Interfaces

Publication details, including instructions for authors and subscription information:
http://pubsonline.informs.org

Neonatal Physician Scheduling at the University of Tennessee Medical Center

Melissa R. Bowers, Charles E. Noon, Wei Wu, J. Kirk Bass

Interfaces
Vol. 46, No. 2, March–April 2016, pp. 168–182
ISSN 0092-2102 (print) | ISSN 1526-551X (online)
http://dx.doi.org/10.1287/inte.2015.0839

Neonatal Physician Scheduling at the University of Tennessee Medical Center

Melissa R. Bowers, Charles E. Noon, Wei Wu
Department of Business Analytics and Statistics, The University of Tennessee, Knoxville, Tennessee 37996
{mrbowers@utk.edu, cnoon@utk.edu, wwu3@vols.utk.edu}

J. Kirk Bass
Regional Neonatal Associates, Knoxville, Tennessee 37920, jkbass@utmck.edu

A common issue faced by physician groups is how to schedule 24-7 coverage for hospital units such as an emergency department. The first step is to determine the shifts to be covered. The second step, assigning physicians to specific shifts, is complicated because shifts vary with respect to duration, day of week, time of day, and desirability. To ensure workload fairness, physician groups often create "equality" schedules in which they evenly divide shifts, by type, among physicians. This problem can be readily modeled and solved via optimization. This paper presents a novel approach that incorporates individual physician shift-type preference and seeks, for each physician, a schedule that is superior to his or her equality schedule. We formulate and solve the problem as a binary, mixed-integer program designed to maximize relative gains in individual physician preference. We describe the methodology and real-world implementation within a neonatal intensive care group, and present a hybrid version of the model that is capable of simultaneously accommodating physicians who prefer an equality schedule and physicians willing to deviate from the equality schedule in pursuit of a schedule that better fits their shift preferences. Increases in schedule preference for the latter group ranged from 6.3 to 8.5 percent.

Keywords: healthcare; integer programming applications; physician scheduling; personnel scheduling.
History: This paper was refereed.

Acute-care medicine is characterized by the need to provide physician coverage sufficient to meet the staffing needs of a specific medical unit. For example, an emergency department may contract with a physician group to provide specified levels of physician staffing around the clock. The physician group must then decide how to provide those staffing levels, typically identifying a collection of shifts which, taken together, will provide the desired staffing levels. The final step is to assign individual physicians to individual shifts over some time horizon. The problem of assigning physicians to shifts is characterized by hard constraints or work rules and work patterns, which are based on quality-of-care and safety issues, and soft constraints, which are based on workload and lifestyle choices.

A number of software products have been developed to address these issues in physician scheduling. Although the products have some differences in interface, platform, and specific features, they all essentially work by capturing the constraints (hard and soft) and then use heuristics to determine the individual schedules. To assign appropriate workloads, some of the packages require the user to specify how many of each type of shift each doctor should work. Some of the packages also have the ability to attach relative values (in the form of points or monetary values) to the various types of shifts and provide workload scores for a given scheduling period. Many physician groups, however, still construct schedules by hand. Often, this responsibility falls to one member of the group and can consume a considerable amount of time depending on the number of shifts requiring coverage and the number of physicians and their constraints.

In addition to creating schedules that are feasible with respect to the myriad of constraints, the issue of workload fairness presents a serious challenge for schedulers producing schedules either manually or by scheduling software. A workload-fair schedule is one that all members of the physician group consider equal in terms of apportioned work. For example, suppose that four physicians need to provide single-physician coverage for a four-week period, and that each day requires coverage for two 12-hour shifts (a day shift and a night shift). Over the course of the four weeks, the four physicians will need to be assigned to the 28 day shifts and 28 night shifts. If each physician is assigned to work seven day shifts and seven night shifts, it is apparent that the schedule would be perceived as being fair. If, however, one physician is assigned to work all of his or her day shifts on Saturdays and Sundays, that would not be perceived as being fair. For this example, a schedule perceived as being fair would be one in which, over the course of the four weeks, each physician works one of each of the 14 possible shifts in a week (e.g., one Saturday day shift, one Saturday night shift, one Sunday day shift, one Sunday night shift, and so on). We refer to this form of a schedule as a pure *equality* schedule in that each physician is working an equal number of each type of shift.

