W_{k+1} \leq \mathbb{E}(X_{\tau(1)} + \cdots + X_{\tau(k)})^+ + \mathbb{E}(X_{\tau(1)} + \cdots + X_{\tau(k-1)})^+.$$

Proof. By Lemma 3.6, W_k is stochastically dominated by $|X_{\tau(1)} + \cdots + X_{\tau(k-1)}|$. As a consequence,

$$\mathbb{E}W_{k+1} = \mathbb{E}(W_k + X_{\tau(k)})^+ \leq \mathbb{E}(|X_{\tau(1)} + \cdots + X_{\tau(k-1)}| + X_{\tau(k)})^+. \quad (16)$$

If Y and Z are independent and both have symmetric distributions, then

$$\mathbb{E}(|Y| + Z)^+ = \mathbb{E}(Y + Z)^+ + \mathbb{E}Y^+.$$

This is easily checked by conditioning on $|Y| = a$ and $|Z| = b$, as then Y is either a or $-a$ with probability $\frac{1}{2}$, and similarly for Z . Applying this result to the upper bound in (16), we find the upper bound in the lemma. \square

Lemma 3.10. *Under Assumption 3.2, for any ℓ ,*

$$\mathbb{E}W_{k+1} \geq \frac{1}{2} \left(\mathbb{E}(X_{\tau(1)} + \cdots + X_{\tau(k)})^+ + \mathbb{E}(X_{\tau(1)} + \cdots + X_{\tau(\ell)})^+ + \mathbb{E}(X_{\tau(\ell+1)} + \cdots + X_{\tau(k)})^+ \right).$$

Proof. Let $S'_\ell = X_{\tau(1)} + X_{\tau(2)} + \dots + X_{\tau(\ell)}$, so that $S'_\ell = S_k - S_{k-\ell}$. As $W_{k+1} = \max\{0, S_1, \dots, S_k\}$, we then have

$$W_{k+1} \geq \max\{0, S_{k-\ell}, S_k\} = \max\{0, S_{k-\ell}, S_{k-\ell} + S'_\ell\} = (S_{k-\ell} + (S'_\ell)^+)^+. \quad (17)$$

If Y and Z are independent and both have symmetric distributions, then

$$\mathbb{E}(Y + Z^+)^+ = \frac{1}{2} (\mathbb{E}(Y + Z)^+ + \mathbb{E}Y^+ + \mathbb{E}Z^+).$$

Again, this is easily checked by conditioning on $|Y| = a$ and $|Z| = b$, as then Y is either a or $-a$ with probability $\frac{1}{2}$, and similarly for Z . Applying this result to the lower bound in (17), we find the lower bound in the lemma. \square

APPENDIX B. PROOFS CORRESPONDING TO SECTION 4

In this appendix we prove Proposition 4.4, Proposition 4.5, and hence deduce Theorem 4.2, building on the ideas presented in Section 4.

Proof of Proposition 4.4. For ease of notation, we assume throughout this proof that patients are renumbered so that $\tau(i) = i$. Note that $\Sigma_k^2 := \sum_{i=1}^k \sigma_{n,i}^2$ is the variance of both S_k and $\tilde{S}_k(a)$. In order to apply Theorem 4.3 we scale all steps, and hence S_k and $\tilde{S}_k(a)$, by a factor $1/\Sigma_k$. For both martingales the squared increments $(X_{n,k-i+1})^2$ are independent of the previous increments, so that after rescaling the conditional variance after k steps equals one for both martingales. Note that we can recognize the q_k from Assumption 4.1 in the upper bound, so we find for any x that

$$\left| \mathbb{P}\left(\frac{S_k}{\Sigma_k} > x\right) - \mathbb{P}(Z > x) \right| \leq C_\delta q_k^{1/(3+\delta)}, \quad (18)$$

$$\left| \mathbb{P}\left(\frac{\tilde{S}_k(a)}{\Sigma_k} > x\right) - \mathbb{P}(Z > x) \right| \leq C_\delta q_k^{1/(3+\delta)}. \quad (19)$$

Using inequality (18) and Chebyshev's inequality, we find

$$\begin{aligned} \frac{1}{\Sigma_k} \int_0^\infty \mathbb{P}(S_k > a) da &= \int_0^\infty \mathbb{P}\left(\frac{S_k}{\Sigma_k} > x\right) dx \\ &\leq \int_0^{1/\sqrt{q_k}^{1/(3+\delta)}} \left(\mathbb{P}(Z > x) + C_\delta q_k^{1/(3+\delta)}\right) dx + \int_{1/\sqrt{q_k}^{1/(3+\delta)}}^\infty \frac{1}{x^2} dx \\ &\leq \mathbb{E}Z^+ + (C_\delta + 1)\sqrt{q_k}^{1/(3+\delta)}. \end{aligned}$$

A similar reasoning using (19) finds the same upper bound for

$$\frac{1}{\Sigma_k} \int_0^\infty \mathbb{P}(\tilde{S}_k(a) > a) da.$$

Finally,

$$\begin{aligned} \mathbb{E}W_{n,k+1} &\leq \int_0^\infty \left(\mathbb{P}(S_k \geq a) + \mathbb{P}(\tilde{S}_k(a) > a) \right) da = \int_0^\infty \left(\mathbb{P}(S_k > a) + \mathbb{P}(\tilde{S}_k(a) > a) \right) da \\ &\leq 2\mathbb{E}Z^+ + 2(C_\delta + 1)q_k^{(3+\delta)/2}. \end{aligned}$$

□

Proof of Proposition 4.5. We will first lay out the technical groundwork that we need for this proof.

Fix any n and $k \leq n$ (later we will take limits). As in the proof above, assume that $\tau(i) = i$. Recall from Section 4 that $(S'_j)_{j \leq k}$ is the random walk with steps $X_{n,i} \mathbb{1}_{X_{n,i} \leq c_{n,k}}$ instead of $X_{n,i}$; and $\widehat{S}'_k(a)$, $\tilde{S}'_k(a)$ are defined based on this random walk in the same way that $\widehat{S}_k(a)$ and $\tilde{S}_k(a)$ were defined based on (S_j) . For now, $c_{n,k}$ can take any positive value. We thus have

$$\mathbb{P}(W'_{n,k+1} \geq a) \geq \mathbb{P}(S'_k > a) + \mathbb{P}(\widehat{S}'_k(a) > a) \geq \mathbb{P}(S'_k > a) + \mathbb{P}(\tilde{S}'_k(a) > a + 2c_{n,k}).$$

Note that the steps no longer have mean zero, and so $\tilde{S}'_j(a)$ is no longer a martingale. To repair this issue, we must know how much the change in steps due to the indicator affects the mean and variance of all the steps. For this we have the following lemma. Define $\Sigma_{n,k}^2 := \text{Var } S_k$, $(\Sigma'_{n,k})^2 := \text{Var } S'_k$ and $\gamma_{n,k} := (\Sigma_{n,k}/c_{n,k})^\delta$.

Lemma B.1.

- (1) $\frac{\sum_{i=1}^k \mathbb{E}[(X_{n,i}) \mathbb{1}_{X_{n,i} > c_{n,k}}]}{\Sigma_{n,k}} \leq q_k \gamma_{n,k}^{(1+\delta)/\delta}$; and
- (2) $\frac{(\Sigma'_{n,k})^2}{\Sigma_{n,k}^2} \geq 1 - 2q_k \gamma_{n,k}$.

Proof. 1 Using that $X_{n,i} > c_{n,k}$ implies $|X_{n,i}|^{1+\delta}/c_{n,k}^{1+\delta} > 1$, we find

$$\mathbb{E}[(X_{n,i}) \mathbb{1}_{X_{n,i} > c_{n,k}}] \leq \mathbb{E}\left[\left(|X_{n,i}|^{2+\delta}/c_{n,k}^{1+\delta} \right) \mathbb{1}_{X_{n,i} > c_{n,k}} \right] \leq \mathbb{E}|X_{n,i}|^{2+\delta}/c_{n,k}^{1+\delta}.$$

Summing over i and dividing both sides by $\Sigma_{n,k}$ we find

$$\frac{\sum_{i=1}^k \mathbb{E}[(X_{n,i}) \mathbb{1}_{X_{n,i} > c_{n,k}}]}{\Sigma_{n,k}} \leq \frac{\sum_{i=1}^n \mathbb{E}|X_{n,i}|^{2+\delta}}{\Sigma_{n,k} c_{n,k}^{1+\delta}} = \frac{\sum_{i=1}^n \mathbb{E}|X_{n,i}|^{2+\delta}}{\Sigma_{n,k}^{2+\delta}} \left(\frac{\Sigma_{n,k}}{c_{n,k}} \right)^{1+\delta} \leq q_k \gamma_{n,k}^{(1+\delta)/\delta},$$

as was claimed.

