Models, algorithms and performance analysis for adaptive operating room scheduling

Guanlian Xiao, Willem van Jaarsveld, Ming Dong & Joris van de Klundert

Published online: 19 Jun 2017.

Submit your article to this journal ⟶

Article views: 3189

View related articles ⟶

View Crossmark data ⟶

Citing articles: 2 View citing articles ⟶

Taylor & Francis
Taylor & Francis Group

Check for updates

Models, algorithms and performance analysis for adaptive operating room scheduling

Guanlian Xiao[a,b], Willem van Jaarsveld[a], Ming Dong[c] and Joris van de Klundert[b*]

[a]*Industrial Engineering & Innovation Sciences, Eindhoven University of Technology, Eindhoven, Netherlands;* [b]*Department of Health Services Management and Organization, Institute of Health Policy and Management, Erasmus University Rotterdam, Rotterdam, The Netherlands;* [c]*Department of Operations Management, Antai College of Economics and Management, Shanghai Jiao Tong University, Shanghai, China*

The complex optimisation problems arising in the scheduling of operating rooms have received considerable attention in recent scientific literature because of their impact on costs, revenues and patient health. For an important part, the complexity stems from the stochastic nature of the problem. In practice, this stochastic nature often leads to schedule adaptations on the day of schedule execution. While operating room performance is thus importantly affected by such adaptations, decision-making on adaptations is hardly addressed in scientific literature. Building on previous literature on adaptive scheduling, we develop adaptive operating room scheduling models and problems, and analyse the performance of corresponding adaptive scheduling policies. As previously proposed (fully) adaptive scheduling models and policies are infeasible in operating room scheduling practice, we extend adaptive scheduling theory by introducing the novel concept of committing. Moreover, the core of the proposed adaptive policies with committing is formed by a new, exact, pseudo-polynomial algorithm to solve a general class of stochastic knapsack problems. Using these theoretical advances, we present performance analysis on practical problems, using data from existing literature as well as real-life data from the largest academic medical centre in The Netherlands. The analysis shows that the practically feasible, basic, 1-level policy already brings substantial and statistically significant improvement over static policies. Moreover, as a rule of thumb, scheduling surgeries with large mean duration or standard deviation early appears good practice.

Keywords: operating room scheduling; adaptive scheduling; stochastic knapsack problem; operations research in healthcare

1. Introduction

Globally, the demand for health services increases as societies advance economically and populations age. This holds particularly true for the demand for hospital services, and hence for the surgical services provided in operating rooms (ORs). The services provided in ORs form a primary source of hospital revenue and drive a large share of hospital costs (Denton, Viapiano, and Vogl 2007; Cardoen, Demeulemeester, and Beliën 2010). As a result, the planning and scheduling of ORs are continuous and high priority activities in many hospitals around the globe. OR planning and scheduling have also received considerable attention in scientific literature, as demonstrated by the review of Cardoen, Demeulemeester, and Beliën (2010), which has identified 246 publications, among which 114 in the first decade of the present millennium. Other comprehensive reviews which address OR scheduling are provided by Erdogan et al. (2011), Gupta (2007), Gupta and Denton (2008) and Cayirli and Veral (2003).

In this work, we consider the scarcely researched common practice of making adaptations to the surgical schedule while schedule execution is in progression (May et al. 2011). Evidence shows that the stochastic nature of the operating room processes causes the effectiveness and efficiency of operating room performance to be importantly impacted by such adaptation decisions (Stepaniak, Mannaerts, et al. 2009). Aiming to contribute to practical performance improvement, we first develop required scientific models and algorithms, and subsequently analyse resulting improvement obtained for real-life problems. From the perspective of the operating room scheduling classification scheme of Cardoen, Demeulemeester, and Beliën (2010), the presented research thus addresses operating room uncertainty.

Strategic planning is commonly considered as the highest level in hierarchical production planning frameworks, as is also the case in hospital production management (Vissers, Bertrand, and De Vries 2001). At the strategic level, hospital management decides which specialties to offer, how many operating rooms to run and on capital intensive equipment. At the next lower level, the tactical level, hospital management divides the resulting operating room capacity, often in the form of periodic rosters which specify slots (or blocks) of operating room time to specialties or individual doctors. For instance,

---

Orthopaedics may get surgery time assigned every Monday from 8 am to 4 pm in room 2. At the subsequent operational level, specialties or doctors select for each slot which patients to operate, in alignment with the operating room planners. Prior to the day of operation, the final operating room schedule typically specifies a sequence in which the selected patients will be treated, as well as expected start and end times for each patient. The stochastic, uncertain nature of surgical operations causes the schedules to often require adaptation while being in execution. This brings us at the executional, bottom, level of the production planning hierarchy.

Our work focuses on making optimal use of adaptivity, at the scheduling level and at the execution level. Because in practice adaptations are often confined to the slots assigned to a specialty or individual doctor, we develop models and methods within the context of a single operating room on a single day. Within such a slot, some of the scheduled surgeries may require much more time, others may take much less time, patients are cancelled as they are unfit for surgery, emergency patients may have to be inserted into the schedule, et cetera. Such events may lead to adaptations in the form of changes in schedule times, order, deletion, insertion, and addition of patients, and of cancelling patients at the end of the day. More precisely classifying the presented research in terms of the classification scheme of Cardoen, Demeulemeester, and Beliën (2010), we consider the problem of scheduling elective and emergent patients (field patient characteristics) assigned to a single operating room (field decision delineation) using exact combinatorial methods (field research methods), and test our methods on two real data-sets of empirical origin (field applicability of research).

The review by Cardoen, Demeulemeester, and Beliën (2010) already demonstrates that much scientific progress has been made at various of the aforementioned planning levels to improve operating room performance. Moreover, further recent research to better cope with uncertainty is provided by Tang and Wang (2015) who consider the tactical planning problem of robustly allocating limited operating room capacity to subspecialties so as to provide timely and accessible treatment. Another tactical problem is considered by Dellaert, Cayiroglu, and Jeunet (2016), who strive to balance waiting time and resource utilisation. Jebali and Diabat (2015) consider a problem of a more operational nature, namely to select elective surgeries, while considering uncertainties related to surgery duration as well as patient length of stay in the ICU and the ward.

Despite adaptations being common practice, it has however received little attention in the scientific literature. Zhou et al. (1999) and Dexter et al. (2004), Dexter, Willemsen-Dunlap, and Lee (2007) address deterministic solution approaches and analyse methods to construct a priori schedules using historic average durations as expected surgery durations. These authors show that solving the resulting static deterministic OR scheduling problems may result in poor OR utilisation and high patient waiting times. A natural subsequent improvement is to consider static stochastic approaches, as proposed in Denton, Viapiano, and Vogl (2007) and Batun et al. (2011). These authors propose one stage programmes, or two stage stochastic programmes with decisions taken in the first stage. Alternatively, Shylo, Prokopyev, and Schaefer (2012) propose a chance-constrained-based approach which minimises idle time while constraining the probability of overtime work. Wang, Tang, and Fung (2014) consider a chance-constrained model for a closely related problem with multiple operating rooms. van Oostrum et al. (2008) and Van Essen et al. (2014) also use chance constraints, and focus on the impact of the OR schedule on ward occupancy and overtime. These stochastic models capture important practically relevant uncertainties of OR scheduling.

All of the aforementioned work however considers the problem as static, or off-line, and disregards the common practice of adapting a priori schedules as the day proceeds (May et al. 2011). A first study which includes such adaptations is the simulation study by Stepaniak et al. (2012), who analyse the impact of risk attitudes of schedulers on their adaptation decision-making and the resulting operating room performance. Another simulation study which analyses adaptive policies is provided by Vermeulen et al. (2009), who consider a planning problem arising in the context of allocating CT scan capacity over different patient groups. Zhang, Xie, and Geng (2014) investigate dynamic allocation of surgeons/surgeries to multiple operating rooms. In contrast to the problem we study however, their model assumes that all selected patients must be treated. More closely related is the recent work of Xiao et al. (2016) who solve a version of the problem where adaptation is allowed only once per day, using a computationally intensive three-stage stochastic recourse model. We also refer to their work for a general discussion of the connection between stochastic OR scheduling and the stochastic scheduling literature. For the OR scheduling literature, a general analytic approach towards adaptivity in operating room schedules is lacking in the scientific literature, despite being common practice.

Our main contributions are twofold. First, we develop adaptive models and solution methods for single day single operating room scheduling problems and present theoretical results on their relative performance. Second, we show that adaptive policies can substantially improve over non-adaptive policies in practice, using two real-life data-sets from Erasmus Medical Center (one of which is available via scientific literature). We reflect on practically relevant insights and guidelines in the discussion and conclusions.

2. Embedding in adaptive scheduling literature

Ilhan, Iravani, and Daskin (2011) consider adaptive policies for a knapsack problem with deterministic weights and stochastic rewards. The single operating room scheduling we consider in this paper differs from their problem as it has stochastic weights and deterministic rewards. Adaptivity has also been widely studied for settings with weights that are unknown until the item is inserted in the knapsack, and where items are added one by one until the capacity is exceeded. In such problems, it is often assumed that all items added until capacity is exceeded contribute to total reward (Dean, Goemans, and Vondrák 2008; Bhalgat, Goel, and Khanna 2011; Levin and Vainer 2013; Balseiro and Brown 2016; Blado, Hu, and Toriello 2016), or that exceeding the capacity of the knapsack results in a final reward of zero (Chen and Ross 2014). In accordance with operating room scheduling practice, we model exceeding capacity using overtime penalties (cost).

Because adaptive problems are of a recursive nature, their solution can be computationally demanding. To help control computation times, Ilhan, Iravani, and Daskin (2011) introduce the concept of l-level adaptivity, and propose and analyse corresponding l-level heuristics. l-level adaptivity is defined by restricting adaptation to the time of completion of the next l jobs (viz. patients). We adopt the l-level adaptivity framework of Ilhan, Iravani, and Daskin (2011), yet translate it to the setting of stochastic weights and costs for overtime, and to the practice of operating room scheduling.

In operating room scheduling, it is not desirable to adapt schedules instantly. Both for technical reasons and for reasons of patient-centeredness, it is considered infeasible to decide on the next patient to schedule immediately after a surgery on a current patient has finished. Typically, the sequence of next patients to be scheduled is committed one or more patients ahead. To model this practical requirement, we introducing the concept of committing, i.e. to consider the schedule for subsequent patients as partly fixed. More specifically, k-level committing is defined as prohibiting adaptations to the sequence of the next k patients in the schedule, after which an l-level policy allows adaptations for the next l decisions.

The core of any l-level adaptive policy for the knapsack problem is formed by an exact algorithm for the corresponding *static* knapsack problem. In our case, the static problem refers to the version of the problem where all patients are selected and sequenced at the beginning of the day, and adaptations are not allowed. Merzifonluoğlu, Geunes, and Romeijn (2012) study static policies for a knapsack problem with stochastic weights and rewards, and a penalty for exceeding the knapsack capacity. Their analysis assumes normally distributed weights. Computational results on stochastic knapsack problems with normally distributed item sizes or rewards are also presented in the seminal work on sample average approximation by Kleywegt, Shapiro, and Homem-de Mello (2002). To support their l-level heuristics, Ilhan, Iravani, and Daskin (2011) develop an algorithm for the static problem with normally distributed rewards and deterministic weights.

In the setting of OR scheduling, evidence shows that the normal distribution fits empirical surgery durations relatively poorly because of its positive skew (Stepaniak, Heij, et al. 2009). In Section 4, we develop a new pseudo-polynomial algorithm to solve the static stochastic version of our problem to optimality and which forms the core of the subsequently analysed adaptive policies. It solves the static stochastic knapsack problem for a general class of problems where the distributions of the weights fall into the class of additive distributions. The normal distribution falls into this class.

The problem objective under consideration (classification field performance measures) is rooted in the work of Dexter, Macario, and Traub (1999), Dexter et al. (1999) and Stepaniak et al. (2012) who aim to balance the advantages and disadvantages of performing surgeries in overtime. More specifically, we consider the problem of maximising the rewards which result from providing surgical services to patients selected from a given set, minus the corresponding cost of working overtime. Rewards and costs are not only to be understood in terms of hospital revenue and operating costs, but may also include health benefits, patient satisfaction and employee satisfaction. As outlined by Stepaniak et al. (2012), overtime work may result in higher labour costs and cause job dissatisfaction and employee turnover among OR staff. On the other hand, delay of surgery to the next day or beyond may result in patient dissatisfaction, patient anxiety, longer recovery (perhaps with worse outcomes) and ultimately, loss of health. Obviously, policies which allow more generous adaptation (l-level adaptivity with large l in combination with low k-level committing) provide more opportunity for cost-efficient solutions and to reduce cancellations. Such policies, however, may increase dissatisfaction and anxiety from patients who have less certainty about their surgery (time). The presented analysis and results do not mean to disregard such dissatisfaction and anxiety, but provide insight in the gains which can be made through adaptive scheduling with committing, where the committing level can be set to control dissatisfaction and anxiety. We explicitly reflect on the practical considerations and implications in the discussion and conclusions.

3. Models for studying adaptive OR scheduling

We now present our models for single day single room adaptive scheduling. The models assume a fixed set of patients from which to select and schedule patients for surgery. Typically this is the set of patients initially selected by the specialty or surgeon, perhaps added with reserve patients and acute recent arrivals. The task is to select patients for surgery from this set, and provide a treatment sequence for the selected patients. The objective function will count a deterministic reward for

each patient treated at the end of the day. The patient rewards can depend on condition, but also on urgency, or for instance be higher for patients which got cancelled the previous day(s). Further, the objective function charges a penalty per time unit overtime. Overtime occurs when the sum of surgery durations of selected patients exceeds the (given) regular working time. Surgery durations are random variables, whose value is unknown until surgery is completed. However, the stochastic distribution functions describing the surgery durations are known beforehand. Hence expected overtime costs for a given selection of patients can be calculated on beforehand.

A **static policy** selects and sequences all patients at the beginning of the day, and does not make adaptations to the selection or sequence as the surgeries are executed and actual surgery durations become known. An **adaptive policy**, by contrast, may adapt the selection and/or sequence as surgeries are completed and surgery durations of treated patients become known. Hence, unlike static policies, adaptive policies can take the actual duration of completed surgeries into account. Our main goal is to study the objective function values attained by various adaptive policies, especially in relation to objective function values attained by static policies. The objective function measures OR performance.

Because it may not always be feasible or desirable in practice to select the next patient at the time of completion of the current surgery, we introduce the concept of committing. An **adaptive policy with k committings** commits k patients for, meaning that no adaptations are allowed to first k patients in the sequence. Furthermore, whenever a surgery is finished, and $k - 1$ committed patients remain, a next patient is added to the committed sequence, so that the committed sequence again consists of k patients. The simplest case of committing is an **adaptive policy with** 1 **committing**. It always selects the next patient at the *start* of the surgery of the current patient, and this patient must thus be treated next, regardless of the surgery duration of the current patient.

