

The single-day surgery scheduling problem: sequential decision-making and threshold-based heuristics

William L. Herring · Jeffrey W. Herrmann

© Springer-Verlag 2011

Abstract Scheduling elective surgeries is a dynamic, sequential decision-making process that must balance the costs of deferring waiting cases and blocking higher-priority cases. Although other surgery scheduling problems have received extensive treatment in the literature, this paper presents the first single-day scheduling problem formulation to capture this aspect of the scheduling process while also incorporating surgical block schedules, block release policies, and waiting lists. Theoretical results for the special case in which all cases have the same duration motivate a range of threshold-based heuristics for the general problem with multiple case durations. Our computational results demonstrate the effectiveness of the proposed heuristics and show how block release dates affect the quality of the scheduling decisions. Based on these results, we propose a new approach to surgery scheduling. In particular, to make more equitable waiting list decisions, operating room (OR) managers should gradually release unused OR time over the course of several days leading up to the day of surgery.

Keywords Healthcare · Surgery scheduling · Dynamic programming · Heuristics

1 Introduction

The health care system has been the focus of increasing attention from operations researchers and management scientists in recent years. The growth in this research

W. L. Herring (✉)
Applied Mathematics and Statistics and Scientific Computation Program,
University of Maryland, College Park, USA
e-mail: wherring@math.umd.edu

J. W. Herrmann
Department of Mechanical Engineering and Institute for Systems Research,
University of Maryland, College Park, USA

19 area is motivated by increasing costs and the rising demand for health care services
20 and is facilitated by the improved quality and availability of data generated by the
21 system. The problem of scheduling surgical patients into operating rooms benefits
22 from a robust literature. This interest in surgery scheduling is intensified by the key
23 role that operating room (OR) scheduling plays in determining hospital occupancy
24 levels and because the OR is the most resource-intensive and profitable unit of a hos-
25 pital (Macario et al. 1995; McManus et al. 2003). At its core, the surgery scheduling
26 problem, in all its variations, involves the allocation of a fixed amount of resources
27 (ORs, hospital staff) under uncertain demand (see Cardoen et al. 2010, for a thorough
28 review). Like other scheduling problems, surgery scheduling approaches hope to make
29 more efficient use of existing resources. This improved efficiency in the context of
30 the health care system comes with the added societal benefit of lowering costs and
31 increasing access to health care.

32 1.1 Literature review

33 The majority of hospitals schedule their OR suites using cyclic master, or block, sur-
34 gery schedules, in which available OR space is assigned to specific surgical specialties,
35 or service lines. For hospitals using block schedules, the literature on elective surgery
36 scheduling describes the problem as consisting of three stages: (1) determining the
37 amount of OR time to allocate to various surgical specialties, (2) creating a block
38 schedule implementing the desired allocations, and (3) scheduling individual patients
39 into available time (Blake and Donald 2002; Santibañez et al. 2007; Testi et al. 2007).
40 First stage decisions typically reflect the long-term strategic goals of hospital manage-
41 ment, such as meeting the demand for surgical specialties' services, achieving desired
42 levels of patient throughput, or maximizing revenue (Blake and Carter 2002; Gupta
43 2007; Santibañez et al. 2007; Testi et al. 2007). The second and third stages repre-
44 sent medium- and short-term operational and tactical decisions, respectively, but differ
45 markedly in their objectives. In recent years, block scheduling models have transitioned
46 from simply implementing desired allocation levels (Blake and Donald 2002) to level-
47 ing hospital bed occupancy and minimizing overcapacity (Beliën and Demeulemeester
48 2007; van Oostrum et al. 2008). Research on individual patient scheduling, including
49 patient selection, room placement, and sequencing, typically aims to minimize patient
50 delays or maximize OR utilization (Denton et al. 2007; Guinet and Chaabane 2003).
51 The large number of stakeholders involved in surgery scheduling has also motivated
52 a number of multi-objective models for both block scheduling and individual patient
53 scheduling (Blake and Carter 2002; Beliën et al. 2009; Cardoen et al. 2009).

54 A fundamental, but understudied, element of the day-to-day job of surgery sched-
55 uling is the transition from the generality of the block schedule (in which OR time
56 is allocated to specialties) to the specificity of a completed schedule for a particular
57 day (in which OR time is assigned to specific cases). As will be discussed below,
58 individual patients are scheduled into blocks of OR time in a dynamic and sequen-
59 tial decision-making process. Block schedules are often incorporated as constraints in
60 individual patient scheduling models (Hans et al. 2008; Pham and Klinkert 2008; Testi
61 et al. 2007), but these models schedule patients all at once rather than sequentially over
62 time. A key component of the schedule evolution is the scheduling of add-on cases

off of surgical waiting lists, but existing papers on this topic only consider decisions made at the tail end of the scheduling process and fail to capture earlier aspects of the process (Dexter et al. 1999; Dexter and Traub 2002; Gerchak et al. 1996). This paper proposes a new model for the single-day surgery scheduling problem that seeks to address these shortcomings.

The model proposed below has much in common with existing research on capacity allocation (Schütz and Kolisch 2010a,b) and, in particular, airline revenue management (Brumelle and Walczak 2003; Lee and Hersh 1993; Subramanian et al. 1999). In these problems, a finite resource (OR time, seats on a flight) with a fixed expiration date (the day of surgery, the departure time for the flight) must be allocated to competing demand classes. The demand from each of the classes arrives over time, and decision-makers must decide how much of the resource to allocate to lower priority classes and how much to reserve for higher priority classes. In these existing models, arriving demand must be accepted or rejected at the moment of its arrival and rejected demand is lost. In the surgery scheduling problem, however, lower priority demand is placed on a waiting list, or request queue, and can be accepted at a number of different decision points leading up to the day of surgery. While the proposed surgery scheduling model displays a threshold behavior similar to the revenue management models (particularly to Lee and Hersh 1993), the introduction of the request queue concept increases the complexity of the state space and complicates the analysis leading to these results.

The problem studied in the remainder of this paper is motivated by and models the dynamic nature of surgery scheduling, in which the OR schedule for a single day evolves over time. The immediate goal is to capture the most important relationships and develop insights into the process, particularly focusing on the often contentious nature of block release and RQ decisions. While the initial results presented in this paper are largely theoretical in nature, we believe that they lay an important foundation from which to design more equitable and transparent block release and RQ policies. Our modeling approach reflects our discussions with collaborators at the University of Maryland Medical Center (UMMC), whose response to our results and their policy implications has been positive. Their feedback informs a number of directions for future collaborative research that will be required before implementation at UMMC or any other particular hospital.

1.2 Dynamics of surgery scheduling

This description and the resulting problem statement are based on a case study of the surgery scheduling system at the UMMC (Herring 2011). Similar systems are used by other hospitals that use block scheduling (Dexter et al. 2003; Dexter and Macario 2004; Ozkarahan 2000).

Block schedules are concerned primarily with elective surgery, and the demand for elective surgery arrives over the course of multiple days before the day of surgery, rather than all at once. For a given day on which an OR suite is open for surgery, the available OR time is allocated to specific surgical specialties according to a block schedule. Blocks of OR time are initially controlled exclusively by these “primary” specialties, and the primary specialties are free to choose and sequence their cases (“primary” cases) within their allocated blocks as they see fit. Specialties or surgeons




Fig. 1 The evolution of Friday's OR schedule and RQ for a single OR over the week before surgery

that do not have allocated OR time on that day, but still wish to perform a surgery, must submit their cases ("secondary" cases) to the surgical request queue (RQ) for that day. If a primary specialty's allocated OR time has been filled, it may also submit its excess cases to the RQ. In the period leading up to the day of surgery, cases accumulate on the RQ and OR managers look to schedule these RQ cases into OR time that has not been filled by the primary specialties. In practice, there is a set day before surgery, referred to as the *block release date*, after which OR managers can begin assigning RQ cases to unfilled blocks. Before the block release date, RQ cases may not be assigned to open times. After the block release date, the *request queue policy* determines how and when RQ cases will be assigned to open times. Together, the block release dates and request queue policies control the transition from the block schedule to the individual patient scheduling stage of surgery scheduling.

RQ decisions are made every day, once a day, before any new cases have arrived. In practice, OR managers must manage the schedule for many future surgery days, and it is impractical for them to spend time reconsidering the RQ cases too frequently. The demand for surgeries is often communicated to the hospital by the specialties once a day, typically after the surgeons have seen their patients in a clinic or made their rounds. It is also important to note that this problem does not consider the process of scheduling urgent (emergent) cases that are generated on the day of surgery. In practice, entire blocks may be reserved for urgent cases, or the schedule may include some slack for these cases. These and other similar approaches affect the scheduled capacity of the OR, which is taken as given in this problem.

The following example illustrates the importance of considering the sequential nature of the OR manager's decisions and how they interact with the demand for surgery. Figure 1 shows the evolution of the schedule for a single OR and the RQ for a given Friday over the course of a week before surgery, as demand for Friday's OR time arrives and the OR manager makes daily RQ decisions. The OR has been allocated to a primary specialty, and the demand for surgery is divided into primary cases and secondary cases. Note that not all cases arrive on the same day. If there is still open time on Friday's schedule, primary cases are added directly to the schedule; otherwise they are placed on the RQ. Secondary cases are always placed on the RQ. Because Wednesday is the block release date, RQ cases may not be added to Friday's schedule before Wednesday. Starting on Wednesday, the OR manager uses RQ policies to decide whether to schedule RQ cases into any time that is still available. In this

141 example, the decision on Wednesday morning to add a RQ case to the schedule means
142 that there is no time available for the primary case that arrives on Wednesday, which
143 is sent to the RQ (as seen on Thursday). Had the OR manager decided not to schedule
144 the RQ case on Wednesday, then the primary case would have filled the OR, and the
145 OR manager would have been prevented from scheduling the secondary case. In this
146 dynamic, stochastic setting, there is no guarantee that the OR manager's decision will
147 be optimal (when viewed ex post), but this paper aims to provide insights that can help
148 the OR manager make better decisions on average.

149 The dynamic, sequential, and stochastic aspect of creating an OR schedule empha-
150 sizes the importance of considering the joint impact of block schedules, block release
151 dates, and request queue policies. However, no existing work in the surgery schedul-
152 ing literature addresses all three of these components simultaneously. As mentioned
153 above, some recent papers use block schedules as constraints in models that schedule
154 batches of individual patients (Hans et al. 2008; Pham and Klinkert 2008; Testi et al.
155 2007). A study of different block release dates concludes that the timing of the block
156 release has little impact on OR efficiency (Dexter et al. 2003; Dexter and Macario
157 2004). However, these papers, which add single RQ cases to existing schedules at
158 different points in time, fail to consider the potential effect of RQ decisions on the
159 evolution of the schedule after the RQ case has been added. The other papers that
160 consider scheduling add-on cases limit their analysis to the day before and the day of
161 surgery, thus neglecting the role of the block release date (Dexter et al. 1999; Gerchak
162 et al. 1996).

