Evaluating the impact of flexible practices on the master surgical scheduling process: an empirical analysis

Evaluating the impact of flexible practices on the master surgical scheduling process: an empirical analysis

FILIPPO VISINTIN
University of Florence
Dipartimento di Ingegneria Industriale
IBIS Lab.
PO Box 50134, Viale Morgagni 40/44, Florence, Italy

PAOLA CAPPANERA
University of Florence
Dipartimento di Ingegneria dell'Informazione
IBIS Lab.
PO Box 50139, Via Santa Marta 3, Florence, Italy

AND

CARLO BANDITORI
University of Florence
Dipartimento di Ingegneria Industriale
IBIS Lab.
PO Box 50134, Viale Morgagni 40/44, Florence, Italy

DOI 10.1007/s10696-014-9208-9

*corresponding Author

Abstract

This study focuses on the *master surgical scheduling* problem and adds two main contributions. First, it presents a novel mixed integer programming model to support the master surgical schedule production. Second, it uses the model to investigate the impact, in terms of scheduled surgeries, of the flexible management of three critical resources, namely *surgical teams*, *operating rooms* and *surgical units*. Our analysis revealed that to maximise the number of surgeries scheduled, it is sufficient to introduce flexibility with respect to surgical teams and ORs. In fact, if both these resources are managed flexibly, then introducing flexibility with respect to surgical units carries no additional advantages. However, if surgical teams or ORs (or both) are not managed flexibly, then managing surgical units flexibly produces significant benefits. In addition, our study shows that if surgical teams cannot be managed flexibly, then introducing flexibility with respect to ORs yields significant benefits. Similarly, it reveals that if ORs cannot be managed flexibly, then introducing flexibility with respect to surgical teams yields significant benefits as well. The work is based on real data from the Meyer University Children's Hospital in Florence.

1 Introduction

The Operating Theatre (OT) is widely acknowledged as the functional area driving most hospitals' costs and revenues (Denton et al., 2007). The surgical scheduling process, i.e. the process by which OT activities are planned, dramatically influences OT performance and, as such, it is the object of growing attention from hospital managers worldwide. Such a process, however, is extremely complex to manage. In fact, it requires the consideration of many resources (operating rooms (ORs), surgical teams, and nursing staff as well as downstream resources, such as surgical units and intensive care units (ICUs)) operating in a context affected by a high variability (Litvak and Long, 2000) and characterised by people - surgeons, patients, hospital managers - with conflicting priorities (Glouberman and Mintzberg, 2001). The complexity of the surgical scheduling process coupled with its significant economic and social impact has thus stimulated, in recent years, intensive research activities as well (Cardoen et al., 2010, Guerriero and Guido, 2011, May et al., 2011). The literature, indeed, abounds with models supporting the scheduling of surgical activities. In particular, the mainstream literature presents the consensus that solving a surgical scheduling problem requires addressing three intertwined sub-problems (Beliën and Demeulemeester, 2007): (i) the *case-mix planning*, i.e. the determination (usually on a yearly basis) of the total amount of OR time to assign to each surgical specialty, (ii) the *master surgical scheduling*, i.e. the determination of the specialty (or specialties) to assign to each OR on each day of the planning horizon (e.g. two weeks or one month) and, in certain cases, the specification of the number and typology of surgeries to be performed each day, and finally (iii) the *selection and sequencing* of patients who have to undergo surgery. Typically, these three sub-problems are solved in cascade; the case-mix determined at the first stage is used in the definition of the master surgical schedule (MSS). The MSS, in turn, is used as input for patient selection and sequencing.

This study focuses on the *master surgical scheduling* sub-problem. In the literature, the models supporting such a sub-problem consider slightly different sets of resources (ORs, surgical units, surgical teams, and the ICU) and make different assumptions about how *flexibly* these critical resources are managed. Some studies propose scenario analysis to assess the effects associated with the flexible management of certain resources, such as surgical teams or surgical units (Banditori et al., 2013, Agnetis et al., 2012). However, despite the fact that flexibility is by no means a new topic (Balasubramanian et al., 2012, Buzacott and Mandelbaum, 2008, Chou et al., 2008, Gupta and Shanthikumar, 2008), to date the literature lacks of contributions that have systematically studied the impact of flexibility on OT performance.

This study addresses this gap by adding two main contributions. First, it presents a novel mixed integer programming model to support MSS production. Second, it uses the model to investigate the main and interaction effects associated with the flexible management of three critical resources: *surgical teams*, *ORs* and *surgical units*.

The model assumes that surgical cases can be organised into homogeneous surgery groups (Santibáñez et al., 2007, Banditori et al., 2013) based on their specialty, their expected surgical time (ST) and their expected length of stay (LoS), that is, based on the extent to which these cases are expected to “consume” the previously mentioned critical resources. The model creates a solution indicating for each *OR session* (i.e. for each day, for each OR and for each session) in the planning horizon the number of surgeries to perform and the surgery group these cases must belong to. The model’s objective function is the maximisation of the number of scheduled surgeries.

In addition to presenting the model, we show how such a model can be modified by acting on its variables, parameters and constraints to incorporate a more or less flexible management of surgical teams, ORs and surgical units. The different versions of the model are then used to carry out an experimental campaign based on a 2^3 experimental design (Montgomery and Runger, 2003). In detail, we consider the way the three critical resources are managed as *factors* and we assume two possible levels for each factor: “high” when the resource is managed in a flexible way and “low” otherwise. More specifically:

1) With respect to surgical teams (“*Teams*” factor), we analyse the case where the assignment of surgical teams to sessions is fixed (fixed surgical teams assignment, low level) and the case where such an assignment can change every time the MSS is produced (variable surgical teams assignment, high level).
2) With respect to ORs (“*ORs*” factor), we distinguish the case where ORs are used to perform, within the same session, either long-stay (LoS>1 day) surgeries or short-stay (LoS=1 day) surgeries (dedicated sessions, low level) and the case where both types of surgeries can be performed within the same session (mixed sessions, high level).

3) With respect to surgical units *("Units"* factor), we distinguish the case where units characterised by the same care setting (in terms of nursing staff, equipment, etc.) are used to host cases of specific specialties only (dedicated units, low level) and the case where these units are pooled to host patients of all specialties (pooled units, high level).

In the remainder of the paper, when a resource is managed flexibly, we will say that the hospital implements a *flexible practice* with respect to such a resource. Combining factors and factor levels, we obtained eight (=2^3) configurations. For each of them we ran the optimisation model in correspondence with 30 randomly generated instances. These instances were obtained starting from real data coming from the Meyer University Children's Hospital (hereinafter Meyer Hospital) a leading Italian hospital. The remainder of the paper is organised as follows: in Section 2, we provide a brief review of the literature. In Section 3, we describe the optimisation models. In Section 4, we illustrate the experimental campaign. In Section 5, we present the empirical results and in Section 6 we discuss them. Subsequently, in Section 7, we draw the conclusions and outline the direction of our future research efforts.

2 Literature review

The master surgical scheduling problem has been the object of several studies (see the reviews of Cardoen et al. (2010), Guerriero and Guido (2011), May et al. (2011)). In Table 1, we review the most important mathematical models supporting the production of MSS that appeared in the literature. Each column of the table (except the last one) represents a resource, while each row represents a model. In each cell, we specify if and how the resource is modelled. When a resource is not explicitly considered in the model, the cell contains "NEC." In the last column of the table, instead, we report the methodology adopted in the relevant study.

