

Scheduling, Revenue Management, and Fairness in an Academic-Hospital Radiology Division

Richard Baum, MD, Dimitris Bertsimas, PhD, Nathan Kallus

Rationale and Objectives: Physician staff of academic hospitals today practice in several geographic locations including their main hospital. This is referred to as the extended campus. With extended campuses expanding, the growing complexity of a single division's schedule means that a naive approach to scheduling compromises revenue. Moreover, it may provide an unfair allocation of individual revenue, desirable or burdensome assignments, and the extent to which the preferences of each individual are met. This has adverse consequences on incentivization and employee satisfaction and is simply against business policy.

Materials and Methods: We identify the daily scheduling of physicians in this context as an operational problem that incorporates scheduling, revenue management, and fairness. Noting previous success of operations research and optimization in each of these disciplines, we propose a simple unified optimization formulation of this scheduling problem using mixed-integer optimization.

Results: Through a study of implementing the approach at the Division of Angiography and Interventional Radiology at the Brigham and Women's Hospital, which is directed by one of the authors, we exemplify the flexibility of the model to adapt to specific applications, the tractability of solving the model in practical settings, and the significant impact of the approach, most notably in increasing revenue by 8.2% over previous operating revenue while adhering strictly to a codified fairness and objectivity.

Conclusions: We found that the investment in implementing such a system is far outweighed by the large potential revenue increase and the other benefits outlined.

Key Words: Revenue management; personnel scheduling; fairness; mixed-integer optimization.

©AUR, 2014

With the introduction of the extended medical campus to today's hospitals, the complexity of physician schedules has dramatically increased as department members who have historically been confined to a single institution and location are now routinely spread across an entire region. At the same time, there are important concerns associated with assigning a physician to a daily assignment. The average revenue yielded from assigning one physician to an assignment may differ somewhat from assigning another physician to the same or different assignment. This affects both the departmental revenue and the physician's own revenue, and therefore, there are immediate concerns for fairness in the income of the department's physicians. It also means that a naively optimal schedule that maximizes total

revenue without considering fairness will also enlarge the disparities in individual revenues and not offer equal opportunities to practice at varied locations. Other concerns include individual site and practice preferences, time-off requests, physician effort and fatigue, and the distribution of on-call assignments. These entail both strict limitations on possible schedules and softer fairness concerns in accommodating all the physicians equally.

Thus, deciding whom to place where in an extended medical campus is a complex exercise that takes many considerations into account. Failure to successfully mitigate all of these concerns will on the one hand overexert the physicians and create unfairness and on the other generate poor revenue for the department. Inequity, in turn, entails disincentivization, lowers employee satisfaction, and is simply unacceptable based on business principles. Where whole departments are worth in the hundreds of millions of dollars of revenue per year to a hospital, missing out on a few percentage points due to improper scheduling amounts to a large loss.

For these reasons, this is a problem of considerable and growing importance for many medical departments where the extended medical campus is expanding. At the same time, this significance is often overlooked, and this scheduling problem has not been addressed from optimal scheduling, revenue management, or quantitative fairness perspectives,

Acad Radiol 2014; 21:1322–1330

Division of Angiography and Interventional Radiology, Brigham and Women's Hospital, Boston, Massachusetts (R.B.); Operations Research Center, Massachusetts Institute of Technology, Cambridge, Massachusetts (D.B., N.K.); and Sloan School of Management, Massachusetts Institute of Technology, One Amherst St, Rm E40-147, Cambridge, MA 02478 (D.B.). Received January 9, 2014; accepted May 20, 2014. Funding Sources: This material is based upon work supported by the National Science Foundation Graduate Research Fellowship under grant number 1122374. **Address correspondence to:** D.B. e-mail: dbertsim@mit.edu

©AUR, 2014

<http://dx.doi.org/10.1016/j.acra.2014.05.009>

traditionally being done somewhat informally by hand, without clear quantifiable objectives and with poorer results as the task increased in sophistication over time.

One of the authors (R.B.) is the Herbert L. Abrams Director of the Division of Angiography and Interventional Radiology at the Brigham and Women's Hospital, Boston, Massachusetts (mentioned as BWH IR henceforth) and has the responsibility of managing the physician staff and the overall revenue of the division. In particular, he is responsible for the daily scheduling of the eight physicians in the division, himself included. With the academic responsibilities of a teaching hospital affiliated with Harvard Medical School and an extended campus that includes three external clinics in Boston suburbs in addition to the facilities at the hospital, BWH IR faces a formidable scheduling task with great implications for the physician staff, their revenue, and the departmental revenue.

The present article is motivated by the need of studying these tasks and their ramifications on revenue and human resources from an operational and quantitative point of view. To our knowledge, this is a novel consideration of the problem in the literature, incorporating aspects of disciplines in which operations research has a proven track record of success, including scheduling, revenue management, and quantifying fairness. As a complex operational planning task with direct impact on revenue and fairness, scheduling the daily assignments of the physician staff of a medical department could benefit greatly from operations research.