163 1.3 Problem statement

164 In our setting, both the demand for elective surgery and the RQ are day specific. Sur-
165 geons submit cases for a specific date in the future on which they would like to perform
166 surgery. Similarly, cases on a given day's RQ that are never scheduled are not auto-
167 matically rolled over to future days' RQs (although surgeons can resubmit their cases
168 for future days). As a result, we can separate the general problem of scheduling the
169 operating rooms into subproblems that consider only one day's schedule. Two types
170 of costs are associated with the process of scheduling an OR suite that uses block
171 scheduling: (1) utilization costs related to under- and over-utilization of the ORs and
172 (2) customer satisfaction costs incurred when cases are forced to wait on the RQ or
173 when primary cases are blocked from their allocated ORs due to RQ placements. The
174 utilization costs are common in the literature, but the satisfaction costs are a distinctive
175 feature of this formulation.

176 Consider again the example presented earlier. On Wednesday in Fig. 1, the OR
177 schedule for Friday has room for one additional case and the RQ contains a case that
178 is waiting to be scheduled. The OR manager has the following choice: either schedule
179 the RQ case into the available time or defer scheduling the case (and consider it again
180 on Thursday if the time is still open). Deferring the case leaves open the possibility
181 that a primary case may arrive to use the remaining time, but is undesirable because
182 the surgeon and the patient associated with the case must wait at least one more day
183 for a decision. This satisfaction cost is defined as a *deferral cost*. If the RQ case is
184 scheduled (thus filling the OR) and another primary case arrives, then the OR manager

185 has blocked the primary service line's access to its allocated room. This satisfaction
186 cost is defined as a *blocking cost*. The balance between deferral costs and potential
187 blocking costs informs the OR manager's decision to schedule or defer RQ cases.
188 While these deferral and blocking costs are primarily satisfaction costs, it is important
189 to note that their relative values may reflect a range of practical concerns, including
190 differences in patient acuity levels, specialized resource needs, and contributions to
191 revenue between surgical specialties.

192 On each day leading up to the day of surgery, the OR manager must choose the num-
193 ber of RQ cases to add to each OR's schedule for the day of surgery. The manager's
194 objective is to minimize the expected total cost of deferral and blocking penalties
195 incurred on the days before surgery and OR underutilization on the day of surgery.
196 Block release dates restrict the days on which RQ cases can be added to the schedule
197 and may differ from room to room. RQ policies dictate which decisions to make in
198 which scenarios.

199 The dynamic relationship between decisions and costs suggests a stochastic
200 dynamic programming (SDP) formulation. Because block release dates serve as con-
201 straints, the optimal policy combination would use no block release dates and use the
202 optimal decisions from the unconstrained SDP as the RQ policy. For this reason, the
203 formulation and analysis that follow consider an SDP unconstrained by block release
204 dates. Because block release dates are used in practice (to protect OR time allocated
205 to primary service lines and to reduce the OR manager's workload by limiting the
206 number of days that must be considered), the added cost of imposing block release
207 dates will be considered.

208 A special case of the single-day surgery scheduling problem for a single OR in
209 which all of the cases have the same duration is formulated as an SDP and analyzed
210 in [Herring \(2010\)](#). A more general formulation for a single OR with multiple case
211 durations is presented in Sect. 2. Section 3 summarizes the results of the special case
212 and analyzes the optimal solutions for the general case. Insights from these results
213 motivate a range of threshold-based heuristics. Section 4 presents the results of com-
214 putational tests that compare these heuristics and study the impact of block release
215 dates. Finally, Sect. 5 offers some concluding remarks and indicates directions for
216 future research.

217 Overall, the results in this paper demonstrate the effectiveness of the proposed
218 threshold-based heuristics and show how block release dates affect the quality of the
219 scheduling decisions. These results provide valuable insights into how OR managers
220 can use threshold-based decision rules to make equitable waiting list decisions. More-
221 over, it is the first step towards modeling, understanding, and optimizing more realistic
222 problems involving more case types, multiple operating rooms, and other complexities
223 that are necessary in order to provide meaningful decision support to an OR manager.

224 2 Stochastic dynamic programming formulation

225 This section presents a stochastic dynamic programming formulation of the single-
226 day surgery scheduling problem (SDSSP) that considers the general case with multiple
227 case durations. The problem focuses on the development of the schedule for a single
228 OR on a single, fixed day of surgery. The OR has been allocated to a primary specialty

229 according to a block schedule. The two sources of demand (primary and secondary)
 230 are assumed to be independent, as are the number of cases that arrive each day. RQ
 231 decisions for the OR must be made once a day on each day leading up to the day of
 232 surgery, and are made before any additional cases for the day of surgery arrive.

233 For modeling purposes, cases will be separated into different types based on their
 234 duration (the amount of OR time allocated to the case). While this approach requires
 235 discretizing the case durations, which may not capture the entire range of possible
 236 durations, there is a precedent in the literature for treating case durations in this man-
 237 ner (see, for instance, [Guinet and Chaabane 2003](#)). In addition to helping maintain
 238 the computational tractability of the SDP, this approach allows the analysis to stay
 239 focused on the principal issue of how OR managers balance potential blocking and
 240 deferral costs.

241 2.1 Formulation

242 The input data is defined as follows:

243 N = number of days before the day of surgery on which surgical demand is generated
 244 C = capacity of the OR in hours
 245 K = number of case types
 246 d_k = duration (in hours) of cases of type k , for $k = 1, \dots, K$

247 The stochastic demand for surgery is given by two classes of random variables associ-
 248 ated with the primary cases and secondary cases. From this point forward, all references
 249 to days refer to the number of days before the day of surgery, with day 0 representing
 250 the day of surgery.

251 T_j^k = number of primary cases of type k that arrive on day j , $j = 1, \dots, N$
 252 R_j^k = number of secondary cases of type k that arrive on day j , $j = 1, \dots, N$

253 The blocking and deferral costs for each day before surgery and the utilization costs
 254 for the day of surgery are similarly defined.

255 h_j^k = deferral cost on day j for cases of type k , $j = 1, \dots, N$
 256 r_j^k = blocking cost on day j for cases of type k , $j = 1, \dots, N$
 257 h_0^k = penalty for unscheduled cases of type k left on RQ on the day of surgery
 258 r_0 = penalty for unused space in the OR on the day of surgery

259 Three classes of state variables are required to represent the state of the OR schedule
 260 at the start of day j (before the day j 's RQ decisions have been made and before any
 261 new cases have arrived). The number of blocking eligible hours in the OR reflects the
 262 presence of RQ cases that were added to the schedule on previous days but have not
 263 yet incurred a blocking penalty. This auxiliary state accounts for the fact that a primary
 264 case might be blocked by a RQ case scheduled on a previous day.

265 C_j = available hours remaining on the OR schedule on day j , $j = 0, \dots, N$
 266 B_j = number of blocking eligible hours on the OR schedule on day j , $j = 0, \dots, N$
 267 W_j^k = number of cases of type k on the RQ on day j , $j = 0, \dots, N$

268 Finally, the decision variables can be defined.

269 x_j^k = number of RQ cases of type k to add to the OR schedule on day j , $j = 0, \dots, N$

270 With the exception of the day of surgery, the costs incurred each day are separated
 271 into deferral costs and blocking costs. Deferral penalties are assessed only up to the
 272 number of RQ placements that are feasible.

273 ND_j^k = number of cases of type k deferred on day j , $j = 1, \dots, N$
 274 $= \min \left(\left\lfloor \frac{C_j}{d_k} \right\rfloor, W_j^k \right) - x_j^k$

275 When the RQ decisions are made on day j , that day's arrivals have not yet occurred.
 276 Therefore, the number of blocking penalties incurred on day j is a random variable that
 277 depends upon the value of the decision variables and on the arrival of primary cases.
 278 The consideration of different case durations raises an important question related to
 279 computing blocking costs. If a longer primary case is blocked and sent to the RQ,
 280 should a blocking penalty be assessed for the entire case or only for the part of the
 281 case that overlaps with the blocking eligible hours on the OR schedule? Suppose, for
 282 example, that the OR schedule has one available hour ($C_j = 1$) and one blocking eligi-
 283 ble hour ($B_j = 1$) and that a primary case with a duration of 2 h arrives ($T_j^k = 1$ for the
 284 case type with $d_k = 2$). Because this case will not fit in the available time, it is blocked
 285 and placed on the RQ. The blocking penalty can be assessed either for the entire 2-h
 286 duration or for just the fraction of the case that overlaps with the blocking eligible hour.
 287 (Both approaches will be considered during computational testing). Computing these
 288 blocking quantities requires an algorithm (as opposed to a closed form expression). The
 289 pseudocode for this algorithm, referred to as *ProcessDay*, is given in the Appendix.

290 NB_j^k = number of cases of type k blocked on day j , $j = 1, \dots, N$

291 FB_j^k = fraction of cases of type k blocked on day j , $j = 1, \dots, N$

292 fit_j^k = number of primary cases of type k added to the OR schedule on day j , $j =$
 293 $1, \dots, N$

294 This is now sufficient information to define the value function, Bellman's optimality
 295 equation, and the day-to-day transition equations. Note that the daily blocking penalty
 296 is a convex combination of the number blocked and the fraction blocked, with the
 297 weighting parameter λ satisfying $0 \leq \lambda \leq 1$. This parameter allows us to address the
 298 question of full or partial blocking, and its impact will be explored during computa-
 299 tional testing. For convenience, some quantities are expressed in vector form in the
 300 optimality equation and subsequent discussion.

301 For $j = 1, \dots, N$:

302 $V_j(C_j, B_j, W_j)$ = minimum expected remaining cost from state (C_j, B_j, W_j) on day j
 303 $= \min_{x_j} \left\{ \sum_{k=1}^K \left(h_j^k N D_j^k + r_j^k E \left[\lambda N B_j^k + (1 - \lambda) F B_j^k \right] \right) \right.$
 304 $\left. + E \left[V_{j-1}(C_{j-1}, B_{j-1}, W_{j-1}) \right] \right\}$

$$\begin{aligned}
 & \text{s.t. } \sum_{k=1}^K d_k x_j^k \leq C_j \\
 & 0 \leq x_j^k \leq W_j^k \text{ for } k = 1, \dots, K \\
 & \text{where } C_{j-1} = C_j - \sum_{k=1}^K d_k (x_j^k + fit_j^k) \\
 & B_{j-1} = B_j + \sum_{k=1}^K d_k (x_j^k - FB_j^k) \\
 & W_{j-1}^k = W_j^k - x_j^k + R_j^k + (T_j^k - fit_j^k)
 \end{aligned}$$

The boundaries for the SDP are day N and day 0 . On day N the system is initialized to empty, and on day 0 the boundary costs come from the unused time on the OR schedule and the cases that remain on the RQ. As with deferral costs, the objective function on day 0 penalizes only those cases on the RQ that could have been feasibly added to the schedule. The boundary value function is defined as follows:

$$\begin{aligned}
 V_0(C_0, B_0, W_0) &= \min_{x_0} \left\{ \sum_{k=1}^K h_0^k \left(\min \left(\left\lfloor \frac{C_0}{d_k} \right\rfloor, W_0^k \right) - x_0^k \right) \right. \\
 & \quad \left. + r_0 \left(C_0 - \sum_{k=1}^K d_k x_0^k \right) \right\} \\
 & \text{s.t. } \sum_{k=1}^K d_k x_0^k \leq C_0 \\
 & 0 \leq x_0^k \leq W_0^k \text{ for } k = 1, \dots, K
 \end{aligned}$$