Table 1 – MSS literature review: resources modelled and operational assumptions

Paper	Surgical teams	OR	Surgical units' beds	Other resources	Type of analysis and solution technique
Our study	Number of sessions per surgical specialty bounded on a daily and on a weekly basis Session assignment performed: - Once and then considered as fixed (low flex) - Every time the MSS is produced, but only limited changes are allowed with respect to a predefined assignment (high flex)	Fully interchangeable ORs Two sessions per day/OR Sessions: - Dedicated (low flex) - Mixed (high flex)	Three types of surgical units (one day surgery unit and two regular units). - All units are dedicated to specific patient types, no mismatch allowed (low flex) - Regular units are pooled (high flex)	NEC	Single criterion exact optimisation, scenario analysis Mixed integer programming
Blake et al. (2002)	Number of sessions per surgical specialty bounded on a daily and on a weekly basis Session assignment performed once and then kept constant in the following period	Partially interchangeable ORs One session per day/OR Mixed sessions	NEC	Medical equipment	Single criterion heuristic optimisation, scenario analysis Mixed integer programming, constructive heuristic
Vissers et al. (2005)	NEC	Fully interchangeable ORs One session per day/OR Mixed sessions	Two types of surgical units (ICU and regular unit) Dedicated units, no mismatch allowed	ICU nursing staff	Single criterion exact optimisation Mixed integer programming
Santibáñez et al. (2007)	Number of sessions per surgical specialty bounded on a daily and on a monthly basis Session assignment performed once and then considered as fixed	Partially interchangeable ORs One or two sessions per day/OR Mixed sessions	Two types of surgical units (SCU and regular unit) Dedicated units, no mismatch allowed	NEC	Single criterion exact optimisation Mixed integer programming
van Oostrum et al. (2008)	NEC	Fully interchangeable ORs One session per day/OR Mixed sessions	Two types of surgical units (ICU and regular unit) Dedicated units, no mismatch allowed	NEC	Multi-criteria exact optimisation, multi-criteria heuristic optimisation Mixed integer programming, column generation, decomposition approach
Beliën et al. (2009)	Number of sessions per surgical specialty bounded on a weekly basis Session assignment performed once and then considered as fixed	Fully interchangeable ORs One or more sessions per day/OR Mixed sessions	Several types of surgical units Dedicated units, no mismatch allowed	NEC	Multi-criteria heuristic optimisation, goal programming Simulated annealing
Tànfani and Testi (2010)	Number of sessions per surgical specialty bounded on a weekly basis Session assignment performed every time MSS is produced	Fully interchangeable ORs One or two sessions per day/OR Mixed sessions	Two types of surgical units (ICU and regular unit) Dedicated units, no mismatch allowed	NEC	Single criterion heuristic optimisation Constructive heuristic
Banditori et al. (2013)	Number of sessions per surgical specialty bounded on a daily basis Session assignments performed every time MSS is produced	Partially interchangeable ORs Two sessions per day/OR Mixed sessions	Three types of surgical units Dedicated units, no mismatch allowed	NEC	Multi-criteria hierarchical exact optimisation, scenario analysis Mixed integer programming, discrete event simulation
Agnetis et al. (2012)	Number of sessions per surgical specialty bounded on a daily and on a weekly basis Session assignment performed: - Once and then considered as fixed (low flex) - Every time the MSS is produced, but only limited changes are allowed with respect to a predefined assignment (medium flex) - Every time the MSS is produced without limiting the changes allowed with respect to a predefined assignment (high flex)	Partially interchangeable ORs One or two sessions per day/OR Dedicated sessions	NEC	NEC	Single criterion exact optimisation, scenario analysis Mixed integer programming

In order to emphasise similarities and differences between our study and the related literature, we have added a row representing our model. When a study proposes both flexible and rigid approaches to manage a resource, we report all the alternatives in the table. Table 1 reveals that most of the authors considered three main critical resources in their models: *surgical teams, ORs* and *surgical units' beds*. Therefore, the remainder of this review will focus on these resources.

Surgical teams, i.e. the teams of surgeons belonging to the same specialty that actually carry out surgeries are considered explicitly in all but two models (i.e. the model of Vissers et al. (2005) and van Oostrum et al. (2008)). In the remaining works, the availability of surgical teams is modelled by limiting the number of sessions that each surgical specialty can perform on a weekly basis and/or daily basis. Based on these constraints, almost all models assign sessions to specialties, thereby identifying *when* a surgical team will potentially operate in the planning horizon (*session assignment*). In addition, some models (Santibáñez et al., 2007, van Oostrum et al., 2008, Banditori et al., 2013) also determine the *type* and/or the *number* of surgeries that surgical teams will execute in each session (*surgery types assignment*). In (Agnetis et al., 2012), instead, one of the proposed models assumes that the *session assignment* has already been done and, consequently, supports the surgery types assignment only. Most studies suggest that the *session assignment* should be carried out once and should not be changed frequently (Guerriero and Guido, 2011). The underlying assumption of these studies is that it is not technically feasible to change the *session assignment* on a monthly (or more frequent) basis because it would make it very complex for surgeons to coordinate their activities inside and outside the OT (van Oostrum et al., 2010). Nonetheless, Agnetis et al. (2012) demonstrate that small and frequent changes in the *session assignment* can yield substantial benefits and that these benefits are higher than those associated with large yet less frequent changes. Therefore, the authors argue that a limited amount of flexibility in managing surgical teams can produce benefits that are higher than the organisational cost of implementing this solution. For that reason, we decided to include this latter case in our study and compare it with the case where the session assignment is considered as already having been performed.

Contrary to surgical teams, ORs are considered as critical in all the reviewed models. However, different authors model these resources in different ways. A first distinction is between *interchangeable* and *partially interchangeable ORs.* The former can host every type of surgery; the latter, instead, can host only a limited subset of surgeries and/or specialties. A second distinction pertains to how OR time is divided into sessions. Some authors consider one session per OR per day (van Oostrum et al., 2008), some consider two (Santibáñez et al., 2007) or more (Beliën et al., 2009) sessions per OR per day, while others allow both daily sessions and shorter sessions (Agnetis et al., 2012). A third distinction concerns the types of surgery that can be performed in the same OR session. For example, Agnetis et al. (2012) distinguishes two macro-types of surgeries: *general surgeries* and *day surgeries*. The former includes all the procedures leading to a LoS of at least two days (one night), and the latter includes those procedures associated with a LoS of just one day. Based on this distinction, Agnetis et al. (2012)'s model allows only *dedicated sessions*, meaning that within the same session it is not possible to execute both *day-surgeries and general surgeries*. Instead, other models (e.g. Banditori et al. (2013)) allow *mixed sessions* where these types of surgeries can coexist. While the interchangeability of an OR depends on the structural characteristics (e.g. the presence of certain equipment) of the OR itself, hospital managers have more degrees of freedom in deciding how to subdivide the OR time. Nonetheless, this decision is influenced by the actual number of surgical teams available for each specialty (Banditori et al., 2013). For example, all-day-long sessions cannot be planned for those specialties relying on less than two surgical teams per day (except in extraordinary cases, one team cannot operate for the entire day). The decision to organise dedicated or mixed sessions, instead, is generally free. The literature suggests that surgeons usually prefer dedicated sessions; surgeons, in fact, can reduce surgery time because of the repetitive nature of their work (Hans et al., 2008). On the other hand, a mixed session makes the scheduling process less constrained and as such, it potentially allows scheduling a greater number of surgeries. In this study, we explore both options.