In the remainder of this section, we provide a detailed description of the above. In Section 3.1, we detail the patient set and reward structure, and in Section 3.2 we detail the various patient selection policies (static, adaptive, adaptive with committings).

3.1 *The patient set and objective function*

The patient set is partitioned into m subsets, called classes. A class corresponds to a single surgical procedure, e.g. hip replacement, or a set of homogenous surgical procedure types, e.g. acute trauma of the lower extremities. We denote by $\hat{n}_i$ the number of patients at the beginning of the day for class i ($i = 1, \ldots, m$) and denote its corresponding vector $\hat{\mathbf{n}} = (\hat{n}_1, \ldots, \hat{n}_m)$. We denote by ξ_{ij} the random surgery duration of patient $j = 1, \ldots, \hat{n}_i$ in class i. As said, we assume the distribution of ξ_{ij} to be known, but the actual duration to be revealed upon completion of surgery on the patient. Each surgical duration is independent of all other surgical durations. The surgical durations of patients within the same class follow the same distribution. Thus within class i, the random variables ξ_{ij} for $j = 1, \ldots, \hat{n}_i$ are independent and identically distributed.

The more patients selected and receive surgery, the higher the hospital revenue and (or) the more patients enjoy the resulting health benefits. This is modelled by a reward r_i contributed for each patient within class i for which surgery is completed by the end of the day. At the same time however, selecting many patients increases the likelihood of *overtime*, which is associated with increased risk on adverse events, patient dissatisfaction, employee dissatisfaction and overtime costs. Overtime occurs when the sum of the durations of completed surgeries exceeds the regular opening time of the OR $\hat{t}$. In the practical context of Erasmus Medical Center for instance, $\hat{t}$ equals 450 min. Overtime cost is calculated using a standard price p per time unit overtime work (as in Denton, Viapiano, and Vogl 2007; Batun et al. 2011).

The yield (or net reward) is defined as the sum of the rewards of the patients for which surgery has been completed at the end of the day (including overtime) minus the overtime costs. Note that overtime is a random variable because surgery times are random variables. Moreover, for adaptive policies the selection of patients for which surgery is completed at the end of the day is dependent on the actual surgery times. Thus yield is also a random variable. As we are interested to maximise yield, and yield is a random variable, we choose the *expected* yield as objective function of our single day single operating room adaptive scheduling problem.

Note that the *worst case* expected overtime costs for treating a patient of type i is $pE(\xi_{ij})$. This worst case occurs when the entire surgery takes place during overtime. Thus, $r_i > pE(\xi_{ij})$ would imply that all class i patients are always selected in any optimal solution. We assume $r_i < pE(\xi_{ij})$ to avoid this uninteresting and unrealistic case.

3.2 *Patient selection*

In this section we discuss the static and adaptive policies in more detail.

3.2.1 *The static problem*

Static policies may take the stochasticity of durations into account, but they disregard making adaptations to the patient selection or sequence as execution of the schedule progresses. Now, let $V_s := V_s(\hat{t}, \hat{\mathbf{n}})$ denote the optimal solution value of the static problem and therefore be the maximum expected yield attainable by a static policy. More generally, for future use, we define $V_s(t, \mathbf{n})$ as the optimal solution value for the static problem in which there are t time units and patient vector $\mathbf{n}$ remaining, $t = 1, \ldots, \hat{t}$, $\mathbf{n} = (n_1, \ldots, n_m)$, $n_i = 0, 1, \ldots, \hat{n}_i$. A formal model for the static problem is now defined by $V_s(t, \mathbf{n}) :=$

$$\max \sum_{i=1}^{m} \sum_{j=1}^{n_i} r_i \cdot x_{ij} - p \cdot y, \tag{1}$$

$$\text{s.t.} \quad y = E_{\boldsymbol{\xi}}\left[(\sum_{i=1}^{m} \sum_{j=1}^{n_i} \xi_{ij} \cdot x_{ij} - t)^+ \right], \tag{2}$$

$$x_{i(j+1)} \geq x_{ij}, \forall i = 1, 2, \ldots, m, \ j = 1, 2, \ldots, n_i - 1, \tag{3}$$

$$x_{ij} \in \{0, 1\}, \forall i = 1, 2, \ldots, m, \ j = 1, 2, \ldots, n_i. \tag{4}$$

Here, $x_{ij} = 1$ if and only if the j-th patient of class i is selected for surgery. Since all patients within a class are identical, we may impose without loss of generality (WLOG) that per class, patients are selected in decreasing order of index. This is enforced by (3).

To address the complexity of the static problem, let us consider the special case where all surgical durations are deterministic (have variance zero), with $\xi_{ij} = d_i$ for $i = 1, \ldots, m$, $j = 1, 2, \ldots, \hat{n}_i$ and $p \to +\infty$. Then any optimal solution for this problem will have $y = 0$. When $\hat{n}_i = 1$, for $i = 1, \ldots, m$, the (decision version of the) problem is then equivalent to the well-known knapsack problem, and therefore (weakly) NP-Complete (see Garey and Johnson (1979) and the references therein). Hence, we may conclude that the static problem is NP-hard.

We next introduce notation for the static policy with side-constraints: some patients *must* be scheduled. This is *not* a situation that will be studied in this paper: We introduce the notation because it will be convenient in subsequent sections. Let $\mathbf{n}' = (n'_1, \ldots, n'_m)$. Then define

$$V_s(t, \mathbf{n}, \mathbf{n}') \tag{5}$$

as the maximum expected yield obtainable under a static schedule with remaining time t and available patient vector $\mathbf{n}$, with the constraint that for each class $i = 1, \ldots, m$ at least n'_i patients must be scheduled.

3.2.2 *The adaptive problem*

We now formalise the model and notation for the adaptive problem. The adaptive problem is most naturally formulated via a dynamic programming formulation. Indeed, observe that after each surgery completion, the state $(t, \mathbf{n})$ of the system consists in the number of time units t remaining ($t = 1, \ldots, \hat{t}$) and the number of untreated patients $\mathbf{n} = (n_1, \ldots, n_m)$ of each type remaining ($n_i = 0, 1, \ldots, \hat{n}_i$). Let $V(t, \mathbf{n})$ denote the maximum expected yield attainable via an adaptive policy in this state.

In each state, the choice is whether to treat another patient, and if so, from which class. Let $[\mathbf{n}] := \{i \mid n_i > 0\}$ denote the classes of patients available for surgery for patient vector $\mathbf{n}$. In each class, we again assume WLOG that patients are selected in decreasing order of index, implying that patient n_i is the candidate patient in class $i \in [\mathbf{n}]$. Thus, when selecting a patient in class $i \in [\mathbf{n}]$, we move from state $(t, \mathbf{n})$ to state $(t - \xi_{in_i}, \mathbf{n} - \mathbf{e}_i)$ and obtain a reward r_i. Here $\mathbf{e}_i$ denotes the i-th standard unit vector. This gives rise to the following recursion:

$$V(t, \mathbf{n}) = \max\left\{ \max_{i \in [\mathbf{n}]} r_i + E_{\xi_{in_i}}[V(t - \xi_{in_i}, \mathbf{n} - \mathbf{e}_i)], -p \cdot (-t)^+ \right\}. \tag{6}$$

That is, in each step, we either select the patient that maximises our expected reward, or we select no more patients. If $[\mathbf{n}] = \emptyset$, then set $\max_{i \in [\mathbf{n}]} f(i) := -\infty$ for any $f(i)$. Note that in (6), instead of incurring overtime costs after completing each patient, overtime costs $-p \cdot (-t)^+$ are incurred for the entire overtime period when we decide to admit no more patients.

For any state $(t, \mathbf{n})$, the optimal solution for the adaptive problem $V(t, \mathbf{n})$ forms an upper bound on the maximum expected yield attainable by any adaptive policy for the corresponding problem. We refer to adaptive policies which are guaranteed to attain this upper bound as *exact adaptive policies*. Such exact adaptive policies have associated yield $V(\hat{t}, \hat{\mathbf{n}})$. Adaptive policies which do not necessarily deliver solutions having expected value $V(t, \mathbf{n})$ are henceforth referred to as *heuristic adaptive policies*.

3.2.3 *The adaptive problem with committing*

The adaptive problem defined in Section 3.2.2 allows to select the next patient to undergo surgery after surgery on the previous patient completes. As argued in Section 1, this is not necessarily realistic. Instead, one may wish to commit one or several patients for surgery in advance. The recursion below formally models the arising adaptive problem where k patients are committed in advance.

To introduce a formulation with committing, we let $\mathbf{c} = (c_1, c_2, \ldots, c_k)$ denote the queue of committed patients. A patient of class c_1 is thus committed to undergo surgery next, and then a patient of class c_2, etc. Let $n_i(\mathbf{c})$ denote the number of patients of class i that are in the queue $\mathbf{c}$, we thus have $\mathbf{n}(\mathbf{c}) = (n_1(\mathbf{c}), n_2(\mathbf{c}), \ldots, n_m(\mathbf{c})) = \sum_{i=1}^{k} \mathbf{e}_{c_i}$. $\mathbf{n} = (n_1, \ldots, n_m)$ denotes the vector of untreated patients: There are $n_i - n_i(\mathbf{c})$ untreated patients in class i that are *not* in queue to be treated. When selecting the next patient i to add to the queue, we must thus select $i \in [\mathbf{n} - \mathbf{n}(\mathbf{c})]$.

We let $V(t, \mathbf{n}, \mathbf{c})$ denote the maximum expected yield for the adaptive problem with k committings, where patient vector $\mathbf{n}$ and t time units are remaining, and the queue of committed patients is $\mathbf{c} = (c, c_2, \ldots, c_k)$. (We will often omit the subscript $_1$ of the first coordinate c of $\mathbf{c}$ to improve legibility in formulas.) The following recursion then characterises the problem:

$$V(t, \mathbf{n}, (c, c_2, \ldots, c_k))$$
$$= \max \left\{ r_c + \max_{i \in [\mathbf{n} - \mathbf{n}(\mathbf{c})]} E_{\xi_{cn_c}}[V(t - \xi_{cn_c}, \mathbf{n} - \mathbf{e}_c, (c_2, c_3, \ldots, c_k, i))], V_s(t, \mathbf{n}(\mathbf{c}), \mathbf{n}(\mathbf{c})) \right\} \quad (7)$$

So, given the state $(t, \mathbf{n}, (c, c_2, \ldots, c_k))$, we either commit a patient i and treat patient c, transiting to state $(t - \xi_{cn_c}, \mathbf{n} - \mathbf{e}_c, (c_2, c_3, \ldots, c_k, i))$ and collecting a reward r_c, or we may decide to commit no further patients, and to finish the surgery on the committed patients $\mathbf{n}(\mathbf{c})$ to obtain a reward $V_s(t, \mathbf{n}(\mathbf{c}), \mathbf{n}(\mathbf{c}))$ (cf. (5) in Section 3.2.1).

Let $V(t, \mathbf{n}, \mathbf{0})$ denote the expected yield of the optimal adaptive policy with k committings. At the beginning of the day, we need to either find the optimal queue $\mathbf{c}$ such that $\mathbf{n}(\mathbf{c}) \leq \mathbf{n}$, or we need to use a static policy with fewer than k patients. We obtain:

$$V(t, \mathbf{n}, \mathbf{0}) = \max \left\{ \max_{\mathbf{c}|\mathbf{n}(\mathbf{c}) \leq \mathbf{n}} V(t, \mathbf{n}, \mathbf{c}), V_s(t, \mathbf{n}) \right\}$$

Note that there is no need to impose that *fewer* than k patients are selected in the static alternative $V_s(t, \mathbf{n})$: If it is optimal to select more than k patients in the static policy, then $\max_{\mathbf{c}|\mathbf{n}(\mathbf{c}) \leq \mathbf{n}} V(t, \mathbf{n}, \mathbf{c}) \geq V_s(t, \mathbf{n})$.

4. Analysis of adaptive policies

In this section, we use the concept of *l-level adaptivity* to develop effective and computationally efficient heuristic adaptive policies (with and without committing). These adaptive policies will rely on efficient solution approaches for the static problem to overcome the computational difficulties of solving the adaptive problems as such. The concept of *l*-level adaptivity was introduced by Ilhan, Iravani, and Daskin (2011) for a knapsack problem with deterministic weights and stochastic rewards. We adapt the concept to our setting with stochastic weights and deterministic rewards and extend it to the case with committing. For this new setting, we derive results along the lines of results in Ilhan, Iravani, and Daskin (2011) and subsequently extend them to the case with committing. Moreover, we derive new monotonicity results on the impact of committing.

The *l*-level adaptive policies are introduced in Section 4.1. Section 4.2 presents monotonicity with respect to the number of committed patients.

4.1 *A tractable adaptive policy*

The *l- level adaptive policy* is defined recursively. The 0-level adaptive problem is the static problem. The 1-level adaptive problem is now informally defined as the problem of maximising expected yield by selecting a first patient, and solving the problem which remains upon completion of this first patient as a 0-level adaptive problem, i.e. as a static problem. Notice that the 1-level problem is adaptive as, by definition, it takes the realisation of the surgery duration time of the first patient into account in the selection and sequence of subsequent patients. The *l*-level adaptive policy refers to the problem of selecting a first patient, and solving the problem which remains upon completion of the first patient as an $(l - 1)$-level problem. More formally, denoting by $V_a^l(t, \mathbf{n})$ the expected yield of the *l*-level adaptive problem in state $(t, \mathbf{n})$, we obtain:

$$V_a^l(t, \mathbf{n}) = \begin{cases} \max \left\{ \max_{i \in [\mathbf{n}]} r_i + E_{\xi_{in_i}}[V_a^{l-1}(t - \xi_{in_i}, \mathbf{n} - \mathbf{e}_i)], -p \cdot (-t)^+ \right\} & \text{if } l \geq 1 \\ V_s(t, \mathbf{n}) & \text{if } l = 0 \end{cases} \quad (8)$$

This can be generalised as follows to the l-level adaptive problem with k committings we obtain:

$$V_a^l(t, \mathbf{n}, \mathbf{c}) = \begin{cases} \max \left\{ r_c + \max_{i \in [\mathbf{n} - \mathbf{n}(\mathbf{c})]} E_{\xi_{cn_c}} [V_a^{l-1}(t - \xi_{cn_c}, \mathbf{n} - \mathbf{e}_c, \tilde{\mathbf{c}})], V_s(t, \mathbf{n}(\mathbf{c}), \mathbf{n}(\mathbf{c})) \right\} & \text{if } l \geq 1 \\ V_s(t, \mathbf{n}, \mathbf{n}(\mathbf{c})) & \text{if } l = 0 \end{cases} \tag{9}$$

where $\mathbf{c} = (c, c_2, \ldots, c_k)$ and $\tilde{\mathbf{c}} = (c_2, c_3, \ldots, c_k, i)$. At the beginning of the day, the reward attained by this latter policy is thus:

$$V_a^l(t, \mathbf{n}, \mathbf{0}) := \max \left\{ \max_{\mathbf{c} | \mathbf{n}(\mathbf{c}) \leq \mathbf{n}} V_a^l(t, \mathbf{n}, \mathbf{c}), V_s(t, \mathbf{n}) \right\} \tag{10}$$

As l-level problems only consider the realisations of surgery duration for the first l patients, they can be viewed as a hybrid between the adaptive problem and the static problem. The following result is along the lines of a similar result in Ilhan, Iravani, and Daskin (2011) and generalises that result to the case with committings:

THEOREM 1 *For any remaining time t, patient vector $\mathbf{n}$, committed patient vector $\mathbf{c}$, and for $l' \geq \sum_{i=1}^{m} n_i - k$, it holds that*

$$V_s(t, \mathbf{n}, \mathbf{n}(\mathbf{c})) \leq V_a^1(t, \mathbf{n}, \mathbf{c}) \leq V_a^2(t, \mathbf{n}, \mathbf{c}) \leq \cdots \leq V_a^{l'}(t, \mathbf{n}, \mathbf{c}) = V_a^{l'+1}(t, \mathbf{n}, \mathbf{c}) = \cdots = V(t, \mathbf{n}, \mathbf{c})$$

and for policies without committing and $l' \geq \sum_{i=1}^{m} n_i$, it holds that

$$V_s(t, \mathbf{n}) \leq V_a^1(t, \mathbf{n}) \leq V_a^2(t, \mathbf{n}) \leq \cdots \leq V_a^{l'}(t, \mathbf{n}) = V_a^{l'+1}(t, \mathbf{n}) = \cdots = V(t, \mathbf{n})$$

For all proofs, see Appendix 1. Theorem 1 shows that l-level adaptive policies outperform $(l-1)$-level adaptive policies. Hence, it is preferable in practice to use the largest l which is computationally feasible.