Some minor rearrangements of the day 0 value function show it to be an integer knapsack problem with some additional constants (with respect to \mathbf{x}_0) in the objective function.

$$\begin{aligned}
 V_0(C_0, B_0, W_0) &= r_0 C_0 + \sum_{k=1}^K h_0^k \min \left(\left\lfloor \frac{C_0}{d_k} \right\rfloor, W_0^k \right) \\
 & \quad - \min_{x_0} \left\{ \sum_{k=1}^K (h_0^k + d_k r_0) x_0^k \right\} \\
 &= r_0 C_0 + \sum_{k=1}^K h_0^k \min \left(\left\lfloor \frac{C_0}{d_k} \right\rfloor, W_0^k \right) \\
 & \quad + \max_{x_0} \left\{ \sum_{k=1}^K (h_0^k + d_k r_0) x_0^k \right\}
 \end{aligned}$$

The final details needed to complete the formulation involve some aspects of the *ProcessDay* algorithm for computing the blocking quantities. The algorithm represents

329 how the system processes a particular realization of the primary case arrival random
 330 variables(see the Appendix for the pseudocode). In order to process arrivals of dif-
 331 ferent types, it is necessary to assume some sort of case type prioritization on behalf
 332 of the primary service line. Because this prioritization is beyond the control of the
 333 OR manager, it is not included as a decision variable in the model. (Note that the
 334 prioritization of the RQ case types is within the OR manager’s control and is reflected
 335 as such in the daily RQ decision variables). We assume throughout that the primary
 336 service line prioritizes its cases by decreasing duration. That is, longer duration cases
 337 receive priority over the shorter duration cases. Other prioritization rules could easily
 338 be considered within the framework of the algorithm.

339 2.2 Implementation issues and solution structure

340 As the problem is currently formulated, there is no upper bound on the size of the
 341 request queue. In order to implement a solution algorithm for the SDP, however, the
 342 RQ must be truncated at some point. Early implementations of code for the SDP
 343 showed that a truncation rule that combined the states with large RQs by ignoring any
 344 values beyond the point of truncation created a “bounce-back” effect that impacted
 345 the optimal solutions. Apparent irregularities in the optimal solutions followed the
 346 location of the truncation, necessitating a more sophisticated technique for limiting
 347 the number of RQ states that would not allow the system to prematurely return from
 348 the point of truncation.

349 Instead of truncating, we combined the states with large RQs (those greater than
 350 a given boundary) into a state in which the RQ is described as “infinite.” In order
 351 to justify that this technique does not produce solution irregularities of its own, two
 352 properties must hold. First, once the RQ state variable for a case type exceeds what
 353 will fit into the remaining time on the schedule, this condition must continue to hold
 354 for all subsequent states. Second, the optimal decisions must be identical for all states
 355 beyond the selected boundary. We claim below that these properties hold if the RQ
 356 boundary for each case type is beyond the range of feasible decisions, which allows
 357 the RQ boundaries to be set as small as possible and greatly reduces the computational
 358 complexity of the SDP implementation.

359 The following definitions will be useful both in the statement and proof of the
 360 desired claims. The first definition applies to $j = 1, \dots, N$

$$\begin{aligned}
 361 \quad F_j(C_j, B_j, \mathbf{W}_j, \mathbf{x}_j) &= \sum_{k=1}^K \left(h_j^k N D_j^k + r_j E \left[\lambda N B_j^k + (1 - \lambda) F B_j^k \right] \right) \\
 362 \quad &+ E \left[V_{j-1}(C_{j-1}, B_{j-1}, \mathbf{W}_{j-1}) \right] \\
 363 \quad F_0(C_0, B_0, \mathbf{W}_0, \mathbf{x}_0) &= \sum_{k=1}^K h_0^k \left(\min \left(\left\lfloor \frac{C_0}{d_k} \right\rfloor, W_0^k \right) - x_0^k \right) \\
 364 \quad &+ r_0 \left(C_0 - \sum_{k=1}^K d_k x_0^k \right)
 \end{aligned}$$

365 The dependence on the decision variables, x_j^k , is not explicitly stated on the right hand
 366 side of the first definition but is implied for all quantities that depend on x_j^k . With these
 367 functions thus defined, an alternative statement of the value function for $j = 0, \dots, N$
 368 can be given, with the minimization still subject to the same constraints described
 369 above. This re-statement of the value function also provides a concise way to refer to
 370 the optimal decision for each feasible state.

$$371 \quad V_j(C_j, B_j, W_j) = \min_{x_j} \{F_j(C_j, B_j, W_j, x_j)\}$$

$$372 \quad x_j(C_j, B_j, W_j) = \arg \min_{x_j} \{F_j(C_j, B_j, W_j, x_j)\}$$

372 Finally, a quantity referred to as the “trimmed RQ state” is defined for each feasible
 373 state that reflects the maximum number of feasible RQ placements for each case type.

$$374 \quad \hat{W}_j^k = \min \left(\left\lfloor \frac{C_j}{d_k} \right\rfloor, W_j^k \right)$$

375 The two claims below represent the two properties described above. The first claim
 376 justifies the notion that once the number of RQ cases exceeds the remaining capacity,
 377 it will continue to do so. The second claim states that the value functions and optimal
 378 decision vectors are identical for all states where the number of a certain case type on
 379 the RQ exceeds what is feasible. Proofs of these claims are given in the Appendix.

380 **Claim 1** *If $W_j^k \geq \left\lfloor \frac{C_j}{d_k} \right\rfloor$, then $W_{j-1}^k \geq \left\lfloor \frac{C_{j-1}}{d_k} \right\rfloor$ for all possible future states.*

381 **Claim 2** *For all feasible states (C_j, B_j, W_j) for all days $j = 0, \dots, N$, the following*
 382 *equalities hold:*

- 383 (i) $F_j(C_j, B_j, W_j, x_j) = F_j(C_j, B_j, \hat{W}_j, x_j)$
- 384 (ii) $V_j(C_j, B_j, W_j) = V_j(C_j, B_j, \hat{W}_j)$
- 385 (iii) $x_j(C_j, B_j, W_j) = x_j(C_j, B_j, \hat{W}_j)$

386 The final two parts of Claim 2 are the justification for combining all states past the
 387 RQ boundary into one state. For case type k , $\left\lfloor \frac{C}{d_k} \right\rfloor$ is used as the boundary because it
 388 equals the maximum number of cases of type k that will fit into an empty OR schedule.
 389 Because the value function and optimal decisions are identical for all states beyond
 390 this boundary (all other things being equal), aggregating these states will not impact
 391 the resulting optimal policy.

392 3 Optimal solution behavior

393 3.1 The special case of unit durations

394 In the special case when all cases have a unit duration (SDSSP1), the optimal decisions
 395 follow a threshold policy (Herring 2010). That is, for each day before the day of surgery,
 396 a fixed number of hours are reserved for future primary cases, and the optimal decisions

Table 1 Optimal thresholds for SDSSP1 across a range of cost ratios and arrival scenarios

Blocking-to-deferral cost ratio	Arrival scenarios														
	Early					Middle					Late				
	Days before surgery														
	4	3	2	1	0	4	3	2	1	0	4	3	2	1	0
1:1	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0
3:1	3	1	1	1	0	2	2	2	1	0	1	2	2	2	0
5:1	4	2	1	1	0	3	3	2	2	0	3	3	3	3	0
7:1	5	3	2	1	0	4	4	3	2	0	4	4	4	4	0

397 take the system as close to this threshold as possible. Furthermore, the optimal thresh-
 398 olds can be computed efficiently (without solving the SDP) from the daily deferral and
 399 blocking costs and the primary case arrival distributions. The existence and optimality
 400 of these thresholds are formally stated in the following claim (proved in Herring 2010).

401 **Claim 3** Assume that all cases have unit duration and the daily costs satisfy $h_j \leq$
 402 r_j for all $j \geq 0$ and $r_{j+1} \geq r_j$ for all $j \geq 1$. Then for all $j \geq 0$ there exists a thresh-
 403 old Y_j such that the optimal decision for all feasible day j states (C_j, B_j, W_j) is given
 404 by $\min(W_j, \max(0, C_j - Y_j))$.

405 In addition to the threshold structure, two other observations stand out in this expres-
 406 sion for the optimal decision. First, the blocking eligible state (B_j) plays no role in
 407 the optimal decision. This state essentially keeps track of the schedule’s RQ history,
 408 thus implying that past RQ decisions should have no bearing on the current optimal
 409 decision. Second, the number of cases on the RQ (W_j) influences the optimal deci-
 410 sion only as an upper bound on the number of RQ cases that can be scheduled. This
 411 indicates that having a very large number of cases on the RQ should not force the OR
 412 manager to place more of them.

413 Table 1 shows the optimal thresholds for SDSSP1 across a range of different input
 414 scenarios for an OR schedule that has a capacity of four cases (for the complete input
 415 data, see Herring 2010). Recall that the thresholds reflect how much space to save
 416 on the OR schedule for future primary cases, so higher thresholds lead to fewer RQ
 417 cases being scheduled. The behavior of the optimal thresholds is fairly intuitive. When
 418 blocking and deferral costs hold equal weight, the result is a greedy RQ policy. How-
 419 ever, as blocking becomes more expensive (relative to deferring), fewer RQ cases are
 420 scheduled. The thresholds are also highly dependent on the timing of the primary
 421 case arrivals. More RQ cases are scheduled after the bulk of the expected demand has
 422 passed, but when primary cases are expected to arrive late, few if any RQ cases are
 423 accepted until the day of surgery.

424 3.2 The general case with multiple durations

425 The simple and intuitive nature of the optimal thresholds for SDSSP1 raises a fun-
 426 damental question about the extent to which this behavior is apparent in the general

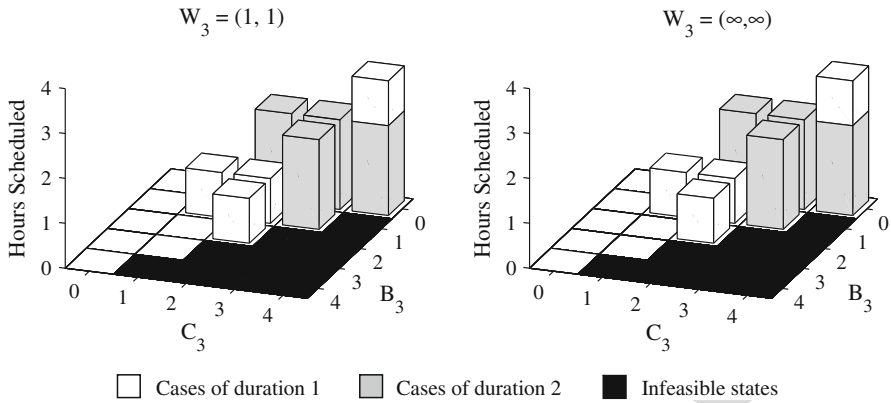


Fig. 2 Policy maps showing the number of hours of cases scheduled by the optimal decisions 3 days before surgery for two RQ states across all feasible hours remaining (C_3) and blocking eligible hours (B_3) states




Fig. 3 Hours remaining on the OR schedule after the optimal decisions 3 days before surgery for a single RQ state across all feasible hours remaining (C_3) and blocking eligible hours (B_3) states

427 case. The policy maps in Fig. 2 show the optimal decisions for a single instance of
 428 SDSSP across a representative selection of states. In this example, the scheduled OR
 429 capacity is 4 h, and there are two case types: 1-h cases and 2-h cases. The rest of the
 430 input data come from one of the test problems described in the next subsection. The
 431 policy maps show the total hours of each case type taken from the RQ and placed
 432 on the OR schedule by the optimal decision for each state. A single map shows all
 433 possible hours remaining (C_j) and blocking eligible hours (B_j) states for a single RQ
 434 state on a particular day. Note that the sum of the hours remaining and the blocking
 435 eligible hours cannot exceed the scheduled capacity of the room (as indicated by the
 436 infeasible states).