Finally, *surgical units*, i.e. the facilities where patients are cared for following surgical procedures, are considered in six out of eight models. These units are usually classified based on the intensity of care required by the hospitalised patients: e.g. ICUs, day-surgery units, regular units. Moreover, these units are characterised by a given capacity that is expressed in terms of the number of beds. Certain hospitals (e.g. Meyer Hospital) allocate patients to the regular units based on the specialty. Such a practice makes it easier and faster for surgeons to control and visit their hospitalised patients. Different models assume different numbers of units and unit types. All the reviewed models except Banditori et al. (2013) constrain each type of patient to be hospitalised into a specific unit. In general, the literature (Vincent et al., 1998) suggests that it is risky to accommodate patients requiring thorough care in units characterised by reduced nursing staff or that are physically located far away from the intensive care unit. Thus, units should be pooled only if they are characterised by similar care settings, which is the flexible practice explored in this study. Banditori et al. (2013)'s model, instead, violates this recommendation and allows bed mismatches whenever they allow increasing the OT throughput.

According to Table 1, it can be argued that flexible practices are considered in several studies. However, no study proposes an analysis that investigates how different flexible practices can interact. With this study, we aim to address this literature gap. In sum, our study (i) proposes a model that considers critical resources that are included in the vast majority of the other studies; (ii) investigates three flexible practices that are reasonable and justified in light of the extant literature but that previous studies have considered only separately or by combining a very limited number of different scenarios (two at maximum); (iii) assesses, in statistical terms, the main and the interaction effects of the mentioned practices and to the best of our knowledge is the only study to do so. These facts ensure that the results presented in the next sections can be of value for a wide audience of practitioners and scholars and also that this study adds a significant contribution to the literature.

3 Model description

In this section, we present the mathematical models we have developed. Specifically, first we present a version of the model that does not implement any flexible practice (hereafter referred to as the "rigid model"). Then we show how such a model can be modified to incorporate flexibility with respect to the management of surgical teams, ORs and surgical units.

All the models presented in this work address a twofold problem: (i) determining the number of cases to assign to each OR session of the planning horizon; (ii) determining the surgery group these cases must belong to. The models consider three critical resources: (i) ORs, whose available time is organised in sessions; (ii) surgical units, which accommodate patients after the surgery; (iii) surgical teams, dedicated to one specialty each, whose availability is defined in terms of number of OR sessions. Cases belonging to the same surgery group require the same specialty, the same amount of OR time and will occupy a surgical unit for the same amount of time.

Let us define the following sets and parameters that are common to the rigid model and to its extensions as follows:

W the set of weeks in the planning horizon, indexed by w

D the set of days in the planning horizon, indexed by d

T the set of sessions, indexed by t

O the set of ORs, indexed by o

S the set of surgical specialties, indexed by s

K the set of surgery groups, indexed by k

M a suitably big constant

H_{odt} the available time of OR o, on day d and session t

F_{bd} the number of beds in the surgical unit b available on day d

L_{sw} the availability of surgical teams with specialty s for week w, expressed in number of OR sessions

s_k the specialty of surgery group k

r_k the typology of surgery group k (short-stay surgery – SS vs. long-stay surgery – LS)

c_k the average surgery duration of surgery group k

b_k, a_k the average number of days of hospitalisation, before and after surgery, required by surgery group k

Y_k the minimum number of procedures of surgery group k to be scheduled.

3.1 Rigid model

In this model, we assume that the *session assignment* has already been done. Consequently, we rely on an *allocation grid G* as an input. Specifically, for each specialty s, day d and session t, G_{sdt} is equal to 1 if specialty s is allocated on day d, session t, and 0 otherwise.

Grid G must respect the following feasibility constraints:

$$\sum_{s\in S} G_{sdt} \le |O| \qquad \forall d \in D, \forall t \in T \qquad (3.1.1)$$

$$\sum_{d=7w-6}^{7w} \sum_{t\in T} G_{sdt} = L_{sw} \qquad \forall s \in S, \forall w \in W \qquad (3.1.2)$$

Constraints (3.1.1) assure that on each day d and session t, the number of specialties assigned to an OR does not exceed the number of available ORs ($|O|$). Constraints (3.1.2) instead control that the number of sessions assigned weekly to a given s is exactly the value resulting from the upstream case-mix planning problem. Then, in the rigid model, an OR o has to be assigned to each triple (s,d,t) for which $G_{sdt} = 1$. For a matter of convenience, we denote this with the following:

$$\bar{G} = \{(s,d,t) \text{ s.t. } s \in S,\, d \in D, t \in T \text{ and } G_{sdt} = 1 \} \quad .$$

The rigid scheduling model takes the following two main decisions:

1. Assign an OR to each triple (s,d,t) in $\bar{G}$
2. Determine, for each surgery group k, the number of procedures to schedule in correspondence with each triple (s,d,t) in $\bar{G}$ where s is the specialty associated with k.

Then let us define the following main and auxiliary variables:

q_{go} binary, 1 if triple $g = (s,d,t)$ in $\bar{G}$ is assigned to OR o, 0 otherwise

y_{kdto} the number of procedures in surgery group k assigned to OR o on day d in time slot t

z_{bd} the number of beds belonging to surgical unit b occupied on day d

u_{odt} binary, 1 if OR o on day d and session t is dedicated to short-stay surgeries, 0 otherwise.

Using these variables and parameters, we can state the rigid model as follows:

$$\max \sum_{\substack{k\in K, o\in O,\\ d\in D, t\in T}} y_{kodt} \qquad (3.1.3)$$

$$\sum_{s\in S: g=(s,d,t)\in \bar{G}} q_{go} \leq 1 \qquad \forall o \in O, \forall d \in D, \forall t \in T \qquad (3.1.4)$$

$$\sum_{o \in O} q_{go} = 1 \qquad \forall g \in \overline{G} \qquad (3.1.5)$$

$$\sum_{k \in K: s_k = s} y_{kdto} \leq M q_{go} \qquad \forall g = (s,d,t) \in \overline{G}, \forall o \in O \qquad (3.1.6)$$

$$\sum_{k \in K} c_k y_{kdto} \leq H_{odt} \qquad \forall o \in O, \forall d \in D, \forall t \in T \qquad (3.1.7)$$

$$\sum_{\substack{k \in K, o \in O, \\ t \in T}} \sum_{d' = \max(1, d - a_k)}^{\min(|D|, d + b_k)} y_{kd'to} = z_{bd} \qquad \forall b \in B, \forall d \in D \qquad (3.1.8)$$

$$z_{bd} \leq F_{bd} \qquad \forall b \in B, \forall d \in D \qquad (3.1.9)$$

$$\sum_{k \in K: r_k = \text{"SS"}} y_{kdto} \leq M u_{odt} \qquad \forall o \in O, \forall d \in D, \forall t \in T \qquad (3.1.10)$$

$$\sum_{k \in K: r_k = \text{"LS"}} y_{kdto} \leq M (1 - u_{odt}) \qquad \forall o \in O, \forall d \in D, \forall t \in T \qquad (3.1.11)$$

$$\sum_{\substack{o \in O, d \in D, \\ t \in T}} y_{kdto} \geq Y_k \qquad \forall k \in K \qquad (3.1.12)$$

$$q_{go} \in \{0,1\} \qquad \forall g \in \overline{G}, \forall o \in O \qquad (3.1.13)$$

$$y_{kdto} \in \mathrm{N} \qquad \forall k \in K, \forall d \in D, \forall t \in T, \forall o \in O \qquad (3.1.14)$$

$$z_{bd} \in \mathrm{N} \qquad \forall b \in B, \forall d \in D \qquad (3.1.15)$$

$$u_{odt} \in \{0,1\} \qquad \forall o \in O, \forall d \in D, \forall t \in T \qquad (3.1.16)$$