An alternative to the equality schedule can be described using a simple example. Suppose two physicians (A and B) must provide single coverage for six day shifts and six night shifts. Each physician perceives a fair schedule as working three day and three night shifts (i.e., an equality schedule). Suppose both physicians prefer day shifts to night shifts; however, they differ with respect to the degree of preference. Physician A believes that working a night shift is equivalent to working 1.3 day shifts, whereas physician B equates a night shift to 1.6 day shifts. The equality schedule (three day shifts and three night shifts for each), based on preference scores, assigns to physician A 6.9 day-shift equivalents (computed as $3\times 1+3\times 1.3$) and to physician B, 7.8 day-shift equivalents. If, however, physician A is assigned five night shifts and physician B is assigned the remaining shifts (six day and one night), then the schedules would score as 6.5 day-shift equivalents for physician A and 7.6 for physician B. The resulting schedules would, therefore, be improvements for each physician compared to his or her respective equality schedule. This simple example illustrates the value of mutually beneficial swaps; however, such swaps are not likely to happen naturally in real-world situations. One reason is that physicians tend to trade only like shifts (i.e., a night shift for a night shift) when they need to be off from a particular shift. Hence, in moving from the equality schedule, the trade of three day shifts for two night shifts is not likely to occur. Furthermore, such a trade assumes each party knows his or her own relative preference and, more importantly, that this preference differs from that of the other individual.

This paper focuses on an approach for developing shift schedules that maintains fairness, yet incorporates individual relative preference for types of shifts. The goal is to create schedules through an objective process that can handle the computational complexity of the constraints while, at the same time, results in schedules that are better aligned with an individual physician's lifestyle or stage of life. Our work is more focused on the model and its implementation than on the computational algorithm for model solution. We describe the application of this method for a nine-physician neonatology group; but, we also test it on larger instances.

Literature Review

Staff scheduling, also known as rostering, is the prescriptive process of characterizing personnel demands over a time horizon and ultimately identifying individual worker schedules to meet that demand. The area has been heavily researched and the contributions in the literature generally consist of either new model formulations, new solution methods, novel real-world applications, or some combination thereof. Ernst et al. (2004b) and Van den Bergh et al. (2013) present reviews and classifications of the staff-scheduling and rostering literature. As a testament to the number of problem variations, applications, and solution methodologies, Ernst et al. (2004a) give a separate, annotated review of approximately 700 papers. Models differ on dimensions such as type of demand (e.g., task or shift based) or type of scheduling (e.g., individual or crew). Application areas include transportation systems, other types of service operations (including healthcare), and manufacturing. Given that most models have a discrete optimization component and many applications involve a large number of combinations, the solution methods run the full spectrum from enumeration to mathematical optimization to heuristics.

Within the application area of healthcare, the vast majority of papers tackle the problem of nurse scheduling. In contrast to physician scheduling, creating nurses' schedules is complicated by union rules, grades of nurses, and a large number of individuals needing to be scheduled. Bekkan (2010), Burke et al. (2004), Cheang et al. (2003), and De Causmaecker and Vanden Berghe (2010) give comprehensive reviews of nurse-scheduling and rostering literature. Erhard et al. (2015) provides a review of physician scheduling within hospital settings.

Scheduling and rostering models in healthcare can also be characterized based on whether the model treats a group of individuals as being interchangeable (e.g., emergency department physicians) or considers attributes and needs of distinct individuals (e.g., a specialty surgeon). With the latter, the data requirements, dimensionality, and complexity increase notably. The prior research in physician scheduling has variations in model specifics and solution approaches for the underlying combinatorial problem. Below, we review the most relevant works with a focus on methods for incorporating individual preference and fairness.

In Gunawan and Lau (2010, 2012), physician assignments in a government hospital are modeled as a bi-objective optimization problem in which the objectives are to (1) Maximize the number of ideal duty assignments, and (2) Minimize the number of unscheduled duties. The problem is then solved via a commercial solver using several formulations for handling the bi-objective aspects. The ideal assignments are represented by a proposed schedule defined by a weekday, a shift (morning or afternoon), and a duty. The model then allocates as many of the ideal assignments as possible, while maximizing the number of duties that are covered. Their approach is tested on a collection of randomly generated instances and a real case study from a surgical department. This approach incorporates individual physician preference, but only in a binary fashion in that an ideal assignment is given a value of 1 in the objective, and an assignment that is not ideal is essentially given a value of 0.

Gendreau et al. (2007), Beaulieu et al. (2000), Rousseau et al. (2002), Bourdais et al. (2003), and Carter and Lapierre (2001) include methods for incorporating individual physician requests for days off or days on, desired work patterns, and work hours per week. In these papers, the authors typically utilize a base optimization model and then introduce or relax as many of the preference constraints as possible to reach a feasible schedule. Several Canadian hospitals use these methods to schedule emergency department physicians.