2 Analogous to the proof of part (1), we can also deduce that

$$\frac{\sum_{i=1}^k \mathbb{E}\left[X_{n,i}^2 \mathbb{1}_{X_{n,i} > c_{n,k}} \right]}{\Sigma_{n,k}^2} \leq q_k \gamma_{n,k}. \quad (20)$$

Using that $\mathbb{E}X_{n,i} = 0$, we find

$$\begin{aligned} \frac{(\Sigma'_{n,k})^2}{\Sigma_{n,k}^2} &= \frac{\sum_{i=1}^n \mathbb{E} \left[X_{n,i}^2 \mathbb{1}_{X_{n,i} \leq c_{n,k}} \right]}{\Sigma_{n,k}^2} - \frac{\sum_{i=1}^n (\mathbb{E} [X_{n,i} \mathbb{1}_{X_{n,i} \leq c_{n,k}}])^2}{\Sigma_{n,k}^2} \\ &= \frac{\Sigma_{n,k}^2 - \sum_{i=1}^n \mathbb{E} \left[X_{n,i}^2 \mathbb{1}_{X_{n,i} > c_{n,k}} \right]}{\Sigma_{n,k}^2} - \frac{\sum_{i=1}^n (\mathbb{E} [X_{n,i} \mathbb{1}_{X_{n,i} > c_{n,k}}])^2}{\Sigma_{n,k}^2}. \end{aligned}$$

Now applying Jensen's inequality to the last term and (20),

$$\frac{(\Sigma'_{n,k})^2}{\Sigma_{n,k}^2} \geq \frac{\Sigma_{n,k}^2 - 2 \sum_{i=1}^n \mathbb{E} \left[X_{n,i}^2 \mathbb{1}_{X_{n,i} > c_{n,k}} \right]}{\Sigma_{n,k}^2} \geq 1 - 2q_k \gamma_{n,k},$$

which is what we wanted to prove. \square

Now we have the following result. Define, with Z a standard normal random variable and C_δ the constant featuring in Theorem 4.3,

$$D_k := \int_0^{1/\sqrt{q_k}^{1/(3+\delta)}} \mathbb{P}(Z > x) dx - C_\delta \sqrt{q_k}^{1/(3+\delta)}.$$

Lemma B.2. *We have*

$$\frac{\mathbb{E}W_{n,k+1}}{\Sigma_{n,k}} \geq 2D_k \sqrt{1 - 2q_k \gamma_{n,k}} - \frac{2c_{n,k}}{\Sigma_{n,k}} - 2q_k \gamma_{n,k}^{(1+\delta)/\delta}.$$

Proof. Again we want to apply Theorem 4.3. This time we not only need to divide by the standard deviation $\Sigma'_{n,k}$ of $S'_{n,k}$, but also subtract the (negative) mean. We then find

$$\begin{aligned} \frac{1}{\Sigma'_{n,k}} \int_0^\infty \mathbb{P}(S'_k > a) da &= \int_{-\mathbb{E}S'_k/\Sigma'_{n,k}}^\infty \mathbb{P} \left(\frac{S'_k - \mathbb{E}S'_k}{\Sigma'_{n,k}} > x \right) dx \\ &\geq \int_0^{1/\sqrt{q_k}^{1/(3+\delta)}} \left(\mathbb{P}(Z > x) - C_\delta q_k^{1/(3+\delta)} \right) dx - \int_0^{-\mathbb{E}S'_k/\Sigma'_{n,k}} \mathbb{P} \left(\frac{S'_k - \mathbb{E}S'_k}{\Sigma'_{n,k}} > x \right) dx \\ &\geq D_k + \frac{\mathbb{E}S'_k}{\Sigma'_{n,k}}, \end{aligned}$$

where we used in the last step that probabilities are bounded by one.

Next, we need a bound for

$$\frac{1}{\Sigma'_{n,k}} \int_0^\infty \mathbb{P}(\tilde{S}'_k(a) > a + 2c_{n,k}) da. \quad (21)$$

Note that $\tilde{S}'_j(a)$ is no longer a martingale, as the mean step size is no longer zero. However, we can remedy this by noting that $\tilde{S}'_k(a)$ is bounded below by

$$\tilde{S}'_k(a) - \sum_{i=1}^k (X_{n,i}) \mathbb{1}_{X_{n,i} > c_{n,k}} = \tilde{S}'_k(a) - S_k + S'_k,$$

and $\tilde{S}_j(a) - S_j + S'_j - \mathbb{E}S'_j$ is again a martingale with mean zero. Applying Theorem 4.3 to this martingale and taking into account the shift $\mathbb{E}S'_k$ and overshoot $2c_{n,k}$, we find that (21) is bounded below by

$$D_k + \frac{\mathbb{E}S'_k}{\Sigma'_{n,k}} - \frac{2c_{n,k}}{\Sigma'_{n,k}}.$$

Now adding up these two lower bounds, we find

$$\frac{\mathbb{E}W'_{n,k+1}}{\Sigma'_{n,k}} \geq 2D_k + \frac{2\mathbb{E}S'_k}{\Sigma'_{n,k}} - \frac{2c_{n,k}}{\Sigma'_{n,k}}.$$

Now note that, due to Lemma B.1 2, $\Sigma'_{n,k}/\Sigma_{n,k} \geq \sqrt{1 - 2q_k\gamma_{n,k}}$. In addition, $\mathbb{E}W_{n,k+1} \geq \mathbb{E}W'_{n,k+1}$. Consequently,

$$\frac{\mathbb{E}W_{n,k+1}}{\Sigma_{n,k}} \geq 2D_k \sqrt{1 - 2q_k\gamma_{n,k}} + \frac{2\mathbb{E}S'_k}{\Sigma_{n,k}} - \frac{2c_{n,k}}{\Sigma_{n,k}}.$$

Applying Lemma B.1 1,

$$\frac{\mathbb{E}W_{n,k+1}}{\Sigma_{n,k}} \geq 2D_k \sqrt{1 - 2q_k\gamma_{n,k}} - 2q_k\gamma_{n,k}^{(1+\delta)/\delta} - \frac{2c_{n,k}}{\Sigma_{n,k}},$$

which is the bound we wanted to prove. \square

We are now finally ready to proceed with the proof of Proposition 4.5. The freedom remains to choose $c_{n,k}$, which we have only assumed to be positive. We would like to have $c_{n,k}/\Sigma_{n,k} \rightarrow 0$, $q_k\gamma_{n,k}^{(1+\delta)/\delta} \rightarrow 0$ and $q_k\gamma_{n,k} \rightarrow 0$ as $k \rightarrow \infty$. A choice that achieves this goal is

$$c_{n,k} := \sqrt{q_k}^{-1/(\delta+1)} \Sigma_{n,k},$$

so that $\gamma_{n,k} = q_k^{-\delta/(2\delta+2)}$. We thus obtain, with $\bar{q}_k := \sqrt{q_k}^{-1/(\delta+1)} + \sqrt{q_k}$,

$$\frac{\mathbb{E}W_{n,k+1}}{\Sigma_{n,k}} \geq 2 \left(\int_0^{1/\sqrt{q_k}^{-1/(3+\delta)}} \mathbb{P}(Z > x) dx - C_\delta \sqrt{q_k}^{-1/(3+\delta)} \right) \sqrt{1 - 2\sqrt{q_k}^{-2-\delta/(\delta+1)}} - 2\bar{q}_k.$$

Note that this converges to $\mathbb{E}|Z|$ as $q_k \rightarrow 0$, so this completes the proof of Proposition 4.5. \square

Proof of Theorem 4.2. We first state an immediate consequence of combining Proposition 4.4 and Proposition 4.5. We use $W_{n,i}^{\text{SVF}}$ and $I_{n,i}^{\text{SVF}}$ to refer specifically to the waiting and idle times for the SVF sequence, and $W_{n,i}^{\text{OPT}}$ and $I_{n,i}^{\text{OPT}}$ for the optimal sequence.