The objective function (3.1.3) maximises the number of procedures scheduled in the planning horizon. Constraints (3.1.4) guarantee that each OR-session can host a specialty at most. Constraints (3.1.5) assure that all the triples (s,d,t) in $\overline{G}$ are assigned to some OR o. Constraints (3.1.6) bind together assignment variables q and variables y: specifically, they state that if the triple (s,d,t) in $\overline{G}$ has not been assigned to OR o, then no procedure belonging to a group characterised by specialty s can be performed in OR o, on day d and session t. In contrast, when the triple g=(s,d,t) in $\overline{G}$ is assigned to OR o (q_{go}=1), then the corresponding constraint is redundant since it imposes that the number of procedures of that specialty scheduled in that OR session does not exceed the suitably big constant M. Specifically, M is set equal to the maximum number of shortest procedures a session can host. Constraints (3.1.7) guarantee that the total duration of the procedures scheduled in an OR session does not exceed the available time of that OR session. Constraints (3.1.8) and (3.1.9) are used to properly manage beds; specifically, for each day d and surgical unit b, they respectively compute the number z_{bd} of beds occupied and limit such a number to the bed availability F_{bd}. To correctly determine the bed occupancy on a given day d, we have to consider all the patients whose stay in the surgical units, before (b_k) and after surgery (a_k), overlaps day d. More specifically, in a given day d, we have to consider the beds occupied by patients who have undergone a surgery before day d and who are still in the hospital on day d as well as all the patients who will undergo surgeries after day d and that have been pre-hospitalised, in addition to the patients that undergo a surgery exactly on day d. Constraints (3.1.10) and (3.1.11) refer to the management of dedicated sessions, and they assure that in a given OR session, long-stay and short-stay surgeries are mutually exclusive. In fact, the binary variable u_{odt} is equal to 1 if OR o on day d and session t is dedicated to short-stay surgeries. In this case, constraints (3.1.11) assure that in that OR session no long-stay surgery is performed. In contrast, when u_{odt} is equal to 0, the corresponding OR session can host only long-stay surgeries. Constraints (3.1.12) relate to target efficiency and they impose that for each surgery group k the number of procedures performed is not smaller than the target value Y_k. Indeed, the MSS must guarantee to schedule a minimum number of surgeries for each surgery group. Such a requisite is set to avoid solutions planning an excessive number of surgeries belonging to easy-to-schedule surgery groups (i.e. groups characterised by short ST and LoS). This method ensures a reasonable waiting time for patients of each group and allows distributing complex-to-schedule surgeries over time.

Finally, constraints (3.1.13), (3.1.14), (3.1.15), and (3.1.16) define the domain of the variables.

In the following section, we describe how to extend/modify the rigid model in order to take into account the flexible practices discussed in the previous section.

3.2 Flexibility with respect to surgical teams

In this scenario, differently from the rigid model, the allocation grid is not an input for the scheduling model. Instead, the grid is the output of the model that decides the specialty to assign to each OR, day and session in the planning horizon. However, only limited variations with respect to the original grid are allowed in order to guarantee that the new grid is still implementable. To this aim, the following variables are defined:

x_{sdto} binary, 1 if specialty s is assigned to OR o on day d and session t, 0 otherwise

$x^{+}{}_{sdt}$ binary, 1 if a swap from 0 to 1 occurs with respect to G_{sdt}, 0 otherwise.

All the constraints that in the rigid model are implicitly satisfied by the pre-defined grid G have now to be explicitly guaranteed through the following set of constraints:

$$\sum_{s\in S} x_{sodt} \le 1 \qquad \forall o \in O, \forall d \in D, \forall t \in T \qquad (3.2.1)$$

$$\sum_{o\in O} x_{sodt} \le 1 \qquad \forall s \in S, \forall d \in D, \forall t \in T \qquad (3.2.2)$$

$$\sum_{o\in O, t\in T} \sum_{d=7w-6}^{7w} x_{sodt} = L_{sw} \qquad \forall s \in S, \forall w \in W \qquad (3.2.3)$$

$$\sum_{k\in K: s_k = s} y_{kodt} \le M x_{sodt} \qquad \forall s \in S, \forall o \in O, \forall d \in D, \forall t \in T \qquad (3.2.4)$$

$$x_{sodt} \in \{0,1\} \qquad \forall s \in S, \forall o \in O, \forall d \in D, \forall t \in T \qquad (3.2.5)$$

Specifically, constraints (3.2.1) assure that each OR on each day and in each session of the planning horizon is assigned to at most one specialty. Constraints (3.2.2) guarantee that each specialty is assigned to at most one OR in each day and session. Constraints (3.2.3) impose that the number of sessions assigned weekly to a given specialty s is exactly the value resulting from the upstream case-mix planning problem. Constraints (3.2.4) bind together assignment (x) and scheduling (y) variables; specifically, they assure that no procedure with specialty s is scheduled in OR o, on day d and session t unless specialty s has been assigned to that OR, on that day and session. Conversely, these constraints become redundant when $x_{sodt}=1$ since they impose that the number of procedures scheduled does not exceed a suitably defined big M. Finally, constraints (3.2.5) define the domain of the assignment variables.

Furthermore, we introduced the following constraints to control the variations with respect to the grid G:

$$\sum_{o \in O} x_{sodt} \leq G_{sdt} + x^{+}_{sdt} \qquad \forall s \in S, \forall d \in D, \forall t \in T \qquad (3.2.6)$$

$$\sum_{d \in D, t \in T} x^{+}_{sdt} \leq \overline{A} \qquad \forall s \in S \qquad (3.2.7)$$

Specifically, constraints (3.2.6) allow that any variation of element G_{sdt} can occur. In particular, if $G_{sdt} = 0$, i.e. if specialty s is not allocated to day d, session t, the new grid defined through variables x may allow that specialty s is assigned to some OR in that day and session. When this variation occurs, variable x^{+}_{sdt} takes value 1 and it accounts for a zero to one swap with respect to G. In addition, x_{sodt} specifies the OR o to which specialty s is assigned in day d, session t.

One to zero swaps, instead, do not need to be explicitly controlled. In fact, since we hypothesize that the number of sessions dedicated to each specialty in the planning period is constant, when a one to zero swap occurs also a zero to one swap takes place and this latter swap is controlled by x^{+}_{sdt} as well. The number of zero to one swaps affecting the specialty s cannot exceed the maximum number $\overline{A}$ of allowed variations (see constraints (3.2.7)).