Now, repeated application of l-level policies entails to solve the problem which arises after each surgery as an l-level problem, rather than using an l-level problem for selecting the first patient, and an $(l-1)$-level problem for the second patient, etc. In this approach, in each state $(t, \mathbf{n})$ a patient is selected from the class i^* that attains the maximum in (8) (or (9) for committing). Repeatedly selecting patients in this fashion is a policy that results in an expected yield for each state $(t, \mathbf{n})$, and we denote this expected yield by $V_{act}^l(t, \mathbf{n})$ (or $V_{act}^l(t, \mathbf{n}, \mathbf{c})$ for committing). The following theorem supports the use of this policy. It generalises a result in Ilhan, Iravani, and Daskin (2011) to the case with committing:

THEOREM 2 *For any remaining time t, patient vector $\mathbf{n}$, committing vector $\mathbf{c}$, and adaptivity level $l \geq 1$, it holds that $V_{act}^l(t, \mathbf{n}, \mathbf{c}) \geq V_a^l(t, \mathbf{n}, \mathbf{c})$ and $V_{act}^l(t, \mathbf{n}) \geq V_a^l(t, \mathbf{n})$.*

4.2 *The impact of committings*

This section regards the maximum expected yield attainable in adaptive problems with committing as a function of the number of committed patients. The following theorem shows that both for the adaptive problem, and for the l-level adaptive problem, the maximum attainable expected yield decreases with the number of committed patients.

THEOREM 3 (COMMITTING DETERIORATES PERFORMANCE) *Let $k, k' \in \mathbb{N}$ with $k < k'$ and consider committing vectors $\mathbf{c}_1 = (c, c_2, \ldots, c_k)$ and $\mathbf{c}_2 = (c, c_2, \ldots, c_k, c_{k+1}, \ldots, c_{k'})$. Then, for any time t, patient vector $\mathbf{n}$ such that $\mathbf{n} \geq \mathbf{n}(\mathbf{c}_2)$, and adaptivity level $l \geq 0$ it holds that $V_a^l(t, \mathbf{n}, \mathbf{c}_1) \geq V_a^l(t, \mathbf{n}, \mathbf{c}_2)$. Moreover, for the optimal policy it holds that: $V(t, \mathbf{n}, \mathbf{c}_1) \geq V(t, \mathbf{n}, \mathbf{c}_2)$.*

From this result, it follows easily that the maximal attainable reward at the beginning of the day decreases in the committing level. Intuitively, these results follow from the fact that the more patients are committed, the less room remains for optimisation in the adaptive selection of subsequent patients.

In summary, this section introduces adaptive policies and showed that in expectation they deliver better solutions than static policies. However, the l-level adaptive policies we developed frequently call an exact static policy. This is only computationally feasible for practical purposes, when the exact static policy is very efficient. The next section develops such an exact static policy for a broad class of surgery duration distribution functions.

5. A pseudo-polynomial algorithm for the static stochastic knapsack problem

As discussed in the introduction, the literature presents various solution approaches to stochastic knapsack problems. Merzifonluoğlu, Geunes, and Romeijn (2012) provide a branch and bound approach for the case with normally distributed weights and deterministic rewards. In the context of OR scheduling, however, Stepaniak, Heij, et al. (2009) have found that the symmetric normal distributions do not fit surgical durations well, and demonstrate that other distributions such as

the lognormal distribution provide a significantly better fit. Below, we present a new, exact, pseudo-polynomial algorithm to solve the (static) stochastic knapsack problem, which applies to a general class of distributions and includes the normal distribution as a special case. The presented approach extends the well-known pseudo-polynomial dynamic programming solution method for the deterministic knapsack problem. As this dynamic programming algorithm is not a main contribution of this paper, we will not present explicit computational analysis in Section 6. However, the computational analysis evaluates the 1-level adaptive policy, which makes frequent calls to the dynamic programming algorithm. These results confirm the computational suitability of our method.

Definition 1   Let random variables $\xi_1, \ldots, \xi_n$ be independently distributed, with ξ_i having probability density function $f(\theta_i)$, then $f(\theta)$ is **closed under addition** iff $\sum_{i=1}^{j} \xi_i \sim f(\sum_{i=1}^{j} \theta_i), \forall j = 1, 2, \ldots, n$.

This assumption captures a broad class of random variables. Unfortunately, the lognormal distribution advocated by Stepaniak, Heij, et al. (2009) for modelling surgery durations is not captured by this assumption. But we show in Section 6.1 and in the appendices that a class of random distributions satisfying 1 can be constructed that fits the empirical data rather well, confirming the flexibility and suitability of the assumption.

Parameters θ_i can be single valued, e.g. specifying the parameter λ in a Poisson distribution, or vector valued, e.g. specifying the mean μ and variance σ^2 for normal distribution, which are both closed under addition. We now present a pseudo-polynomial solution method for the static problem $V_s(t, \mathbf{n})$ for the case where the probability density functions are closed under addition and all parameter values are integral. The integral requirement can be relaxed by scaling techniques as used in Section 6.

Let ξ_{ij} represent the service time of the jth patient in class i and $\xi_{ij} \sim f(\theta_i)$. Let $f(\theta)$ be defined as in Definition 1. For any given parameter θ, let $R(\theta)$ be the maximum reward over all subsets $\Omega_{\mathbf{n}}$ of the set of patients $\mathbf{n}$ such that $\sum_{(i,j) \in \Omega_{\mathbf{n}}} \theta_i = \theta$. Let $OPT(\theta)$ denote the similarly defined objective function value. Then formally $R(\theta)$ and $OPT(\theta)$ can be defined as follows:

$$R(\theta) = \max \sum_{i=1}^{m} \sum_{j=1}^{n_i} r_i x_{ij} \tag{11}$$

$$\text{s.t.} \sum_{i=1}^{m} \sum_{j=1}^{n_i} \theta_i x_{ij} = \theta, x_{ij} \in \{0, 1\}, \forall i = 1, \ldots, m, j = 1, \ldots, n_i$$

$$OPT(\theta) = \max \sum_{i=1}^{m} \sum_{j=1}^{n_i} r_i x_{ij} - p \cdot E_\xi[(\xi - t)^+] \tag{12}$$

$$\text{s.t.} \sum_{i=1}^{m} \sum_{j=1}^{n_i} \theta_i x_{ij} = \theta, x_{ij} \in \{0, 1\}, \forall i = 1, \ldots, m, j = 1, \ldots, n_i$$

where $\xi = \sum_{i=1}^{m} \sum_{j=1}^{n_i} \xi_{ij} x_{ij} \sim f(\theta)$, if there is no feasible solution to formula (11), define $R(\theta) = -\infty$.

Notice that $R(\theta)$ is a multidimensional Knapsack problem because of being closed under addition. We now present a dynamic programming formulation for $R(\theta)$ and $OPT(\theta)$.

THEOREM 4   *For any parameter θ, and state $(t, \mathbf{n})$, $OPT(\theta) = R(\theta) - p \cdot E_\zeta[(\zeta - t)^+]$, with $\zeta \sim f(\theta)$.*

Theorem 4 implies that for any θ, $OPT(\theta)$ can be found by solving the multidimensional Knapsack problem $R(\theta)$. We now briefly address how to solve $R(\theta)$ for the problem under consideration where there are m patient classes with n_i patients $i = 1 \ldots, m$.

Definition 2   For any θ, let $R(\mathbf{n}, \theta)$ be the maximum reward over all subsets of patients $\mathbf{n}$ such that $\sum_{i=0}^{m} \sum_{j=1}^{n_i} \theta_i x_{ij} = \theta$, $x_{ij} \in \{0, 1\} \forall i = 1, \ldots, m, j = 1, \ldots, n_i$. Moreover, we let $R(\mathbf{n}, \theta) = -\infty$ if $\sum_{i=0}^{m} \sum_{j=1}^{n_i} \theta_i x_{ij} \neq \theta$ for all subsets of $\mathbf{n}$.

Introducing ℓ_n as the largest element in set $[\mathbf{n}]$, which indicates the class with largest index in $\mathbf{n}$, $R(\mathbf{n}, \theta)$ can now be recursively defined as follows:

THEOREM 5   *For any $\mathbf{n}$ and θ,*

$$R(\mathbf{n}, \theta) = \max_{x_{\ell_n n_{\ell_n}} \in \{0,1\}} R(\mathbf{n} - x_{\ell_n n_{\ell_n}} e_{\ell_n} - n_{\ell_n}(1 - x_{\ell_n n_{\ell_n}})e_{\ell_n}, \theta - \theta_{\ell_n} x_{\ell_n n_{\ell_n}}) + r_{\ell_n} x_{\ell_n n_{\ell_n}} \tag{13}$$
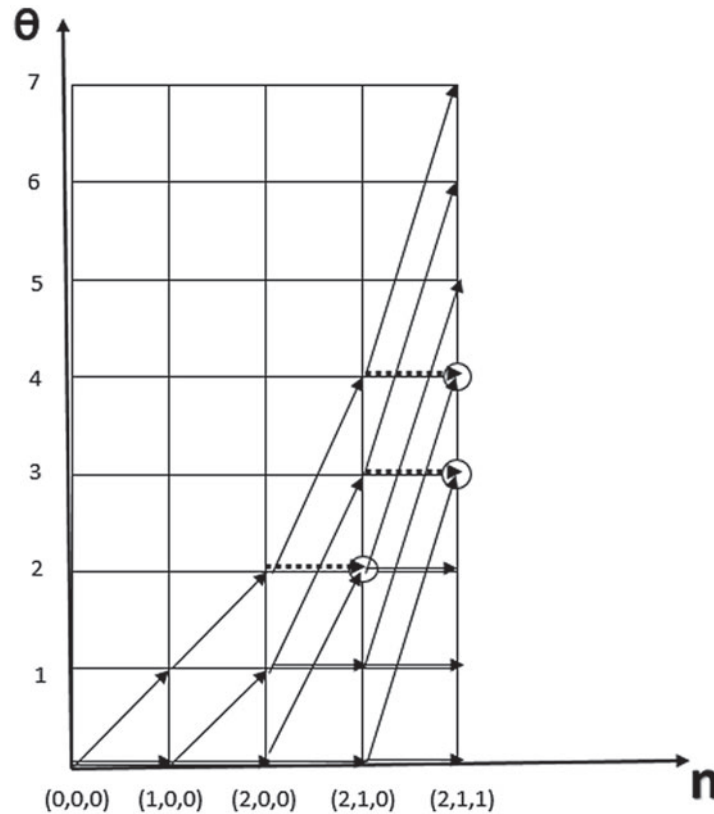
Figure 1. Dynamic procedures.

The recursion in the theorem can be interpreted to repeatedly consider selection of the patient with the largest index in class ℓ_n and proceed to state $R(\mathbf{n} - e_{\ell_n}, \theta - \theta_{\ell_n})$ if the patient is scheduled or to state $R(\mathbf{n} - n_{\ell_n} e_{\ell_n}, \theta)$ otherwise. The recursion ensures the sequence assumption in Section 3.

Figure 1 illustrates the dynamic programming recursion. It considers an example: there are three classes of patients with parameters $\mathbf{n} = (2, 1, 1)$, $\theta_1 = 1$, $\theta_2 = 2$, $\theta_3 = 3$ and the rewards are $r_1 = 1$, $r_2 = 3$, $r_3 = 5$, respectively. The node corresponding to state $((0, 0, 0), 0)$ forms the origin of the state space, and the transitions are depicted by arcs. If a state space node has two incoming arcs, the solid arc represents the arc maximising the reward.

From Figure 1, we can also understand how to find $OPT(\theta)$. For each state $(\hat{\mathbf{n}}, \theta)$, $R(\hat{\mathbf{n}}, \theta)$ and the corresponding optimal subset $\Omega_{\hat{\mathbf{n}}}^*$ of patients can be determined using dynamic programming. Moreover, for this subset of patients $\Omega_{\hat{\mathbf{n}}}^*$, $p \cdot E_\xi[(\xi - t)^+]$, where $\xi = \sum_{(i,j) \in \Omega_{\hat{\mathbf{n}}}^*} \xi_{ij} x_{ij} \sim f(\theta)$ can be calculated. The optimal solution can then be found by maximising over all θ, $R(\hat{\mathbf{n}}, \theta) - p \cdot E_\xi[(\xi - t)^+]$. This procedure can be extended to the case of committing, as the committed patients $\mathbf{c}$ can simply be accounted for by setting $\xi \sim f(\theta + \sum_{i=1}^k \theta_{c_i})$. Through this an optimal track is derived to arrive at state $(\hat{\mathbf{n}}, \theta)$ if it is feasible and to get maximum objective value $R(\hat{\mathbf{n}}, \theta)$

Further, we notice that the presented approach is polynomial in the number of patients, but not in the input size of the problem if it is efficiently encoded. This is due to the fact that the presented problem with multiple patients per patient class is in fact a high multiplicity scheduling problem, see for instance Hochbaum and Shamir (1990, 1991). Hence, policies which iteratively select patients by applying a pseudo-polynomial algorithm can be classified as a pseudo-polynomial list generating algorithm (Brauner et al. 2005).

Table 1. Average coefficient of variation (CV) and number of patient classes for each case.