Lemma B.3. *Under Assumption 4.1, for any $\varepsilon > 0$ there exists a K depending on ε only, such that for all $k \geq K$ and for all $n \geq k - 1$,*

$$\mathbb{E}W_{n,k}^{\text{SVF}} \leq (1 + \varepsilon) \mathbb{E}W_{n,k}^{\text{OPT}}.$$

Proof. By Proposition 4.5, for any $\varepsilon > 0$, we can choose k sufficiently large such that

$$\mathbb{E}W_{n,k+1} \geq (1 - \varepsilon) \sqrt{\sum_{i=1}^k \sigma_{n,\tau(i)}^2} \cdot \mathbb{E}|Z| \geq (1 - \varepsilon) \sqrt{\sum_{i=1}^k \sigma_{n,i}^2} \cdot \mathbb{E}|Z|$$

for any sequence τ , in particular for the optimal sequence. Here we also used that $\sigma_{n,1}^2, \dots, \sigma_{n,k}^2$ are the k smallest variances. For the SVF sequence, Proposition 4.4 gives us that for sufficiently large k we have

$$\mathbb{E}W_{n,k+1}^{\text{SVF}} \leq (1 + \varepsilon) \sqrt{\sum_{i=1}^k \sigma_{n,i}^2} \cdot \mathbb{E}|Z|.$$

Combining these two bounds completes the proof. \square

By the Lindley recursion we have

$$\mathbb{E}W_{n,k+1} = \mathbb{E}(W_{n,k} + X_{n,\tau(k)})^+ \geq \mathbb{E}(W_{n,k} + X_{n,\tau(k)}) = \mathbb{E}W_{n,k},$$

so, for any n and τ , $\mathbb{E}W_{n,k}$ is increasing in k . With K as in Lemma B.3, we consequently have

$$\sum_{k=1}^n \mathbb{E}W_{n,k}^{\text{SVF}} \leq K \mathbb{E}W_{n,K}^{\text{SVF}} + \sum_{k=K+1}^n \mathbb{E}W_{n,k}^{\text{SVF}} \leq (1 + \varepsilon) \left(K \mathbb{E}W_{n,K}^{\text{OPT}} + \sum_{k=K+1}^n \mathbb{E}W_{n,k}^{\text{OPT}} \right).$$

Again using that $\mathbb{E}W_{n,k}$ is increasing in k , we also have

$$K \mathbb{E}W_{n,K}^{\text{OPT}} / \sum_{k=K+1}^n \mathbb{E}W_{n,k}^{\text{OPT}} \leq \frac{K}{n - K},$$

which is smaller than ε for sufficiently large n . Then we have

$$\begin{aligned} \sum_{k=1}^n \mathbb{E}W_{n,k}^{\text{SVF}} &\leq (1 + \varepsilon) \left(K \mathbb{E}W_{n,K}^{\text{OPT}} + \sum_{k=K+1}^n \mathbb{E}W_{n,k}^{\text{OPT}} \right) \\ &\leq (1 + \varepsilon)^2 \sum_{k=K+1}^n \mathbb{E}W_{n,k}^{\text{OPT}} \leq (1 + \varepsilon)^2 \sum_{k=1}^n \mathbb{E}W_{n,k}^{\text{OPT}}. \end{aligned}$$

Recall that the total expected idle time is equal to $\mathbb{E}W_{n,n}$. Since also $\mathbb{E}W_{n,n}^{\text{SVF}} \leq (1 + \varepsilon) \mathbb{E}W_{n,n}^{\text{OPT}}$ for $n \geq K$,

$$\omega \mathbb{E}W_{n,n}^{\text{SVF}} + (1 - \omega) \sum_{k=1}^n \mathbb{E}W_{n,k}^{\text{SVF}} \leq (1 + \varepsilon)^2 \left(\omega \mathbb{E}W_{n,n}^{\text{OPT}} + (1 - \omega) \sum_{k=1}^n \mathbb{E}W_{n,k}^{\text{OPT}} \right)$$

for n sufficiently large (independent of ω). As this holds for any $\varepsilon > 0$, we find $\varrho_\omega(\mathbf{B}_n) \rightarrow 1$ as $n \rightarrow \infty$, and Theorem 4.2 is proved. \square

APPENDIX C. PROOFS CORRESPONDING TO SECTION 5

C.1. Proofs corresponding to Section 5.1 and 5.3. Here we prove Proposition 5.4, Proposition 5.5, and Theorem 5.6 from Section 5. In order to prove Proposition 5.4, we need a number of lemmas. We will denote the waiting times and idle times under the SVF sequence with schedule $\mathbf{x} = \boldsymbol{\mu} + \alpha\boldsymbol{\sigma}$ by W_i^{SVF} and I_i^{SVF} respectively, and we will denote the waiting times and idle times under the optimal combination of sequence and schedule by W_i^{OPT} and I_i^{OPT} .

Lemma C.1. *Let M be the maximum of a random walk with steps Y_1, \dots, Y_k . Now let $i \in \{1, \dots, k\}$, and let $c \geq 1$. Let M' be defined as the maximum of a random walk with steps $Y_1, \dots, Y_i, cY_{i+1}, \dots, cY_k$. Then $M \leq M'$.*

Proof. Suppose that the maximum of the first random walk is attained at time j , that is, $M = Y_1 + \dots + Y_j$. If $j \leq i$, then the second random walk also attains the value $Y_1 + \dots + Y_j$, so $M \leq M'$. If $j > i$, then the second random walk attains the value $Y_1 + \dots + Y_i + cY_{i+1} + \dots + cY_j$. Now $Y_{i+1} + \dots + Y_j \geq 0$, as otherwise $Y_1 + \dots + Y_i > Y_1 + \dots + Y_j$, in contradiction with M being the maximum. From this and $c \geq 1$ it follows that

$$Y_1 + \dots + Y_i + cY_{i+1} + \dots + cY_j \geq Y_1 + \dots + Y_j,$$

so $M \leq M'$. □

Lemma C.2. *Suppose Assumption 5.2 holds, and we use the schedule $\mathbf{x} = \boldsymbol{\mu} + \alpha\boldsymbol{\sigma}$. Let M_k be the all-time maximum of a random walk with i.i.d. steps distributed as $\sigma_k(B - \alpha)$. Then W_{k+1}^{SVF} is stochastically dominated by M_k , for all k .*

Proof. By (3), we know that W_{k+1}^{SVF} is the maximum of a random walk with steps $B_k - x_k, B_{k-1} - x_{k-1}, \dots, B_1 - x_1$. Now note that $B_i \stackrel{d}{=} \mu_i + \sigma_i B$ and $x_i = \mu_i + \alpha\sigma_i$, hence $B_i - x_i \stackrel{d}{=} \sigma_i(B - \alpha)$. So W_{k+1}^{SVF} can be represented as the maximum of a random walk with steps distributed as $\sigma_k(B - \alpha), \dots, \sigma_1(B - \alpha)$.

We first multiply the last step of this random walk σ_2/σ_1 . By Lemma C.1, we then see that W_{k+1} is stochastically dominated by the maximum of a random walk with steps distributed as $\sigma_k(B - \alpha), \dots, \sigma_2(B - \alpha), \sigma_2(B - \alpha)$. The next step is to multiply the last two steps with σ_3/σ_2 . Again by Lemma C.1, W_{k+1}^{SVF} is stochastically dominated by the maximum of a random walk with steps $\sigma_k(B - \alpha), \dots, \sigma_3(B - \alpha), \sigma_3(B - \alpha), \sigma_3(B - \alpha)$. Continuing in this way, we find that W_{k+1}^{SVF} is stochastically dominated by the maximum of a random walk with k steps distributed as $\sigma_k(B - \alpha)$.

Now note that adding extra steps to a random walk can only increase the maximum. Therefore, we now find that W_{k+1}^{SVF} is stochastically dominated by M_k . \square

Lemma C.3. *Under Assumption 5.2, when using the schedule $\mathbf{x} = \boldsymbol{\mu} + \alpha\boldsymbol{\sigma}$, we have for all k that*

$$\mathbb{E}W_{k+1}^{\text{SVF}} \leq \frac{\sigma_k}{2\alpha}.$$

Proof. We need the following bound by [21]. Let M be the all-time maximum of a random walk with i.i.d. steps distributed as Y , with $\mathbb{E}Y < 0$. Then

$$\mathbb{E}M \leq \frac{\text{Var}(Y)}{2|\mathbb{E}Y|}.$$

Let M_k be as in Lemma C.2. Then, by Lemma C.3 and Kingman's bound, we have

$$\mathbb{E}W_{k+1}^{\text{SVF}} \leq \mathbb{E}M_k \leq \frac{\text{Var}(\sigma_k(B - \alpha))}{2|\mathbb{E}(\sigma_k(B - \alpha))|} = \frac{\sigma_k^2}{2\sigma_k\alpha} = \frac{\sigma_k}{2\alpha},$$

as claimed. \square

Now we are ready to prove Proposition 5.4.