3.3 Flexibility with respect to ORs

To implement this type of flexibility, it is sufficient to remove constraints (3.1.10) and (3.1.11), thus enlarging the feasibility region and allowing both short-stay and long-stay surgeries to be scheduled in the same session.

3.4 Flexibility with respect to surgical units

Each procedure is associated with a surgical unit. If surgical units are managed flexibly, then they are pooled. With this method, patients can be hospitalised in units that differ from the one originally assigned to them. To model this practice, we introduce the following variables:

$v_{bb'd}$ the number of beds of surgical unit *b'* used in place of beds of surgical unit *b* on day *d*.

Constraints (3.1.9) in the rigid model are then updated with constraints (3.4.1). These constraints allow that on a given day for a given surgical unit the number of beds occupied in that unit exceeds capacity. Moreover, we add constraints (3.4.2), which limit the overall number of beds occupied to the overall bed availability.

$$z_{bd} \leq F_{bd} + \sum_{b' \in B: b' \neq b} v_{bb'd} \qquad \forall b \in B, \forall d \in D \qquad (3.4.1)$$

$$\sum_{b \in B} z_{bd} + \sum_{b,b' \in B: b \neq b'} v_{bb'd} \leq \sum_{b \in B} F_{db} \qquad \forall d \in D \qquad (3.4.2)$$

4 Methodology

To assess the effects of the implementation of flexible practices on MSS efficiency, we use a 2^3 factorial design comprising the following:

- Three factors: *Teams, ORs, Units***.** Each factor corresponds to one of the critical resources incorporated in the model.

- Two possible levels for each factor: *high* when the resource is managed in a flexible way and *low* otherwise.
- One response variable, i.e. the number of surgeries scheduled.

Factors and factor levels are reported in Table 2, and the experimental design is illustrated in Figure 1.

Table 2 - Factorial design

Symbol	Factor Name	Low level	High level
A	Teams	Fixed surgical teams assignment. The allocation grid is fixed.	Variable surgical teams assignment. At maximum, one swap per specialty is allowed with respect to a predefined allocation grid.
B	ORs	Dedicated sessions. OR can host either short-stay or long-stay surgeries.	Mixed sessions. OR can host both short-stay and long-stay surgeries.
C	Units	Dedicated units. Unit 1 and Unit 2 are dedicated to different types of long-stay patients.	Pooled units. Unit 1 and Unit 2 are used interchangeably.

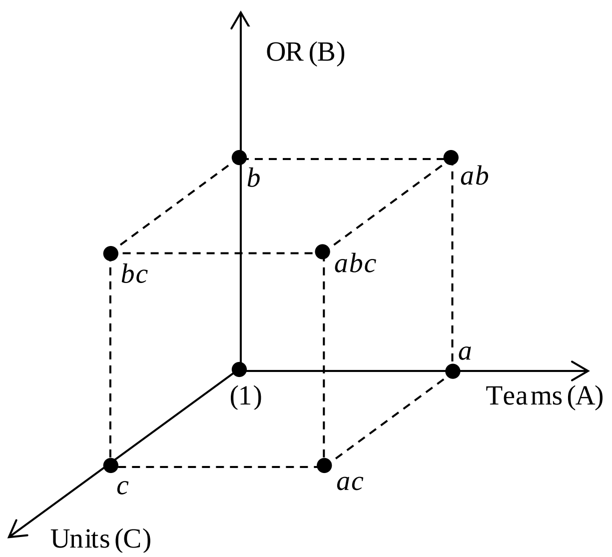

Treatment	Factors		
	Teams (A)	**ORs (B)**	**Units (C)**
(1)	Low	Low	Low
a	High	Low	Low
b	Low	High	Low
c	Low	Low	High
ab	High	High	Low
ac	High	Low	High
bc	Low	High	High
abc	High	High	High

Figure 1 Experimental design

Each factor is associated with an uppercase letter (A, B, C). Each vertex of the cube represents a treatment. Treatments are labelled according to the Montgomery and Runger's (2003, p.524) notation. According to this notation, a treatment combination is represented by a series of lowercase letters. If a letter is present, the corresponding factor is run at the high level in that treatment combination; if it is absent, the factor is run at its low level. The treatment combination with all the factors at the low level is represented by (1).

To implement the different treatments, the optimisation model is extended as described in Section 3. For each treatment, we have analysed the result of the optimisation model in correspondence of 30 randomly generated instances. These instances differ in each other's in terms of allocation grid G.

We coded the optimisation models in AMPL and solved them through the IBM ILOG Cplex Solver (version 12.4) running on a personal computer equipped with an Intel Core i7 processor and 8 GB of RAM. For each optimisation run, we bound the computational time to 1 hour. The results of our experimental campaign are presented in the next section.

5 Empirical results

In this section, we present the data we used to run the optimisation model(s) and the results of the experiments.

5.1 Input data

As we pointed out in the introduction, our study was inspired by Meyer Hospital. Such a hospital is characterised by the following features:

(i) *12 surgical specialties*. Each surgical specialty is associated with surgical teams that can cover a certain number of sessions per week.

(ii) *38 surgery groups*. Surgery groups have been created following Banditori et al.'s (2013) methodology. For each surgery group (k), we calculated the mean value of LoS and ST and used these values to set the parameters a_k and c_k of the optimisation models, respectively.

(iii) A *planning horizon* of two weeks.

(iv) Defined lower bounds (Y_k) for the number of surgeries to schedule within the planning horizon for each surgery group k. These lower bounds are fixed by the hospital's top management on a yearly basis.

(v) 3 *surgical units*: a day surgery unit and two regular units (Unit 1 and 2). The day surgery unit contains 14 beds, and Unit 1 and Unit 2 contain 19 and 14 beds, respectively.

(vi) The day surgery unit can host only short-stay patients, i.e. patients whose expected LoS is one day (no night), regardless of the speciality. In contrast, Units 1 and 2 can accommodate long-stay patients only for certain specialties. Long-stay patients can be hospitalised either in Unit 1 or in Unit 2, and mismatches are not allowed.

(vii) *4 interchangeable ORs* dedicated to elective patients. Each OR is open 10 hours a day, 5 days per week. The OR time is subdivided into two sessions, morning and afternoon. Additional OR sessions and beds are allocated to non-elective patients (emergencies and urgencies).

(viii) *OR* sessions are "dedicated," i.e. in a session where long-stay surgeries are performed, no short-stay surgery can be scheduled and vice versa. In addition, afternoon sessions can host only long-stay surgeries, while morning sessions can host both long-stay and short-stay surgeries.

(ix) An allocation grid G that fixes the specialty to assign to each OR session.

(x) No deviation from the allocation grid G is tolerated.

Features (i, ii, iii, iv, v, and vii) do not change across treatments and instances. Features (vi, viii and x) change depending on the treatment, as described in Table 2. Feature (ix) changes according to the instance, which is randomly generated. The Meyer Hospital case corresponds to the treatment (1) in Table 2.

5.2 Optimisation output

Table 3 shows the results of the optimisation models. It displays the mean values, calculated across instances, of the scheduled surgeries and of the optimality gap. In addition, for each treatment, the table shows the number of instances for which the optimisation model found the optimal solution and the minimum and the maximum optimality gaps across the 30 instances.