RELEVANT LITERATURE

Personnel scheduling is a well-studied problem in the literature and one in which operations research, and particularly optimization, has offered success. There are a myriad of special problems, solutions, and studies. We refer the interested reader to Ernst et al. (10) for an exhaustive review and annotated bibliography of studies of personnel scheduling. There is in particular a significant amount of work on scheduling nurses (9,10,14) and on scheduling physicians in an emergency room (2,7) that also focuses on lower level structures such as the timing of shifts. Notable examples of applying mixed-integer optimization (MIO) methods to personnel scheduling problems in other domains are Stojkovic et al. (16) and Caprara et al. (6).

In addition, there have been studies of fairness in scheduling, mostly in the context of computer process scheduling and scheduling transmissions over communication networks. A few notable examples are Hahne (11), Liua et al. (13), and Zaharia et al. (17). Fairness has also been studied in other operational applications ranging from travel arrangements for National Basketball Association games (1) to allocations of kidney transplantations (4).

Fairness has also been studied abstractly in the generic context of utility allocation by a central decision maker. In the study by Bertsimas et al. (3), the authors seek to bound the inefficiency of fairness, or the relative difference between

the social optimum and different fair allocations. For a thorough review of the philosophical and mathematical studies of fairness over the years, we refer the reader to the references in Bertsimas et al. (3) under section 3.

CONTRIBUTIONS AND STRUCTURE OF ARTICLE

We here argue that this overall scheduling and revenue management problem is amenable to solution by mathematical optimization. Quantifying the above considerations and nontangible partialities as variables and constraints using existing theory allows for a computable mathematical solution that not only generates feasible schedules but also chooses the most desired one in some sense, such as maximal revenue and minimal mismatch with individual scheduling preferences. We propose a simple unified optimization formulation using MIO that is flexible, efficiently solvable in practical settings, and impactful.

These qualities are exemplified through the motivating case for this study, BWH IR, where the optimization approach is implemented. In particular, the scheduling procedure is successfully implemented and incorporated into the workings of the division. Scheduling optimally in this way offered an increase of 8.2% in projected departmental revenue, amounting to a considerable improvement given the significant revenue already generated by the division. The schedule adheres strictly to a codified fairness, and we observe empirically an improvement in various measures of inequity. We find that compared to the initial fixed cost of implementing such a scheduling system including the costs of solver software, the clear revenue gains are far greater, not to mention the potential gains in fairness and employee incentivization and satisfaction.

The structure of the article is as follows. In the next section, we present the problem as motivated by the case of BWH IR. We then proceed to review the Materials and Methods used in addressing it. In Fairness Considerations we discuss the fairness considerations in the problem and discuss the mathematization of fairness and its implications. In The Extended Campus at the Division, we describe the structure of the extended campus at BWH IR as it pertains to the scheduling problem. In Data and Inputs, we discuss the data we use in the case of BWH IR. In Mathematical Formulation section, we present the generic unified formulation of the problem as an MIO. In Implementation at BWH IR, we discuss the implementation of the proposed methodology to the specific case of BWH IR, discussing both adapting the optimization model to this case and also practical considerations of actually carrying out the implementation. We then discuss the results we observed in this application in Results. Finally, we offer some concluding remarks in Discussion.

THE EXTENDED CAMPUS AT THE DIVISION

The Department of Radiology at BWH consists of 176 physicians practicing in nine locations including the main

campus in the Longwood medical center of Boston (8). Within it is the Division of Angiography and Interventional Radiology, responsible for all vascular and interventional radiologic procedures including related imaging and clinical work. Committed to teaching and research, the division also undertakes all related teaching responsibilities and is an active participant in the BWH Biomedical Research Institute. The division employs eight physicians, including its director. The division is responsible for a significant yearly revenue for the hospital. This has great ramifications both for the hospital and the division and for the individual physicians.

Schedules are planned on a quarterly basis to accommodate changing demands at various practicing locations. Fairness and objectivity is a very conscious concern. In fact this is one of the main reasons that the schedules of each physician are available to the whole staff. The natural assignment cycle at BWH IR is a 2-week one, where every such block should resemble more or less the others. We generically refer to the length of such a scheduling cycle as a scheduling block, which could be shorter or longer in different applications. To maintain a high standard of care and in consideration of staff, a physician may not be consecutively overexerted. Some assignments are naturally more taxing than others.

The eight physicians at BWH IR have many responsibilities. The hospital's four operating rooms must be manned by two physicians, 5 days a week. Each physician must hold a clinic at the hospital at least as often as semiweekly to see his or her own patients and schedule operations. The physician must then also have sufficient time at the operating room to treat his or her own patients. The clinic itself brings in very little revenue but is critical to the functioning of the department and indirectly brings in revenue through operations, which are only scheduled through a clinical visit. A day in the operating room is more taxing than a day in the clinic.

Physicians are also committed to academic excellence as members of the Harvard Medical School. At least 1 day weekly is allotted for academic activities for each of the physicians. Because academic days leave the physician with a more flexible schedule, there are many personal preferences regarding the day of the week on which one is assigned to academic duties. It is also less taxing than other activities. Other academic obligations include scheduled absence due to conference attendance.