Proposition 5.4. *Suppose α is given by (15). Under Assumption 5.2,*

$$C_\omega(\text{id}, \boldsymbol{\mu} + \alpha\boldsymbol{\sigma}) \leq \sqrt{2\omega} \sum_{i=1}^{n-1} \sigma_i.$$

Proof. We already have a bound on the mean waiting time in Lemma C.3, so we proceed by considering the mean idle time. Taking expectations in (5) and using the fact that $x_i = \mu_i + \alpha\sigma_i$,

$$\sum_{i=1}^n \mathbb{E}I_i^{\text{SVF}} + \sum_{i=1}^n \mu_i = \sum_{i=1}^n \mu_i + \alpha \sum_{i=1}^{n-1} \sigma_i + \mathbb{E}W_n^{\text{SVF}}.$$

Hence, by virtue of Lemma C.3,

$$\sum_{i=1}^n \mathbb{E}I_i^{\text{SVF}} = \alpha \sum_{i=1}^{n-1} \sigma_i + \mathbb{E}W_n^{\text{SVF}} \leq \alpha \sum_{i=1}^{n-1} \sigma_i + \frac{\sigma_{n-1}}{2\alpha}.$$

Upon combining the bounds for the mean waiting times and the total mean idle time, we find that

$$C_\omega(\text{id}, \boldsymbol{\mu} + \alpha\boldsymbol{\sigma}) \leq \alpha\omega \sum_{i=1}^{n-1} \sigma_i + \frac{\omega}{2\alpha} \sigma_{n-1} + \frac{1-\omega}{2\alpha} \sum_{i=1}^{n-1} \sigma_i.$$

Through standard calculus we find that this upper bound is minimized for the α given in (15). The corresponding upper bound for this α is

$$\sqrt{2\omega(1-\omega)} \sqrt{\sum_{i=1}^{n-1} \sigma_i \left(\sum_{i=1}^{n-1} \sigma_i + \frac{\omega}{1-\omega} \sigma_{n-1} \right)}.$$

Using the fact that $\sigma_{n-1} \leq \sum_{i=1}^{n-1} \sigma_i$, we find the upper bound in Proposition 5.4. \square

In order to prove Proposition 5.5, we need the following lemma, which is known from the classical newsvendor problem (see e.g. [20]).

Lemma C.4. *Let X be a random variable, with Q_X its quantile function. Then for $\omega \in (0, 1)$,*

$$\omega \mathbb{E}(X - c)^- + (1 - \omega) \mathbb{E}(X - c)^+$$

is minimal for $c = Q_X(1 - \omega) = \inf\{x : 1 - \omega \leq \mathbb{P}(X \leq x)\}$.

Proposition 5.5. *Under Assumption 5.2, for any sequence and schedule,*

$$C_\omega(\tau, \mathbf{x}) \geq [\omega \mathbb{E}B(\omega)^- + (1 - \omega) \mathbb{E}B(\omega)^+] \sum_{i=1}^{n-1} \sigma_i.$$

Proof. Consider some arbitrary sequence τ and schedule \mathbf{x} . Recall that the idle and waiting times satisfy the recursions in (1). We consider

$$\omega \mathbb{E}I_{k+1} + (1 - \omega) \mathbb{E}W_{k+1} = \omega \mathbb{E}(W_k + B_{\tau(k)} - x_{\tau(k)})^- + (1 - \omega) \mathbb{E}(W_k + B_{\tau(k)} - x_{\tau(k)})^+. \quad (22)$$

Now note that $W_k + B_{\tau(k)} - x_{\tau(k)} = B_{\tau(k)} - (x_{\tau(k)} - W_k)$, so by minimizing (22) over all possible values of $x_{\tau(k)} - W_k$, we find by Lemma C.4 that

$$\omega \mathbb{E}I_{k+1} + (1 - \omega) \mathbb{E}W_{k+1} \geq \omega \mathbb{E}(B_{\tau(k)} - Q_{B_{\tau(k)}}(1 - \omega))^- + (1 - \omega) \mathbb{E}(B_{\tau(k)} - Q_{B_{\tau(k)}}(1 - \omega))^+.$$

Because $B_{\tau(k)} \stackrel{d}{=} \mu_{\tau(k)} + \sigma_{\tau(k)} B$, it follows that $Q_{B_{\tau(k)}}(1 - \omega) = \mu_{\tau(k)} + \sigma_{\tau(k)} Q_B(1 - \omega)$.

Then

$$B_{\tau(k)} - Q_{B_{\tau(k)}}(1 - \omega) \stackrel{d}{=} \mu_{\tau(k)} + \sigma_{\tau(k)} B - (\mu_{\tau(k)} + \sigma_{\tau(k)} Q_B(1 - \omega)) = \sigma_{\tau(k)} (B - Q_B(1 - \omega)),$$

which, recalling that $B(\omega) = B - Q_B(1 - \omega)$, leads to

$$\omega \mathbb{E}I_{k+1} + (1 - \omega) \mathbb{E}W_{k+1} \geq \sigma_{\tau(k)} [\omega \mathbb{E}B(\omega)^- + (1 - \omega) \mathbb{E}B(\omega)^+].$$

Note that $\sum_{i=1}^{n-1} \sigma_{\tau(i)} \geq \sum_{i=1}^{n-1} \sigma_i$, as the σ_i were put in increasing order. Now summing over k we find the lower bound of Proposition 5.5. \square

To prove Theorem 5.6 we need the following lemma, which fulfills the same role for lognormally distributed service times as Lemma C.2 does for location-scale families. We then prove the upper and lower bounds needed for Theorem 5.6 in two propositions.

Lemma C.5. *Suppose the B_i are lognormally distributed with $m_1 \leq \dots \leq m_n$ and $s_1^2 \leq \dots \leq s_n^2$, and that we use the schedule $\mathbf{x} = (1 + \alpha)\boldsymbol{\mu}$. Let M_k be the all-time maximum of a random walk with i.i.d. steps distributed as $B_k - x_k$. Then $\mathbb{E}W_{k+1}^{\text{SVF}} \leq \mathbb{E}M_k$.*

Proof. Let $M_k^{(i)}$ be the maximum of a random walk with steps distributed as $B_k - x_k, B_{k-1} - x_{k-1}, \dots, B_{i+1} - x_{i+1}$ followed by i steps distributed as $B_i - x_i$. We will prove that $\mathbb{E}M_k^{(i)} \leq \mathbb{E}M_k^{(i+1)}$ for $i = 1, \dots, k-1$. As $\mathbb{E}W_{k+1}^{\text{SVF}} = \mathbb{E}M_k^{(1)}$ and $\mathbb{E}M_k^{(k)} \leq \mathbb{E}M_k$ (adding extra steps increases the maximum), it then follows that $\mathbb{E}W_{k+1}^{\text{SVF}} \leq \mathbb{E}M_k$.

Consider the random walk with steps distributed as $B_k - x_k, B_{k-1} - x_{k-1}, \dots, B_{i+1} - x_{i+1}$, followed by i steps distributed as $B_i - x_i$. Let Z be a standard normal random variable. Note that

$$\begin{aligned} B_i - x_i &= B_i - (1 + \alpha)\mu_i \stackrel{\text{d}}{=} \exp(m_i + s_i Z) - (1 + \alpha) \exp(m_i + s_i^2/2) \\ &= \exp(m_i) [\exp(s_i Z) - (1 + \alpha) \exp(s_i^2/2)]. \end{aligned}$$

Use Lemma C.1 to show we get an upper bound on $M_k^{(i)}$ by replacing all steps distributed as $B_i - x_i$ by steps distributed as

$$X'_i := \exp(m_{i+1} + (s_{i+1}^2 - s_i^2)/2) [\exp(s_i Z) - (1 + \alpha) \exp(s_i^2/2)].$$

Let Z' be another standard normal random variable independent of Z . Then

$$B_{i+1} - x_{i+1} \stackrel{\text{d}}{=} \exp(m_{i+1} + s_i Z + \epsilon Z') - (1 + \alpha) \exp(m_{i+1} + (s_i^2 + \epsilon^2)/2)$$

with $\epsilon := \sqrt{s_{i+1}^2 - s_i^2}$, and

$$\begin{aligned} &\mathbb{E}[\exp(m_{i+1} + s_i Z + \epsilon Z') - (1 + \alpha) \exp(m_{i+1} + (s_i^2 + \epsilon^2)/2) | X'_i] \\ &= \mathbb{E}[\exp(m_{i+1} + s_i Z + \epsilon Z') - (1 + \alpha) \exp(m_{i+1} + (s_i^2 + \epsilon^2)/2) | Z] \\ &= \exp(m_i + \epsilon^2/2) [\exp(s_i Z) - (1 + \alpha) \exp(s_i^2/2)] = X'_i. \end{aligned}$$