Low-dimensional cases			High-dimensional cases		
Case name	Avg. CV	# of patient classes	Case name	Avg. CV	# of patient classes
vHOHWK-HCV	0.86	3	Ophthalmology	0.31	6
vHOHWK-LCV	0.76	3	Gynaecology	0.35	9
EMC-HCV	0.31	3	Orthopedics	0.29	10
EMC-LCV	0.19	3	Urology	0.24	12

6. Computational results

In this section, we investigate various policies. We introduce acronyms below:

StS   Exact static policy (Section 3.2.1);
ANC   Exact adaptive policy (Section 3.2.2);
AOC   Exact adaptive policy with one committing (Section 3.2.3);
1-ANC   Repeated exact 1-level adaptive policy (Section 4.1);
1-AOC   Repeated exact 1-level adaptive policy with one committing (Section 4.1).

We let $V(\cdot)$ denote the corresponding expected yield. E.g. $V(\text{StS})$ denotes the expected yield provided by an exact static policy. Because the 1-level adaptive policies are computed along a sample path, the expected yields $V(1\text{-ANC})$ and $V(1\text{-AOC})$ are estimated using Monte Carlo simulation with 2000 samples. We perform benchmarks using data from Van Houdenhoven et al. (2007) and from recently collected data, both originating from Erasmus Medical Center. We refer to the data Van Houdenhoven et al. (2007) as vHOHWK, and to the recently collected data as EMC. The benchmarks compare:

- The improvement of ANC and AOC over StS;
- The performance of 1-ANC and 1-AOC, when compared to ANC and AOC;
- The computational efficiency of 1-ANC and 1-AOC compared to ANC and AOC;
- Typical sequences in which patients from various classes are scheduled;
- The values of 1-ANC and 1-AOC in practice, using data from Erasmus Medical Center.

6.1 Description of test cases

The experiments fit surgery time using a class of finite parameter compound Poisson distributions (see Appendix 2 in the electronic companion appendix). The distributions are additive (as in Definition 1), asymmetric and non-negative, unlike the normal distribution. Moreover, they fit closer to the empirical data than chi-squared and fixed-scale Gamma distributions. We construct four low-dimensional cases consisting of three patient classes: Two based on the data from vHOWK and two based on the more recent data obtained directly from Erasmus Medical Centre. Additionally, we construct 4 high-dimensional cases based on the more recent data from Erasmus Medical Centre. Summarising statistics are given in Table 1. The table includes information on the average coefficient of variation (CV) of the surgery distributions in the class. The CV is defined as the standard deviation divided by the mean. For detailed information on the surgery time distributions for each patient we refer to Appendix B.2. A verification of the goodness of fit of the surgery time distributions with respect to the Erasmus Medical Center data is presented in Appendix B.3.

For all cases, daily regular capacity is set at 450 min (7.5 h). W.l.o.g., we may rescale the unit overtime cost p to 1. We let rewards r_i depend on mean μ_i and standard deviation σ_i of the surgery time distribution for class i by setting $r_i = \alpha\mu_i + \beta\sigma_i$. We investigate a total of 15 combinations of (α, β) to determine sensitivity: each $\alpha \in \{0.1, 0.3, 0.5, 0.75\}$ and $\beta \in \{0, 0.1, 0.2, 0.3\}$, but with $(\alpha, \beta) = (0.75, 0.3)$ omitted because for some of the patient classes in the test cases, this combination does not satisfy $r_i \leq p\mu_i$.

6.2 The performance of the 1-level adaptive policies

In the following, we investigate the performance of the 1-level adaptive policies 1-ANC and 1-AOC. As benchmarks, we use the optimal adaptive policies ANC and AOC, respectively. These latter policies are computationally tractable only for small numbers of patient classes, and we thus use the low-dimensional cases (see Table 1) to perform the benchmarks. The parameters (α, β) are varied over 15 combinations, see Section 6.1. We consider $p \in \{2, \ldots, 7\}$ patients per class, giving $3p$

Table 2. The average and standard deviation of the relative benefit of adaptivity (rBoA) computed for all 90 instances corresponding to each of the 4 low-dimensional cases of Table 1.

Case name	Avg. CV	rBoA of ANC	rBoA of AOC
vHOHWK-HCV	0.86	$28.5\% \pm 12.5\%$	$18.3\% \pm 5.1\%$
vHOHWK-LCV	0.76	$27.9\% \pm 17.1\%$	$17.7\% \pm 9.1\%$
EMC-HCV	0.31	$9.9\% \pm 5.7\%$	$8.1\% \pm 4.3\%$
EMC-LCV	0.19	$2.8\% \pm 0.7\%$	$0.1\% \pm 0.1\%$
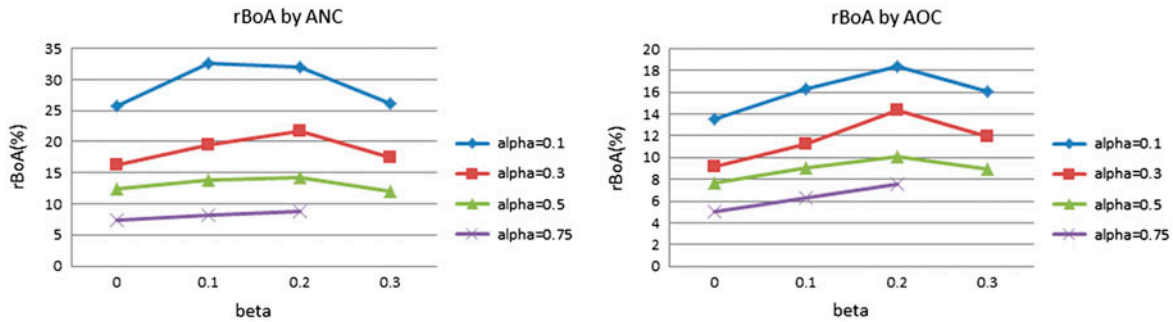
Figure 2. The dependence of the rBoA on α and β. Each point represents the average of 24 instances that use that specific value of (α, β).

patients in total. We perform a full-factorial experiment, resulting in a total of (4 cases) $\times$ (15 values for α, β) $\times$ (6 value for p) = 360 test instances.

The quantities $V(\text{ANC}) - V(\text{StS})$ and $V(\text{AOC}) - V(\text{StS})$ have been termed *benefit of adaptivity* (Dean, Goemans, and Vondrák 2008). Define the relative benefit of adaptivity (rBoA) as $[V(\text{ANC}) - V(\text{StS})]/V(\text{StS})$ (no committing), and as $[V(\text{AOC}) - V(\text{StS})]/V(\text{StS})$ (one committing). Because $V(\text{ANC}) \geq V(\text{1-ANC}) \geq V(\text{StS})$ and $V(\text{AOC}) \geq V(\text{1-AOC}) \geq V(\text{StS})$ by Theorem 1, the 1-level adaptive policies capture a fraction of the maximal benefit of adaptivity. We now first investigate the total benefit of adaptivity for our test instances, and then assess which fraction of this benefit is captured by the heuristic 1-level adaptive policies. Table 2 gives the benefit of adaptivity for each of the low-dimensional cases, averaged over the corresponding 90 instances. The benefits are especially substantial for the $vHOHWK$ instances and increasing with the average CV. The benefits are statistically significant at the 5% level for 3 out of the 8 cases. Obviously the variances are large because of the different values for the parameters α and β. The AOC for the EMC-LCV case forms an exception and gives very little benefit. This is also the only case where the 1-level committing appears to cost more than half of the benefit without committing. The dependence of the rBoA on the parameters α and β is plotted in Figure 2. The figure shows that the rBoA increases significantly with decreasing α, i.e. if the reward for selecting patients with higher variation in surgery duration is higher relative to overtime costs. The rBoA is less sensitive to β, and dependence on β is not monotonic. The figure also shows that in the problem with committing, the benefit from adaptivity is less for these small-scale instances.

We now investigate the fraction of the maximum benefit of adaptivity captured by the 1-level adaptive policies: $[V(\text{1-ANC}) - V(\text{StS})]/[V(\text{ANC}) - V(\text{StS})]$ (without committing) and $[V(\text{1-AOC}) - V(\text{StS})]/[V(\text{AOC}) - V(\text{StS})]$ (with committing). This ratio takes on values in [0%, 100%], and the closer it is to 100%, the greater proportion of the benefit is captured by the 1-level adaptive policies. Due to the existence of the sampling error of our Monte Carlo estimator for 1-ANC and 1-AOC however, the ratio can occasionally exceed 1. Table 3 shows the results for each of the four low-dimensional cases, averaged over the 90 instances for each case. The results of 1-AOC for EMC-LCV are omitted from the table, as the denominator $V(\text{AOC}) - V(\text{StS})$ is close to 0. The table shows that 1-ANC has excellent performance, mostly capturing (well) above 90% of the maximum benefit of adaptivity. 1-AOC has a decent performance, capturing $60 - 90\%$ of the benefit of adaptivity. Figure 3 shows that the performance of the heuristic is robust over a wide range of values for α and β.

6.3 Insights from the schedules

We inspect sample paths of the 1-ANC and 1-AOC schedules, to obtain insights into the order in which patients from various patient classes should be sequenced. Patient classes for the low-dimensional instances are numbered in order of increasing standard deviation, see Table B1 in Appendix B.2. Table 4 shows the frequencies of the most common sequences. The results

Table 3. The fraction of the benefit of adaptivity gained by the heuristics for the four low-dimensional cases. Reported numbers correspond to the average and standard deviation over the 90 instances for that case.

Case name	ANC	AOC
vHOHWK-HCV	95.6% ± 11.0%	79.8% ± 16.9%
vHOHWK-LCV	100.0% ± 6.3%	69.7% ± 24.7%
EMC-HCV	90.2% ± 7.4%	68.6% ± 13.0%
EMC-LCV	89.7% ± 15.4%	Omitted

Table 4. Most common scheduling sequence under 1-ANC and corresponding fraction of sample paths in which this sequence is used for 1-ANC and 1-AOC.

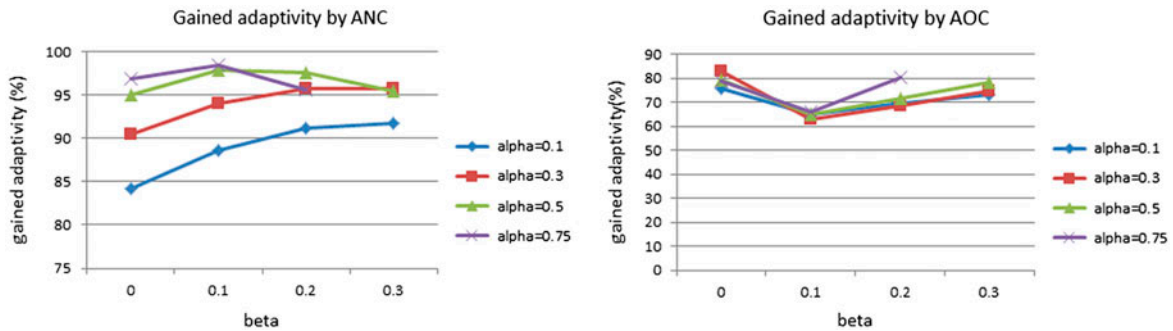| Case	Most common sequence	Fraction (1-ANC) (%)	Fraction (1-AOC) (%)
vHOHWK-HCV	(3,2,1)	98.73	66.23
vHOHWK-LCV	(3,2,1)	84.90	60.86
EMC-HCV	(3,2,1)	99.78	18.25
EMC-LCV	(3,2,1)	75.26	57.51

Figure 3. The fraction of the benefit of adaptivity captured by the heuristics for various values of α and β. Each point represents the average of 24 instances that use that specific value of α, β, and the average of 18 instances for the 1-AOC.
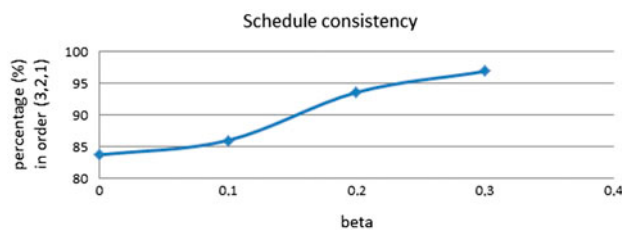

Figure 4. Dependence of schedule consistency, i.e. fraction of schedules in the order $(3, 2, 1)$, on β.

in the table represent the average over 90 instances per case (see Section 4), and 2000 sample paths per instance. The table shows that the 1-ANC policy typically sequences patients in order of decreasing standard deviation. To a lesser degree, this also holds for 1-AOC. The degree by which sequences adhere to this order increases with β, as shown in Figure 4.
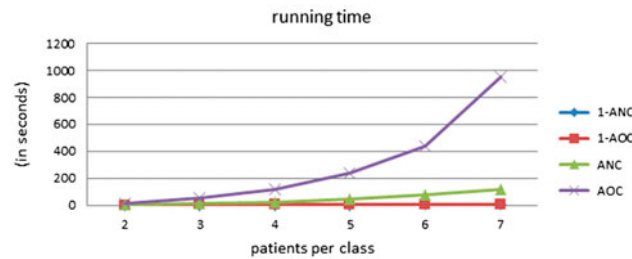
Figure 5. The running time of the various adaptive policies for cases with 3 patient classes. Each data point corresponds to the average of 60 instances that share the same number of patients per class.

Table 5. Running time of various adaptive policies for high-dimensional problems.

Case	# patients per class	1-ANC	1-AOC	ANC	AOC
Ophthalmology	(2,2,2,2,2,2)	0.2 min	0.2 min	5.4 min	57 min
	(3,3,3,3,3,3)	0.4 min	0.2 min	> 24 h	> 24 h
Gynaecology	(1,1,1,1,1,1,1,1,1)	0.1 min	0.2 min	6.6 min	161 min
	(2,2,2,2,2,2,2,2,2)	1.2 min	0.6 min	> 24 h	> 24 h

6.4 *Computational efficiency*

We investigate the computational efficiency of the 1-level adaptive policies and the optimal adaptive policies. An important factor governing running time is the number of patient classes. We first test the computational efficiency using the 360 low-dimensional problem instances described in Section 6.2. These instances have only three patient classes. Our experiments indicate that the number of patients per class is the main factor determining running time for these instances. We therefore plot the running time against this parameter in Figure 5. Computations were performed using a laptop with an i5-3210M CPU and 8GB RAM. The figure shows that while the 1-level adaptive policies computes solutions nearly instantly, the computation times for the optimal adaptive policy increase rapidly as the number of patients per class grows.

To investigate the scalability of the policies for cases with more patient classes, we perform additional tests for the cases Ophthalmology and Gynaecology (see Table 1). We determine the average computation time for 15 combinations of the parameters (α, β) (see Section 6.1). The results are summarised in Table 5. The table shows that the computation times for the 1-level adaptive policies are shorter than 1.2 min: short enough for application in practice. However, the computation times for the optimal adaptive policy are quite significant: for the cases with 18 patients over 6 or 9 classes, the computation time even exceeds 24 h. We conclude that the optimal schedules cannot be computed in a reasonable timescale for cases with moderate numbers of patients and patient classes.