In any one work day and over the weekend, there must be one physician (and a physician's assistant) on call to attend to emergencies. Being on call is highly taxing and undesirable—even more so over the weekend—and carries important fairness concerns. Should a physician feel he or she is on call more than another, he or she would be highly dissatisfied and rightly so. Being on call at nights and over the weekend means less recovery from fatigue and this too must be considered in scheduling.

At times that all assignments are filled or assigning is otherwise infeasible, a physician is usually assigned to "float" meaning that he or she can work wherever needed during the day.

There are also days in which administrative tasks are required of some of the physician staff.

In addition to responsibilities at the Hospital, the physicians treat patients at three Vascular and Vein Care Centers in Newton Corner, East Bridgewater, and Foxborough, MA. These clinics are a major source of revenue, but because here revenue is greatest, it is also the most variable. Because of appointment scheduling concerns at the external offices, the days of the week at which these centers are open are fixed.

MATERIALS AND METHODS

In the following, we review the materials and methods we employ to address the problem. Methods include MIO and quantifications of fairness. Materials include data and the practical implementation of the MIO model.

Fairness Considerations

As discussed, fairness in revenue is a central concern in scheduling the physicians. At BWH IR, revenue discrepancies between physicians at the same assignment stem mostly from varying levels of productivity and individual referral-based practice development.

It is important the physicians feel that they are being treated fairly as it is important to incentivize them to succeed and generate revenue. Besides fairness in scheduling, further incentivization is attained through a novel revenue structure that is unique to the division and was not reported on previously. Of the revenue generated by each physician that is allotted for compensation, a percentage is given to that physician as compensation and the rest is combined in a common pool. The common pool is divided equally among the physicians as additional compensation. This shifts the incentives of the individual physician toward departmental success and teamwork while still encouraging individual performance, and by having the direct effect of shrinking the spread of the distribution of individual incomes, it improves fairness in the most direct of ways.

It is also important to treat revenue fairness through scheduling to further reduce this spread and to ensure that no physician is left with all the low-revenue assignments and each has the opportunity to build up their practice. Because some physicians generate more revenue than others at certain assignments, maximizing departmental revenue naively does not overlap with providing equal opportunities to the physicians to practice and to generate revenue. Therefore, a more refined approach to the optimization must be taken.

In this study, we formulate fairness as described by Rawls (15) where we maximize the prospects of the least well-off. This is also known as *maximin* fairness in optimization contexts. Under this choice, an assignment of utilities is deemed fair if it maximizes the utility of the party that is assigned least utility. If two assignments are both equally fair in this sense but one has greater total utility assigned to all parties, clearly that one is preferable and will be the one we choose.

Another important fairness concern is in accommodating personal scheduling preferences. Every physician has his or her own scheduling preferences. In making a schedule, these should be met whenever possible, but physicians should also be accommodated equally. We again take the Rawlsian approach and model unfairness as the worst level of accommodation among physicians. By limiting the worst level, we ensure a fair control on all levels.

In addition to these, there are fairness concerns with the most desirable and undesirable assignments. For example, the undesirable on-call assignments need to be evenly assigned among physicians. On the other hand, the float assignment is a desirable one and must also be assigned evenly, considered separately. These assignments must be evenly spread among the staff and are not exchangeable to revenue or to one another. Therefore, here we take the approach of limiting the spread of number of assignments. For each highly desirable or undesirable assignment, the difference between the maximal number of such assignments made to any one physician and the minimal number is constrained to be small.

Finally, besides the constraints on a feasible schedule, to be fair, the scheduling process must be objective. Of course, being a computer-automated process for which an individual is nothing more than the relevant data, the optimization model is an objective scheduler free of human prejudice and favoritism.

Mixed-integer Optimization

MIO is a very general and powerful modeling framework for decision problems that involve both discrete and continuous decisions. Abstractly, for n binary decisions x_1, \dots, x_n and k continuous decisions y_1, \dots, y_k , the general binary MIO model can be written as:

$$\begin{aligned}
 &\text{minimize } \sum_{i=1}^n c_i x_i + \sum_{i=1}^k d_i y_i \\
 &\text{such that } x_i \in \{0, 1\} \quad i = 1, \dots, n \\
 &\quad y_i \geq 0 \quad i = 1, \dots, k \\
 &\quad \sum_{i=1}^n a_{ji} x_i + \sum_{i=1}^k b_{ji} y_i \leq g_j \quad j = 1, \dots, m,
 \end{aligned} \tag{MIO}$$

where c_i, d_i, a_{ji}, b_{ji} , and g_j are given data. The model (MIO) subsumes many decision problems. As a concrete example, consider the problem of forming a team of at most p professionals to complete N operations or tasks as part of a project. The potential pool of team members is numbered $1, \dots, n$. For each, data are provided for the daily rate they charge for work (d_i , \$/day), the daily rate at which they complete operations (r_i , operations/day), and the minimal (g_i) and maximal (G_i) amount of days they wish to work if hired for the project. This problem involves binary decisions, $x_i = 1$ if i is hired and 0 otherwise, as well as continuous decisions, y_i the (possibly fractional) number of days for which i is hired to work if hired. The problem of finding the team to complete the project with the least expenditures on salary is directly representable in the form of (MIO):

$$\begin{aligned}
 &\text{minimize } \sum_{i=1}^k d_i y_i \\
 &\text{such that } x_i \in \{0, 1\}, \quad y_i \geq 0 \quad i = 1, \dots, n \\
 &\quad \sum_{i=1}^n x_i \leq p, \quad \sum_{i=1}^n (-r_i) y_i \leq -N \\
 &\quad y_i \leq G_i x_i, \quad -y_i \leq -g_i x_i \quad i = 1, \dots, n.
 \end{aligned}$$

We refer the interested reader to Bertsimas and Weismantel (5) for a thorough discussion of MIO, its theory, solution methods, applications, and, in particular, for much more on how to model various types of problems, constraints, and objectives using MIO.