437 The fact that the policy map does not change when the number of cases on the RQ
 438 is increased confirms the notion (as indicated by Claims 1 and 2 in Sect. 2) that having
 439 more cases on the RQ does not pressure the optimal decision into taking more cases.
 440 The maps also immediately show that certain behaviors from SDSSP1 do not extend
 441 to the general case. In particular, a change in the number of blocking eligible hours can
 442 change the optimal decision. This change is evident between $B_3 = 0$ and $B_3 = 1$ for
 443 both maps. Figure 3 depicts the number of available hours on the OR schedule after

Table 2 Input data for a set of 162 test problems for the general case of SDSSP

Parameter	Notation	Range of values
Blocking parameter	λ	0, 0.5, 1
Deferral cost ratio	$h_j^2: h_j^1$	1:1, 2:1
Blocking cost ratio	$r_j^2: r_j^1$	1:1, 2:1, 3:1
Blocking-to-deferral cost ratio	$r_j^1: h_j^1$	1:1, 3:1, 5:1
Case type arrival rate ratio	$E[T_j^2]: E[T_j^1]$	1:2, 1:1, 2:1

444 the optimal decisions and clearly illustrates the fact that the decisions do not exactly
 445 follow a threshold pattern. However, it does appear that the decisions are guided by
 446 a “target” threshold. In this case, the OR manager generally seeks to leave 1 h avail-
 447 able in the OR, and deviates from this target only by small amounts and in a few
 448 situations.

449 3.3 Test problems

450 The results in Table 1 show that the optimal thresholds for SDSSP1 depend on the
 451 primary case arrival distribution and on the ratio of blocking costs to deferral costs.
 452 Two other potential factors merit exploration: the ratio of the deferral costs and block-
 453 ing costs between the case types and the weighting parameter (λ) that balances the
 454 number blocked (NB_j^k) and the fraction blocked (FB_j^k).

455 Table 2 shows the ranges of the input data used to create an initial set of 162 test
 456 problems. The remaining problem parameters are held constant for all the instances.
 457 The scheduled capacity of the OR is set to 4 h ($C = 4$), and there are two case types
 458 with durations of 1 and 2 h, respectively ($d_1 = 1, d_2 = 2$). The problems are solved
 459 over a period of 4 days before the day of surgery ($N = 4$). The blocking and deferral
 460 costs are held constant from day to day for each instance (following the ratios in
 461 Table 2), and the cost of unused OR time is set to ten ($r_0 = 10$). The case arrivals on
 462 each day follow Poisson distributions, and a total of 6 h of primary cases and 6 h of
 463 secondary cases are expected over the time horizon. For the secondary cases, an equal
 464 number of each case type is expected ($E[R_j^1] = E[R_j^2] = 0.5$), and the expected
 465 number of primary cases is weighted between the two case types as indicated by the
 466 ratios in Table 2. The expected number of primary arrivals is the same every day,
 467 implying that this initial set of test problems will not capture the influence of the
 468 arrival timing on the optimal solutions. This factor and the impact of block release
 469 dates will be explored in Sect. 4.

470 3.4 Solution behavior

471 A thorough analysis of policy maps similar to those in Figs. 2 and 3 for each of the
 472 test problems reveals insight into three fundamental aspects of the solution behavior:
 473 (1) the extent to which the solutions follow an approximate threshold pattern, (2) the
 474 nature of the target thresholds (relative to the input data), and (3) the factors that drive

475 the selection of one case type over another. While it is not feasible to present the entire
 476 set of policy maps here (see [Herring 2011](#) for a representative selection), the results
 477 of the analysis can be summarized as follows.

478 The most striking and important trend is that the optimal decisions exhibit an
 479 approximate threshold behavior like the example shown in Figs. 2 and 3. In many
 480 cases the thresholds are exact, but even when the optimal solutions deviate from the
 481 apparent threshold they do so in only a small number of states and rarely by more than
 482 a single hour in either direction.

483 Of the five factors listed in Table 2, two have no significant impact on the nature
 484 of the optimal decisions and apparent target thresholds. First, the impact of the block-
 485 ing weight parameter (λ) is minimal. Smaller λ values effectively lower the blocking
 486 costs, and, as a result, more hours of RQ cases tend to be scheduled when $\lambda = 0$ than
 487 when $\lambda = 1$. However, the change rarely amounts to more than an hour's difference
 488 for a handful of states and never affects the apparent threshold guiding the optimal
 489 decisions. Similarly, the relative prevalence of one primary case type over another
 490 ($E[T_j^2] : E[T_j^1]$) has very little influence on the optimal decisions. This suggests that
 491 the total hours of cases expected each day (which is held constant throughout) is more
 492 important in determining the target threshold than the prevalence of a particular case
 493 type. Finally, neither factor has any impact on the optimal solutions' tendency to prefer
 494 one case type over another (i.e. taking two 1-h cases over one 2-h case, or vice versa).

495 These observations leave the deferral and blocking cost structure as the primary
 496 determinant of the target threshold and the selection between the two case types. As in
 497 SDSSP1, the ratio of blocking cost to deferral cost (in the form of the $r_j^1 : h_j^1$ ratio) has
 498 the most influence on the approximate thresholds for the general SDSSP. Furthermore,
 499 the apparent thresholds for the three ratios listed in Table 2 (1:1, or *greedy*; 3:1, or
 500 *low*; and 5:1, or *high*) exhibit threshold trends much like those observed in Table 1 for
 501 SDSSP1. When the two costs are weighted equally, the optimal decisions universally
 502 try to fill the OR schedule as much as possible (thus the 'greedy' label). When the ratio
 503 is 'low,' the thresholds roughly aim to preserve enough open time for the total hours
 504 of cases expected each day. When the ratio is 'high,' the thresholds move to preserve
 505 enough open time for the total cumulative remaining hours of expected cases.

506 The other cost ratios included in this initial set of test problems (the blocking and
 507 deferral cost ratios between the case types) also impact the optimal solution behavior.
 508 Recall that the cost ratio discussed above is set for the first type ($r_j^1 : h_j^1$), but not for
 509 the second case type ($r_j^2 : h_j^2$). However, the between-case-type ratios ($h_j^2 : h_j^1$ and
 510 $r_j^2 : r_j^1$) affect the ratio of the blocking cost to the deferral cost for the second case type.
 511 When the between-case-type ratios combine to make the cost ratio for the second
 512 case type higher than the cost ratio for the first case type, the approximate thresholds
 513 increase by as much as an hour. This pattern is reversed when the between-case-type
 514 ratios combine to make the cost ratio lower for the second case type. While these
 515 trends in the target thresholds are perhaps not as clean as those observed for SDSSP1,
 516 they provide evidence that the optimal solutions to SDSSP2 exhibit threshold behavior
 517 similar to that observed for SDSSP1.

518 In some cases, the optimal solutions prefer one case type over another. In particu-
 519 lar, when the deferral costs are the same for both case types ($h_j^2 : h_j^1 = 1 : 1$), the

solutions uniformly prefer the shorter case type. When the deferral costs are proportional to the case durations ($h_j^2 : h_j^1 = 2 : 1$), the prioritization becomes more difficult to decipher. However, a weak preference seems to exist for the type with the higher ratio of blocking cost to duration. Together, these observations suggest a prioritization that first considers the ratio of deferral cost to duration and then considers the ratio of blocking cost to duration. Another critical property, particularly in light of the approximate thresholds discussed above, is an overall tendency to aim for the threshold rather than prioritize one case type over another.

These three sets of observations suggest a framework for threshold-based heuristics. Stochastic dynamic programs like the SDSSP are known to be constrained by the ‘curse of dimensionality.’ Without pursuing a full analysis of the complexity of SDSSP, it is clear that the problem complexity grows as the initial OR capacity and the number of case types increase. This complexity makes heuristic approaches necessary for solving realistically sized problems in which the scheduled OR capacity may be as large as 8 h and the case durations may exceed 4 h.

4 Threshold-based heuristics

The proposed heuristics for the single-day surgery scheduling problem have two components: (1) a threshold for each day and (2) a case type prioritization rule. The thresholds identify how much open time to preserve for future primary cases and, for each state, define a target number of hours to take off the RQ. The case type prioritization rule then indicates how to reach this target.

4.1 Threshold determination

The trends in the approximate threshold behavior motivate three simple methods for determining thresholds and a fourth method based on the optimal threshold algorithm from SDSSP1. Each of these methods is described below:

- *Greedy*: Set each threshold to 0 h.
- *Day-to-Day*: Set each threshold to the expected number of hours of cases arriving that day.
- *Cumulative*: Set each threshold to the expected cumulative remaining number of hours of cases.
- *Smart*: Use the optimal threshold algorithm from SDSSP1.

The final method requires some manipulation of the input data in order to apply the SDSSP1 algorithm. Specifically, the multiple case types must be combined into a single case type. The arrival distributions for primary cases are combined to form a single distribution for the hours of arrivals on each day. The deferral and blocking costs for each case type are combined to form weighted costs, with the weights determined by the expected number of arrivals of each case type. The following definitions create the input data for the threshold algorithm.

$$\begin{aligned} \tilde{T}_j &= \sum_{k=1}^K d_k T_j^k \\ \tilde{h}_j &= \frac{\sum_{k=1}^K h_j^k E[T_j^k]}{E[\tilde{T}_j]}, \quad j = 0, \dots, N \\ \tilde{r}_j &= \frac{\sum_{k=1}^K r_j^k E[T_j^k]}{E[\tilde{T}_j]}, \quad j = 1, \dots, N \end{aligned}$$

For more details on the specific algorithm used to find the optimal thresholds, see the proof of the threshold optimality claim in [Herring \(2010\)](#).

The first three methods, which are based on the observed behavior for specific problem scenarios, should work well for some problems and poorly for others. Specifically, the *Greedy* threshold rule should work well for problems with “greedy” cost ratios. Similarly, the *Day-to-Day* and *Cumulative* threshold rules should work well for problems with the “low” and “high” cost ratios, respectively. The final threshold method can adapt to the input data and should perform well in all scenarios.

4.2 Case type prioritization

Three different case type prioritization schemes are based on the insights gleaned from the optimal solution behavior. The last two schemes described below are for days $1, \dots, N$ only. Because the deferral and blocking cost structure is different on the day of surgery (a single penalty, r_0 , for OR hours left empty rather than multiple blocking costs, r_j^k), these ratio dependent prioritizations take advantage of the knapsack nature of the day 0 costs (as described in Sect. 2).