Stolletz and Brunner (2012) imbed the notion of fairness by pursuing an even distribution of duties, especially for those shifts that are less desirable. The model tries to satisfy individual preferences for specific days off and/or desired shifts. The paper also explores the trade-offs between the dual objectives of pursuing cost minimization or fairness maximization. In the approach presented in Fügener et al. (2015), the same notion of fairness is applied, but across shift demands characterized by variation in the required levels of experience. The model presented in Baum et al. (2014) includes the notion of revenue fairness among physicians, while also trying to generate schedules that maximize overall revenue for the organization. Physician preference is incorporated by including each individual's ideal block schedule within the formulation and then trying to minimize the deviations from it.

Two *Interfaces* papers on physician scheduling provide valuable modeling and implementation insights. In the first paper (Cohn et al. 2009), the authors develop a methodology for creating full-year schedules for residents with various shift types across three hospitals. The method iterates between a progressively constrained optimization model and a review for feasibility and fairness by chief residents. The authors also provide some insightful lessons learned concerning the process of successfully applying operations research techniques. The second paper (Ferrand et al. 2011) describes a model and its application for creating a cyclic schedule for emergency physicians. The result is an eight-week schedule that is fair and ergonomically agreeable so that it can be repeated every eight weeks indefinitely. The results were well received, and provide further evidence that a physician group can converge on the concepts of fairness and work patterns.

In our review of prior research, the notion of individual preference applies mainly to the idea of incorporating individual requests for days off (or days on) or work patterns. We are not aware of any approaches that allow an individual's preference for specific shift types to be incorporated within a physician-rostering model designed to maximize relative gains in individual physician preference.

The Motivation and Problem Statement

Regional Neonatal Associates is a nine-physician group that provides in-house coverage for the 63-bed Neonatal Intensive Care Unit (NICU) at the University of Tennessee Medical Center in Knoxville, Tennessee. The group also provides call coverage for high-risk deliveries at two other hospitals in the Knoxville area. Several members of the group contacted the authors about methods for physician scheduling. Over the last several decades, one of the group's physicians manually constructed the group's schedules. Prior to that physician's retirement, the group had grown concerned about who would invest the necessary time and possess the expertise to create the schedules in the future. Additionally, the physicians perceived the resulting schedules varied with respect to workload, but that was difficult to establish because of the lack of a quantitative scheduling system. A meeting with management science faculty resulted in the formulation of the following problem statement: *Given a set of shifts that must be covered by a group of physicians and a collection of work rules, determine a schedule that is fair with respect to workload and, if possible, incorporate individual preferences for shift types such that each physician's schedule is better than or equal to a simple equal division of shift types.*

Shift Types and Relative Preference Scoring

The first step was to clearly identify the shifts requiring coverage and then group them into like shift types. The group provides 24-7 on-site coverage at the University of Tennessee Medical Center (UTMC); this coverage consists of two physicians during the weekday days from 8 AM to 5 PM and one physician on overnights and weekends. The group also provides 24-7 on-call coverage for the two off-site local hospitals. Although the off-site coverage typically involves relatively few call-ins during a week, it requires an exclusive physician assignment; that is, the physician assigned to cover the off-site hospitals cannot be assigned to any other shift during the week. In discussions with us, the physician group identified seven shift types (Table 1). The last column shows the number of each type of shift to be covered each week. For example, each weekday, two physicians cover UTMC (sides A and B of the unit); hence, 10 *WeekDay* day shifts must be covered per week.

In Table 1, the last shift type is considered a single shift although it requires the physician to work both Saturday and Sunday day shifts, which are six hours each. In total, 19 physician shifts must be covered each week. Within a shift type, individual shifts are considered equivalent; that is, a day shift on Tuesday is considered equivalent to a day shift on Friday. In contrast, a night shift on Monday is not considered equivalent to a night shift on Friday, because each is a different type of shift. Figures 1(a) and 1(b) illustrate the 19 shifts (and their indices for formulation purposes) and an example set of shift assignments. The assignments are feasible in that they violate no obvious constraints, such as a night-shift assignment immediately followed by a day-shift assignment, and they possess desirable qualities, such as continuity of patient care delivered when the same physician is assigned to cover two or more weekday shifts in a given week.

With nine physicians, an equality schedule would span a nine-week period. Within that nine-week period,

Shift-type name	Shift-type code	Description	Location	Days	Period	No. of shifts each week
WeekLong	T_7	Full-week call	Two off-site local hospitals	All week (Monday 8 AM thru Monday 8 AM)	24 hours per day	1
WeekDay	T_1	Weekday day	UTMC	Monday thru Friday	8 AM–5 PM	10
WeekNight	T_2	Mon-Thurs overnight	UTMC	Monday thru Thursday	4 PM–9 AM	4
FriNight	T_3	Friday overnight	UTMC	Friday	4 PM–10 AM	1
SatNight	T_5	Saturday overnight	UTMC	Saturday	12 PM–10 AM	1
SunNight	T_6	Sunday overnight	UTMC	Sunday	12 PM–9 AM	1
SatSunDay	T_4	Weekend days	UTMC	Saturday and Sunday	7 AM–1 PM	1

Table 1: The physician group identified seven shift types, and we assigned a code to each.