Importantly, there are powerful commercial software packages almost solely dedicated to the solution of problems of the form of (MIO), including IBM CPLEX (used here), Gurobi, XPRESS, Mosek, and Microsoft Solver, which is often bundled with and can be called from the popular Microsoft Excel. Open-source solutions include GNU LPK, SCIP, and Coin-OR CBC.

We will formulate our scheduling problem in the form of (MIO) and use IBM CPLEX along with its modeling language OPL to implement and solve it.

Data and Inputs

The optimization model takes as input various data and tuning parameters. In particular, we design the model around the data that are available. For our study, BWH IR provided raw data regarding historical revenue of different physicians at different assignments. We also surveyed physicians' scheduling preferences. Moreover feedback regarding whether a schedule conformed to their practices allowed for tuning of parameters.

BWH IR provided both historical daily schedules and daily revenues for each physician for one financial year. With hardly any variability in revenue with the time of day, we model a physician's expected revenue at an assignment as a weighted average of his or her historical revenue at that assignment, assigning more weight to more recent revenue. Although this served well for the purposes of BWH IR, it is simple to incorporate additional structure such as day-of-week dependence on the revenue where needed.

To understand physicians' individual scheduling preferences, each was asked to fill in their own ideal one-block schedule. The participation of the physicians in this process also reinforced trust in the scheduling procedure as the results resembled the inputs of the physicians.

Parameters on structural constraints of the schedule were determined in two ways. Some of the constraints are physically describable, and their parameters were provided by discussion on what constitutes a feasible schedule. For example, the number of staff needed for each assignment on each day—such as at least two physicians in the four operating rooms on each weekday—and the number of slots available for each assignment are quantities that are known and fixed. In addition, some assignments are previously fixed such as vacation and conferences. Other examples include identifying

which are the assignments that must be distributed evenly (such as on-call duty) and how large their spread can be and what are the assignments to which each physician must be assigned at least a certain number of times each block (such as academic duties).

Some of the constraints were not immediately quantifiable and had to be tuned. An example is in the exertion limits, which, although clearly necessary, are less palpable. In this case, a manager can determine whether a schedule is feasible but clear quantifiable rules are difficult to determine without experiment. For this purpose, we provide as tuning parameters effort ratings on an arbitrary scale for each assignment and tuned limits on effort over any two consecutive days and on effort over any one block. An initial guess is made by the manager, and then, the values are tuned so to produce schedules that are acceptable in their effort requirements as deemed by the manager. Another tuning parameter is the exchange rate between discrepancy in meeting preferences and revenue.

We summarize and provide mathematical notation for the inputs in [Appendix A](#).

Mathematical Formulation

We now present the formulation of the problem as an MIO. We discuss briefly the decision variables and objective function but defer the full formulation with a discussion of the constraints to [Appendix B](#).

Given a set S of staff members to be scheduled, a set A of daily assignments, the blocks to be scheduled $B = \{1, \dots, B\}$, and block days $D = \{1, \dots, D\}$, we introduce the following decision variables for each scheduling block b , block day d , assignment a , and staff member s :

$$x_{bdas} = \begin{cases} 1, & \text{if } s \text{ is assigned to activity } a \text{ on day } d \text{ of block } b, \\ 0, & \text{otherwise.} \end{cases}$$

Given data R_{as} , the expected revenue in dollars generated when assigning staff member s to assignment a , and P_{das} , equal to 1 if assignment a is in staff member s 's preferred block schedule on day d and otherwise 0, we introduce two key quantities as functions of the decision variables:

- The per-day average revenue of staff member s under schedule x :

$$R_s = \frac{1}{BD} \sum_{b=1}^B \sum_{d=1}^D \sum_{a \in A} R_{as} x_{bdas}.$$

- The average number of assignment preferences unmet for staff member s in block b under schedule x :

$$P_{bs} = \frac{1}{D} \sum_{d=1}^D \sum_{a \in A} \max\{P_{das} - x_{bdas}, 0\}.$$

Our primary objective is the fairness objective: to maximize the least individual revenue of the staff, $\min_{s \in S} R_s$, and to minimize the block average of the largest discrepancies in meeting scheduling preferences by staff member, $\frac{1}{B} \sum_{b=1}^B \max_{s \in S} P_{bs}$. We combine these two with a trade-off parameter λ :

$$\min_{s \in S} R_s - \frac{\lambda}{B} \sum_{b=1}^B \max_{s \in S} P_{bs}.$$

To ensure that among fair schedules we choose an efficient one, we also consider an efficiency objective: to maximize total revenue, $\sum_{s \in S} R_s$, and to minimize all scheduling preference discrepancies, $\frac{1}{B} \sum_{b=1}^B \sum_{s \in S} P_{bs}$. We combine this with the fairness objective using positive trade-off parameters ϵ_1 , for the total revenue, and ϵ_2 , for the total sum of discrepancies.