- *Duration*: Case types are prioritized by decreasing duration.
- *Ratios*: Case types are prioritized first by decreasing deferral cost-to-duration ratio, and second by decreasing blocking cost-to-duration ratio. Remaining ties are ordered by decreasing duration.
- *Threshold First*: Cases are prioritized as in the *Ratios* rule, but the prioritization is subjugated to first reaching the target threshold.

This final prioritization scheme is motivated by the observation that the optimal decisions often place a higher priority on reaching the threshold than favoring one case type over another. However, implementing this final rule requires a simple integer program (IP). In addition to input data on the state of the RQ (W_j^k) and the case types (d_k), the following definitions are needed for the IP:

D = target number of hours to take from the RQ and place on the OR schedule
 p_k = priority of case type k

The priority values amount to a permutation of the set of case types $\{1, \dots, K\}$, where a lower priority score denotes a preferred case type. Assuming that the case durations

591 and the scheduled OR capacity are integers, the IP can be defined as follows:

$$\begin{aligned}
 & \min \sum_{k=1}^K (C+1)^{p_k} x_j^k + (C+1)^{K+1} z \\
 592 \quad & \text{s.t.} \quad \sum_{k=1}^K d_k x_j^k + z = D \\
 & \quad 0 \leq x_j^k \leq W_j^k, \quad x_j^k \text{ integer}, \quad k = 1, \dots, K \\
 & \quad z \geq 0
 \end{aligned}$$

593 The slack variable z is given sufficient weight in the objective function to ensure that
 594 the target is reached if at all possible. If there is more than one way to hit the target, the
 595 RQ variables are weighted according to their priorities to ensure that more preferred
 596 cases are selected first. Note that the combination of base and exponent for the objec-
 597 tive function weights reflects a kind of base $(C+1)$ arithmetic that ensures that no
 598 combination of less-preferred cases will be chosen over a more preferred case type.

599 As mentioned above, the *Ratio* and *Threshold First* prioritization schemes rely on
 600 the knapsack structure of the day 0 costs for the day of surgery. The *Ratio* scheme uses
 601 a greedy knapsack heuristic, ordering the case types by decreasing value-to-weight
 602 ratio and greedily taking as many of each type as possible. Because the *Threshold First*
 603 method already relies on the availability of an IP solver, the day 0 knapsack problem
 604 is solved to optimality with this prioritization scheme.

605 4.3 Heuristic quality

606 Combining the threshold determination rules and the case type prioritization schemes
 607 yields twelve heuristics. The problem instances explored in Sect. 3 are used to test
 608 these heuristics, and their solutions are compared with the optimal solution values.
 609 Notice that the blocking weight parameter (λ) and the primary arrivals case type prev-
 610 alence ($E[T_j^2] : E[T_j^1]$), which have little influence on the optimal solutions, play
 611 similarly limited roles in the heuristics. For this reason, the heuristics are tested on the
 612 original test instances that have $\lambda = 1$ and $E[T_j^2] : E[T_j^1] = 1 : 1$. This leaves a set
 613 of 18 test problems that vary only in their deferral and blocking cost structure.

614 An effective way to test the quality of the various heuristics is to imbed them in a
 615 simulation environment. The evolution of an OR schedule with specified input data
 616 and a RQ decision-making policy can easily be simulated to compute the total cost
 617 associated with a RQ policy and a particular realization of the arrival random vari-
 618 ables. Each of the heuristics is simulated for 10,000 replications for each of the selected
 619 test problems, and the costs associated with each problem are expressed as percent
 620 deviations from the expected value of the optimal solution. Table 3 shows the mean
 621 deviation across all replications for each of the heuristics applied to the selected test
 622 problems. The 95% confidence interval (CI) half-widths for each of the values shown
 623 in Table 3 are given in Table 6 in the Appendix. Note that for all negative entries in the
 624 table, zero (representing the expected value of the optimal solution) falls within the
 625 CI. Recall that the three simple thresholds (*Greedy*, *Day-to-Day*, and *Cumulative*) are

Table 3 Average percent deviation from the optimal cost for twelve threshold-based heuristics applied to test problems representing a range of deferral and blocking cost structures

$h_j^2 : h_j^1$	1:1									Mean deviation	
	1:1			3:1			5:1			Target	All
$r_j^1 : h_j^1$	1:1			3:1			5:1				
$r_j^2 : r_j^1$	1:1	2:1	3:1	1:1	2:1	3:1	1:1	2:1	3:1		
Greedy											
Duration	1.3	5.2	7.9	19.6	50.3	87.2	61.5	129.2	194.9	4.8	61.9
Ratios	2.8	7.7	9.4	19.8	51.0	88.0	59.9	128.2	194.8	6.6	62.4
Threshold First	18.0	18.8	17.5	24.6	53.2	88.6	62.4	126.9	191.3	18.1	66.8
Day-to-Day											
Duration	50.0	29.2	15.1	4.4	7.3	15.5	11.4	31.8	48.2	9.1	23.7
Ratios	60.1	36.2	21.2	9.1	9.9	17.3	14.7	33.4	47.4	12.1	27.7
Threshold First	47.4	27.2	13.2	2.9	5.6	14.1	10.0	30.6	46.9	7.5	22.0
Cumulative											
Duration	96.7	62.2	39.0	11.6	4.4	3.7	2.9	6.3	8.8	6.0	26.2
Ratios	110.6	73.5	47.6	20.1	11.0	9.2	10.7	12.3	16.0	13.0	34.6
Threshold First	94.6	60.3	37.5	10.4	3.2	2.5	1.8	5.3	7.8	5.0	24.8
Smart											
Duration	1.3	31.7	20.8	8.3	4.4	3.7	2.9	4.2	3.8	n/a	9.0
Ratios	2.8	31.5	23.0	16.2	11.0	10.5	10.7	11.5	11.8	n/a	14.3
Threshold First	18.0	36.4	20.0	7.2	3.2	2.2	1.8	2.6	2.3	n/a	10.4
$h_j^2 : h_j^1$	2:1									Mean deviation	
$r_j^1 : h_j^1$	1:1			3:1			5:1			Target	All
$r_j^2 : r_j^1$	1:1	2:1	3:1	1:1	2:1	3:1	1:1	2:1	3:1		
Greedy											
Duration	-0.2	-1.1	5.2	3.0	24.4	46.8	31.7	75.3	126.6	1.3	34.6
Ratios	7.8	-1.1	5.2	5.1	24.4	46.8	33.2	75.3	126.6	4.0	35.9
Threshold First	25.3	17.1	18.5	11.3	29.5	50.2	35.4	77.4	127.0	20.3	43.5
Day-to-Day											
Duration	81.7	48.6	35.5	6.0	2.1	4.5	4.5	13.0	28.6	4.2	24.9
Ratios	89.0	48.6	35.5	8.4	2.1	4.5	4.8	13.0	28.6	5.0	26.1
Threshold First	81.7	48.6	35.5	5.9	2.1	4.5	4.4	13.0	28.6	4.2	24.9
Cumulative											
Duration	155.1	98.5	74.1	24.6	10.2	4.3	6.4	3.5	6.1	5.3	42.5
Ratios	156.8	98.5	74.1	25.2	10.2	4.3	6.9	3.5	6.1	5.5	42.8
Threshold First	155.1	98.5	74.1	24.6	10.2	4.3	6.4	3.5	6.1	5.3	42.5

Table 3 continued

$h_j^2 : h_j^1$	2:1									Mean deviation	
	1:1			3:1			5:1			Target	All
$r_j^1 : h_j^1$	1:1			2:1			3:1				
$r_j^2 : r_j^1$	1:1	2:1	3:1	1:1	2:1	3:1	1:1	2:1	3:1		
Smart											
Duration	-0.2	-1.1	37.3	9.6	9.1	4.3	5.9	3.5	3.2	n/a	8.0
Ratios	7.8	-1.1	37.3	12.1	9.1	4.3	6.4	3.5	3.2	n/a	9.2
Threshold First	25.3	17.1	40.0	9.7	9.1	4.3	6.0	3.5	3.2	n/a	13.1

Averages across problems are shown both for targeted instances (indicated by italics) and all instances

626 expected to perform better for some instances than others. These targeted instances
 627 are shown in italics in Table 3, and the average performance of the heuristics across
 628 problem instances is shown for both the targeted instances and for all instances.

629 Several observations can be made from these results. First, the scenario-specific
 630 thresholds perform well when they should. In most cases, the cost of the solutions is
 631 less than 10% away from the optimal cost regardless of how the case types are prior-
 632 itized. As might be expected, the quality of these thresholds deteriorates rapidly for
 633 other cost structures. The *Smart* threshold adapts to the different cost structures, with
 634 deviations typically falling at or below 10% and only exceeding 20% on a few of the
 635 “greedy” problems (those with $r_j^1 : h_j^1 = 1 : 1$). The *Smart* threshold outperforms the
 636 other thresholds as a general-purpose heuristic. Furthermore, its flexibility suggests
 637 that it will be able to conform and adapt to an even broader range of cost structures
 638 than those tested here.

639 None of the three prioritization schemes significantly outperforms the others across
 640 the entire range of problems. The *Threshold First* prioritization scheme is worse than
 641 the other prioritization schemes when paired with the *Greedy* and *Smart* thresholds
 642 and applied to the “greedy” problems. A greedy approach is appropriate for these
 643 problems, as observed from the policy maps, and this discrepancy between the prior-
 644 itization schemes confirms the intuition that case type preference matters more when
 645 larger numbers of RQ cases are being placed. For the other subsets of problems, the
 646 *Duration* and *Threshold First* prioritization schemes slightly outperform the *Ratios*
 647 rule. Note that for the given case durations (1 and 2 h) the (decreasing) *Duration* rule
 648 will always reach the target if it is possible. This confirms the earlier observation that
 649 reaching the threshold is often more important than a particular case type preference.

650 4.4 Impact of arrival timing

651 The results for SDSSP1 (see Table 1) suggest that the timing of the primary case arriv-
 652 als also impacts the target thresholds. In order to study how this affects the performance
 653 of the proposed heuristics, a second set of test problems is needed. The full range of
 654 inputs used in Table 3 (fixed blocking weight parameter, equal numbers of each case
 655 type expected over the course of the problem, and varying cost ratios) is combined

Table 4 Average percent deviation from the optimal solution for three threshold-based heuristics applied to test problems representing a range of primary service line arrival patterns and cost structures

$r_j^1 : h_j^1$	Early, Early			Middle, Middle			Late, Late			Early, Late			Late, Early		
	1:1	3:1	5:1	1:1	3:1	5:1	1:1	3:1	5:1	1:1	3:1	5:1	1:1	3:1	5:1
Heuristic															
Greedy + Duration	1.9	15.8	40.0	-1.1	24.4	75.3	2.3	21.4	77.0	1.8	17.6	60.1	0.4	15.0	61.5
Smart + Duration	1.9	20.1	12.0	-1.1	9.1	3.5	2.3	6.4	0.4	1.8	5.5	3.5	0.4	1.3	3.8
Smart + Threshold First	14.7	20.3	12.1	17.1	9.1	3.5	28.8	6.4	0.4	18.3	5.5	3.5	19.3	1.3	3.8

with four new arrival patterns. Each of the two case types is set to arrive either “Early” or “Late,” and all possible combinations are tested. The arrival random variables are still assumed to follow Poisson distributions for each day, with the “Early” arrival rates set to (1, 0.5, 0.25, 0.25, 0) and the “Late” arrival rates set to (0.25, 0.25, 0.5, 1, 0). This new test set consists of 72 problem instances.