6.5 *The benefits of adaptive scheduling at Erasmus Medical Center*

For the four realistic high-dimensional cases given in Table 1, the optimal adaptive policies ANC and AOC are computationally intractable, as shown in Table 5. However, our efficient 1-level adaptive policies 1-ANC and 1-AOC can be applied to these instances. We use one patient per class for Orthopaedics and Urology, and consider two options for the patients per class for Gynaecology and Ophthalmology. The results are given in Table 6, which reports the estimated relative improvement of 1-AOC and 1-ANC in comparison to StS. For reward parameters, we set $\alpha = 0.1, \beta = 0, \alpha = 0.75, \beta = 0$ and $\alpha = 0.75, \beta = 0.2$. The table shows a substantial improvement for the 1-level adaptive policy without committing as well as for the 1-level adaptive policy with committing, be it to a smaller extend. We can clearly identify that the relative improvement decreases in α, but no explicit tendency in β. It also tends to increase in the average CV. The two percentages in bold face are low because all the mean surgical times of Orthopaedics are relatively long, and overtime cost is relatively high compared with reward. As a consequence, only two to three surgeries are selected by our heuristic during the day, which limits its potential. Interestingly, one of these orthopaedics cases forms the only exception where the improvement over the optimal static scheduling is not statistically significant at the 5% level.

Table 6. The relative improvement of 1-ANC and 1-ANC compared to StS for 4 specialism and various numbers of patient per class, estimated using Monte Carlo.

Case	Avg. CV	# of patients	1-ANC		
			$\alpha = 0.1, \beta = 0$	$\alpha = 0.75, \beta = 0$	$\alpha = 0.75, \beta = 0.2$
Ophthalmology	0.31	(1,2,1,2,1,2)	15.9% ± 0.6%	3.7% ± 0.5%	2.8% ± 0.4%
		(2,2,2,2,2,2)	13.3% ± 0.6%	3.5% ± 0.3%	2.6% ± 0.4%
Gynaecology	0.35	(1,1, …,1)	17.5% ± 1.1%	5.2% ± 0.6%	5.5% ± 0.9%
		(2,2, …,2)	16.1% ± 1.1%	4.4% ± 0.5%	6.5% ± 1.0%
Orthopedics	0.29	(1,1, …,1)	**6.1% ± 0.6%**	3.3% ± 0.6 %	5.8% ± 0.6%
Urology	0.24	(1,1, …,1)	13.4% ± 0.5%	3.2% ± 0.5 %	3.0% ± 0.4%
			1-AOC		
Case	Avg. CV	# of patients	$\alpha = 0.1, \beta = 0$	$\alpha = 0.75, \beta = 0$	$\alpha = 0.75, \beta = 0.2$
Ophthalmology	0.31	(1,2,1,2,1,2)	7.5% ± 0.8%	2.4% ± 0.5%	2.3% ± 0.4%
		(2,2,2,2,2,2)	8.4% ± 0.6%	2.5% ± 0.3%	2.0% ± 0.3%
Gynaecology	0.35	(1,1, …,1)	6.6% ± 1.1%	3.3% ± 0.6%	2.8% ± 0.7%
		(2,2, …,2)	6.5% ± 1.1%	3.3% ± 0.5%	2.2% ± 0.7%
Orthopedics	0.29	(1,1, …,1)	**0.3% ± 0.4%**	2.4% ± 0.6%	3.2% ± 0.5%
Urology	0.24	(1,1, …,1)	9.6% ± 0.4%	2.0% ± 0.5%	1.8% ± 0.3%

7. Discussion and conclusion

Adaptive scheduling has received considerable attention from the operations research community since the seminal work of Dean, Goemans, and Vondrák (2008), but has so far not been applied to operating room production research, where adaptive scheduling is common practice May et al. (2011). Our work is the first to introduce formal adaptive models and solution methods to the field of operating room scheduling. For the purpose of making the models more realistic and practically applicable to operating room planning, we introduce the concept of committing, thus extending the existing adaptive models to adaptive models with committing. The presented analysis applies the l-level adaptive policies proposed by Ilhan, Iravani, and Daskin (2011) and extensions of these policies to adaptive policies with committing. Our analytical results demonstrate that committing deteriorates performance, while adaptivity enhances performance. Moreover, we show that l-level committing adaptive policies outperform the static policy in expectation. For practical purposes, the concept of adaptive scheduling with committing can be further refined to better match the reality of operating room scheduling. For instance, the level of committing can decrease as the day progresses, and extensions to incorporate arrival of acute patients are of interest.

For practical implementation, l-level adaptive policies require an efficient exact method for the static problem. We develop such an algorithm in the form of a new pseudo-polynomial dynamic programming algorithm that is valid for the broad class of additive distributions, which includes the normal distribution. This algorithm for the static problem therefore is considerably more general than existing approaches, which typically assume normal distributions. The computational analysis demonstrates the efficiency of the dynamic programming algorithm.

The performance analysis on practical problems shows that the 1-level adaptive policy tends to capture most of the gap between the static and the optimal adaptive policy. For instances based on real-life data, the 1-level adaptive policy gives a statistically significant performance benefit of 2 to 15%. The 1-level 1-committing adaptive policy, which may be preferred in practice because it does not make last-minute changes, typically gives significant performance benefits in the order of 2 to 8%.

Now let us turn to the practical implications and derive some practical guidelines from the analysis and results. We observe from our computational experiments that adaptivity brings substantial and significant performance gains over static policies, thus explaining why it is commonly practised in the first place. As shown by Stepaniak, Mannaerts, et al. (2009), Stepaniak et al. (2012) however, the adaptation decisions made by human planners may differ considerably and lead to considerable performance differences. Hence, the proposed models and methods may be of value by enabling consistent performance improvement. The 1-level adaptive algorithms in particular are computationally efficient and provide near optimal solutions.

To use these algorithms in practice, it is required to set an appropriate level of committing, so as to balance inconveniences from adaptivity for the patients with other performance attributes such as risk of cancellation, overtime costs, hospital revenue and employee dissatisfaction from overtime work. Our analysis does not provide optimal levels of committing, these need to be set by management. Analysing these trade-offs can also be viewed as an area for further research. Practical application

will also require that the rewards of treating patients are quantified relative to overtime costs. This also requires managerial judgement, as these rewards and costs also include constituents which are not trivially quantified (e.g. health outcomes, and employee dissatisfaction). Future research to learn these parameter values from empirical decisions is called for.

Finally, our results show that, as a rule, adaptive policies tend to sequence long surgeries and surgeries with relatively large standard deviations early. Thus, until decision support for adaptive scheduling is available, these rules of thumb might already guide practical adaptive operating room scheduling.

Disclosure statement

No potential conflict of interest was reported by the authors.

Funding

References

Balseiro, S. R., and D. B. Brown. 2016. *Approximations to Stochastic Dynamic Programs via Information Relaxation Duality*. Tech. rep., Working Paper. http://faculty.fuqua.duke.edu/~dbbrown/bio/papers/balseiro_brown_approximations_16.pdf .

Batun, S., B. T. Denton, T. R. Huschka, and A. J. Schaefer. 2011. "Operating Room Pooling and Parallel Surgery Processing under Uncertainty." *INFORMS Journal on Computing* 23 (2): 220–237.

Bhalgat, A., A. Goel, and S. Khanna. 2011. "Improved Approximation Results for Stochastic Knapsack Problems." In *Proceedings of the Twenty-second Annual ACM-SIAM Symposium on Discrete Algorithms*, 1647–1665. SIAM.

Blado, D., W. Hu, and A. Toriello. 2016. "Semi-infinite Relaxations for the Dynamic Knapsack Problem with Stochastic Item Sizes." *SIAM Journal on Optimization* 26 (3): 1625–1648.

Brauner, N., Y. Crama, A. Grigoriev, and J. van de Klundert. 2005. "A Framework for the Complexity of High-multiplicity Scheduling Problems." *Journal of Combinatorial Optimization* 9 (3): 313–323.

Cardoen, B., E. Demeulemeester, and J. Beliën. 2010. "Operating Room Planning and Scheduling: A Literature Review." *European Journal of Operational Research* 201 (3): 921–932.

Cayirli, T., and E. Veral. 2003. "Outpatient Scheduling in Health Care: A Review of Literature." *Production and Operations Management* 12 (4): 519–549.

Chen, K., and S. M. Ross. 2014. "An Adaptive Stochastic Knapsack Problem." *European Journal of Operational Research* 239 (3): 625–635.

Dean, B. C., M. X. Goemans, and J. Vondrák. 2008. "Approximating the Stochastic Knapsack Problem: The Benefit of Adaptivity." *Mathematics of Operations Research* 33 (4): 945–964.

Dellaert, N., E. Cayiroglu, and J. Jeunet. 2016. "Assessing and Controlling the Impact of Hospital Capacity Planning on the Waiting Time." *International Journal of Production Research* 54 (8): 2203–2214.

Denton, B., J. Viapiano, and A. Vogl. 2007. "Optimization of Surgery Sequencing and Scheduling Decisions under Uncertainty." *Health Care Management Science* 10 (1): 13–24.

Dexter, F., R. H. Epstein, R. D. Traub, and Y. Xiao. 2004. "Making Management Decisions on the Day of Surgery Based on Operating Room Efficiency and Patient Waiting Times." *Anesthesiology* 101 (6): 1444–1453.

Dexter, F., A. Macario, and R. D. Traub. 1999. "Which Algorithm for Scheduling Add-on Elective Cases Maximizes Operating Room Utilization? Use of Bin Packing Algorithms and Fuzzy Constraints in Operating Room Management." *Anesthesiology* 91 (5): 1491–1500.

Dexter, F., A. Macario, R. D. Traub, M. Hopwood, and D. A. Lubarsky. 1999. "An Operating Room Scheduling Strategy to Maximize the Use of Operating Room Block Time: Computer Simulation of Patient Scheduling and Survey of Patients' Preferences for Surgical Waiting Time." *Anesthesia & Analgesia* 89 (1): 7–20.

Dexter, F., A. Willemsen-Dunlap, and J. D. Lee. 2007. "Operating Room Managerial Decision-making on the Day of Surgery with and without Computer Recommendations and Status Displays." *Anesthesia & Analgesia* 105 (2): 419–429.

Erdogan, S. A., B. T. Denton, J. Cochran, L. Cox, P. Keskinocak, J. Kharoufeh, and J. Smith. 2011. *Surgery Planning and Scheduling*. Wiley Encyclopedia of Operations Research and Management Science. http://ca.wiley.com/WileyCDA/Section/id-380199.html.

Garey, M. R., and D. S. Johnson. 1979. *Computers and Intractability: A Guide to the Theory of NP-completeness*. San Francisco, CA: WH Freeman & Co.

Gupta, D. 2007. "Surgical Suites' Operations Management." *Production and Operations Management* 16 (6): 689–700.

Gupta, D., and B. Denton. 2008. "Appointment Scheduling in Health Care: Challenges and Opportunities." *IIE Transactions* 40 (9): 800–819.

Hochbaum, D. S., and R. Shamir. 1990. "Minimizing the Number of Tardy Job Units under Release Time Constraints." *Discrete Applied Mathematics* 28 (1): 45–57.

Hochbaum, D. S., and R. Shamir. 1991. "Strongly Polynomial Algorithms for the High Multiplicity Scheduling Problem." *Operations Research* 39 (4): 648–653.

Ilhan, T., S. M. Iravani, and M. S. Daskin. 2011. "Technical Note – The Adaptive Knapsack Problem with Stochastic Rewards." *Operations research* 59 (1): 242–248.

Jebali, A., and A. Diabat. 2015. "A Stochastic Model for Operating Room Planning under Capacity Constraints." *International Journal of Production Research* 53 (24): 7252–7270.

Kleywegt, A. J., A. Shapiro, and T. Homem-de Mello. 2002. "The Sample Average Approximation Method for Stochastic Discrete Optimization." *SIAM Journal on Optimization* 12 (2): 479–502.

Levin, A., and A. Vainer. 2013. "The Benefit of Adaptivity in the Stochastic Knapsack Problem with Dependence on the State of Nature." *Discrete Optimization* 10 (2): 147–154.

May, J. H., W. E. Spangler, D. P. Strum, and L. G. Vargas. 2011. "The Surgical Scheduling Problem: Current Research and Future Opportunities." *Production and Operations Management* 20 (3): 392–405.

Merzifonluoğlu, Y., J. Geunes, and H. E. Romeijn. 2012. "The Static Stochastic Knapsack Problem with Normally Distributed Item Sizes." *Mathematical Programming* 134 (2): 459–489.

van Oostrum, J. M., M. Van Houdenhoven, J. L. Hurink, E. W. Hans, G. Wullink, and G. Kazemier. 2008. "A Master Surgical Scheduling Approach for Cyclic Scheduling in Operating Room Departments." *OR Spectrum* 30 (2): 355–374.

Shylo, O. V., O. A. Prokopyev, and A. J. Schaefer. 2012. "Stochastic Operating Room Scheduling for High-volume Specialties under Block Booking." *INFORMS Journal on Computing* 25 (4): 682–692.

Stepaniak, P., C. Heij, G. Mannaerts, M. de Quelerij, and G. De Vries. 2009. "Modeling Procedure and Surgical Times for Current Procedural Terminology-anesthesia-surgeon Combinations and Evaluation in Terms of Case-duration Prediction and Operating Room Efficiency: A Multicenter Study." *Anesthesia & Analgesia* 109 (4): 1232–1245.

Stepaniak, P. S., G. H. H. Mannaerts, M. de Quelerij, and G. de Vries. 2009. "The Effect of the Operating Room Coordinator's Risk Appreciation on Operating Room Efficiency." *Anesthesia & Analgesia* 108 (4): 1249–1256.

Stepaniak, P., R. van der Velden, J. van de Klundert, and A. Wagelmans. 2012. "Human and Artificial Scheduling System for Operating Rooms." In *Handbook of Healthcare System Scheduling*, 155–175. Springer.

Tang, J., and Y. Wang. 2015. "An Adjustable Robust Optimisation Method for Elective and Emergency Surgery Capacity Allocation with Demand Uncertainty." *International Journal of Production Research* 53 (24): 7317–7328.

Van Essen, J. T., J. M. Bosch, E. W. Hans, M. van Houdenhoven, and J. L. Hurink. 2014. "Reducing the Number of Required Beds by Rearranging the OR-schedule." *OR spectrum* 36 (3): 585–605.

Van Houdenhoven, M., J. M. van Oostrum, E. W. Hans, G. Wullink, and G. Kazemier. 2007. "Improving Operating Room Efficiency by Applying Bin-packing and Portfolio Techniques to Surgical Case Scheduling." *Anesthesia & Analgesia* 105 (3): 707–714.

Vermeulen, I. B., S. M. Bohte, S. G. Elkhuizen, H. Lameris, P. J. Bakker, and H. La Poutré. 2009. "Adaptive Resource Allocation for Efficient Patient Scheduling." *Artificial Intelligence in Medicine* 46 (1): 67–80.

Vissers, J. M., J. Bertrand, and G. De Vries. 2001. "A Framework for Production Control in Health Care Organizations." *Production Planning & Control* 12 (6): 591–604.