It follows by Lemma 2.5 that $X'_i \leq_{\text{cx}} B_{i+1} - x_{i+1}$. As the maximum of a random walk is a convex function in each of the individual stepsizes, we can replace each step distributed as X'_i by one distributed as $B_{i+1} - x_{i+1}$. Therefore, $\mathbb{E}M_k^{(i)} \leq \mathbb{E}M_k^{(i+1)}$, which completes the proof. \square

Proposition C.6. *Suppose the B_i are lognormally distributed with $m_1 \leq \dots \leq m_n$ and $s_1^2 \leq \dots \leq s_n^2$. Suppose α is given by*

$$\alpha = \frac{1}{\sqrt{2\omega}} \sqrt{(\exp(s_{n-1}^2) - 1)}.$$

Then

$$C_\omega(\text{id}, (1 + \alpha)\boldsymbol{\mu}) \leq 2\omega\alpha \sum_{i=1}^{n-1} \exp(m_i + s_i^2/2).$$

Proof. With M_k as in Lemma C.5, we can now apply Kingman's bound to find

$$\begin{aligned} \mathbb{E}W_{k+1}^{\text{SVF}} &\leq \mathbb{E}M_k \leq \frac{\text{Var}(B_k)}{2|\mathbb{E}B_k - (1 + \alpha)\mathbb{E}B_k|} = \frac{1}{2\alpha} (\exp(s_k^2) - 1) \exp(m_k + s_k^2/2) \\ &\leq \frac{1}{2\alpha} (\exp(s_{n-1}^2) - 1) \exp(m_k + s_k^2/2). \end{aligned}$$

By equation (5) we also have

$$\begin{aligned} \sum_{i=1}^n \mathbb{E}I_i^{\text{SVF}} &= \alpha \sum_{i=1}^{n-1} \mathbb{E}B_i + \mathbb{E}W_n^{\text{SVF}} \\ &\leq \alpha \sum_{i=1}^{n-1} \exp(m_i + s_i^2/2) + \frac{1}{2\alpha} (\exp(s_{n-1}^2) - 1) \exp(m_{n-1} + s_{n-1}^2/2) \\ &\leq \alpha \sum_{i=1}^{n-1} \exp(m_i + s_i^2/2) + \frac{1}{2\alpha} (\exp(s_{n-1}^2) - 1) \sum_{i=1}^{n-1} \exp(m_i + s_i^2/2). \end{aligned}$$

In total, we then find

$$C_\omega(\text{id}, (1 + \alpha)\boldsymbol{x}) \leq \omega\alpha \sum_{i=1}^{n-1} \exp(m_i + s_i^2/2) + \frac{1}{2\alpha} (\exp(s_{n-1}^2) - 1) \sum_{i=1}^{n-1} \exp(m_i + s_i^2/2).$$

Minimizing this upper bound over α , we obtain the result. \square

Proposition C.7. *Suppose the B_i are lognormally distributed with $m_1 \leq \dots \leq m_n$ and $s_1^2 \leq \dots \leq s_n^2$. Then, for any sequence and schedule,*

$$C_\omega(\tau, \boldsymbol{x}) \geq [(1 - \omega)\mathbb{P}(Z \geq Q_Z(1 - \omega) - s_1) - \omega\mathbb{P}(Z \leq Q_Z(1 - \omega) - s_1)] \sum_{i=1}^{n-1} \exp(m_i + s_i^2/2).$$

Proof. Similar as in the location-scale family case, we can use Lemma C.4 to find

$$\begin{aligned} \omega\mathbb{E}I_{k+1} + (1 - \omega)\mathbb{E}W_{k+1} &= \omega\mathbb{E}(B_{\tau(k)} - (x_{\tau(k)} - W_k))^- + (1 - \omega)\mathbb{E}(B_{\tau(k)} - (x_{\tau(k)} - W_k))^+ \\ &\geq \omega\mathbb{E}(B_{\tau(k)} - Q_{B_{\tau(k)}}(1 - \omega))^- + (1 - \omega)\mathbb{E}(B_{\tau(k)} - Q_{B_{\tau(k)}}(1 - \omega))^+. \end{aligned}$$

Computing this lower bound, we find with Z being a standard normal random variable that

$$\omega\mathbb{E}I_{k+1} + (1 - \omega)\mathbb{E}W_{k+1}$$

$$\geq \exp(m_{\tau(k)} + s_{\tau(k)}^2/2) \left[(1 - \omega)\mathbb{P}(Z \geq Q_Z(1 - \omega) - s_{\tau(k)}) - \omega\mathbb{P}(Z \leq Q_Z(1 - \omega) - s_{\tau(k)}) \right].$$

It can be easily seen that

$$(1 - \omega)\mathbb{P}(Z \geq Q_Z(1 - \omega) - s_{\tau(k)}) - \omega\mathbb{P}(Z \leq Q_Z(1 - \omega) - s_{\tau(k)})$$

is an increasing function in $s_{\tau(k)}$, and so is minimal for $s_{\tau(k)} = s_1$. Using this and summing over k , we find the lower bound of the proposition. \square

Proof of Theorem 5.6. This follows directly from Proposition C.6 and Proposition C.7. \square

C.2. Proofs corresponding to Section 5.4. We proceed by providing a detailed argumentation behind the results presented in Section 5.4.

Lemma 5.7. *For the SVF sequence in the limit $c \rightarrow \infty$, the optimal interarrival times for groups 1 and 2 are $x^{(1)} = \frac{1}{2}\sqrt{2}$ and $x^{(2)} = a$ respectively. Patients in group 1 incur an average expected waiting time of $\mathbb{E}W^{(1)} = \frac{1}{2}\sqrt{2}$, whereas patients in group 2 incur no waiting time. This yields an average total cost of $\frac{1}{2}(a + \sqrt{2})$.*

Proof. First, consider group 1. Crucially, Kingman's bound is known to be tight for this distribution of $B^{(1)}$ as $c \rightarrow \infty$ ([9]). Thus

$$\lim_{c \rightarrow \infty} \mathbb{E}W^{(1)} = \frac{\text{Var } B^{(1)}}{2x^{(1)}} = \frac{1}{2x^{(1)}}.$$

The average cost for patients in group 1 is thus

$$\mathbb{E}I^{(1)} + \mathbb{E}W^{(1)} = x^{(1)} + \frac{1}{2x^{(1)}},$$

which is minimized by the choice $x^{(1)} = \frac{1}{2}\sqrt{2}$.

Now consider group 2. Note that $W^{(2)}$ is the maximum of a random walk with steps distributed as $B^{(2)} - x^{(2)}$. As a lower bound, we take the value of the random walk right before the step where $B^{(2)}$ takes value $-a$ for the first time. The number of steps we take then has a geometric distribution with probability $\frac{1}{2}$ (i.e. the probability of taking k steps is $(1/2)^{k+1}$ for $k = 0, 1, 2, \dots$), which has mean 1, and each of these steps has value $a - x^{(2)}$. Therefore, we find

$$\mathbb{E}I^{(2)} + \mathbb{E}W^{(2)} \geq x^{(2)} + (a - x^{(2)}) = a.$$

For $x^{(2)} = a$, this lower bound is achieved, as all steps $B^{(2)} - a$ can never be positive and so the waiting time is zero. In conclusion, the minimal average cost as $c \rightarrow \infty$ is $\frac{1}{2}(a + \sqrt{2})$, attained with $x^{(1)} = \frac{1}{2}\sqrt{2}$ and $x^{(2)} = a$. \square

Lemma 5.8. $\lim_{c \rightarrow \infty} \mathbb{E}V^{(1)} = \lim_{c \rightarrow \infty} \mathbb{E}V^{(2)}$.

Proof. By the Lindley recursion we have $\mathbb{E}V^{(2)} = \mathbb{E}(V^{(1)} + B^{(1)} - x^{(1)})^+$, and conditioning on the two values that $B^{(1)}$ can take we then find

$$\begin{aligned} \lim_{c \rightarrow \infty} \mathbb{E}V^{(2)} &= \lim_{c \rightarrow \infty} (1-p)\mathbb{E}\left(V^{(1)} - \frac{1}{c} - x^{(1)}\right)^+ + p\mathbb{E}\left(V^{(1)} + c - x^{(2)}\right)^+ \\ &\leq \lim_{c \rightarrow \infty} \mathbb{E}V^{(1)} + pc = \lim_{c \rightarrow \infty} \mathbb{E}V^{(1)} + \frac{c}{1+c^2} = \lim_{c \rightarrow \infty} \mathbb{E}V^{(1)}. \end{aligned}$$

We also have $\mathbb{E}V^{(2)} \geq \mathbb{E}V^{(1)}$, since $V^{(1)}$ results from the same random walk but with an extra step $B^{(2)} - x^{(2)} \leq 0$ in the beginning. \square

Lemma 5.9. $\lim_{c \rightarrow \infty} \mathbb{E}V^{(1)} \leq \lim_{c \rightarrow \infty} \mathbb{E}\tilde{V}$.