The overall objective we consider is

$$\begin{aligned} &\text{maximize} \left(\min_{s \in S} R_s + \epsilon_1 \sum_{s \in S} R_s - \frac{\lambda}{B} \sum_{b=1}^B \right. \\ &\quad \left. \times \left(\max_{s \in S} P_{bs} + \epsilon_2 \sum_{s \in S} P_{bs} \right) \right) \end{aligned}$$

In our modeling framework, we give priority to the fairness objective and therefore choose very small values for ϵ_1 and ϵ_2 . This choice ensures that fairness, as we define it, is maximized and that among all equally fair allocations, we select the one that leads to the highest overall revenue and the least total scheduling preference discrepancy. The existence of positive, but small, ϵ_1 and ϵ_2 that ensure exactly this is a consequence of Iancu and Trichakis (12).

In the next section, we discuss the implementation of this optimization model in the case of BWH IR.

Implementation at BWH IR

We implement the scheduling as an optimization over a quarterly horizon. At the beginning of each fiscal quarter, the model is solved to (near) optimality with conference dates and time-off requests fixed by the manager and are not allowed to vary. Should any fixed assignments change, such as vacation or conferences being scheduled, the model is re-optimized with the part of the schedule that is already past set as fixed as well. Note that the quarterly horizon for scheduling also restricts the horizon and memory of fairness considerations. We add some application-specific constraints to the generic model, which we review in [Appendix C](#). The ability to introduce application-specific constraints exemplifies the model's flexibility.

A quarterly MIO problem has approximately 160,000 variables and 140,000 constraints and is solved to high precision in well less than 10 minutes on a personal computer using IBM ILOG CPLEX 12.3. The running time, in fact, as long as it is less than 1 day, is really quite irrelevant because of the infrequency of solving the optimization model.

Sep 3 Mon		Sep 4 Tue		Sep 5 Wed		Sep 6 Thu		Sep 7 Fri		Sep 8 Sat	Sep 9 Sun
#1	Albert	#1	Emma	#1	Herbert	#1	Carrol	#1	Daniel	OnCallWkd	Emma
#2	Barbara	#2	Albert	#2	Gillian	#2	Herbert	#2	Herbert	-	Frank
Clinic	Carrol	Float	Frank	Clinic	Frank	Clinic	Daniel	Clinic	Barbara	-	Barbara
Float	Daniel	Academic	Gillian	Float	Emma	Float	Frank	Academic	Frank	-	Daniel
Academic	Emma	Academic	Herbert	Float	Carrol	Academic	Barbara	Academic	Carrol	-	Albert
NCIR	Frank	NCIR	Carrol	Academic	Daniel	Academic	Albert	Academic	Gillian	-	Carrol
EBIR	Gillian	EBIR	Daniel	NCIR	Albert	NCIR	Emma	NCIR	Emma	-	Gillian
FBIR	Herbert	FBIR	Barbara	EBIR	Barbara	EBIR	Gillian	FBIR	Albert	-	Herbert
OnCall	Herbert	OnCall	Frank	OnCall	Barbara	OnCall	Carrol	OnCall	Emma		
Sep 10 Mon		Sep 11 Tue		Sep 12 Wed		Sep 13 Thu		Sep 14 Fri		Sep 15 Sat	Sep 16 Sun
#1	Daniel	#1	Barbara	#1	Carrol	#1	Barbara	#1	Emma	OnCallWkd	Carrol
#2	Barbara	#2	Emma	#2	Herbert	#2	Carrol	#2	Daniel	-	Frank
Clinic	Albert	Float	Frank	Clinic	Emma	Clinic	Gillian	Clinic	Herbert	-	Emma
Academic	Emma	Academic	Albert	Float	Frank	Float	Daniel	Float	Albert	-	Barbara
Academic	Carrol	Academic	Herbert	Float	Gillian	Float	Albert	Academic	Frank	-	Daniel
NCIR	Frank	NCIR	Carrol	Academic	Daniel	Admin	Frank	Academic	Barbara	-	Albert
EBIR	Gillian	EBIR	Daniel	NCIR	Albert	NCIR	Emma	NCIR	Carrol	-	Gillian
FBIR	Herbert	FBIR	Gillian	EBIR	Barbara	EBIR	Herbert	FBIR	Gillian	-	Herbert
OnCall	Albert	OnCall	Daniel	OnCall	Barbara	OnCall	Herbert	OnCall	Carrol		

Figure 1. An example of a fragment of the Web-based interface displaying one block optimally scheduled. The physicians' names are obscured using alphabetical pseudonyms. NCIR, EBIR, and FBIR refer to the external clinics in Newton Corner, East Bridgewater, and Foxborough, MA, respectively. #1 and #2 refer to the operating room assignments. The schedule display is designed to resemble the format of existing schedules and share its idiosyncrasies.