Wang, Y., J. Tang, and R. Y. Fung. 2014. "A Column-generation-based Heuristic Algorithm for Solving Operating Theater Planning Problem under Stochastic Demand and Surgery Cancellation Risk." *International Journal of Production Economics* 158: 28–36.

Xiao, G., W. van Jaarsveld, M. Dong, and J. van de Klundert. 2016. "Stochastic Programming Analysis and Solutions to Schedule Overcrowded Operating Rooms in China." *Computers & Operations Research* 74: 78–91.

Zhang, Z., X. Xie, and N. Geng. 2014. "Dynamic Surgery Assignment of Multiple Operating Rooms with Planned Surgeon Arrival Times." *IEEE Transactions on Automation Science and Engineering* 11 (3): 680–691.

Zhou, J., F. Dexter, A. Macario, and D. A. Lubarsky. 1999. "Relying Solely on Historical Surgical Times to Estimate Accurately Future Surgical Times is Unlikely to Reduce the Average Length of Time Cases Finish Late." *Journal of Clinical Anesthesia* 11 (7): 601–605.

Appendix 1. Proofs of theorems

Recall that $V_s(t, \mathbf{n}, \mathbf{n}')$ is defined as the optimal solution to the static problem, with the additional restriction that n_i' patients of each class i must be scheduled. In particular, if we consider the constraints

$$x_{ij} = 1, \quad \forall i = 1, \ldots, m, j = n_i - n_i' + 1, \ldots, n_i \tag{A1}$$

then define $V_s(t, \mathbf{n}, \mathbf{n}')$ as (1-4+A1). Note that this definition is in line with our definition of $V_s(t, \mathbf{n}, \mathbf{n}')$ in Section 3.2.1. Moreover, let $V_s(t, \mathbf{n}, \mathbf{c}) := V_s(t, \mathbf{n}, \mathbf{n}(\mathbf{c}))$.

We first present some results concerning V_s that will be useful in the proofs of Theorems 1–3.

LEMMA 1   *For any t, $\mathbf{n}$, $\mathbf{n}'$ such that $\mathbf{n}' \leq \mathbf{n}$, and $i \in [\mathbf{n} - \mathbf{n}']$, it holds that*

$$V_s(t, \mathbf{n}, \mathbf{n}' + \mathbf{e}_i) \leq r_i + E_{\xi_{in_i}}[V_s(t - \xi_{in_i}, \mathbf{n} - \mathbf{e}_i, \mathbf{n}')]$$

Proof.   Denote any feasible solution of $V_s(t, \mathbf{n}, \mathbf{n}'+\mathbf{e}_i)$ as $x_{jk}, j = 1, \ldots, m, k = 1, \ldots, n_j$. According to the definition of $V_s(t, \mathbf{n}, \mathbf{n}'+\mathbf{e}_i)$, we can derive that $x_{jk} = 1$ for $j = 1, \ldots, m, k = n_j - n_j' + 1, \ldots, n_j$ and $x_{i,n_i-n_i'} = 1$. Moreover, $n_i - n_i' > 0$ due to $i \in [\mathbf{n} - \mathbf{n}']$. Let x_{jk}' equal x_{jk} excluding element x_{in_i} which is 1 for sure, then x_{jk}' is also a feasible solution of $V_s(t - \xi_{in_i}, \mathbf{n} - \mathbf{e}_i, \mathbf{n}')$ for any realisation

of ξ_{in_i}. Its objective value for $r_i + V_s(t - \xi_{in_i}, \mathbf{n} - \mathbf{e}_i, \mathbf{n}')$ is definitely no bigger than the optimal value of $r_i + V_s(t - \xi_{in_i}, \mathbf{n} - \mathbf{e}_i, \mathbf{n}')$. But, taking the expectation, this objective value equals $V_s(t, \mathbf{n}, \mathbf{n}' + e_i)$ with solution x_{jk}. This proves the lemma. $\square$

LEMMA 2 *For any t, $\mathbf{n}$, $\mathbf{n}'$ such that $\mathbf{n}' \leq \mathbf{n}$, it holds that*

$$V_s(t, \mathbf{n}, \mathbf{n}') = \max\left\{ \max_{i \in [\mathbf{n}-\mathbf{n}']}[V_s(t, \mathbf{n}, \mathbf{n}' + \mathbf{e}_i)], V_s(t, \mathbf{n}', \mathbf{n}') \right\}$$

Proof. We first note that $V_s(t, \mathbf{n}, \mathbf{n}')$ is a relaxation of $V_s(t, \mathbf{n}, \mathbf{n}' + \mathbf{e}_i)$ for all i and of $V_s(t, \mathbf{n}', \mathbf{n}')$, so $V_s(t, \mathbf{n}, \mathbf{n}') \geq V_s(t, \mathbf{n}, \mathbf{n}' + \mathbf{e}_i)$ and $\forall i, V_s(t, \mathbf{n}, \mathbf{n}') \geq V_s(t, \mathbf{n}', \mathbf{n}')$. This proves

$$V_s(t, \mathbf{n}, \mathbf{n}') \geq \max\left\{ \max_{i \in [\mathbf{n}-\mathbf{n}']}[V_s(t, \mathbf{n}, \mathbf{n}' + \mathbf{e}_i)], V_s(t, \mathbf{n}', \mathbf{n}') \right\}.$$

Next we will prove the reverse inequality. Let $x_{ij}, i \in \{1, \ldots, m\}, j \in \{1, \ldots, n_i\}$ denote the optimal solution to $V_s(t, \mathbf{n}, \mathbf{n}')$. The one of the following two must be true:

- if $\exists i$, s.t. $x_{i, n_i - n_i'} = 1$, then we can add this constraint to $V_s(t, \mathbf{n}, \mathbf{n}')$ without changing the objective function to obtain $V_s(t, \mathbf{n}, \mathbf{n}' + \mathbf{e}_i)$. So $V_s(t, \mathbf{n}, \mathbf{n}') \leq \max_{i \in [\mathbf{n}-\mathbf{n}']} V_s(t, \mathbf{n}, \mathbf{n}' + \mathbf{e}_i)$.
- if $x_{ij} = 0$ for $i = 1, \ldots, m, j = 1, \ldots, n_i - n_i'$, then $V_s(t, \mathbf{n}, \mathbf{n}') = V_s(t, \mathbf{n}', \mathbf{n}')$.

This proves the reverse inequality, which completes the proof. $\square$

THEOREM 1 *For any remaining time t, patient vector $\mathbf{n}$, committed patient vector $\mathbf{c}$, and for $l' \geq \sum_{i=1}^{m} n_i - k$, it holds that*

$$V_s(t, \mathbf{n}, \mathbf{n}(\mathbf{c})) \leq V_a^1(t, \mathbf{n}, \mathbf{c}) \leq V_a^2(t, \mathbf{n}, \mathbf{c}) \leq \cdots \leq V_a^{l'}(t, \mathbf{n}, \mathbf{c}) = V_a^{l'+1}(t, \mathbf{n}, \mathbf{c}) = \cdots = V(t, \mathbf{n}, \mathbf{c})$$

and for policies without committing and $l' \geq \sum_{i=1}^{m} n_i$, it holds that

$$V_s(t, \mathbf{n}) \leq V_a^1(t, \mathbf{n}) \leq V_a^2(t, \mathbf{n}) \leq \cdots \leq V_a^{l'}(t, \mathbf{n}) = V_a^{l'+1}(t, \mathbf{n}) = \cdots = V(t, \mathbf{n})$$

Proof. We will first prove $V_s(t, \mathbf{n}, \mathbf{n}(\mathbf{c})) \leq V_a^1(t, \mathbf{n}, \mathbf{c})$. We have

$$V_s(t, \mathbf{n}, \mathbf{n}(\mathbf{c})) = \max\left\{ \max_{i \in [\mathbf{n}-\mathbf{n}(\mathbf{c})]}\{V_s(t, \mathbf{n}, \mathbf{n}(\mathbf{c}) + \mathbf{e}_i)\}, V_s(t, \mathbf{n}(\mathbf{c}), \mathbf{n}(\mathbf{c})) \right\},$$

$$\leq \max\left\{ \max_{i \in [\mathbf{n}-\mathbf{n}(\mathbf{c})]}\{r_c + E_{\xi_{cn_c}}[V_s(t - \xi_{cn_c}, \mathbf{n} - \mathbf{e}_c, \mathbf{n}(\tilde{\mathbf{c}}))]\}, V_s(t, \mathbf{n}(\mathbf{c}), \mathbf{n}(\mathbf{c})) \right\},$$

$$= V_a^1(t, \mathbf{n}, \mathbf{c}).$$

Here, and throughout this appendix, we let $\mathbf{c} = (c, c_2, \ldots, c_k)$ and $\tilde{\mathbf{c}} = (c_2, \ldots, c_k, i)$. (The subscript $_1$ of the first coordinate c of $\mathbf{c}$ is omitted to improve legibility.) Here, the first equality follows from Lemma 2 with $\mathbf{n}' = \mathbf{n}(\mathbf{c})$, and the inequality from Lemma 1 with $i = c$. The final equality is a specialisation of (9) to $l = 1$. The result $V_s(t, \mathbf{n}) \leq V_a^1(t, \mathbf{n})$ follows likewise, by noting that $V_s(t, \mathbf{n}) = V_s(t, \mathbf{n}, \mathbf{0})$, with $\mathbf{0}$ the zero vector.

Now, we prove $V_a^l(t, \mathbf{n}, \mathbf{c}) \leq V_a^{l+1}(t, \mathbf{n}, \mathbf{c})$ for all $l \geq 0$ by induction on l. We just proved $l = 0$. Assume the inequality holds for $l - 1$. If $\mathbf{c} \neq \mathbf{0}$, we have

$$V_a^l(t, \mathbf{n}, \mathbf{c}) = \max\left\{ r_c + \max_{i \in [\mathbf{n}-\mathbf{n}(\mathbf{c})]} E_{\xi_{cn_c}}[V_a^{l-1}(t - \xi_{cn_c}, \mathbf{n} - \mathbf{e}_c, \tilde{\mathbf{c}})], V_s(t, \mathbf{n}(\mathbf{c}), \mathbf{n}(\mathbf{c})) \right\} \quad (A2)$$

$$V_a^{l+1}(t, \mathbf{n}, \mathbf{c}) = \max\left\{ r_c + \max_{i \in [\mathbf{n}-\mathbf{n}(\mathbf{c})]} E_{\xi_{cn_c}}[V_a^{l}(t - \xi_{cn_c}, \mathbf{n} - \mathbf{e}_c, \tilde{\mathbf{c}})], V_s(t, \mathbf{n}(\mathbf{c}), \mathbf{n}(\mathbf{c})) \right\}$$

From this, the induction step follows readily. More precisely, if $V_s(t, \mathbf{n}(\mathbf{c}), \mathbf{c})$ attains the outer maximum in (A2), then the result is immediate. On the other hand, if $r_c + \max_{i \in [\mathbf{n}-\mathbf{n}(\mathbf{c})]} E_{\xi_{cn_c}}[V_a^{l-1}(t - \xi_{cn_c}, \mathbf{n} - \mathbf{e}_c, \tilde{\mathbf{c}})]$ attains the maximum, let i^* attain the inner maximisation and $\tilde{\mathbf{c}}^* = (c_2, \ldots, c_k, i^*)$. Then,

$$V_a^l(t, \mathbf{n}, \mathbf{c}) = r_c + E_{\xi_{cn_c}}[V_a^{l-1}(t - \xi_{cn_c}, \mathbf{n} - \mathbf{e}_c, \tilde{\mathbf{c}}^*)],$$

$$\leq r_c + E_{\xi_{cn_c}}[V_a^{l}(t - \xi_{cn_c}, \mathbf{n} - \mathbf{e}_c, \tilde{\mathbf{c}}^*)],$$

$$\leq r_c + \max_{i \in [\mathbf{n}-\mathbf{n}(\mathbf{c})]} E_{\xi_{cn_c}}[V_a^{l}(t - \xi_{cn_c}, \mathbf{n} - \mathbf{e}_c, \tilde{\mathbf{c}})],$$

$$= V_a^{l+1}(t, \mathbf{n}, \mathbf{c}),$$

where the first inequality follows from induction hypothesis. If $\mathbf{c} = \mathbf{0}$, it also holds using their definitions and results of $V_a^l(t, \mathbf{n}, \mathbf{c}) \leq V_a^{l+1}(t, \mathbf{n}, \mathbf{c})$ when $\mathbf{c} \neq \mathbf{0}$. This completes the proof by induction. The proof of $V_a^l(t, \mathbf{n}) \leq V_a^{l+1}(t, \mathbf{n})$ goes in the same fashion.

Finally, we need to prove that for any l' and any $\mathbf{n}, \mathbf{c}$ with $l' \geq \sum_{i=1}^{m} n_i - k$, it holds that $V_a^{l'}(t, \mathbf{n}, \mathbf{c}) = V(t, \mathbf{n}, \mathbf{c})$. For ease of exposition, we define $\bar{O}_{\mathbf{n}, \mathbf{c}} := \sum_{i=1}^{m} n_i - k$. We proceed by induction on l'. For $l' = 0$, we have $V(t, \mathbf{n}, \mathbf{c}) = V(t, \mathbf{n}(\mathbf{c}), \mathbf{c}) = V_s(t, \mathbf{n}(\mathbf{c}), \mathbf{c})$, because

$\mathbf{c}$ are all the available patients since $\mathbf{n} = \mathbf{n}(\mathbf{c})$. Now, assume the statement holds for some l', and let $\mathbf{n}^+$ be any patient vector such that $\bar{O}_{\mathbf{n}^+, \mathbf{c}} \leq l' + 1$. If $\mathbf{c} \neq \mathbf{0}$, we then have

$$V_a^{l'+1}(t, \mathbf{n}^+, \mathbf{c}) := \max \left\{ r_c + \max_{i \in [\mathbf{n}^+ - \mathbf{n}(\mathbf{c})]} E_{\xi_{cn_c}} [V_a^{l'}(t - \xi_{cn_c}, \mathbf{n}^+ - \mathbf{e}_c, \tilde{\mathbf{c}})], V_s(t, \mathbf{n}^+(\mathbf{c}), \mathbf{c}) \right\}$$

$$= V(t, \mathbf{n}^+, \mathbf{c}) \tag{A3}$$

because $\bar{O}_{\mathbf{n}^+ - \mathbf{e}_c, \tilde{\mathbf{c}}} = \bar{O}_{\mathbf{n}^+, \mathbf{c}} - 1 \leq l'$, implying that $V_a^{l'}(t - \xi_{cn_c}, \mathbf{n}^+ - \mathbf{e}_c, \tilde{\mathbf{c}}) = V(t - \xi_{cn_c}, \mathbf{n}^+ - \mathbf{e}_c, \tilde{\mathbf{c}})$ by induction hypothesis. This completes the inductive argument.