Proof. Let $M^{(1)}$ be the maximum of $S^{(1)}$ over the first $2c$ steps only. We will first show that $\mathbb{E}[V^{(1)} - M^{(1)}] \rightarrow 0$ as $c \rightarrow \infty$. Recall that by Kingman's bound, we have the rough estimate $\mathbb{E}V^{(1)} \leq K$, where $K = (1+a^2)/(2a+\sqrt{2})$. Let E be the event that $S_{2c}^{(1)} \leq -2c$, and let \bar{E} be its complement. By Chebyshev's inequality,

$$\mathbb{P}(\bar{E}) = \mathbb{P}\left(S_{2c}^{(1)} - \mathbb{E}S_{2c}^{(1)} > c(2a + \sqrt{2} - 2)\right) \leq \frac{2c(1+a^2)}{(c(2a + \sqrt{2} - 2))^2} = O(c^{-1}).$$

Let Z be the maximum of the random walk obtained by removing the first $2c$ steps from $(S^{(1)})_{j \geq 0}$. Clearly Z has the same law as $V^{(1)}$, and $V^{(1)} - M^{(1)} \leq Z$, so that

$$\mathbb{E}\left[V^{(1)} - M^{(1)} \mid \bar{E}\right] \leq \mathbb{E}Z \leq K.$$

Condition now on E occurring. Then $V^{(1)} - M^{(1)} \leq (Z - 2c)^+$. Let R be the event that $Z \geq c$. When the random walk defining Z first crosses level c , it will still be below level $2c$. As the random walk after crossing level c has the same law as the random walk defining Z , we find that

$$\mathbb{E}\left[(Z - 2c)^+\right] = \mathbb{E}\left[(Z - 2c)^+ \mid R\right] \mathbb{P}(R) \leq \mathbb{E}[Z] \mathbb{P}(R) \leq K \mathbb{P}(R).$$

By Markov's inequality, $\mathbb{P}(R) \leq \mathbb{E}Z/c \leq K/c$. Combining the above,

$$\mathbb{E}\left[V^{(1)} - M^{(1)}\right] \leq K\mathbb{P}(\bar{E}) + K\mathbb{P}(R)\mathbb{P}(E) = O(c^{-1}),$$

which vanishes as $c \rightarrow \infty$, as required.

Let \tilde{M} be the maximum of \tilde{S} over the first $2c$ steps. It now suffices to show that $\mathbb{E}[M^{(1)} - \tilde{M}] \rightarrow 0$ as $c \rightarrow \infty$. Let F be the event that there is at least one upwards jump in

the first $2c$ steps. Then,

$$\mathbb{P}(F) = 1 - (1 - p)^{2c} = 2pc + O((pc)^2) = \frac{2}{c} + O(c^{-2}).$$

If F does not occur, then $M^{(1)} = \tilde{M} = 0$. We therefore concentrate on the scenario that F does apply. Note that the increments $B_i^{(2)}$ are independent of F . We have

$$S_j^{(1)} - \tilde{S}_j = \sum_{i=1}^j B_i^{(2)}.$$

By Hoeffding's inequality, $\mathbb{P}(S_j^{(1)} - \tilde{S}_j > t) \leq e^{-t^2/(ja)}$, and so by a union bound,

$$\mathbb{P}(M^{(1)} - \tilde{M} > c^{2/3}) \leq 2c \cdot e^{-c^{1/3}/(2a)}.$$

Since it is always the case that $M^{(1)} - \tilde{M} \leq 2ca$, we have that

$$\mathbb{E}[M^{(1)} - \tilde{M} | F] \leq 4c^2 a e^{-c^{1/3}/(2a)} + c^{2/3} = O(c^{2/3}).$$

We thus obtain that $\mathbb{E}[M^{(1)} - \tilde{M}] = \mathbb{E}[M^{(1)} - \tilde{M} | F] \mathbb{P}(F) = O(c^{-1/3})$, which vanishes as $c \rightarrow \infty$, as required. \square

APPENDIX D. LOGNORMAL DISTRIBUTIONS THAT SATISFY THE DILATION ORDER

Proposition D.1. *Suppose that A and B are lognormally distributed random variables, such that $\mathbb{E}[\ln A] \leq \mathbb{E}[\ln B]$ and $\text{Var}(\ln A) \leq \text{Var}(\ln B)$. Then $A \leq_{\text{dil}} B$.*

Proof. Let $m_1 = \mathbb{E}[\ln A]$, $m_2 = \mathbb{E}[\ln B]$, $s_1^2 = \text{Var}(\ln A)$ and $s_2^2 = \text{Var}(\ln B)$. Let Z and Z' be two independent standard normal random variables, and let $\epsilon = \sqrt{s_2^2 - s_1^2}$. Then

$$\begin{aligned} A - \mathbb{E}A &\stackrel{d}{=} \hat{X} := \exp(m_1 + s_1 Z) - \exp(m_1 + s_1^2/2), \\ B - \mathbb{E}B &\stackrel{d}{=} \hat{Y} := \exp(m_2 + s_1 Z + \epsilon Z') - \exp(m_2 + (s_1^2 + \epsilon^2)/2). \end{aligned}$$

Now note that

$$\begin{aligned} \mathbb{E}[\hat{Y} | \hat{X}] &= \mathbb{E}[\hat{Y} | Z] = \mathbb{E}[\exp(m_2 + s_1 Z + \epsilon Z') - \exp(m_2 + (s_1^2 + \epsilon^2)/2) | Z] \\ &= \exp(m_2 + \epsilon^2/2) (\exp(s_1 Z) - \exp(s_1^2/2)) = \exp(m_2 - m_1 + \epsilon^2/2) \hat{X}. \end{aligned}$$

Thus by Lemma 2.5,

$$\exp(m_2 - m_1 + \epsilon^2/2)(A - \mathbb{E}A) \leq_{\text{cx}} B - \mathbb{E}B.$$

(This is similar to the argument that $X'_i \leq_{\text{cx}} B_{i+1} - x_{i+1}$ in the proof of Lemma C.5.) Now we can apply Theorem 3.A.18 from [34], that says that $X \leq_{\text{dil}} aX$ for $a \geq 1$, to see that

$$A - \mathbb{E}A \leq_{\text{cx}} \exp(m_2 - m_1 + \epsilon^2/2)(A - \mathbb{E}A) \leq_{\text{cx}} B - \mathbb{E}B.$$

This proves that $A \leq_{\text{dil}} B$. □

APPENDIX E. NUMERICAL EXPERIMENTS FOR MEAN-BASED SCHEDULES

In this appendix we describe numerical experiments for exponential and lognormal service time distributions, in which the approximation ratio ϱ_ω is computed. The goal is to gain insight into what performance can be expected for the SVF sequence in practice. As the mean-based schedule is in place, (5) implies that the total expected idle time is equal to $\mathbb{E}W_n$, the expected waiting time of the last patient. The SVF rule is expected to be good for the patients in the beginning of the sequence, as then the previous patient is less variable, but might not be as good for patients at the end of the sequence, as they have high variance patients directly before them. Therefore, we expect the idle time to suffer more than the waiting times when the SVF sequence is used instead of the optimal sequence. This explains why we set $\omega = 1$ unless stated otherwise.

This choice of ω is further motivated by the results in [24]. They constructed examples where SVF is not optimal when mean-based schedules are used, for an objective function that is a weighted average of expected waiting time and expected overtime. When waiting time and overtime had equal weight, typically more than 50 patients were needed for their counterexamples (refuting SVF being optimal, that is). When the overtime cost was much larger than the waiting time cost, much fewer patients were needed in creating counterexamples.

The following proposition reduces the complexity of finding the optimal sequence.