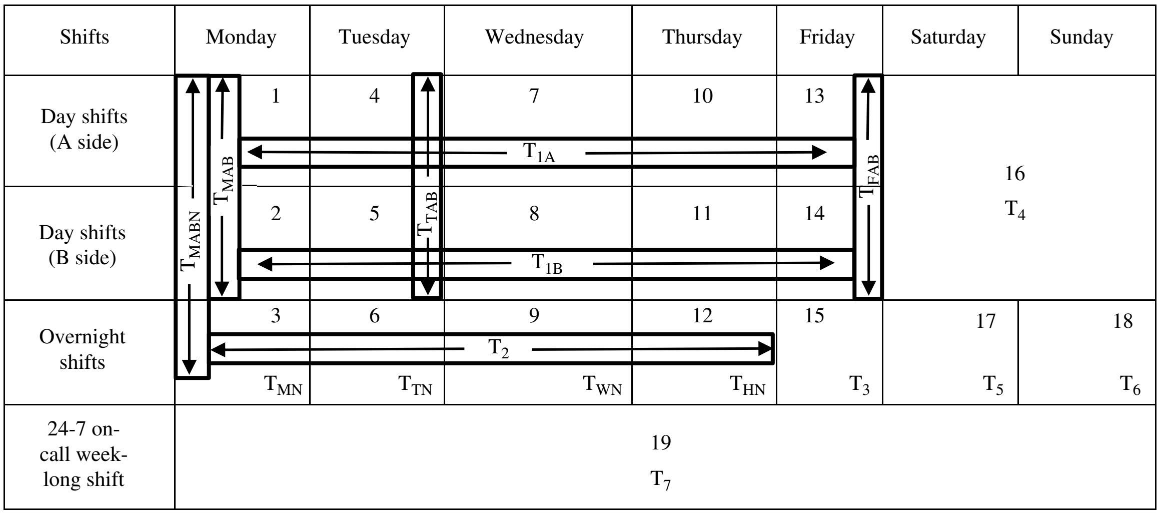

Shifts	Monday	Tuesday	Wednesday	Thursday	Friday	Saturday	Sunday
Day shifts (A side)	Doc2	Doc2	Doc4	Doc4	Doc4	Doc7	
Day shifts (B side)	Doc7	Doc1	Doc1	Doc7	Doc7		
Overnight shifts	Doc6	Doc5	Doc9	Doc6	Doc2	Doc5	Doc2
24-7 on-call week-long shift	Doc8						

Figure 1: (a) (top panel): Each of the distinct 19 weekly shifts and all shift-type subsets are listed. (b) (bottom panel): The one-week schedule presents a feasible assignment of physicians to shifts.

each physician would work 19 shifts, consisting of 10 *WeekDay* shifts, four *WeekNight* shifts, and one each of the remaining five shift types.

The next step was for each physician to articulate his (her) individual relative preferences for each shift type. All physicians agreed that a Monday-to-Friday day (*WeekDay*) shift was the most preferred shift. To provide a base measure for the preference scores, a *WeekDay* shift was assigned a preference score of one preference unit. To determine a preference value for the remaining six shift types, each physician then estimated how many *WeekDay* shifts he or she would consider as a fair trade (in either direction) for each of the remaining six shift types.

The initial physician preference values were unconstrained in that the sum of the values within an individual's vector could add up to any number. The group reviewed the initial individual values and was able to reach agreement on a single set of relative workload values for the group. The first row of Table 2 gives the group's resulting relative workload values (henceforth, referred to as the *workload vector*). The workload values would serve as a basis for comparing workloads between physicians and ensuring workload balance through constraints.

With the relative workload values agreed upon, the physicians were then allowed to adjust their individual preference vectors in the context of the workload vector. One requirement stipulated that their vector of

	Shift type							
	WeekDay	*WeekNight*	*FriNight*	*SatSunDay*	*SatNight*	*SunNight*	*WeekLong*	Vector total
Group workload vector	1	1.5	1.75	2.5	2.5	2.5	4	15.75
Individual physician preference vectors								
Doc1	1	1.5	1.75	2.5	2.5	2.5	4	15.75
Doc2	1	1.4	1.65	2.4	2.65	2.65	4	15.75
Doc3	1	1.5	1.75	2.5	2.5	2.5	4	15.75
Doc4	1	1.5	1.75	2.5	2.5	2.5	4	15.75
Doc5	1	1.5	1.75	2.5	2.5	2.5	4	15.75
Doc6	1	1.5	1.75	2.5	2.5	2.5	4	15.75
Doc7	1	1	2	2	3.75	2	4	15.75
Doc8	1	1.5	1.75	2.5	2.5	2.5	4	15.75
Doc9	1	1.5	2	2.25	2.5	2.5	4	15.75

Table 2: The physicians adjusted the set of individual preference vectors.

preference values had to sum to 15.75 (the sum of the group workload vector). Table 2 shows the adjusted set of individual preference vectors. Six members of the group elected to use the group's workload vector as their individual preference vector; three members (highlighted) chose preference vectors that were distinct.