For the proof of $V_a^{l'}(t, \mathbf{n}) = V(t, \mathbf{n})$, induction starts at $l' = 0$, for which the result is immediate because $V_a^l(t, \mathbf{n}) = V(t, \mathbf{n})$ since $\bar{O}_{\mathbf{n}} = 0$ implies there are no patients. The induction step goes similar to the induction step for the case with committing. $\square$

THEOREM 2   *For any remaining time t, patient vector $\mathbf{n}$, committing vector $\mathbf{c}$, s.t. $\mathbf{n}(\mathbf{c}) \leq \mathbf{n}$, and adaptivity level $l \geq 1$, it holds that* $V_{act}^l(t, \mathbf{n}, \mathbf{c}) \geq V_a^l(t, \mathbf{n}, \mathbf{c})$ *and* $V_{act}^l(t, \mathbf{n}) \geq V_a^l(t, \mathbf{n})$.

Proof. We show that for any o', any $\mathbf{n}$ and $\mathbf{c}$ with $\bar{O}_{\mathbf{n}, \mathbf{c}} = \sum_{i=1}^m n_i - k \leq o'$ and any $\mathbf{n}(\mathbf{c}) \leq \mathbf{n}$ and t it holds that $V_{act}^l(t, \mathbf{n}, \mathbf{c}) \geq V_a^l(t, \mathbf{n}, \mathbf{c})$. We proceed by induction on o' when $\mathbf{c} \neq \mathbf{0}$. For $o' = 0$, $V_{act}^l(t, \mathbf{n}, \mathbf{c}) = V_s(t, \mathbf{n}, \mathbf{c}) = V_a^l(t, \mathbf{n}, \mathbf{c})$, since only patients $\mathbf{c}$ are left. Now, assume the statement holds for o'. For any $t, l \geq 1$, $\mathbf{n}$, $\mathbf{c}$ such that $\bar{O}_{\mathbf{n}, \mathbf{c}} \leq o' + 1$, and $\mathbf{n}(\mathbf{c}) \leq \mathbf{n}$, we have:

$$V_a^l(t, \mathbf{n}, \mathbf{c}) = \max \left\{ r_c + \max_{i \in [\mathbf{n} - \mathbf{n}(\mathbf{c})]} E_{\xi_{cn_c}} [V_a^{l-1}(t - \xi_{cn_c}, \mathbf{n} - \mathbf{e}_c, \tilde{\mathbf{c}})], V_s(t, \mathbf{n}(\mathbf{c}), \mathbf{c}) \right\}$$

where $\tilde{\mathbf{c}} = (c_2, \ldots, c_k, i)$.

Case 1: $V_s(t, \mathbf{n}(\mathbf{c}), \mathbf{c})$ attains the maximum. In that case, $V_{act}^l(t, \mathbf{n}, \mathbf{c}) = V_a^l(t, \mathbf{n}, \mathbf{c})$. Case 2: $r_c + \max_{i \in [\mathbf{n} - \mathbf{n}(\mathbf{c})]} E_{\xi_{cn_c}} [V_a^{l-1}(t - \xi_{cn_c}, \mathbf{n} - \mathbf{e}_c, \tilde{\mathbf{c}})]$ attains the maximum. In that case, let i^* denote the patient class that attains the inner maximum and $\tilde{\mathbf{c}}^* = (c_2, \ldots, c_k, i^*)$. We thus obtain

$$V_a^l(t, \mathbf{n}, \mathbf{c}) = r_c + E_{\xi_{cn_c}} [V_a^{l-1}(t - \xi_{cn_c}, \mathbf{n} - \mathbf{e}_c, \tilde{\mathbf{c}}^*)],$$

$$\leq r_c + E_{\xi_{cn_c}} [V_a^l(t - \xi_{cn_c}, \mathbf{n} - \mathbf{e}_c, \tilde{\mathbf{c}}^*)],$$

$$\leq r_c + E_{\xi_{cn_c}} [V_{act}^l(t - \xi_{cn_c}, \mathbf{n} - \mathbf{e}_c, \tilde{\mathbf{c}}^*)],$$

$$= V_{act}^l(t, \mathbf{n}, \mathbf{c}).$$

Here, the first inequality is a consequence of Theorem 1. The second inequality is the induction hypothesis, which may be applied because $\bar{O}_{\mathbf{n} - \mathbf{e}_i, \tilde{\mathbf{c}}^*} = \bar{O}_{\mathbf{n}, \mathbf{c}} - 1 \leq o'$. The final equality in fact corresponds to the definition of $V_{act}^l(t, \mathbf{n}, \mathbf{c})$, because by assumption, patient class j was optimal for $V_a^l(t, \mathbf{n}, \mathbf{c})$. This completes the induction step. Using this result and definitions of $V_a^l(t, \mathbf{n}, \mathbf{0})$ and $V_{act}^l(t, \mathbf{n}, \mathbf{0})$ we can deduce that $V_a^l(t, \mathbf{n}, \mathbf{0}) \leq V_{act}^l(t, \mathbf{n}, \mathbf{0})$. The proof of $V_{act}^l(t, \mathbf{n}) \geq V_a^l(t, \mathbf{n})$ goes likewise. $\square$

THEOREM 3   (COMMITTING DETERIORATES PERFORMANCE) *Let $k, k' \in \mathbb{N}$ with $k < k'$ and consider committing vectors $\mathbf{c}_1 = (c, c_2, \ldots, c_k)$ and $\mathbf{c}_2 = (c, c_2, \ldots, c_k, c_{k+1}, \ldots, c_{k'})$. Then, for any time t, patient vector $\mathbf{n}$ such that $\mathbf{n} \geq \mathbf{n}(\mathbf{c}_2)$, and adaptivity level $l \geq 0$ it holds that*

$$V_a^l(t, \mathbf{n}, \mathbf{c}_1) \geq V_a^l(t, \mathbf{n}, \mathbf{c}_2)$$

and that:

$$V(t, \mathbf{n}, \mathbf{c}_1) \geq V(t, \mathbf{n}, \mathbf{c}_2).$$

Proof. We first consider the case $k' = k + 1$: let $\mathbf{c}_1 = (c, c_2, \ldots, c_k)$ and $\mathbf{c}_2 = (c, c_2, \ldots, c_k, c_{k+1})$. We will prove that $V_a^l(t, \mathbf{n}, \mathbf{c}_1) \geq V_a^l(t, \mathbf{n}, \mathbf{c}_2)$ by induction on l. For $l = 0$, the result follows from $V_s(t, \mathbf{n}, \mathbf{n}(\mathbf{c}_1)) \geq V_s(t, \mathbf{n}, \mathbf{n}(\mathbf{c}_2))$, which follows readily from Lemma 1 since $\mathbf{n}(\mathbf{c}_1) \leq \mathbf{n}(\mathbf{c}_2)$ by assumption. Now, let the induction hypothesis be that the statement holds for some l. Denote $\tilde{\mathbf{c}}_2 = (c_2, c_3, \ldots, c_k, c_{k+1}, i)$. We have

$$V_a^{l+1}(t, \mathbf{n}, \mathbf{c}_2) = \max \left\{ r_c + \max_{i \in [\mathbf{n} - \mathbf{n}(\mathbf{c}_2)]} E_{\xi_{cn_c}} [V_a^l(t - \xi_{cn_c}, \mathbf{n} - \mathbf{e}_c, \tilde{\mathbf{c}}_2)], V_s(t, \mathbf{n}(\mathbf{c}_2), \mathbf{n}(\mathbf{c}_2)) \right\}.$$

Case 1: suppose that $V_s(t, \mathbf{n}(\mathbf{c}_2), \mathbf{n}(\mathbf{c}_2))$ attains the maximum. Then

$$V_a^{l+1}(t, \mathbf{n}, \mathbf{c}_1) \geq V_s(t, \mathbf{n}, \mathbf{n}(\mathbf{c}_1)) \geq V_s(t, \mathbf{n}, \mathbf{n}(\mathbf{c}_2)) \geq V_s(t, \mathbf{n}(\mathbf{c}_2), \mathbf{n}(\mathbf{c}_2))$$

so we obtain $V_a^{l+1}(t, \mathbf{n}, \mathbf{c}_1) \geq V_a^{l+1}(t, \mathbf{n}, \mathbf{c}_2)$.

Case 2: Suppose:

$$V_a^{l+1}(t, \mathbf{n}, \mathbf{c}_2) = r_c + \max_{i \in [\mathbf{n} - \mathbf{n}(\mathbf{c}_2)]} E_{\xi_{cn_c}} [V_a^l(t - \xi_{cn_c}, \mathbf{n} - \mathbf{e}_c, \tilde{\mathbf{c}}_2)].$$

Denote $\tilde{\mathbf{c}}_1 = (c_2, c_3, \ldots, c_k, i')$. Then

$$V_a^{l+1}(t, \mathbf{n}, \mathbf{c}_2) \leq r_c + E_{\xi_{cn_c}}[V_a^l(t - \xi_{cn_c}, \mathbf{n} - \mathbf{e}_c, (c_2, c_3, \ldots c_k, c_{k+1}))],$$

$$\leq r_c + \max_{i' \in [\mathbf{n} - \mathbf{n}(\mathbf{c}_1)]} E_{\xi_{cn_c}}[V_a^l(t - \xi_{cn_c}, \mathbf{n} - \mathbf{e}_c, \tilde{\mathbf{c}}_1)],$$

$$\leq V_a^{l+1}(t, \mathbf{n}, \mathbf{c}_1),$$

where the first inequality follows from induction hypothesis. This completes the induction step, which completes the proof of the case $k' = k + 1$. The general case where $k' > k$ follows by repeatedly applying the case $k' = k + 1$.

The second claim follows from the first claim by noting that Theorem 1 guarantees that $V_a^l(t, \mathbf{n}, \mathbf{c}) = V(t, \mathbf{n}, \mathbf{c})$ and $V_a^l(t, \mathbf{n}) = V(t, \mathbf{n})$ for $l \geq \bar{O}_{\mathbf{n}, \mathbf{c}}$. $\square$

THEOREM 4 *For any parameter θ, and state $(t, \mathbf{n})$, $OPT(\theta) = R(\theta) - p \cdot E_\zeta[(\zeta - t)^+]$, with $\zeta \sim f(\theta)$.*

Proof. According to the definitions in (11) and (12), let x_{ij}^* $(i = 1, \ldots, m, j = 1, \ldots, n_i)$ be the optimal solution to $R(\theta)$, it is also a feasible solution to $OPT(\theta)$, it is easy to deduce that $OBJ(\theta) \leq OPT(\theta)$. On the contrary, with fixed θ, the second term in $OPT(\theta)$ is fixed, maximise $OPT(\theta)$ is equivalent to maximise $R(\theta)$, and if no feasible solution exist, both objective functions are $-\infty$. Therefore, $OPT(\theta) = OBJ(\theta)$ for any group $(t, \mathbf{n})$ and θ. $\square$

THEOREM 5 *For any $\mathbf{n}$ and θ, prove*

$$R(\mathbf{n}, \theta) = \max_{x_{\ell_n n_{\ell_n}} \in \{0,1\}} R(\mathbf{n} - x_{\ell_n n_{\ell_n}} e_{\ell_n} - n_{\ell_n}(1 - x_{\ell_n n_{\ell_n}})e_{\ell_n}, \theta - \theta_{\ell_n} x_{\ell_n n_{\ell_n}}) + r_{\ell_n} x_{\ell_n n_{\ell_n}} \tag{A4}$$

Proof. We will give the proof by induction.

- For $\ell_n = n_{\ell_n} = 1$, if $\theta_{\ell_n} = \theta$, by letting $x_{\ell_n n_{\ell_n}} = 1$, both sides of (A4) get value r_{ℓ_n}, else both equal to $-\infty$. Therefore the theorem is true for $\ell_n = n_{\ell_n} = 1$.
- Assume the recursion is correct for $\mathbf{n}_0 = \mathbf{n} - \mathbf{e}_j$ and $\ell_{n_0} \geq 1$, next we will prove it is also correct for $\mathbf{n}$.

  ○ if $\exists \mathbf{x}$, s.t. $\sum_{i=1}^{\ell_n} \sum_{j=1}^{n_i} \theta_i x_{ij} = \theta$, $x_{ij} \in \{0, 1\}$, then $R(\mathbf{n}, \theta) > -\infty$, denoting its optimal solution as $\mathbf{x}^* = (x_{11}^*, \ldots, x_{1n_i}^*, \ldots, x_{\ell_n 1}^*, \ldots, x_{\ell_n n_{\ell_n}}^*)$, then $\sum_{i=1}^{\ell_n} \sum_{j=1}^{n_i} \theta_i x_{ij}^* = \theta$.
  For the right hand side of expression (A4), we let $x_{\ell_n n_{\ell_n}} = x_{\ell_n n_{\ell_n}}^*$,

    – if $x_{\ell_n n_{\ell_n}}^* = 1$, then $(x_{11}^*, \ldots, x_{1n_i}^*, \ldots, x_{\ell_n 1}^*, \ldots, x_{\ell_n n_{\ell_n}-1}^*)$ is a feasible solution of $R(\mathbf{n} - e_{\ell_n}, \theta - \theta_{\ell_n})$
    – if $x_{\ell_n n_{\ell_n}}^* = 0$, then $(x_{11}^*, \ldots, x_{1n_i}^*, \ldots, x_{\ell_n-11}^*, \ldots, x_{\ell_n-1n_{\ell_n-1}}^*)$ is a feasible solution of $R(\mathbf{n} - n_{\ell_n} e_{\ell_n}, \theta)$

  thus:

  $$R(\mathbf{n}, \theta) \leq R(\mathbf{n} - x_{\ell_n n_{\ell_n}} e_{\ell_n} - n_{\ell_n}(1 - x_{\ell_n n_{\ell_n}})e_{\ell_n}, \theta - \theta_{\ell_n} x_{\ell_n n_{\ell_n}}) + r_{\ell_n} x_{\ell_n n_{\ell_n}}$$

  and vice versa.
  ○ if θ is not feasible, i.e. $\mathbf{n}exists\mathbf{x}$, s.t. $\sum_{i=1}^{\ell_n} \sum_{j=1}^{n_i} \theta_i x_{ij} = \theta$, $x_{ij} \in \{0, 1\}$, based on definition (2), we have $R(\mathbf{n}, \theta) = -\infty$.
  For the right-hand side of (A4), if there exists a $x_{\ell_n n_{\ell_n}} \in \{0, 1\}$,
  $s.t.$ $R(\mathbf{n} - x_{\ell_n n_{\ell_n}} e_{\ell_n} - n_{\ell_n}(1 - x_{\ell_n n_{\ell_n}})e_{\ell_n}, \theta - \theta_{\ell_n} x_{\ell_n n_{\ell_n}}) + r_{\ell_n} \cdot x_{x_{\ell_n n n_{\ell_n}}} > -\infty$.