TABLE 1. Per-Physician Changes in the Optimized Schedule

Physicians (Pseudonyms)	Change in Revenue, %	Change in Revenue Rank	Preferences Accommodated (Old Schedule), %	Preferences Accommodated (New Schedule), %
Albert	+19.10	0	56.25	75.00
Barbara	+0.23	-1	58.33	79.12
Carroll	+17.41	0	33.33	77.08
Daniel	-1.57	0	47.92	81.25
Emma	+5.04	0	50.00	75.00
Frank	+9.74	+1	47.92	77.08
Gillian	+15.74	0	43.75	77.08
Herbert	+9.14	0	33.33	75.00

The solving software is connected with a relational database backend and a Web-based interface. Department administration need only interact with it to provide such parameters as which assignments need to be fixed and ask whether the solver optimizes the model given these. An example of a resulting scheduled block as it is displayed on the Web-based interface is presented in Figure 1. The database is updated with real revenues as the schedule is executed, and these are in turn used in the weighted average of next quarter's optimization, creating a closed-loop adaptive control system.

RESULTS

The most significant result of optimizing the schedule was a projected 8.2% increase in revenue department wide. This amounts to a sizable increase in revenue yearly. At the same time, the overall structure of the optimized schedule was similar to the schedule the department staff members were used to. This similarity, of course, is a result of the structural

constraints on the schedule with parameters that were tuned by administration and staff to produce such schedules exactly.

Individual revenues were also increased (with one exception) and more preferences were accommodated. We report the per-physician changes in Table 1. As can be seen, the overall ordering of revenues was hardly affected, and nearly all revenues increased substantially. Preference accommodation was both increased and made more equal.

Scheduling fairness, which before was an informal and perhaps uncomfortable issue, was now codified and executed by an objective actor. Besides the 8.2% increase in total revenue, we also see a substantive increase in the minimum revenue and a reduction in the overall spread of revenues. We report these summary statistics in Table 2. We report similar statistics for the fractions of preferences accommodated per physician, which exhibit a marked improvement. Moreover whereas some hand-made schedules were not strictly fairness feasible with regard to spread of (un)desirable assignment, the optimized schedule was necessarily fairness feasible.

TABLE 2. Summary Statistics for Changes in the Optimized Schedule

Increase in total (or average) revenue	8.20%
Reduction in coefficient of variation of revenue *	12.87%
Increase in minimum revenue *	5.04%
Increase in total (or average) preference accommodation	66.29%
Reduction in coefficient of variation of preference accommodation	85.60%
Increase in minimum preference accommodation	125.00%

The collections revenue of the division head is omitted in statistics marked with * for both the old and new schedules, because it is necessarily a low outlier due to additional administrative responsibilities.

Where before scheduling was a task that consumed some days of administration work each month, the optimized schedule is given by a computer-automated process that completes in less than 10 minutes, produces familiar-looking schedules, enforces fairness among physicians in revenue and preference accommodation, all while increasing departmental revenue considerably.

DISCUSSION

In this article, we addressed the problem of determining daily schedules for physicians in a hospital division. The problem is phrased in terms not before considered in the literature to our knowledge, incorporating aspects of scheduling, revenue management, and fairness. All of these are areas where operations research has a proven track record of success both in studying theoretical underpinnings and in practical applications. After describing the problem, we propose a solution by a unified MIO formulation.

The formulation is flexible enough to accommodate additional case-specific constraints, and its solution is computable in a reasonable amount of time in practical settings. In the case of BWH IR, which was the motivating case for the study herein, revenue increased, fairness was codified and enforced, and, past the implementation stage, the scheduling process is made easy.

The method is automated and can be used without a great deal of specialized knowledge. With a relational database-backed system with a Web-based interface, an academic medical department may easily use this scheduling process and reap its benefits. We therefore find that the investment in implementing such a system is far outweighed by the large potential revenue increase and other benefits outlined here.