There are various ways to represent the equivalence between shifts. Our method of grouping the individual shifts into shift types and fixing the *WeekDay* preference value to 1.0 had the advantage that the physicians had to express equivalence values on relatively few (six) other shift types. One value for *WeekNight* would essentially represent four shifts, and the requirement that the vector totals had to sum to 15.75 allowed a change in that particular value to have a relatively greater impact than the other singleton shift types. However, the simplicity and the perceived fairness of having to declare a preference vector relative to the workload vector resonated well with the physicians.

The workload vector was used to retrospectively analyze previous schedules for workload balance under the assumption that administrative work (i.e., administrative duties, research, and committee work) was distributed equally among all physicians. We analyzed three consecutive nine-week periods (spanning January 3, 2011 through July 10, 2011) and compared them to an equality schedule in terms of counts of shifts by type and average nine-week workload. The analysis confirmed that the workloads corresponding to the prior schedules were far from equal. One physician's workload was at least 14 percent more than a perfectly equal distribution of shifts and over 30 percent more than the workload of another physician in the group. Even when using the preference vectors in comparisons, the same magnitude of imbalances remained.

Development of the Hybrid Preference Scheduling Model

With a clear opportunity to improve workload balance, we formulated a mixed-integer optimization model to generate nine-week schedules. The first use of the model was to generate a feasible equality schedule. In this model, we added constraints to ensure that each physician was scheduled to work the same number of each type shift. Next, we modified the optimization model to incorporate the workload values and the individual preference values. The objective was constructed to maximize the minimum percentage improvement in individual total preference value relative to an equality schedule.

Both models' results were reviewed and a discussion ensued on the notion of schedule fairness. Although the relative workload between any two physicians in the preference-based model was essentially equivalent, the number of shifts worked over the nine-week period varied considerably and became the focus of the fairness discussion. Ultimately, a subset of physicians maintained that for a schedule to be fair, each physician must be required to work the same number of each type of shift over the scheduling horizon. The remaining physicians maintained their predilection for a preference schedule as they shaped and accepted a preference-based view of fairness from a relative workload perspective as previously outlined.

Clearly, the physicians segmented into two groups—those who preferred an equality assignment and those who were willing to accept a nonequality schedule that incorporated their individual preferences. In our particular instance, six physicians chose the equality schedule and the remaining three physicians opted for the nonequality schedule based on individual preference.

To accommodate both sets of physicians, we introduced the Hybrid Preference Scheduling Model (HPSM). The model is deemed *hybrid* because it is designed to concurrently accommodate two different scheduling philosophies to address the physician scheduling problem (the equality and the preference approach). The physicians that preferred a true equality schedule could easily be scheduled to work the equivalent of one entire week's workload distributed throughout the nine-week horizon. The total workload associated with the remaining physicians could then be assigned based on physician preference among the remaining pool of doctors who maintained their predilection for a true preference schedule. In essence, the hybrid model facilitates multiple-party, Pareto-improving swaps among the subset willing to accept a nonequality schedule.

To model preference improvement for each physician in the preference subset, we calculate a baseline measure of comparison, the individual physician's preference score for an equality schedule, by summing the corresponding physician's shift-type preferences across all shifts in a one-week period. Once this (equality) baseline is established, the mixed-integer programming model seeks to maximize the minimum percentage improvement in baseline preference across the physicians in the preference subset.

The model design allows easy generation of schedules based on pure equality, pure preference, or a truly hybrid schedule. Constraints are classified into three high-level sets: (1) workload constraints for an equality subset of physicians, (2) workload constraints for a preference subset of physicians, and (3) detailed schedule constraints. Creating an equality, preference, or hybrid schedule involves simply combining the appropriate sets of constraints with the corresponding objective function. The appendix includes the HPSM formulation.