Rather than apply all twelve of the proposed heuristics to this new problem set, three representative heuristics are chosen: *Greedy + Duration*, *Smart + Duration*, and *Smart + Threshold First*. The *Greedy + Duration* heuristic is chosen because it is the most naïve of the twelve and serves as a sort of worst case choice. The other two are chosen because they perform well over the entire initial problem set and most closely represent the trends observed in the policy maps. Table 4 shows the average performance of these heuristics over 10,000 simulated replications on the new set of test problems, again shown as the percent deviation from the optimal solution value. The evenly spread out arrival scenario in Table 3 is referred to as “Middle” and is included in the table for comparison’s sake. In the interest of space, the results are only shown for single values of the between-case-type deferral and blocking ratios. Specifically, these ratios in Table 4 are fixed as $h_j^2 : h_j^1 = r_j^2 : r_j^1 = 2 : 1$ (the case duration ratio). While the exact magnitude of the deviations differs for other deferral and blocking cost structures, the observations and patterns discussed here remain constant across other cost structures. The 95% confidence interval half-widths for the results in Table 4 are given in Table 7 in the Appendix.

Many of the insights from the initial set of test problems are still evident in this second set of tests. As expected, the *Smart* thresholds tend to outperform the *Greedy + Duration* heuristic in all arrival scenarios except the “greedy” problems where the *Greedy* threshold is expected to perform well. On average, the *Smart + Duration* combination still slightly outperforms the *Smart + Threshold First* combination. Most of difference between these two heuristics originates from the “greedy” problems, while their performance is nearly identical for the other cost ratios.

It is also clear from the results in Table 4 that the timing of the primary service line’s arrivals does have some impact on the performance of the heuristics. The *Greedy + Duration* heuristic performs better overall when more of the cases arrive early, while the other two heuristics perform worst in the (Early, Early) scenario. The early arrivals leave fewer arrivals closer to the day of surgery, which reduces the time that must be preserved for remaining cases and makes greedier policies more

691 appropriate. As observed with the first set of test problems, the *Smart* threshold suffers
 692 in scenarios where greedier policies are appropriate and more importance is placed
 693 on the case type preference. On the other hand, the *Smart* threshold paired with either
 694 case type preference does well when lower thresholds are required. In these scenarios,
 695 the case type preference plays a less significant role because fewer hours of cases
 696 are taken. While these observations appear at first to implicate the case type pref-
 697 erences as the weak step of the proposed heuristics, it may in fact be the case that
 698 the threshold behavior itself becomes less precise as the optimal solutions become
 699 “greedier.”

700 4.5 Impact of block release dates

701 From the perspective of the single-day scheduling problem, block release dates act as
 702 constraints on RQ decisions and should lead to lower quality solutions. However, the
 703 extent of the impact will depend upon the nature of the optimal decisions. In particu-
 704 lar, forcing some decisions to zero will have less impact on problems with high target
 705 thresholds and will have more impact on problems with low target thresholds.

706 In order to study the impact of block release dates, the 90 test problems presented
 707 in Tables 3 and 4 are combined to form a single set of problem instances. The feasible
 708 block release dates for these problems range from 3 days before surgery (day 3) to the
 709 day of surgery (day 0). Setting the block release date to day 3 does not constrain the
 710 solutions in any way because day 3 is the first day on which a RQ decision needs to
 711 be made (the system is empty on day 4). Once a block release date has been chosen, a
 712 RQ policy is needed to make decisions after the block is released. Three different RQ
 713 policies are combined with the block release dates: (1) the optimal decisions from the
 714 fully solved SDPs, (2) the *Greedy + Duration* heuristic, and (3) the *Smart + Duration*
 715 heuristic. Each of these twelve combinations is simulated over 10,000 replications
 716 for each problem instance and the average results are shown in Table 5. The results
 717 are again given as percent deviations from the unconstrained optimal solutions. As in
 718 Table 4, the results are only shown for single values of the between-case-type deferral
 719 and blocking ratios ($h_j^2 : h_j^1 = r_j^2 : r_j^1 = 2 : 1$). The exact magnitude of the deviations
 720 for other deferral and blocking cost structures differs some, but the observations and
 721 patterns discussed here remain constant across other cost structures. The 95% con-
 722 fidence interval half-widths for the means in Table 5 are given in the Table 8 in the
 723 Appendix.

724 When paired with the optimal decisions, the day 3 results are consistently within 3
 725 percent of the optimal expected value (with the 95% CI half-widths always including
 726 zero), verifying the simulation’s ability to accurately gauge the performance of a given
 727 policy. When paired with the heuristic RQ policies, the day 3 results are directly in
 728 line with the unconstrained results presented in Tables 3 and 4, confirming for these
 729 instances that setting the block release date to day 3 is equivalent to having no block
 730 release date.

731 As the block release date moves closer to the day of surgery, its impact on solution
 732 quality becomes evident. Note that when the block release date is set to day 0, the choice
 733 of RQ policy loses significance because all the policies become greedy on the day of

Table 5 Average percent deviation from the unconstrained optimal solution for three RQ policies paired with all possible block release dates

$r_j^1 : h_j^1$	Early, Early			Middle, Middle			Late, Late			Early, Late			Late, Early		
	1:1	3:1	5:1	1:1	3:1	5:1	1:1	3:1	5:1	1:1	3:1	5:1	1:1	3:1	5:1
<i>RQ policy</i>															
Block release															
<i>Optimal decisions</i>															
Day 3	2.9	-0.3	-1.6	-1.8	1.4	2.4	-0.9	0.5	0.6	2.6	1.4	-1.0	-1.1	-2.1	-0.6
Day 2	42.1	5.4	-1.3	38.0	6.1	2.5	47.5	8.5	0.8	44.6	7.0	-0.6	44.3	4.3	-0.3
Day 1	93.1	19.9	3.3	81.2	9.8	2.6	100.4	10.9	0.8	92.5	12.6	-0.5	100.6	11.7	1.3
Day 0	143.0	35.3	7.1	115.0	12.8	4.2	124.9	11.1	0.8	123.3	16.7	0.2	141.5	17.3	3.9
<i>Greedy + Duration</i>															
Day 3	5.3	15.2	39.8	0.7	23.6	75.2	1.6	22.5	77.0	8.2	19.9	58.2	-0.5	14.9	59.8
Day 2	43.4	8.2	8.9	40.6	21.4	52.0	50.8	37.8	83.3	52.5	27.9	52.1	44.8	15.1	41.4
Day 1	94.5	21.4	8.1	83.2	16.0	25.6	103.6	36.0	54.7	97.8	28.3	31.5	101.5	23.4	31.3
Day 0	143.0	35.3	7.1	115.0	12.8	4.2	124.9	11.1	0.8	123.3	16.7	0.2	141.5	17.3	3.9
<i>Smart + Duration</i>															
Day 3	5.3	17.1	9.2	0.7	8.2	5.5	1.6	8.8	0.8	8.2	8.1	2.7	-0.5	-0.6	1.5
Day 2	43.4	25.7	8.9	40.6	9.3	5.5	50.8	11.1	0.8	52.5	18.4	2.7	44.8	5.9	1.5
Day 1	94.5	34.9	12.6	83.2	9.8	3.5	103.6	11.1	0.8	97.8	18.4	2.7	101.5	13.8	2.2
Day 0	143.0	35.3	7.1	115.0	12.8	4.2	124.9	11.1	0.8	123.3	16.7	0.2	141.5	17.3	3.9

surgery. For the problems where greedier solutions are appropriate ($r_j^1 : h_j^1 = 1 : 1$), the solutions deteriorate quickly as the block release policy becomes more restrictive. For other blocking-to-deferral cost ratios, the added costs of the block release dates depend heavily on the RQ policy they are paired with. When the block release dates are paired with good RQ policies (the optimal decisions or *Smart + Duration*), the added costs increase slightly as the block release policy becomes more restrictive, but rarely exceed 20% deviation from the optimal. In contrast, when the block release dates are paired with a poor RQ policy (*Greedy + Duration*) the solutions at times improve as the block release policy becomes more restrictive. This is most noticeable when the blocking-to-deferral cost ratio is high ($r_j^1 : h_j^1 = 5 : 1$), which is exactly the scenario where the *Greedy* threshold performs poorly. This ability of block release dates to mitigate the effect of poor RQ policies adds a level of theoretical justification to their use in practice, where the overall quality of RQ policies may be unknown and difficult to ascertain.

5 Conclusions

This paper demonstrates that the threshold behavior that was shown to be optimal for the dynamic single-day surgery scheduling problem with unit case durations can be successfully extended to yield high quality solutions to the scheduling problem with

multiple case durations. Showing that the optimal solutions to SDSSP remain constant for large numbers of cases on the request queue maintains the computational tractability of the SDP and allows a diverse range of test problems to be solved to optimality. A detailed analysis of the optimal policy maps for these test problems reveals that the solutions follow an approximate threshold behavior and gives insight into how good thresholds and case type prioritization rules can be combined to form threshold-based heuristics.

A set of threshold-based heuristics is proposed, and the computational results show that heuristics using the optimal threshold algorithm from SDSSP1 consistently outperform heuristics using other threshold rules. While none of the proposed case type prioritization schemes dominates the others, the results favor schemes that focus on hitting the target thresholds rather than enforcing strict case type preferences. The nature of the target thresholds varies with both the deferral and blocking cost structure and the timing of the primary case arrivals. The results indicate that the threshold-based heuristics perform best when the target thresholds are higher (indicating that fewer RQ cases will be scheduled) and deteriorate somewhat when “greedy” (lower) target thresholds are appropriate.

Block release dates are shown to be increasingly costly for problems with “greedy” target thresholds. However, in scenarios with higher target thresholds, imposing block release dates is not costly if good RQ decisions are made after the block is released. Block release dates can improve solution quality when paired with poor RQ decisions (that is, the resulting solutions are not as bad as they would be otherwise). This observation provides practical justification for the use of traditional block release dates in practice.

This study of the single-day surgery scheduling problem is the first to model the sequential nature of an OR manager’s daily block release and RQ decisions and to investigate the dynamic interaction of these decisions with the block schedule and primary case arrivals. The proposed threshold-based decision rules suggest a new way for hospitals and OR managers to look at block release and RQ policies. Rather than setting a block release date after which all unused OR time is released (as is done in practice), our threshold-based rules would gradually release unused OR time over the course of several days leading up to the day of surgery. These model-produced thresholds have the ability to adapt to differences in priority levels, contributions to revenue, and demand arrival patterns between surgical specialties and would help provide clear and empirical justification for an OR manager’s decisions during the often contentious development of single-day surgery schedules.