Proposition E.1. *Under Assumption 2.4 (dilation ordering assumption) and when mean-based schedules are used, there is an optimal sequence τ that satisfies $\tau(n) = n$, i.e. the largest variance should always be sequenced last.*

Proof. Suppose τ is an optimal sequence. Note that $W_{k+1} = \max\{0, S_1, \dots, S_k\}$ is a convex function in each of the X_i . If $\tau(n) = n$ we are done. If $\tau(n) \neq n$, then under Assumption 2.4 we have $X_{\tau(n)} \leq_{\text{cx}} X_n$, and so switching the positions of patient n and patient $\tau(n)$ can only decrease $\mathbb{E}W_{k+1}$. As this is true for all $k = 1, \dots, n-1$, and the total idle time equals $\mathbb{E}W_n$, switching patients $\tau(n)$ and n can only decrease the cost. The resulting sequence after the switch is therefore also optimal, and we are done. □

n	optimal sequence	optimal cost	cost for SVF	approximation ratio
3	1,2,3	0.2646	0.2646	1
4	1,2,3,4	0.3098	0.3098	1
5	2,1,3,4,5	0.3388	0.3389	1.0003
6	3,1,2,4,5,6	0.3588	0.3590	1.0007
7	4,2,1,3,5,6,7	0.3735	0.3739	1.0011
8	5,3,1,2,4,6,7,8	0.3847	0.3853	1.0014
9	6,4,2,1,3,5,7,8,9	0.3936	0.3943	1.0017
10	7,5,3,1,2,4,6,8,9,10	0.4008	0.4015	1.0019
11	8,5,3,1,2,4,6,7,9,10,11	0.4067	0.4076	1.0021

Table 4. The optimal sequence, optimal cost, cost of the SVF sequence, and approximation ratio ϱ_1 when patient i has an exponentially distributed service time with rate $n + 1 - i$, and mean-based schedules are used.

This proposition still leaves $(n - 1)!$ potential candidates for the optimal sequence. We are not aware of any other generally applicable structural result that can help to further reduce this number. As computing ϱ_ω thus requires checking all $(n - 1)!$ potential candidates to find the optimal sequence, this severely limits the numbers of n for which we can do this.

E.1. Exponential distributions. We first look at the case where the service times are exponentially distributed, as this makes it possible to efficiently calculate the cost for a given sequence and schedule [38]. We set the service rate of patient i to $n + 1 - i$, i.e., $\mathbb{E}B_i = 1/(n + 1 - i)$. This way, the service rates are very different from one another, so that not using the optimal sequence could potentially have a large impact. The optimal sequence is found by complete enumeration over the $(n - 1)!$ candidates. This allowed us to compute the approximation ratio ϱ_1 for problem instances with $n = 3, \dots, 11$ (within a reasonable amount of time). The results are given in Table 4. Choosing ω smaller than 1 resulted in smaller approximation ratios, as suspected, with SVF being optimal for all considered problem instances when $\omega = 0.9$.

For $n = 9$ patients we also computed the approximation ratio for 100 random problem instances, where each service time was exponentially distributed with a service rate drawn independently from a uniform distribution on $[0, 1]$. The largest approximation ratio found this way was 1.0034. For exponentially distributed service times it thus seems that the SVF rule will typically perform within 0.5% of optimal.

From Table 4 we notice that the optimal sequence for each problem instance is “V-shaped”: the variances first decrease, and later increase. We also checked that the optimal sequences for our 100 random problem instances were all V-shaped. Restricting to V-shaped schedules allows us to consider substantially larger values of n . For instances with service

n	optimal sequence	optimal cost	cost for SVF	approximation ratio
3	1,2,3	18.0131	18.0131	1
4	2,1,3,4	32.6526	32.5799	1.0022
5	3,1,2,4,5	50.3769	50.1484	1.0046
6	4,2,1,3,5,6	70.8629	70.4453	1.0059
7	5,3,1,2,4,6,7	93.8700	93.2174	1.0070
8	6,4,2,1,3,5,7,8	119.210	118.302	1.0077
9	7,5,3,1,2,4,6,8,9	146.730	145.538	1.0082
10	8,5,3,1,2,4,6,7,9,10	176.305	174.803	1.0086

Table 5. The optimal sequence, optimal cost, cost of the SVF sequence, and approximation ratio ρ_1 when patient i has a service time that is a discrete approximation to a lognormal distribution with $m_i = \ln(50) + \ln(i)$ and $s_i = 0.33$, and mean-based schedules are used.

times $n + 1 - i$ for patient i , we obtained a bound of 1.0038 for $n = 23$, and for randomly chosen rates, found nothing worse than 1.0048 for $n = 17$. However, we are aware of examples with other distributions for which the optimal schedule is not V-shaped, and so the true value of ρ_ω may be slightly larger.

E.2. Lognormal distributions. Lognormal service times are intensively used in health care [6, 22]. For these distributions there is no clear efficient way to exactly compute the cost. We therefore work with its discrete counterpart, where service times are only allowed to take integer values. The lognormal distribution for patient i has parameters $m_i := \mathbb{E} \ln B_i$ and $s_i^2 := \text{Var} \ln B_i$. Motivated by the two lognormal distributions found to fit real data in [6] (having s_i equal to 0.3169 and 0.3480), in our experiments we have fixed $s_i = 0.33$. We set $m_i = \ln(50) + \ln(i)$, where the 50 is just a scaling parameter meant to improve the discrete approximation. If the time unit would be minutes, this means we round to full minutes while all mean service times are over 50 minutes, in line with the accuracy one expects to be found in practice.

For these discrete random variables, we found the optimal sequence for problem instances with $n = 3, 4, \dots, 10$. The results are reported in Table 5. For each of these problem instances the approximation ratio was below 1.01. We also performed the same experiments for $s_i = 0.75$ for all patients (instead of 0.33), for $n = 3, \dots, 10$. The approximation ratios were even smaller than in the $s_i = 0.33$ case, with the largest approximation ratio being 1.0021. When the same experiments were repeated for $\omega = 0.9$, we found that the SVF sequence is optimal for all problem instances. The experiments are indicative of SVF performing quite well for lognormal distributions, but not as well as for exponential distributions.

To gain insight into the effect of the service-time distribution having a tail on either the left or the right, we also assessed the effect of replacing the service times B_i by the flipped

distribution $-B_i$. For $n = 10$ and $s_i = 0.33$ we then found an approximation ratio of 1.0345, which, while still very close to 1, is much larger than the 1.0086 reported for the non-flipped lognormal distributions.

We also performed experiments based on the lognormal distributions fitted to data in [6], distinguishing between “new” and “return” patients. Lognormal distributions with parameters indicated in Table 2 were found to yield a good fit to the data for these two groups. In line with [6], we consider a session consisting of 10 patients. For our numerical experiments we used discrete approximations to these two lognormal distributions, with service times being rounded to full minutes. We computed the approximation ratio ρ_1 for 11 problem instances, each with 10 patients in total with 0 up to 10 of them corresponding to new patients (and the others to return patients). The largest approximation ratio we found was $\rho_1 = 1.0023$ for 7 new patients and 3 return patients. We thus conclude that the SVF rule performs well for the practical distributions proposed in [6].

APPENDIX F. NUMERICAL EXPERIMENTS FOR OPTIMALLY-SPACED SCHEDULES

Here we consider some numerical experiments when an optimally-spaced schedule is used for each sequence. Similar to Appendix E, we will consider exponential service-time distributions and a discrete approximation to lognormal distributions. In order to find the optimal schedule, we note that each waiting time is a convex function in each of the x_i , as by (3) it is the maximum of linear functions in x_i . Note that the case $\omega = 1$ is no longer interesting, as this results in an optimal schedule with $x_i = 0$ for all i , resulting in a cost of zero for any sequence. Instead, we consider the parameter values $\omega = 0.5$, $\omega = 0.8$ and $\omega = 0.9$.

Again we have, similar to Proposition E.1, the result that the largest variance always should be sequenced last, as stated in the next proposition. This again leaves us with $(n - 1)!$ candidates, and again it seems no more structural results are known, so we can only find the optimal sequence by enumerating over all $(n - 1)!$ candidates.