We solved the first instance of the HPSM model, and the physicians reviewed the resulting schedule. Although the schedule did not include vacation requests or individual shift-pattern preferences, all the physicians deemed it reasonable. Table 3 shows an HPSM schedule depicting shift counts, workload, and preference totals. The subset of those electing preference consisted of physicians 2, 7, and 9. The remaining six physicians constitute the equality subset. The schedules for each of the three physicians who selected preference exhibited preference totals (column 10) that were at least 6 percent less (better) than an equality schedule evaluated using their preference values (column 11). With the group, this established the concept of incorporating individual preference and set the stage for greater model detail and customization. Having built

(1)	(2) *WeekDay*	(3) *WeekNight*	(4) *FriNight*	(5) *SatSunDay*	(6) *SatNight*	(7) *SunNight*	(8) *WeekLong*	(9) Total shift count	(10) Solution evaluated using preference values	(11) Equality evaluated using preference values	(12) Preference difference between solution and equality (%)	(13) Schedule evaluated using workload vector
Doc1	10	4	1	1	1	1	1	19	29.25	29.25		
Doc2	14	1	3	0	1	0	1	20	27	28.95	6.7	27.25
Doc3	10	4	1	1	1	1	1	19	29.25	29.25		
Doc4	10	4	1	1	1	1	1	19	29.25	29.25		
Doc5	10	4	1	1	1	1	1	19	29.25	29.25		
Doc6	10	4	1	1	1	1	1	19	29.25	29.25		
Doc7	5	11	0	0	0	3	1	20	26	27.75	6.3	33
Doc8	10	4	1	1	1	1	1	19	29.25	29.25		
Doc9	11	0	0	3	2	0	1	17	26.75	29.25	8.5	27.5

Table 3: The hybrid model resulted in an improvement for each physician in the preference set.

trust and confidence in the concepts, we moved toward including more types of group and individual constraints by repeated iterations of schedule generation and review.

A model created to meet the group's complex and unique scheduling constraints required many lengthy discussions with the physicians and the generation of numerous schedule prototypes to provide answers to what-if types of questions. The schedule visualization tool shown in Figure 2 served to efficiently communicate responses to such inquiries in an easy-to-understand format. The process of schedule discussion, followed by what-if types of questions, followed by scenario analysis and the visualization of the model response, facilitated the elucidation of a comprehensive set of model constraints. Ultimately, agreement was reached on a representative set of group and individual constraints and guidelines for requesting vacation periods or days off.

We reviewed an actual schedule from a previous nine-week period to verify whether such constraints were being honored in the manually generated schedules and found a large number of violations of the newly agreed upon constraints. Hence, the model-generated schedules would promote improved clinical continuity and healthier rest patterns for the physicians relative to the manually generated schedules.

Model Implementation

The group decided to take a conservative approach regarding implementation of the first optimization-based schedule by agreeing to launch with a pure equality schedule in February 2013. The implementation of the second nine-week optimization-based schedule increased the physicians' levels of confidence in the computer-based scheduling process. As a result, several physicians self-selected back into a preference-based subset for a specific set of shifts. In fall 2013, the group conducted, reviewed, and implemented the first run of the HPSM with the complete set of constraints and time-off requests.

In practice, the scheduling model is designed for implementation every nine weeks on a rolling horizon basis, as we describe above. Prior to running a nine-week schedule, each physician submits requests for extended vacation days, professional days off, personal requests for one or two consecutive days off, or for specific shifts off. Physicians may request "true vacation days" in the form of a one- or two-week block of time during a nine-week scheduling period four times each year. A two-week block of vacation time is rare and must be approved by the group months in advance. The number of requests for days off and shifts off per physician varies significantly from zero to as many as 22 days off (only seven consecutive "true vacation days" in this case) per nine-week period. The more consecutive days off requested, the more compressed the particular physician's schedule becomes because the equivalent nine-week workload must then be scheduled within the remaining days. Additionally, when several physicians request time off during the same period, the remaining physicians are negatively impacted because they must cover all shifts. As such, excessive vacation requests may render the instance infeasible. If this occurs, subsets of the detailed schedule constraints must be relaxed iteratively to reach feasibility. The physician group agreed upon a formal process to relax such constraints in the following order: (1) noncritical quality of care, (2) "no day shift two days after a night shift" (lifestyle constraints), and (3) "at most one night shift among three consecutive night shifts" (lifestyle constraints). The constraints were relaxed according to seniority among the physicians. The constraint sets are described in detail in the appendix as constraint sets 32–36, 27–29, and 21–26, respectively. In an attempt to minimize the need to manually relax constraint sets in pursuit of schedule feasibility, the group agreed on guidelines to limit simultaneous vacation requests. No more than two doctors may request an extended vacation (i.e., longer than seven days) concurrently and no more than three doctors may request to be off on the same day or during the same shift. The medical director must approve all exceptions. Although the group agreed on the vacation limits in spirit, evidence suggested that, in practice, violations were commonplace. Root cause analysis revealed a lack of visibility of the entire group's vacation requests as the reason. Physicians agreed to post all vacation requests to an online calendar on a continuous basis to increase visibility. The next nine-week schedule is generated two to three weeks before the end of the active schedule. In the last week of the current schedule, the authors formulate transition constraints that incorporate vacation