REFERENCES

1. Bean JC, Birge JR. Reducing travelling costs and player fatigue in the National Basketball Association. *Interfaces* 1980; 10(3):98–102.
2. Beaulieu H, Ferland J, Gendron B, et al. A mathematical programming approach for scheduling physicians in the emergency room. *Health Care Management Science* 2000; 3:193–200.
3. Bertsimas D, Farias V, Trichakis N. The price of fairness. *Operations Research* 2011; 59(1):17–31.
4. Bertsimas D, Farias V, Trichakis N. Fairness, efficiency and flexibility in organ allocation for kidney transplantation. *Operations Research* 2011; 61(1):73–87.
5. Bertsimas D, Weismantel R. *Optimization over integers*. Belmont, MA: Dynamic Ideas, 2005.
6. Caprara A, Fischetti M, Toth P, et al. Algorithms for railway crew management. *Mathematical Programming* 1997; 79(1-3):125–141.
7. Carter M, Lapierre S. Scheduling emergency room physicians. *Health Care Management Science* 2001; 4:347–360.
8. Department of Radiology. Available at: www.brighamandwomens.org/Departments_and_Services/radiology. Accessed June 12, 2013.
9. Dowsland K. Nurse scheduling with tabu search and strategic oscillation. *European Journal of Operational Research* 1998; 106(2):393–407.
10. Ernst AT, Jiang H, Krishnamoorthy M, et al. An annotated bibliography of personnel scheduling and rostering. *Annals of Operations Research* 2004; 127(1-4):21–144.
11. Hahne EL. Round-robin scheduling for max-min fairness in data networks. *IEEE Journal on Selected Areas in Communications* 1991; 9(7):1024–1039.
12. Iancu D, Trichakis N. Pareto efficiency in robust optimization. *Management Science* 2013. Articles in Advance.
13. Liu X, Chong EKP, Shroff NB. A framework for opportunistic scheduling in wireless networks. *Computer Networks* 2003; 41:451–474.
14. Miller H, Pierskalla W, Rath G. Nurse scheduling using mathematical programming. *Operations Research* 1976; 24(5):857–870.
15. Rawls J. *A theory of justice*. Cambridge, MA: Harvard University Press, 1971.
16. Stojković M, Soumis D, Desrosiers J. The operational airline crew scheduling problem. *Transportation Science* 1998; 32(3):232–245.
17. Zaharia M, Borthakur D, Sarma JS, et al. Delay scheduling: a simple technique for achieving locality and fairness in cluster scheduling. In: *Proceedings of the 5th European conference on Computer systems*; 2010; 265–278.

APPENDIX A. INPUTS AND NOTATION

Here, we summarize and provide mathematical notation for the inputs and data discussed in Data and Inputs section.

- S : The set of staff members to be scheduled.
- A : The set of distinct daily assignments.
- $B = \{1, \dots, B\}$: The range and number of blocks to be scheduled.
- $D = \{1, \dots, D\}$: The range and number of days in a schedule block.
- $X \subset 2^A$: Assignment exclusivity groups. A collection of subsets of assignments in which each subset represents assignments that cannot be scheduled at the same time. For example, vacation and operating room are exclusive but operating room and on call are not.
- $F_{bdas} \in \{0, 1\}$: If assignment a is fixed for staff member s on day d of block b then this takes value 1, otherwise 0.
- $P_{das} \in \{0, 1\}$: If assignment a is in staff member s 's preferred block schedule on block day d then this takes value 1, otherwise 0.
- R_{as} : The expected revenue in dollars generated when assigning staff member s to assignment a .
- $N_{ad}^{\text{Day}}, C_{ad}^{\text{Day}}$: The daily necessity (possibly 0) and capacity (possibly infinite if no capacity constraint) of staff for assignment a on block day d .
- $N_a^{\text{Block}}, C_a^{\text{Block}}$: The minimal and maximal number of assignments to a for each staff member in every block.
- T_a : The maximal spread among staff for assignment a . Only finite for those highly desirable and undesirable assignments that must be spread equally.
- E_a : The effort level rating of assignment a .
- $L^{\text{Consecutive}}$: The limit on the total effort over any two consecutive days.
- L^{Block} : The limit on the total effort over any one block.
- λ : The exchange rate between discrepancy in meeting preferences and revenue.

APPENDIX B. FULL FORMULATION OF THE GENERIC PROBLEM

We present in this appendix the full formulation of the generic scheduling problem and then explain the significance of the various constraints. We refer readers to Bertsimas and Weismantel (5) for more information on modeling with and solving integer-optimization problems. We use the convention that given a block b and block day d then $d' = d + 1$ and $b' = b$ if $d < D$ and otherwise $d' = 1$ and $b' = \min(b + 1, B)$. That is, b' and d' represent the next day if there is one.

$$\max \left(\min_{s \in S} R_s + \epsilon_1 \sum_{s \in S} R_s - \frac{\lambda}{B} \sum_{b=1}^B \left(\max_{s \in S} P_{bs} + \epsilon_2 \sum_{s \in S} P_{bs} \right) \right)$$

$$\text{s.t. } x \in \{0, 1\}^{B \times D \times A \times S}$$

$$R_s = \frac{1}{BD} \sum_{b=1}^B \sum_{d=1}^D \sum_{a \in A} R_{as} x_{bdas} \quad \forall s \in S$$

$$P_{bs} \geq \frac{1}{D} \sum_{d=1}^D \sum_{a \in A} |x_{bdas} - P_{das}| \quad \forall b \in B, s \in S \quad (3)$$

$$x_{bdas} \geq F_{bdas} \quad \forall b \in B, d \in D, a \in A, s \in S \quad (4)$$

$$\sum_{a \in A} x_{bdas} \geq 1 \quad \forall b \in B, d \in D, s \in S \quad (5)$$

$$\sum_{a \in X} x_{bdas} \leq 1 \quad \forall X \in X, b \in B, d \in D, s \in S \quad (6)$$

$$N_{ad}^{\text{day}} \leq \sum_{s \in S} x_{bdas} \leq C_{ad}^{\text{day}} \quad \forall b \in B, d \in D, a \in A \quad (7)$$

$$N_a^{\text{block}} \leq \sum_{d=1}^D x_{bdas} \leq C_a^{\text{block}} \quad \forall b \in B, a \in A, s \in S \quad (8)$$

$$\max_{s \in S} \sum_{b=1}^B \sum_{d=1}^D x_{bdas} - \min_{s \in S} \sum_{b=1}^B \sum_{d=1}^D x_{bdas} \leq T_a \quad \forall a \in A \quad (9)$$

$$\sum_{a \in A} (x_{bdas} + x_{b'd'as}) E_a \leq L^{\text{Consecutive}} \quad \forall b \in B, d \in D, s \in S \quad (10)$$

$$\sum_{d=1}^D \sum_{a \in A} x_{bdas} E_a \leq L^{\text{Block}} \quad \forall b \in B, s \in S \quad (11)$$