Our modeling framework emerged from our discussions with collaborators in the peri-operative services department at UMMC, who have validated the importance of blocking and deferral costs in RQ decision-making, reviewed these results, and endorsed the concept of threshold-based policies. In order to develop specific threshold policies for this facility, we plan to use data from the hospital’s surgical scheduling system databases to estimate the distributions of surgical demand arrivals for their OR suite. Additionally, we will continue to work with the appropriate OR stakeholders to develop reliable methods for determining accurate deferral and blocking cost values. These values have the potential to depend on differences in patient acuity levels, specialized resource needs, and revenue contributions between surgical specialties, and as

798 such will involve a combination of expert elicitation and data analysis. Finally, contin-
 799 ued research in this area will use data from UMMC to compare schedules developed
 800 using our threshold-based policies with actual historical schedules on more traditional
 801 metrics such as OR utilization and throughput.

802 **Appendix**

803 See Tables 6, 7, 8.

Table 6 95% confidence interval half-widths for results presented in Table 3

$h_j^2 : h_j^1$	1:1								
$r_j^1 : h_j^1$	1:1			3:1			5:1		
$r_j^2 : r_j^1$	1:1	2:1	3:1	1:1	2:1	3:1	1:1	2:1	3:1
Greedy									
Duration	3.6	3.2	2.9	2.7	3.1	3.9	3.1	4.5	6.1
Ratios	3.8	3.4	3.0	2.7	3.2	4.0	3.1	4.5	6.2
Threshold First	3.8	3.4	3.0	2.7	3.2	4.0	3.2	4.5	6.1
Day-to-Day									
Duration	4.0	3.4	2.9	2.5	2.5	2.8	2.5	3.1	3.9
Ratios	4.5	3.7	3.2	2.7	2.7	2.9	2.7	3.2	3.9
Threshold First	4.0	3.4	2.9	2.5	2.5	2.8	2.5	3.1	3.9
Cumulative									
Duration	4.6	3.8	3.3	2.7	2.5	2.5	2.4	2.6	2.8
Ratios	5.1	4.3	3.6	3.0	2.7	2.7	2.7	2.8	3.1
Threshold First	4.6	3.8	3.2	2.6	2.5	2.4	2.4	2.6	2.8
Smart									
Duration	3.6	3.6	3.2	2.6	2.5	2.4	2.4	2.5	2.5
Ratios	3.8	3.7	3.3	2.9	2.7	2.7	2.7	2.8	2.8
Threshold First	3.8	3.7	3.2	2.6	2.5	2.4	2.4	2.4	2.5
$h_j^2 : h_j^1$	2:1								
$r_j^1 : h_j^1$	1:1			3:1			5:1		
$r_j^2 : r_j^1$	1:1	2:1	3:1	1:1	2:1	3:1	1:1	2:1	3:1
Greedy									
Duration	3.5	2.9	2.9	2.3	2.5	3.0	2.5	3.4	4.7
Ratios	3.7	2.9	2.9	2.4	2.5	3.0	2.6	3.4	4.7
Threshold First	3.8	3.1	3.0	2.4	2.6	3.1	2.6	3.5	4.7

Table 6 continued

$h_j^2 : h_j^1$	2:1								
$r_j^1 : h_j^1$	1:1			3:1			5:1		
$r_j^2 : r_j^1$	1:1	2:1	3:1	1:1	2:1	3:1	1:1	2:1	3:1
Day-to-Day									
Duration	4.4	3.5	3.2	2.3	2.2	2.3	2.2	2.5	3.1
Ratios	4.6	3.5	3.2	2.4	2.2	2.3	2.2	2.5	3.1
Threshold First	4.4	3.5	3.2	2.3	2.2	2.3	2.2	2.5	3.1
Cumulative									
Duration	5.5	4.3	3.8	2.7	2.4	2.3	2.3	2.3	2.5
Ratios	5.5	4.3	3.8	2.8	2.4	2.3	2.3	2.3	2.5
Threshold First	5.5	4.3	3.8	2.7	2.4	2.3	2.3	2.3	2.5
Smart									
Duration	3.5	2.9	3.3	2.5	2.3	2.3	2.2	2.3	2.3
Ratios	3.7	2.9	3.3	2.6	2.3	2.3	2.3	2.3	2.3
Threshold First	3.8	3.1	3.3	2.5	2.3	2.3	2.2	2.3	2.3

Table 7 95% confidence interval half-widths for results presented in Table 4

	Early, Early			Middle, Middle			Late, Late			Early, Late			Late, Early		
$r_j^1 : h_j^1$	1:1	3:1	5:1	1:1	3:1	5:1	1:1	3:1	5:1	1:1	3:1	5:1	1:1	3:1	5:1
Heuristic															
Greedy + Duration	4.8	3.8	4.3	2.9	2.5	3.4	2.2	1.8	2.5	2.9	2.4	3.0	3.5	2.8	3.6
Smart + Duration	4.8	3.9	3.7	3.5	2.3	2.3	2.2	1.7	1.6	2.9	2.1	2.1	3.5	2.6	2.6
Smart + Threshold First	5.0	3.9	3.7	3.8	2.3	2.3	2.4	1.7	1.6	3.1	2.1	2.1	3.7	2.6	2.6

Table 8 95% confidence interval half-widths for results presented in Table 5

	Early, Early			Middle, Middle			Late, Late			Early, Late			Late, Early		
$r_j^1 : h_j^1$	1:1	3:1	5:1	1:1	3:1	5:1	1:1	3:1	5:1	1:1	3:1	5:1	1:1	3:1	5:1
RQ policy															
Block release															
Optimal decisions															
Day 3	4.6	3.4	3.2	2.9	2.1	2.2	2.2	1.6	1.6	3.0	2.1	2.0	3.5	2.4	2.6
Day 2	5.1	3.5	3.2	3.3	2.3	2.2	2.6	1.8	1.6	3.3	2.2	2.0	4.1	2.6	2.6
Day 1	6.1	3.8	3.3	4.0	2.4	2.2	3.3	1.8	1.6	3.9	2.2	2.0	5.1	2.8	2.6
Day 0	7.3	4.1	3.3	4.7	2.4	2.3	3.7	1.8	1.6	4.3	2.3	2.0	6.1	2.9	2.6

Table 8 continued

$r_j^1 : h_j^1$	Early, Early			Middle, Middle			Late, Late			Early, Late			Late, Early		
	1:1	3:1	5:1	1:1	3:1	5:1	1:1	3:1	5:1	1:1	3:1	5:1	1:1	3:1	5:1
<i>Greedy + Duration</i>															
Day 3	4.8	3.8	4.3	3.1	2.5	3.4	2.3	1.9	2.5	3.2	2.4	3.0	3.5	2.7	3.6
Day 2	5.2	3.7	3.7	3.5	2.6	3.3	2.7	2.1	2.7	3.5	2.6	3.1	4.1	2.8	3.6
Day 1	6.2	3.9	3.6	4.1	2.6	3.0	3.4	2.3	2.7	4.0	2.6	2.9	5.1	3.1	3.5
Day 0	7.3	4.1	3.3	4.7	2.4	2.3	3.7	1.8	1.6	4.3	2.3	2.0	6.1	2.9	2.6
<i>Smart + Duration</i>															
Day 3	4.8	3.9	3.6	3.1	2.3	2.4	2.3	1.8	1.6	3.2	2.1	2.1	3.5	2.5	2.6
Day 2	5.2	4.1	3.6	3.5	2.4	2.4	2.7	1.8	1.6	3.5	2.3	2.1	4.1	2.6	2.6
Day 1	6.2	4.3	3.7	4.1	2.4	2.3	3.4	1.8	1.6	4.0	2.3	2.1	5.1	2.8	2.6
Day 0	7.3	4.1	3.3	4.7	2.4	2.3	3.7	1.8	1.6	4.3	2.3	2.0	6.1	2.9	2.6

Proofs for Claims 1 and 2

See Sect. 2 for the definition of all variables, parameters, and functions used below.

Claim 1 *If, for a given state, $W_j^k \geq \lfloor \frac{C_j}{d_k} \rfloor$, then $W_{j-1}^k \geq \lfloor \frac{C_{j-1}}{d_k} \rfloor$ for all possible future states.*

Proof:

$$\begin{aligned}
 W_{j-1}^k - \left\lfloor \frac{C_{j-1}}{d_k} \right\rfloor &= \left(W_j^k - x_j^k + R_j^k + (T_j^k - fit_j^k) \right) - \left\lfloor \frac{C_j - \sum_{i=1}^K d_i (x_j^i + fit_j^i)}{d_k} \right\rfloor \\
 &= W_j^k + (R_j^k + T_j^k) - (x_j^k + fit_j^k) \\
 &\quad - \left\lfloor \frac{C_j - \sum_{i=1, i \neq k}^K d_i (x_j^i + fit_j^i)}{d_k} - (x_j^k + fit_j^k) \right\rfloor \\
 &= W_j^k + (R_j^k + T_j^k) - \left\lfloor \frac{C_j - \sum_{i=1, i \neq k}^K d_i (x_j^i + fit_j^i)}{d_k} \right\rfloor \\
 &\geq W_j^k - \left\lfloor \frac{C_j}{d_k} \right\rfloor \\
 &\geq 0
 \end{aligned}$$

The first several steps are simply manipulations of the transition equations defined above. The second to last step uses the fact that the arrival random variables are

816 non-negative, as well as an observation on the nature of the floor function. The final
 817 step reflects the initial assumption that the desired inequality holds for day j . \square

818 **Claim 2** For all feasible states (C_j, B_j, \mathbf{W}_j) for all days $j = 0, \dots, N$, the following
 819 equalities hold:

- 820 (i) $F_j(C_j, B_j, \mathbf{W}_j, \mathbf{x}_j) = F_j(C_j, B_j, \widehat{\mathbf{W}}_j, \mathbf{x}_j)$
- 821 (ii) $V_j(C_j, B_j, \mathbf{W}_j) = V_j(C_j, B_j, \widehat{\mathbf{W}}_j)$
- 822 (iii) $\mathbf{x}_j(C_j, B_j, \mathbf{W}_j) = \mathbf{x}_j(C_j, B_j, \widehat{\mathbf{W}}_j)$

823 *Proof:* The proof proceeds by induction on j .

824 *Base Case ($j = 0$):* The first equality relies directly on the definition of $\widehat{\mathbf{W}}_0^k$.

$$\begin{aligned}
 825 \quad F_0(C_0, B_0, \mathbf{W}_0, \mathbf{x}_0) &= \sum_{k=1}^K h_0^k \left(\min \left(\left\lfloor \frac{C_0}{d_k} \right\rfloor, W_0^k \right) - x_0^k \right) \\
 826 \quad &+ r_0 \left(C_0 - \sum_{k=1}^K d_k x_0^k \right) \\
 827 \quad &= \sum_{k=1}^K h_0^k \left(\widehat{W}_0^k - x_0^k \right) + r_0 \left(C_0 - \sum_{k=1}^K d_k x_0^k \right) \\
 828 \quad &= F_0(C_0, B_0, \widehat{\mathbf{W}}_0, \mathbf{x}_0)
 \end{aligned}$$

829 Note that the feasible regions for the decision vector are identical for states
 830 (C_0, B_0, \mathbf{W}_0) and $(C_0, B_0, \widehat{\mathbf{W}}_0)$. Both sides of equalities (ii) and (iii) represent minimi-
 831 zation problems of identical functions over identical feasible regions, clearly implying
 832 that these equalities must hold.