Week1	Mon	Tue	Wed	Thu	Fri	Sat	Sun
DayA	DOC2	DOC8	DOC8	DOC2	DOC9	DOC9	DOC9
DayB	DOC5	DOC5	DOC9	DOC5	DOC5		
Night	DOC3	DOC4	DOC6	DOC3	DOC1	DOC6	DOC1
24-7	DOC7	DOC7	DOC7	DOC7	DOC7	DOC7	DOC7
Week2	Mon	Tue	Wed	Thu	Fri	Sat	Sun
DayA	DOC2	DOC9	DOC2	DOC7	DOC7	DOC8	DOC8
DayB	DOC8	DOC4	DOC9	DOC4	DOC8		
Night	DOC7	DOC5	DOC1	DOC3	DOC2	DOC5	DOC7
24-7	DOC6	DOC6	DOC6	DOC6	DOC6	DOC6	DOC6
Week3	Mon	Tue	Wed	Thu	Fri	Sat	Sun
DayA	DOC4	DOC8	DOC9	DOC2	DOC9	DOC9	DOC9
DayB	DOC2	DOC3	DOC4	DOC8	DOC3		
Night	DOC1	DOC6	DOC7	DOC1	DOC2	DOC8	DOC7
24-7	DOC5	DOC5	DOC5	DOC5	DOC5	DOC5	DOC5
Week4	Mon	Tue	Wed	Thu	Fri	Sat	Sun
DayA	DOC4	DOC3	DOC3	DOC6	DOC1	DOC1	DOC1
DayB	DOC2	DOC1	DOC4	DOC2	DOC6		
Night	DOC6	DOC5	DOC7	DOC8	DOC4	DOC3	DOC4
24-7	DOC9	DOC9	DOC9	DOC9	DOC9	DOC9	DOC9
Week5	Mon	Tue	Wed	Thu	Fri	Sat	Sun
DayA	DOC8	DOC6	DOC2	DOC2	DOC4	DOC4	DOC4
DayB	DOC6	DOC8	DOC4	DOC6	DOC6		
Night	DOC7	DOC1	DOC5	DOC7	DOC8	DOC9	DOC8
24-7	DOC3	DOC3	DOC3	DOC3	DOC3	DOC3	DOC3
Week6	Mon	Tue	Wed	Thu	Fri	Sat	Sun
DayA	DOC6	DOC1	DOC9	DOC9	DOC1	DOC6	DOC6
DayB	DOC5	DOC9	DOC5	DOC1	DOC6		
Night	DOC7	DOC4	DOC8	DOC7	DOC3	DOC1	DOC3
24-7	DOC2	DOC2	DOC2	DOC2	DOC2	DOC2	DOC2
Week7	Mon	Tue	Wed	Thu	Fri	Sat	Sun
DayA	DOC8	DOC2	DOC5	DOC2	DOC8	DOC3	DOC3
DayB	DOC5	DOC1	DOC1	DOC3	DOC3		
Night	DOC7	DOC8	DOC6	DOC7	DOC5	DOC9	DOC5
24-7	DOC4	DOC4	DOC4	DOC4	DOC4	DOC4	DOC4
Week8	Mon	Tue	Wed	Thu	Fri	Sat	Sun
DayA	DOC7	DOC7	DOC9	DOC1	DOC3	DOC9	DOC9
DayB	DOC1	DOC1	DOC7	DOC3	DOC9		
Night	DOC2	DOC4	DOC5	DOC7	DOC6	DOC2	DOC6
24-7	DOC8	DOC8	DOC8	DOC8	DOC8	DOC8	DOC8
Week9	Mon	Tue	Wed	Thu	Fri	Sat	Sun
DayA	DOC5	DOC3	DOC2	DOC6	DOC5	DOC5	DOC5
DayB	DOC4	DOC4	DOC3	DOC2	DOC6		
Night	DOC8	DOC7	DOC4	DOC3	DOC2	DOC4	DOC7
24-7	DOC1	DOC1	DOC1	DOC1	DOC1	DOC1	DOC1

Figure 2: (Color online) The visualization tool facilitated agreement on constraints by efficiently illustrating a candidate schedule.

requests and the impact of the assignments on the first week of the next nine-week schedule. If the schedule is infeasible, constraints are relaxed as described above. The model is formulated and solved every nine weeks using the AMPL math programming language and Gurobi 5.1.0 math programming solver. An Excel VBA module then transforms the output from the MIP into a visual format easily understood by the physicians. The group's liaison physician then reviews the schedule and shares it with the entire group. Additional constraints are added to accommodate change requests or omissions and the model is re-solved. As the schedule is executed, if additional issues arise (e.g., an illness), the liaison physician can manually facilitate amenable swaps to the visual schedule through the Excel interface and then communicate those changes to the group as a whole. Note that the proposed scheduling model assumes that administrative work is distributed equally among all physicians. The model can be easily modified to accommodate less-than-fulltime clinical loads as a result of increased administrative duties by accordingly adjusting workload constraints.