Proposition F.1. *Under Assumption 2.4 (dilation ordering assumption), there is an optimal combination of sequence τ and schedule \mathbf{x} that satisfies $\tau(n) = n$, i.e. the largest variance should always be sequenced last.*

Proof. Suppose τ and \mathbf{x} are an optimal combination of sequence and schedule. If $\tau(n) = n$ we are done. If $\tau(n) \neq n$, then under Assumption 2.4 we have $B_{\tau(n)} - \mathbb{E}B_{\tau(n)} \leq_{\text{cx}} B_n - \mathbb{E}B_n$. We then propose a new sequence τ' with patients $\tau(n)$ and n switched (so now $\tau'(n) = n$), and a new schedule \mathbf{y} with $y_{\tau(n)} = x_n + \mathbb{E}B_{\tau(n)} - \mathbb{E}B_n$ and $y_i = x_i$ for $i \neq \tau(n)$. In the

expression (3) for W_{k+1} this change effectively makes it so that each term $B_n - \mathbb{E}B_n$ gets replaced by $B_{\tau(n)} - \mathbb{E}B_{\tau(n)}$, and as W_{k+1} is a convex function in this expression we see that this change cannot increase the $\mathbb{E}W_{k+1}$. By expression (5) we see that the expected idle time only changes through the term $\mathbb{E}W_n$, and so the total cost will not be increased after the change. Therefore the sequence τ' with schedule \mathbf{y} is also optimal, and as $\tau'(n) = n$ we are done. \square

F.1. Exponential distributions. Suppose that, as in Appendix E, the service time of patient i is exponentially distributed with rate $n + 1 - i$. The exponential service times allow us to efficiently compute the cost for any given schedule, and the convexity of the cost in each of the interarrival times x_i then allows us to compute the optimal schedule. Due to the extra complexity involved in computing the optimal schedule for each possible sequence, we only computed the approximation ratio for $n = 3, \dots, 8$, for each of mentioned parameter values of ω . For each of these problem instances, we found that the SVF sequence is optimal.

F.2. Lognormal distributions. For lognormal distributions, we will use the same discrete approximation as was used in Appendix E. For patient i , the underlying lognormal distribution is set to have $m_i := \mathbb{E} \ln B_i = \ln(20) + \ln(i)$, the 20 again being a scaling parameter. We need to restrict ourselves to schedules with integer-valued interarrival times in order to be able to effectively compute the cost function for a given schedule. Fortunately, this class of schedules is known to contain the optimal schedule, and the cost function of the resulting discrete optimization problem is L-convex [2]. We can thus use the ODICON solver [36], that can minimize L-convex functions, to find the optimal schedule for each sequence. For $s_i = 0.33$ (s_i being the standard deviation of $\ln B_i$), we computed the approximation ratio for $n = 3, \dots, 7$, for each of the mentioned values of ω . When $s_i = 0.75$, we did the same for $n = 3, \dots, 6$. For all considered problem instances, we found that the SVF sequence is optimal.

Remark F.2. The method to find the optimal interarrival times described here, with the use of ODICON, works for any discrete service-time distribution. It was also used to find the optimal interarrival times in Example 5.1.

REFERENCES

- [1] Ahmadi-Javid A, Jalali Z, Klassen K (2017) Outpatient appointment systems in healthcare: A review of optimization studies. *European Journal of Operational Research* 258(1):3–34.
- [2] Begen MA, Queyranne M (2011) Appointment scheduling with discrete random durations. *Mathematics of Operations Research* 36(2):240–257.

- [3] Berg B, Denton B, Erdogan S, Rohleder T, Huschka T (2014) Optimal booking and scheduling in outpatient procedure centers. *Computers & Operations Research* 50:24–37.
- [4] Billingsley P (1995) *Probability and Measure*, Third edition (Wiley, New York).
- [5] Çayırılı T, Veral E (2003) Outpatient scheduling in health care: a review of literature. *Production and Operations Management* 12(4):519–549.
- [6] Çayırılı T, Veral E, Rosen H (2006) Designing appointment scheduling systems for ambulatory care services. *Health Care Management Science* 9(1):47–58.
- [7] Çayırılı T, Veral E, Rosen H (2008) Assessment of patient classification in appointment system design. *Production and Operations Management* 17(3):338–353.
- [8] Charnetski J (1984) Scheduling operating room surgical procedures with early and late completion penalty costs. *Journal of Operations Management* 5(1):91–102.
- [9] Daley DJ (1997) Some results for the mean waiting-time and workload in GI/GI/k queue. *Dshalalow JH (ed.) Frontiers in Queueing: Models and Applications in Science and Engineering* (CRC Press, Boca Raton), 35–59.
- [10] Denton B, Gupta D (2003) A sequential bounding approach for optimal appointment scheduling. *IIE Transactions* 35(11):1003–1016.
- [11] Denton B, Viapiano J, Vogl A (2007) Optimization of surgery sequencing and scheduling decisions under uncertainty. *Health Care Management Science* 10(1):13–24.
- [12] Erdogan A, Gose A, Denton B (2015) Online appointment sequencing and scheduling. *IIE Transactions* 47(11):1267–1286.
- [13] Guda H, Dawande M, Janakiraman G, Jung KS (2016) Optimal policy for a stochastic scheduling problem with applications to surgical scheduling. *Production and Operations Management* 25(7):1194–1202.
- [14] Gupta D (2007) Surgical suites’ operations management. *Production and Operations Management* 16(6):689–700.
- [15] Gupta D, Denton B (2008) Appointment scheduling in health care: Challenges and opportunities. *IIE Transactions* 40(9):800–819.
- [16] Gupta A, Kumar A, Nagarajan V, Shen X (2018) Stochastic load balancing on unrelated machines. *Proceedings of the SODA 2018 Conference*, 1274–1285.
- [17] Gupta V, Moseley B, Uetz M, Xie Q (2017) Stochastic online scheduling on unrelated machines. *Integer Programming and Combinatorial Optimization* (Springer International Publishing), 228–240.
- [18] Heyde CC, Brown BM (1970) On the departure from normality of a certain class of martingales. *Annals of Mathematical Statistics* 41(6):2161–2165.

- [19] Kemper B, Klaassen CAJ, Mandjes M (2014) Optimized appointment scheduling. *European Journal of Operational Research* 239(1):243–255.
- [20] Khouja M (1999) The single-period (news-vendor) problem: literature review and suggestions for future research. *Omega* 27:537–553.
- [21] Kingman JFC (1962) Some inequalities for the queue GI/G/1. *Biometrika* 49(3-4):315–324.
- [22] Klassen K, Rohleder T (1996) Scheduling outpatient appointments in a dynamic environment. *Journal of Operations Management* 14(2):83–101.
- [23] Kong Q, Lee C, Teo C, Zheng Z (2013) Scheduling arrivals to a stochastic service delivery system using copositive cones. *Operations Research* 61(3):711–726.
- [24] Kong Q, Lee C, Teo C, Zheng Z (2016) Appointment sequencing: Why the smallest-variance-first rule may not be optimal. *European Journal of Operational Research* 255(3):809–821.
- [25] Kuiper A, Kemper B, Mandjes M (2015) A computational approach to optimized appointment scheduling. *Queueing Systems* 79(1):5–36.
- [26] Lindley DV (1952) The theory of queues with a single server. *Mathematical Proceedings of the Cambridge Philosophical Society* 48(2):227–289.
- [27] Mak H, Rong Y, Zhang J (2014) Sequencing appointments for service systems using inventory approximations. *Manufacturing & Service Operations Management* 16(2):251–262.
- [28] Mak H, Rong Y, Zhang J (2015) Appointment scheduling with limited distributional information. *Management Science* 61(2):316–334.
- [29] Mancilla C, Storer R (2012) A sample average approximation approach to stochastic appointment sequencing and scheduling. *IIE Transactions* 44(8):655–670.
- [30] Mittal S, Schulz AS, Stiller S (2014) Robust appointment scheduling. *Proceedings of APPROX/RANDOM 2014*, 356–370.
- [31] Pinedo M (2012) *Scheduling: Theory, Algorithms, and Systems* (Springer, New York).
- [32] Robinson LW, Chen RR (2011) Estimating the implied value of the customer’s waiting time. *Manufacturing & Service Operations Management* 13(1):53–57.
- [33] Rohleder T, Klassen K (2000) Using client-variance information to improve dynamic appointment scheduling performance. *Omega* 28(3):293–302.
- [34] Shaked M, Shanthikumar JG (2007) *Stochastic Orders* (Springer, New York).
- [35] Skutella M, Sviridenko M, Uetz M (2016) Stochastic scheduling on unrelated machines. *Mathematics of Operations Research* 41(3):851–864.

- [36] Tsuchimura N, Moriguchi S, Murota K (2013) Discrete convex optimization solvers and demonstration softwares. *Transactions of the Japan Society for Industrial and Applied Mathematics* 23(2):233–252.
- [37] Vanden Bosch PM, Dietz DC (2000) Minimizing expected waiting in a medical appointment system. *IIE Transactions* 32(9):841–848.
- [38] Wang PP (1998) Sequencing and scheduling N customers for a stochastic server. *European Journal of Operational Research* 255(3):809–821.
- [39] Weiss EN (1990) Models for determining estimated start times and case orderings in hospital operating rooms. *IIE Transactions* 22(2):143–150.