833 *Inductive Step:* Assume that all the desired equalities hold for day $j - 1$. Note that
 834 moving from state (C_j, B_j, \mathbf{W}_j) to $(C_j, B_j, \widehat{\mathbf{W}}_j)$ has no impact on the deferral and
 835 blocking penalties (ND_j^k , NB_j^k , and FB_j^k .) provided that the decision variables are
 836 held constant. The deferrals are computed based on what is feasible, the blocking
 837 penalties do not rely on the RQ outside of determining feasible decisions, and, as in
 838 the base case, the feasible regions are identical for both states. Therefore, looking at
 839 the difference between $F_j(C_j, B_j, \mathbf{W}_j, \mathbf{x}_j)$ and $F_j(C_j, B_j, \widehat{\mathbf{W}}_j, \mathbf{x}_j)$ comes down
 840 to looking at the difference between the subsequent day $j - 1$ states. The next day's
 841 available space (C_{j-1}) and blocking eligible hours (B_{j-1}) will also be identical,
 842 again provided that the decision variables are held constant. Therefore, the focus lies
 843 on potential differences on day $j - 1$ between the RQ states. The goal is to show that
 844 the trimmed day $j - 1$ RQ states (limited to the day $j - 1$ decisions that are feasible)
 845 are equal starting from either W_j^k or \widehat{W}_j^k on day j . It will be helpful to express the RQ
 846 day j to day $j - 1$ transition equation as a function of the starting RQ state, with all
 847 other variables being held constant.

848 $w(W_j^k) =$ subsequent day $j - 1$ RQ state for case type k from state (C_j, B_j, \mathbf{W}_j)

849 $w(\widehat{W}_j^k)$ = subsequent day $j - 1$ RQ state for case type k from state $(C_j, B_j, \widehat{W}_j)$

850 Showing that these two RQ states have equal trimmed RQ states now amounts to
 851 showing that the following equality holds.

$$852 \quad \min\left(\left\lfloor \frac{C_{j-1}}{d_k} \right\rfloor, w(W_j^k)\right) = \min\left(\left\lfloor \frac{C_{j-1}}{d_k} \right\rfloor, w(\widehat{W}_j^k)\right)$$

853 If $W_j^k = \widehat{W}_j^k$, then equality trivially holds. If $W_j^k \neq \widehat{W}_j^k$, then $W_j^k > \left\lfloor \frac{C_j}{d_k} \right\rfloor$ and $\widehat{W}_j^k =$
 854 $\left\lfloor \frac{C_j}{d_k} \right\rfloor$. Then by the results of Claim 1, $w(W_j^k) \geq \left\lfloor \frac{C_{j-1}}{d_k} \right\rfloor$ and $w(\widehat{W}_j^k) \geq \left\lfloor \frac{C_{j-1}}{d_k} \right\rfloor$. As
 855 a result, both sides of the equality reduce to $\left\lfloor \frac{C_{j-1}}{d_k} \right\rfloor$.

856 Returning to equality (i), the difference in question becomes:

$$\begin{aligned} 857 \quad & F_j(C_j, B_j, \mathbf{W}_j, \mathbf{x}_j) - F_j(C_j, B_j, \widehat{\mathbf{W}}_j, \mathbf{x}_j) \\ 858 \quad & = E[V_{j-1}(C_{j-1}, B_{j-1}, \mathbf{w}(\mathbf{W}_j)) - V_{j-1}(C_{j-1}, B_{j-1}, \mathbf{w}(\widehat{\mathbf{W}}_j))] \\ 859 \quad & = E[V_{j-1}(C_{j-1}, B_{j-1}, \mathbf{w}(\widehat{\mathbf{W}}_j)) - V_{j-1}(C_{j-1}, B_{j-1}, \mathbf{w}(\widehat{\mathbf{W}}_j))] \\ 860 \quad & = 0 \end{aligned}$$

861 The first step uses equality (ii) of the induction assumption, and second uses the
 862 result above that the trimmed day $j - 1$ RQ states will be equal for both original day
 863 j RQ states. As in the base case, equalities (ii) and (iii) follow directly from equality
 864 (i). □

865 **Pseudocode for *ProcessDay* algorithm (from Sect. 2)**

866 INPUT: $C_j, B_j, \mathbf{x}_j, T_j, \mathbf{d}$

867 DEFINE:

868 fit_j^k = number of primary arrivals of type k placed into the room on day j
 869 $blfit_j^k$ = number of primary arrivals of type k that would have fit in the room on
 870 day j if the blocking eligible hours were available

871 INITIALIZE:

872 $\tilde{C}_j = C_j - \sum_{k=1}^K d_k x_j^k$ # space available in room after RQ decisions
 873 $\tilde{B}_j = B_j + \sum_{k=1}^K d_k x_j^k$ # blocking eligible hours after RQ decisions
 874 $\tilde{D}_j = \tilde{C}_j + \tilde{B}_j$ # space available if blocking eligible hours were free

875 FOR $k = K$ TO 1 # reflects prioritization of primary arrivals

876 # determine cases of type k that either fit or would have fit into open space

$$877 \quad tmpfit = \frac{\tilde{C}_j}{d_k}$$

$$878 \quad fit_j^k = \min \left(T_j^k, [tmpfit] \right)$$

$$879 \quad blfit_j^k = \min \left(T_j^k, \left\lfloor \frac{\tilde{D}_j}{d_k} \right\rfloor \right)$$

880 # compute blocking penalties

$$881 \quad NB_j^k = blfit_j^k - fit_j^k$$

$$882 \quad FB_j^k = \max \left(0, blfit_j^k - tmpfit \right)$$

883 # compute remaining space for next iteration

$$884 \quad \tilde{C}_j = \tilde{C}_j - d_k fit_j^k$$

$$885 \quad \tilde{D}_j = \tilde{D}_j - d_k blfit_j^k$$

886 END

887 RETURN: fit_j, NB_j, FB_j

888 References

- 889 Beliën J, Demeulemeester E (2007) Building cyclic master surgery schedules with leveled resulting bed
890 occupancy. *Eur J Oper Res* 176:1185–1204
- 891 Beliën J, Demeulemeester E, Cardoen B (2009) A decision support system for cyclic master surgery sched-
892 uling with multiple objectives. *J Sched* 12(2):147–161
- 893 Blake J, Carter M (2002) A goal programming approach to strategic resource allocation in acute care hos-
894 pitals. *Eur J Oper Res* 140:541–561
- 895 Blake J, Donald J (2002) Mount Sinai Hospital uses integer programming to allocate operating room time.
896 *Interfaces* 32(2):63–73
- 897 Brumelle S, Walczak D (2003) Dynamic airline revenue management with multiple semi-Markov demand.
898 *Oper Res* 51:137–148
- 899 Cardoen B, Demeulemeester E, Beliën J (2009) Sequencing surgical cases in a day-care environment: an
900 exact branch-and-price approach. *Comput Oper Res* 36(9):2660–2669
- 901 Cardoen B, Demeulemeester E, Beliën J (2010) Operating room planning and scheduling: a literature
902 review. *Eur J Oper Res* 201:921–932
- 903 Denton B, Viapiano J, Vogl A (2007) Optimization of surgery sequencing and scheduling decisions under
904 uncertainty. *Health Care Manag Sci* 10:13–24
- 905 Dexter F, Macario A, Traub RD (1999) Which algorithm for scheduling elective add-on cases maximizes
906 operating room utilization? Use of bin packing algorithms and fuzzy constraints in operating room
907 management. *Anesthesiology* 91:1491–1500
- 908 Dexter F, Traub R (2002) How to schedule elective surgical cases into specific operating rooms to maximize
909 the efficiency of use of operating room time. *Anesth Analg* 94:933–942
- 910 Dexter F, Traub R, Macario A (2003) How to release allocated operating room time to increase efficiency:
911 predicting which surgical service will have the most under-utilized operating room time. *Anesth Analg*
912 96:507–512

- 913 Dexter F, Macario A (2004) When to release allocated operating room time to increase operating room
 914 efficiency. *Anesth Analg* 98:758–762
- 915 Gerchak Y, Gupta D, Henig M (1996) Reservation planning for elective surgery under uncertain demand
 916 for emergency surgery. *Management Science*. 42(3):321–334
- 917 Guinet A, Chaabane S (2003) Operating theatre planning. *Int J Prod Econ* 85:69–81
- 918 Gupta D (2007) Surgical suites' operations management. *Prod Oper Manag* 16(6):689–700
- 919 Hans E, Wullink G, van Houdenhoven M, Kazemier G (2008) Robust surgery loading. *Eur J Oper Res*
 920 185:1038–1050
- 921 Herring WL (2010) A stochastic dynamic program for the single-day surgery scheduling problem. In: Pro-
 922 ceedings of the 2010 IIE Society for Health Systems Annual Conference, Atlanta, 25–28 Feb 2010
- 923 Herring WL (2011) Prioritizing patients: stochastic dynamic programming for surgery scheduling and mass
 924 casualty incident triage. Ph.D. Dissertation, University of Maryland, College Park, MD. Available at
 925 <http://hdl.handle.net/1903/11689>
- 926 Lee TC, Hersh M (1993) A model for dynamic airline seat inventory control with multiple seat bookings.
 927 *Transp Sci* 27(3):252–265
- 928 Macario A, Vitez T, Dunn B, McDonald T (1995) Where are the costs in perioperative care? Analysis of
 929 hospital costs and charges for inpatient surgical care. *Anesthesiology* 83(6):1138–1144
- 930 McManus M, Long M, Cooper A, Mandell J, Berwick D, Pagano M, Litvak E (2003) Variability in surgical
 931 caseload and access to intensive care services. *Anesthesiology* 98(6):1491–1496
- 932 Ozkarahan I (2000) Allocation of surgeries to operating rooms by goal programming. *J Med Syst* 24(6):339–
 933 378
- 934 Pham D, Klinkert A (2008) Surgical case scheduling as a generalized job shop scheduling problem. *Eur J*
 935 *Oper Res* 185:1011–1025
- 936 Santibañez P, Begun M, Atkins D (2007) Surgical block scheduling in a system of hospitals: an application
 937 to resource and wait list management in a British Columbia health authority. *Health Care Manag Sci*
 938 10:269–282
- 939 Schütz H-J, Kolisch R (2010a) Approximate dynamic programming for capacity allocation in the service
 940 industry. Working paper, available at SSRN: <http://ssrn.com/abstract=1618315>
- 941 Schütz H-J, Kolisch R (2010b) Capacity allocation for demand of different customer–product–combina-
 942 tions with cancellation, no–shows, and overbooking when there is a sequential delivery of service.
 943 Working paper, available at SSRN: <http://ssrn.com/abstract=1618313>
- 944 Subramanian J, Stidham S Jr, Lautenbacher CJ (1999) Airline yield management with overbooking, can-
 945 cellations, and no-shows. *Transp Sci* 33(2):147–167
- 946 Testi A, Tanfani E, Torre G (2007) A three-phase approach for operating theatre schedules. *Health Care*
 947 *Manag Sci* 10:163–172
- 948 van Oostrum JM, van Houdenhoven M, Hurink JL, Hans EW, Wullink G, Kazemier G (2008) A mas-
 949 ter surgical scheduling approach for cyclic scheduling in operating room departments. *OR Spectr*
 950 30(2):355–374