Table 3 Optimisation Output

Treatment	Mean of scheduled surgeries	Mean of optimality gap	Optimal solutions found	Min of optimality gap	Max of optimality gap
(1)	272.1	0.0%	30/30	0.0%	0.0%
a	280.2	2.7%	0/30	1.8%	3.6%
b	278.2	0.0%	30/30	0.0%	0.0%
c	274.7	0.0%	30/30	0.0%	0.0%
ab	286.7	0.5%	7/30	0.0%	1.4%
ac	281.9	2.6%	0/30	1.4%	3.6%
bc	279.2	0.0%	30/30	0.0%	0.0%
abc	287.0	0.5%	5/30	0.0%	1.4%

As can be seen for some treatments and instances, it was not possible to find an optimal solution within the fixed time limit. Nonetheless, the mean optimality gap associated with each treatment never exceeds 3.6%. The table shows that when moving from treatment (1) to treatment *abc*, the number of surgeries scheduled increases by 14.9. Therefore, implementing all the mentioned flexible practices yields, on average, a monthly increase of around 30 surgeries. To interpret these results, we performed an analysis of variance (ANOVA) and assessed the statistical significance and the magnitude of all main and interaction effects. Moreover, we carried out several Tukey's post-hoc tests to compare treatments with each other and rank them in terms of scheduled surgeries while controlling the familywise error rate (Field, 2005, p.310) to a 0.05 level. These statistical analyses are presented in the next section.

5.3 Statistical analysis

Table 4 displays the complete ANOVA table including the magnitude of the estimated effects and their level of significance. The ANOVA analysis included an accurate check of the assumptions of normality of error terms and homogeneity of variance. More specifically, we carried out a Ryan-Joiner test and failed to reject ($p=0.099$) the null hypothesis of normally distributed errors. Similarly, we performed the Levene's test and failed to reject the null hypothesis of the variances being equal ($p=0.087$).

Table 4 Analysis of variance and effects summary table for scheduled surgeries

	DF	Sum of squares	Mean squares	F	Effect	p
Teams	1	3744.6	3744.6	2378.4	7.9	0.000 (*)
Ors	1	1837.1	1837.1	1166.8	5.5	0.000 (*)
Units	1	117.6	117.6	74.7	1.4	0.000 (*)
Teams*Ors	1	4.8	4.8	3.1	0.3	0.080
Teams*Units	1	8.8	8.8	5.6	-0.4	0.020 (*)
Ors*Units	1	33.8	33.8	21.4	-0.8	0.000 (*)
Teams*Ors*Units	1	0.1	0.1	0.0	0.0	0.837
Error	232.0	365.3	2.1			
Total	239.0	6112.0				

(*) significant at the $\alpha = 0.05$ level

Table 4 shows that, assuming an $\alpha = 0.05$ significance level, all the main effects are significant ($p<0.05$). Similarly there is a significant, yet negative, interaction effect between Teams factor (A) and the Units factor (C) ($p=0.020$) and between the ORs factor (B) and Units factor (C) ($p=0.000$). Other 2-way and 3-way interaction effects, instead, are not statistically significant ($p>0.05$). The main and interaction effects are plotted in Figure 2 and Figure 3, respectively.

Looking at the main effects (Figure 2), it can be noted that, on average, an increase in the level of each factor leads to an increase in the number of surgeries scheduled. For example, when the Teams factor (A) is run at a high level (i.e. treatments a, ab, ac, abc), the model schedules, on average, 283.9 surgeries. Instead, when the Teams factor (A) is run at its low level (i.e. treatments b, c, bc, (1)), the model schedules, on average, 276 surgeries (in fact, main effect (A) =283.9-276= 7.9)

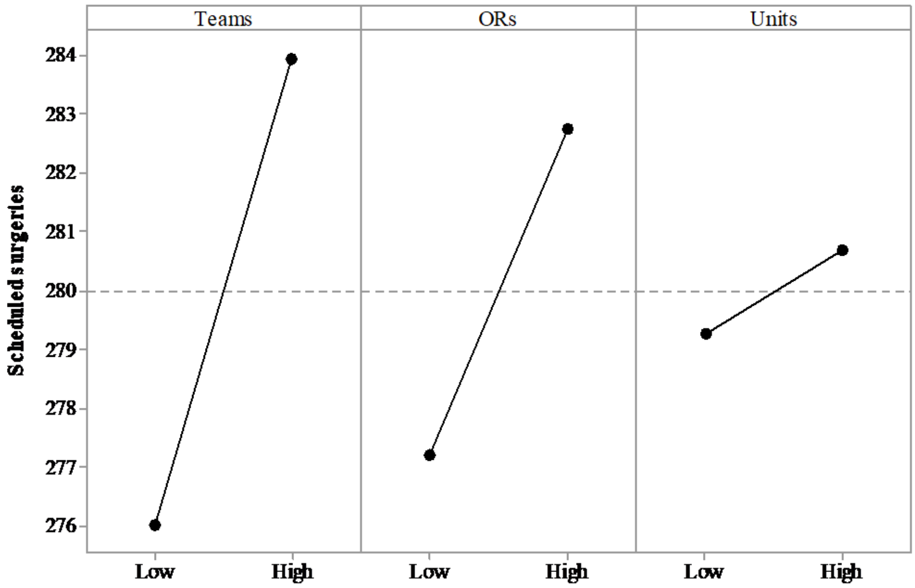

Figure 2 Main effects for scheduled surgeries, mean values

However, as in our case, when one or more significant interaction effects are present, the interpretation of the main effects can be incomplete or misleading. In fact, when an interaction factor is significant, the impact of one factor depends on the level of another factor. For example, in our case, the significant interaction between B and C factors implies that the effect on scheduling surgeries (dependent variable) of B depends on the level of C and vice versa. In particular, since the interaction effect is negative, increasing B when C is at a high level leads, on average, to a variation in terms of the number of scheduled surgeries that is significantly smaller than the variation obtained by increasing B when C is at a low level. If this latter variation were negative, i.e. if increasing B when C is low would lead to a decrease in the surgeries scheduled, then the interpretation of the main effects would be completely misleading. In this latter case, in fact, increasing B from low to high in the presence of a high level of C would determine a decrease of the surgeries scheduled, which is the opposite of what one would expect looking at the main effects of B and C. To prevent misleading interpretations of the main effects, however, it is sufficient to observe the interaction graphs in Figure 3.