    – if $x_{\ell_n n_{\ell_n}} = 1$, and $(x_{11}, \ldots, x_{1n_i}, \ldots, x_{\ell_n 1}, \ldots, x_{\ell_n n_{\ell_n}-1})$ is a feasible solution of $R(\mathbf{n} - e_{\ell_n}, \theta - \theta_{\ell_n})$, then $\sum_{i=1}^{\ell_n} \sum_{j=1}^{n_i} \theta_i x_{ij} = \theta$.
    – if $x_{\ell_n n_{\ell_n}} = 0$, and $(x_{11}, \ldots, x_{1n_i}, \ldots, x_{\ell_n-11}, \ldots, x_{\ell_n-1n_{\ell_n-1}})$ is a feasible solution of $R(\mathbf{n} - n_{\ell_n} e_{\ell_n}, \theta)$, then $\sum_{i=1}^{\ell_n} \sum_{j=1}^{n_i} \theta_i x_{ij} = \theta$ by letting $x_{\ell_n j} = 0$, $j = 1, \ldots, n_{\ell_n}$.

  Then it contradicts to our assumption. So the right-hand side of (A4) is also $-\infty$, and the theorem is true.

  $\square$

Appendix 2. Surgical time distributions

B.1 *Finite parameter compound Poisson distributions*

We use a finite-parameter compound Poisson distribution to fit the surgical time. A Poisson random variable with rate l^θ will be denoted by $p(l^\theta)$.

To construct the finite-parameter class with m parameters, we can freely choose constants $q_1, q_2, \ldots, q_m$ as coefficients of a series of Poisson random variables. After these constants are fixed, each distribution in the class will be fully characterised by the rates $l_k^\theta \in R^+$ of Poisson random variables for $k = 1, 2, \ldots, m$. If for class i we have $l_k^\theta = l_{ik}^\theta$, then the surgical time ξ_{ij} for this class is distributed as follows:

Table B1. Parameters of surgical times and fitting coefficients for low dimensional instances

Groups	Van Houdenhoven et al. (2007) (vHOHWK)			Erasmus medical center		
	Class #	(mean,SD)	$(l_1^\theta, l_2^\theta, l_3^\theta)$	Class #	(mean,SD)	$(l_1^\theta, l_2^\theta, l_3^\theta)$
High CV	1	(80, 65)	(3.838,0, 0.149)	1	(62,18)	(12.371, 0.001, 0.003)
	2	(119, 107)	(3.159,0, 0.477)	2	(97,32)	(17.349, 0.002, 0.168)
	3	(192, 165)	(0.777,0, 1.202)	3	(181,57)	(27.914, 0.004, 0.692)
	Avg. CV=0.86		$(q_1, q_2, q_3) = (15, 60, 150)$	Avg. CV=0.31		$(q_1, q_2, q_3) = (5, 15, 60)$
Low CV	1	(135, 52)	(8.714,0, 0.024)	1	(187,33)	(34.805, 0.861, 0)
	2	(100, 68)	(5.363,0, 0.111)	2	(215,41)	(35.298, 2.225, 0.081)
	3	(102, 125)	(0.926,0, 0.503)	3	(276,53)	(37.960, 4.942, 0.205)
	Avg. CV=0.76		$(q_1, q_2, q_3) = (15, 60, 175)$	Avg. CV=0.19		$(q_1, q_2, q_3) = (5, 15, 60)$

$$\xi_{ij} \sim \sum_{k=1}^{m} q_k p(l_{ik}^\theta)$$

Because $p(l^\theta) + p(l'^\theta) \sim p(l^\theta + l'^\theta)$, this class of compounding distributions is closed under addition. It will be denoted as $f(\theta), \theta = (l_1^\theta, l_2^\theta, \dots, l_m^\theta)$. For our experiments, we will use three parameters $(m = 3)$.

B.2 *Surgical time distributions for the benchmark cases*

In this appendix, we give the parameters (μ, σ) and $(l_1^\theta, l_2^\theta, l_3^\theta)$ for each of the patient classes in each of the cases that are used in our numerical experiments. Two low-dimensional cases were based on mean and variance of surgery distributions from Van Houdenhoven et al. (2007), and two cases were based on data from Erasmus medical centre. Data for these cases are summarised in Table B1. Four high-dimensional cases correspond to the patient classes in four specialisms at Erasmus medical centre. Data are summarised in Table B2.

B.3 *Compound Poisson goodness of fit*

We test the goodness of fit of the three-parameter compound Poisson distribution for fitting the surgery time data obtained from Erasmus medical centre. The data that we fit includes a set-up time of 20 min for cleaning and sterilisation, which is added to each surgical time before distribution estimation.

For our tests, we focus on procedures with the largest number of samples in each of the four specialism. Throughout, we use $q_1 = 5, q_2 = 15, q_3 = 60$, and we fit the parameters $(l_1^\theta, l_2^\theta, l_3^\theta)$ based on mean and variance of the samples. The goodness of fit of the resulting distributions is analysed by Quantile–Quantile Plot (qq-plot) and a chi-square test. In all the graphs below, the solid line with $y = x$ works as a benchmark to reflect a perfect fit. The dashed line is the quantile values corresponding to the theoretical compound Poisson distribution, and dots are our samples. If two distributions are similar, the points will approximately lie on the green line; if there exists an approximate affine relation, the points are approximately on a straight line. From Figures B1–B4, we observe that the distributions fit the data quite well. A chi-squared goodness of fit test with 9 bins confirms this hypothesis at the significance level $p = 0.05$.
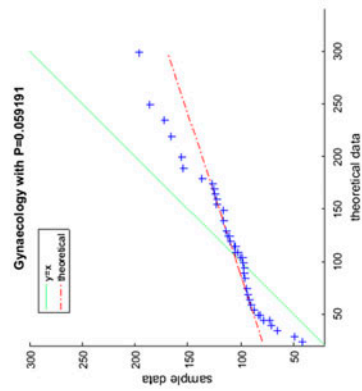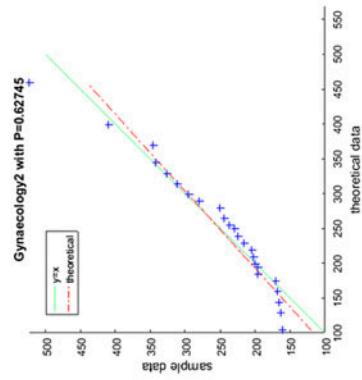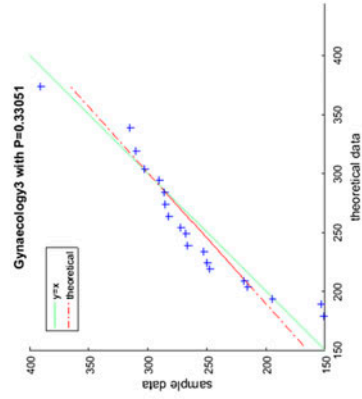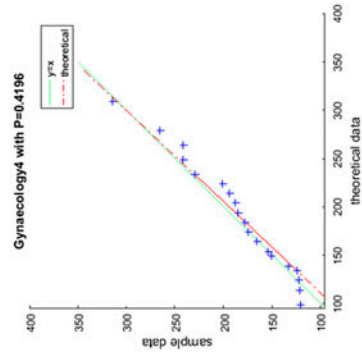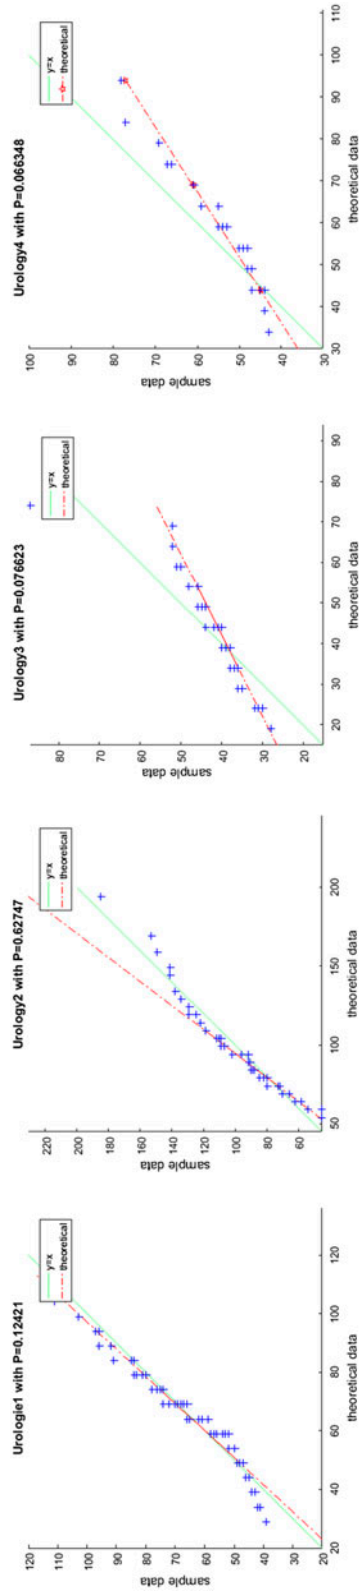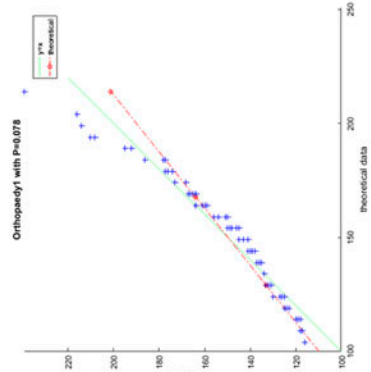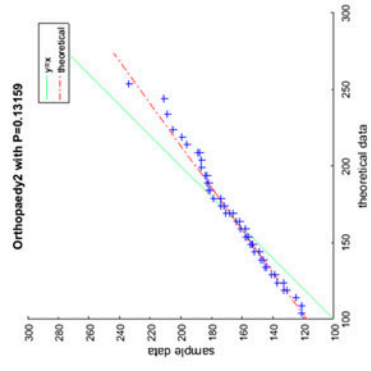
Figure B1. Gynaecology surgery.
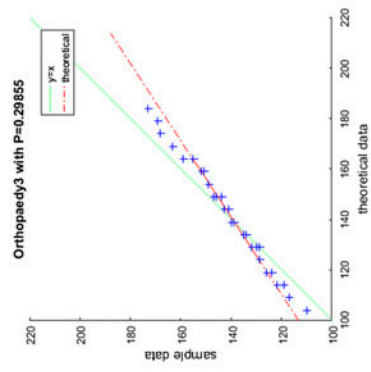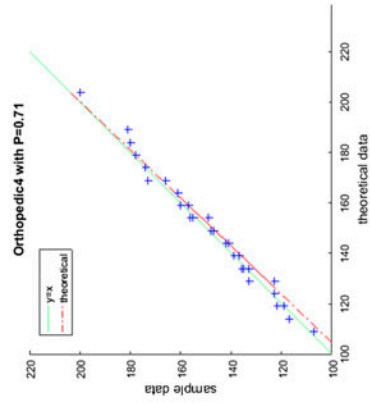
Figure B2.   Urology surgery.
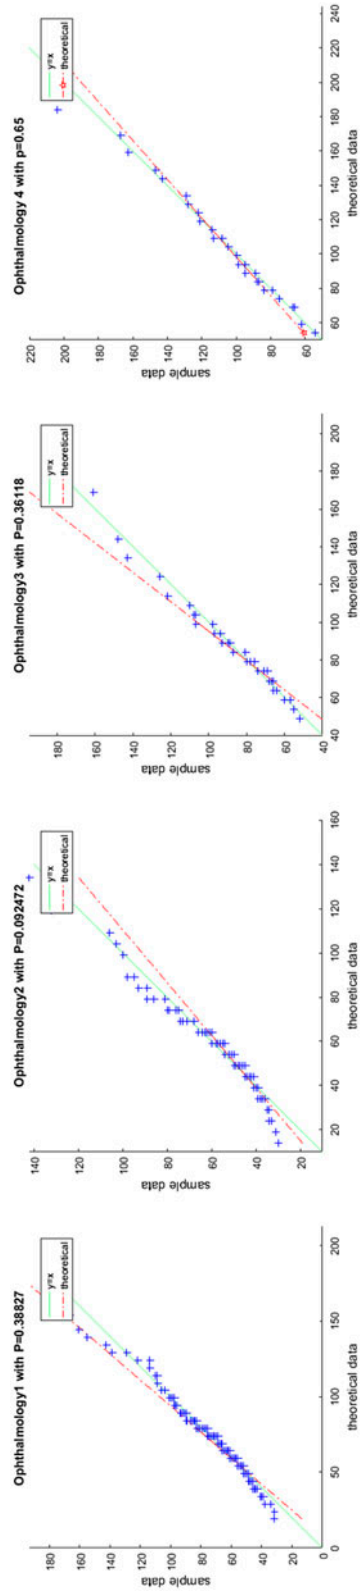
Figure B3. Orthopaedy surgery.

Figure B4. Ophthalmology surgery.

Table B2. Surgical time mean, standard deviation, and fitting parameters obtained for procedures in various categories, from the data of Erasmus medical centre.

Case	(mean,SD)	$(l_1^\theta, l_2^\theta, l_3^\theta)$
Ophthalmology	(85.31,12.95)	(18.249,0,0)
	(55.35,17.92)	(10.904,0.001,0.013)
	(110.76,38.85)	(18.670,0.002,0.289)
	(87.31,27.99)	(16.194, 0.002, 0.105)
	(53.82,19.92)	(10.296, 0.001, 0.038)
	(71.03,24.32)	(13.340, 0.002, 0.071)
Gynaecology	(135.53, 44.85)	(22.249, 0.003, 0.403)
	(150.27,39.74)	(27.033,0.004,0.250)
	(121.81,38.32)	(21.229,0.003,0.260)
	(73.13,29.69)	(12.744,0.002,0.156)
	(87.57,22.67)	(17.231,0.002,0.022)
	(190.14,81.25)	(17.473,0.002,1.712)
	(249,65.65)	(34.694,1.614,0.855)
	(253.04,88.47)	(26.736,0.004,1.988)
	(117.02,64.55)	(10.376,0.001,1.085)
Urology surgery	(49.6, 12.44)	(10.349,0,0)
	(276.25, 52.91)	(37.960,4.942,0.205)
	(187, 32.66)	(34.805,0.861,0)
	(79.47, 17.35)	(16.337,0,0)
	(66.38, 15.69)	(13.671,0,0)
	(214.75, 40.93)	(35.298,2.225,0.081)
	(181.2, 56.51)	(27.914,0.004,0.692)
	(121.55, 27.55)	(23.751,0.003,0.045)
	(53.27, 10.39)	(11.382,0,0)
	(42.03, 10.02)	(8.911,0,0)
	(96.87, 32.24)	(17.349,0.002,0.168)
	(62.09, 17.94)	(12.371,0.001,0.003)
Orthopedic surgery	(77.93, 21.75)	(15.278,0.002,0.025)
	(143.11, 37.79)	(26.022,0.003,0.215)
	(159.89, 38.72)	(29.422,0.004,0.211)
	(77.15, 56.21)	(5.339,0,0.840)
	(235.58, 88.46)	(22.934,0.003,2.014)
	(258.68, 61.49)	(35.852,2.786,0.627)
	(144.77, 25.16)	(29.369,0,0)
	(144.25, 28.46)	(28.517,0.004,0.026)
	(167.56, 40.01)	(30.727,0.004,0.230)
	(149.14, 26.51)	(30.024,0,0)