Constraints (4–8) describe the logical restrictions on a feasible schedule. Constraint (4) incorporates assignments fixed by the scheduler. Constraint (5) incorporates assignments fixed by the scheduler. Constraint (5) constrains each staff member to be assigned to at least one assignment each day. We require this while having an assignment of “unscheduled” so that we may attach to it (possibly negative) effort ratings or even revenue if it makes sense for the application. Constraint (6) ensures that conflicting assignments are not given to a single staff member on a single day. For example, a staff member can only do one workday activity whether it be staffing the clinic, staffing the operation room, being on vacation, and so forth. These are all exclusive activities. But a staff member who is in the operating room can also be scheduled to be on call on the same day—those are nonexclusive activities. A staff member who is on vacation, however, cannot be scheduled to be on call—these too are exclusive.

Constraint (7) describes the daily necessities of assignments to keep the division operational and the capacities of the division to accommodate different activities. For example, the

operating rooms must be manned, so there is a necessity of two for the operating room assignment. On the other hand, no more than two physicians may be assigned to operating room and no more than one to the clinic. Constraint (8) describes what assignments are required of each staff member over a schedule block. For example, the each staff member is to have at least two days dedicated to academic duties per schedule block.

Constraint (9) describes the fair distribution of the highly desirable and undesirable assignments. As mentioned previously, such include academic and on call as examples of highly desirable and undesirable assignments, respectively.

Finally, constraints (10) and (11) describe the limits on exerting each staff member over any two consecutive days and over any block. The parameters involved are derived from tuning by experimentation.

In formulating the optimization problem explicitly as a MIO, modeling the nonlinear constraint (3) requires $BD|A||S|$ auxiliary continuous variables and twice as many additional constraints. Modeling the nonlinear constraint (9) requires $2|A|$ auxiliary continuous variables and $2|S||A|$ additional constraints.

APPENDIX C. ADDITIONAL CONSTRAINTS SPECIFIC TO THE APPLICATION AT (ANONYMIZED) IR

At BWH IR, the natural schedule cycle is 2 weeks. A quarterly schedule has 6.5 blocks, rounded up to 7 whole blocks. Therefore, we have $B = 7$ and $D = 14$, starting with Monday. There are eight staff physicians and about 15 distinct assignments. Conference dates and time-off requests are fixed by the manager and are not allowed to vary.

At BWH IR, the assignment to be on call over the weekend includes being on call Friday, Saturday, and Sunday. Therefore, we add the linear constraints that for each staff member,

for the on-call assignment, the value of x over those 3 days is the same:

$$\begin{aligned} x_{b, 5, \text{ On Call}, s} &= x_{b, 6, \text{ On Call}, s} \\ &= x_{b, 7, \text{ On Call}, s} \quad \forall b \in B, s \in S \end{aligned}$$

$$\begin{aligned} x_{b, 12, \text{ On Call}, s} &= x_{b, 13, \text{ On Call}, s} \\ &= x_{b, 14, \text{ On Call}, s} \quad \forall b \in B, s \in S \end{aligned}$$

Additionally as a business practice, BWH IR does not schedule a physician as on call over two consecutive week-ends. This too is a simple linear constraint where for each staff member, for the on-call assignment, the sum of x over consecutive weekends is bounded by one:

$$x_{b, 7, \text{ On Call}, s} + x_{b, 14, \text{ On Call}, s} \leq 1 \quad \forall b \in B, s \in S$$

$$x_{b, 14, \text{ On Call}, s} + x_{b+1, 7, \text{ On Call}, s} \leq 1 \quad \forall b \in B, s \in S$$

In addition, to meet the varying demands for each physician in the external clinics while somewhat fixing the day of the week that he or she is there to accommodate appointment scheduling there, we require that a certain minimal fraction of block days be assigned to specific physicians at different locations over the long term. For example, at a certain external clinic, on the second Tuesday of each block, at least one time over three blocks should be allotted to a certain physician. For this, we take as input $\mathcal{E} \subset \mathcal{D} \times \mathcal{A} \times \mathcal{S} \times \mathbb{N} \times \mathbb{N}$ and add the constraint

$$\sum_{b'=b}^{b+l} x_{b' das} \geq p, \quad \forall (d, a, s, l, p) \in \mathcal{E}, b \in \mathcal{B}$$

where l represents the horizon of the rotation and p/l the fraction of assignment a allotted to physician s on block days d .

These are all application-specific constraints that can easily be incorporated into the generic MIO model.