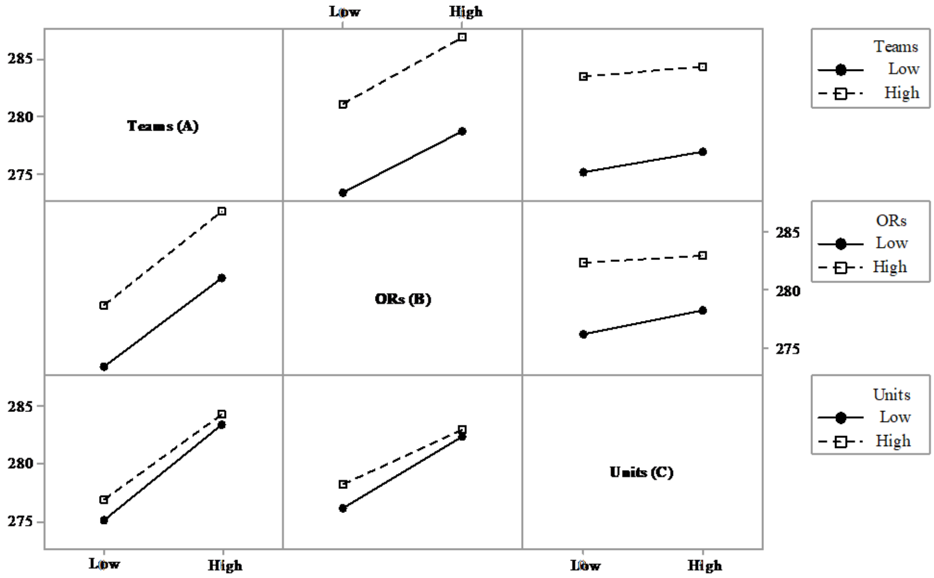

Figure 3 Interaction plot for scheduled surgeries, data means

In Figure 3, the lines in each cell do not cross. Therefore, for each factor, the number of surgeries scheduled is, on average, higher when the factor is high than when the factor is low, regardless of the level of the other factors. Therefore, moving a factor from low to high leads to a benefit in terms of scheduled surgeries, regardless of the levels of the other factors.

To compare treatments with each other and rank them, we used the Tukey's post hoc procedure. This procedure allows us to compare all different combinations of the treatment groups and to control the familywise error rate without sacrificing the statistical power. For each pairwise comparison, we assigned the same rank (1, 2, etc.) to those treatment groups for which the post-hoc test did not allow for the identification of a significant ($p>0.05$) difference between the number of scheduled surgeries. The results of these tests are shown in Table 5 and will be discussed in the next section along with their practical implications.

Table 5 Pairwise comparisons, grouping information using Tukey's method and 95.0% confidence level

Comparisons	Treatment group code	Treatment groups	N	Mean	Rank
1	(1.1)	a, ab, ac, abc	120	283.9	1
	(1.2)	b, c, bc, (1)	120	276	2
2	(2.1)	ab, abc, bc, b	120	282.8	1
	(2.2)	ac, a, c, (1)	120	277.2	2
3	(3.1)	abc, ac, bc, c	120	280.7	1
	(3.2)	ab, a , b, (1)	120	279.3	2
4	(4.1)	ab, abc	60	286.9	1
	(4.2)	ac, a	60	281	2
	(4.3)	bc, b	60	278.7	3
	(4.4)	c, (1)	60	273.4	4
5	(5.1)	abc, ac	60	284.5	1
	(5.2)	ab, a	60	283.4	2
	(5.3)	bc, c	60	276.9	3
	(5.4)	b, (1)	60	275.1	4
6	(6.1)	abc, bc	60	283.1	1
	(6.2)	ab, b	60	282.4	2
	(6.3)	ac, c	60	278.3	3
	(6.4)	a, (1)	60	276.1	4
7	(7.1)	abc	30	287	1
	(7.2)	ab	30	286.7	1
	(7.3)	ac	30	281.9	2
	(7.4)	a	30	280.2	3
	(7.5)	bc	30	279.2	4
	(7.6)	b	30	278.2	5
	(7.7)	c	30	274.7	6
	(7.8)	(1)	30	272.1	7

6 Discussion

Looking at pairwise comparisons 1 to 3 in Table 5, emerges that, between the three investigated flexible practices, the one that, *on average*, leads to the largest increase in the number of surgeries scheduled is the variable surgical teams assignment (groups 1.1 vs. 1.2). This practice is followed by the introduction of mixed session (groups 2.1 vs. 2.2) and by the surgical units pooling (groups 3.1 vs. 3.2).

The pairwise comparisons 4 in Table 5, in their turn, reveal that, *on average*, the introduction of a variable surgical teams assignment significantly increases the number of surgeries scheduled both when ORs can host mixed sessions (groups 4.1 vs. 4.3) and when ORs are organised into dedicated sessions (groups 4.2 vs. 4.4). Similarly, they reveal that introducing mixed sessions increases the number of surgeries scheduled in the presence of both a variable surgical teams assignment (groups 4.1 vs. 4.2) and fixed surgical teams assignment (groups 4.3 vs. 4.4). However, the increase that can be obtained introducing a variable surgical teams assignment is larger than the one that can be obtained introducing mixed sessions (groups 4.2 vs. 4.3).

Similarly, the pairwise comparisons 5 in Table 5, show that, *on average*, introducing a variable surgical teams assignment increases the number of surgeries scheduled both when surgical units are pooled (groups 5.1 vs. 5.3) and when they are not (groups 5.2 vs. 5.4). Similarly, pooling surgical units increases the number of surgeries scheduled both in presence of a variable (groups 5.1 vs. 5.2) and a fixed surgical teams assignment (groups 5.3 vs. 5.4). The increase that can be obtained introducing a variable surgical teams assignment is larger than the one that can be obtained by pooling surgical units (groups 5.2 vs. 5.3).

In the same way, the pairwise comparisons 6 in Table 5, show that, *on average*, introducing mixed sessions increases the number of surgeries scheduled, both when surgical units are pooled (groups 6.1 vs. 6.3) and when they are not pooled (groups 6.2 vs. 6.4). Similarly, pooling surgical units increases the number of surgeries scheduled, both in the presence of dedicated sessions (groups 6.1 vs. 6.2) and in presence of mixed sessions (groups 6.3 vs. 6.4). The increase that can be obtained introducing mixed sessions is larger than those that can be obtained by pooling surgical units (groups 6.2 vs. 6.3)

Finally, from the pairwise comparisons 7 in Table 5 emerges that for hospitals where no flexible practices are implemented, the best results in terms of surgeries scheduled can be achieved by introducing flexibility with respect to surgical teams and ORs (groups 7.2 vs. 7.8). In fact, once these two flexible practices are implemented, pooling surgical units does not yield any significant additional advantage (groups 7.1 vs. 7.2). On the other hand, if mixed session cannot be implemented, then pooling surgical units significantly increases the number of surgeries scheduled both when surgical teams are managed flexibly (groups 7.3 vs. 7.4) and when they are not (groups 7.7 vs. 7.8). Equivalently, if surgical teams cannot be managed flexibly, then pooling surgical units significantly increases the number of surgeries scheduled both when ORs are managed flexibly (groups 7.5 vs. 7.6) and when they are not (groups 7.6 vs. 7.8). Finally, for hospitals where no flexible practice is implemented, introducing flexibility with respect to surgical teams leads to an increase in the scheduled surgeries that is statistically larger than the one that can be obtained by introducing flexibility with respect to ORs and surgical units (groups 7.4 vs. 7.5).

As a final remark, it is worth mentioning that the *statistical significance* of an effect does not necessarily imply that such an effect is also *practically relevant*. The post-hoc test, in fact, reveals if the difference between the mean number of surgeries associated with two treatment groups is statistically different from zero. A difference greater than zero (say, one surgery in two weeks) is not necessarily practically relevant and does not necessarily imply that the associated flexible practice deserves to be implemented. Indeed, the benefits that are possible to obtain with a flexible practice should always be traded off with the costs of implementation. For example, the sessions assignment is often the output of a lengthy and complex negotiation process between stakeholders (surgeons, management, nursing staff) with different priorities and needs. Thus to avoid conflicts, a hospital could also decide to renounce the potential benefits of implementing a variable surgical teams assignment.

7 Conclusion and future research

In this study, we presented a novel mixed integer programming model to address the master surgical scheduling problem. In addition, we evaluated the impact in terms of scheduled surgeries of the implementation of different combinations of three flexible practices: (i) variable surgical teams assignment, (ii) mixed sessions and (iii) pooled surgical units.

Our analysis revealed that to maximise the number of scheduled surgeries it is sufficient to introduce a variable surgical teams assignment *and* mixed sessions. In fact, if both these practices are implemented, pooling surgical units carries no additional advantages. However, if only one of these flexible practices (or none) is implemented, then pooling surgical units produces significant benefits. Moreover, the analysis showed that, if a hospital cannot implement a variable surgical teams assignment, then it can still improve its efficiency by introducing mixed sessions and, similarly, if it cannot implement mixed sessions, it can improve its efficiency by introducing a variable surgical teams assignment.

This study considers hospital features that are included in the vast majority of the contributions available in the master surgical scheduling literature and explores flexible practices that are reasonable according to such a literature. Moreover, it is the first study to propose a systematic analysis of the effect of the implementation of these practices. As such, both the presented model and the implications of the analysis can be of interest for a wide audience of practitioners and scholars.

Of course, this study is not without limitations. First, we investigated only a limited number of hospital settings. For example, we neglected to factor in certain hospital resources (e.g. ICU, electro-medical devices) that are not considered critical at Meyer Hospital but that may be highly critical in other hospitals. Second, we have not investigated how the MIP model would perform in terms of computational time if the problem dimension increases, e.g. if the planning horizon is extended to one month or if the number of ORs and beds increases. Finally, we only considered elective patients. Nonetheless, it might be interesting to investigate how the implementation of flexible practices could help improve hospital performance in presence of emergencies, urgencies and no-shows (Stuart and Kozan, 2012). The extension of the computational campaign to other hospital settings, the analysis of the optimisation model scalability, the design of ad-hoc methodologies to cope with large scale instances and the incorporation of non-elective patients in the analysis will certainly be the object of our future research efforts.

Acknowledgements

The authors are grateful to Meyer University Children's Hospital, and in particular to General Director Dr. Tommaso Langiano, and to the Medical Director Dr. Monica Frassineti for supporting the research project that has inspired this study.

The authors would also express their sincere gratitude to the anonymous referees for their insightful critical comments and suggestions.

References

AGNETIS, A., COPPI, A., CORSINI, M., DELLINO, G., MELONI, C. & PRANZO, M. 2012. Long term evaluation of operating theater planning policies. *Operations Research for Health Care,* 1**,** 95-104.

BALASUBRAMANIAN, H., MURIEL, A. & WANG, L. 2012. The impact of provider flexibility and capacity allocation on the performance of primary care practices. *Flexible Services and Manufacturing Journal,* 24**,** 422-447.

BANDITORI, C., CAPPANERA, P. & VISINTIN, F. 2013. A combined optimization–simulation approach to the master surgical scheduling problem. *IMA Journal of Management Mathematics,* 24**,** 155-187.

BELIËN, J. & DEMEULEMEESTER, E. 2007. Building cyclic master surgery schedules with leveled resulting bed occupancy. *European Journal of Operational Research,* 176**,** 1185-1204.

BELIËN, J., DEMEULEMEESTER, E. & CARDOEN, B. 2009. A decision support system for cyclic master surgery scheduling with multiple objectives. *Journal of Scheduling,* 12**,** 147-161.

BLAKE, J. T., DEXTER, F. & DONALD, J. 2002. Operating room managers' use of integer programming for assigning block time to surgical groups: A case study. *Anesthesia & Analgesia,* 94**,** 143-148.

BUZACOTT, J. A. & MANDELBAUM, M. 2008. Flexibility in manufacturing and services: achievements, insights and challenges. *Flexible Services and Manufacturing Journal,* 20**,** 13-58.

CARDOEN, B., DEMEULEMEESTER, E. & BELIËN, J. 2010. Operating room planning and scheduling: A literature review. *European Journal of Operational Research,* 201**,** 921-932.

CHOU, M. C., TEO, C.-P. & ZHENG, H. 2008. Process flexibility: design, evaluation, and applications. *Flexible Services and Manufacturing Journal,* 20**,** 59-94.

DENTON, B., VIAPIANO, J. & VOGL, A. 2007. Optimization of surgery sequencing and scheduling decisions under uncertainty. *Health Care Management Science,* 10**,** 13-24.

FIELD, A. 2005. *Discovering statistics using SPSS,* London, SAGE Publications.

GLOUBERMAN, S. & MINTZBERG, H. 2001. Managing the care of health and the cure of disease-Part I: Differentiation. *Health care management review,* 26**,** 56-69.

GUERRIERO, F. & GUIDO, R. 2011. Operational research in the management of the operating theatre: a survey. *Health Care Management Science,* 14**,** 89-114.

GUPTA, D. & SHANTHIKUMAR, G. 2008. Editorial. *Flexible Services and Manufacturing Journal,* 20**,** 1-10.

HANS, E., WULLINK, G., VAN HOUDENHOVEN, M. & KAZEMIER, G. 2008. Robust surgery loading. *European Journal of Operational Research,* 185**,** 1038-1050.

LITVAK, E. & LONG, M. C. 2000. Cost and quality under managed care: Irreconcilable differences. *Am J Manag Care,* 6**,** 305-12.

MAY, J. H., SPANGLER, W. E., STRUM, D. P. & VARGAS, L. G. 2011. The surgical scheduling problem: Current research and future opportunities. *Production and Operations Management,* 20**,** 392-405.

MONTGOMERY, D. C. & RUNGER, G. C. 2003. *Applied statistics and probability for engineers,* John Wiley and Sons.
SANTIBÁÑEZ, P., BEGEN, M. & ATKINS, D. 2007. Surgical block scheduling in a system of hospitals: an application to resource and wait list management in a British Columbia health authority. *Health Care Management Science,* 10**,** 269-282.
STUART, K. & KOZAN, E. 2012. Reactive scheduling model for the operating theatre. *Flexible Services and Manufacturing Journal,* 24**,** 400-421.
TÀNFANI, E. & TESTI, A. 2010. A pre-assignment heuristic algorithm for the Master Surgical Schedule Problem (MSSP). *Annals of Operations Research,* 178**,** 105-119.
VAN OOSTRUM, J. M., BREDENHOFF, E. & HANS, E. W. 2010. Suitability and managerial implications of a master surgical scheduling approach. *Annals of Operations Research,* 178**,** 91-104.
VAN OOSTRUM, J. M., VAN HOUDENHOVEN, M., HURINK, J., HANS, E., WULLINK, G. & KAZEMIER, G. 2008. A master surgical scheduling approach for cyclic scheduling in operating room departments. *OR Spectrum,* 30**,** 355-374.
VINCENT, C., TAYLOR-ADAMS, S. & STANHOPE, N. 1998. Framework for analysing risk and safety in clinical medicine. *BMJ: British Medical Journal,* 316**,** 1154.
VISSERS, J., ADAN, I. & BEKKERS, J. 2005. Patient mix optimization in tactical cardiothoracic surgery planning: a case study. *IMA Journal of Management Mathematics,* 16**,** 281-304.