Model Testing on Larger and More Diverse Practices

Although the neonatal physician group we studied at UTMC is small, it is the second largest NICU in the state. The two largest UTMC hospital-based groups providing 24-7 coverage are the anesthesia and hospitalist groups, which employ 27 and 24 physicians, respectively. We explored the applicability of the scheduling model to larger and more demographically diverse practices. A test data set was generated to compare model performance among practices comprised of 9, 18, and 27 physicians among general neonatal and hospitalist practices since hospitalist and neonatal coverage-duty requirements are quite similar. Based on actual physician preference data, in combination with physician expertise, five demographically based physician profiles and corresponding shift preferences were defined (Table 4). As above, each vector of preferences sums to 15.75.

Demographic physician data, available from three medical professional organizations (Frintner 2013, Smart 2013, American Board of Pediatrics 2014) were used to generate six data sets, three to represent neonatal practice groups with 9, 18, and 27 physicians each and three to represent hospitalist groups with 9, 18, and 27 physicians, respectively. A pure preference version of the scheduling model was tested using each of the six data sets described with no limit on the number of 24-7 call shifts worked in a given scheduling horizon. Table 5 shows the results for each of the six test scenarios. In each case, the table includes the maximum and minimum percentage improvement in physician preference over that of an equality schedule as a range and the physician count by profile type.

In general, more variation in the mix of physicians and preferences results in a better schedule in terms of overall percentage improvement in preference over an equality schedule; the maximum improvement increases as the number of physicians increases.

Conclusions, Limitations, and Future Research

The presented method of incorporating individual relative shift preferences into a physician rostering model has not been explored by previous researchers and is not present in commercial software packages. The method is general and can be applied within any rostering problem in which individual lifestyles may evoke different relative preferences among shift types.

Physician profile	*WeekDay*	*WeekNight*	*FriNight*	*SatSunDay*	*SatNight*	*SunNight*	*WeekLong*
A Young–married/with partner, no children	1	1.3	2.25	2	3.1	2.1	4
B Parent of young children	1	1.5	2	3	3	2	3.25
C Young–single, no children	1	1.5	2.1	2	2.75	2.4	4
D Parent of older children	1	1.25	1.5	3	2.5	1.5	5
E Nearing retirement	1	1.6	1.8	2	2.8	2.55	4

Table 4: Representative physician profiles and corresponding shift-type preferences are defined.

	General neonatal practices		General hospitalist practices	
Total no. of physicians	Physician count by A, B, C, D, E profile type, respectively	Range of improvement over equality (all physicians) (%)	Physician count by A, B, C, D, E profile type, respectively	Range of improvement over equality (all physicians) (%)
9	1, 1, 1, 3, 3	6.60–7.89	1, 2, 1, 4, 1	7.12–7.89
18	2, 3, 1, 5, 7	5.26–7.95	3, 5, 1, 7, 2	6.60–9.65
27	3, 4, 1, 9, 10	5.98–9.95	4, 8, 1, 11, 3	5.98–11.4

Table 5: Variations in the mix of physicians and preferences improve schedule results.

The optimization-based approach facilitates complex multiple-party trades and creates schedules that are individually preferred over culturally acceptable equality schedules, which are often the default in equally partnered physician groups.

One limitation of this study is that all physicians were considered equally qualified to work any of the shift types. In some settings, consideration must be given with respect to clinical expertise (board certifications) or productivity. Although these exceptions can be handled with additional constraints and objectives, they were not explored in the current model. Another characteristic of the neonatology group was that physician pay was not schedule based. Some physician groups pay differential rates depending on the type of shift or the number of hours worked. This scenario creates a situation in which the physicians may have to consider income impacts when declaring their preference vectors. This problem can, however, be handled using the workload vector as a guide for the differential rates.

Future enhancements include testing the model in other types of hospital-coverage practices (e.g., emergency medicine, anesthesiology, hospital medicine, radiology) and incorporating more formal methods for constraint relaxation when initial models are constrained to the point of being either infeasible or very limited with respect to individual improvement relative to an equality schedule. We might also try to generate additional improvements by constraining to insure minimum gains and then maximizing the total improvement. Although this might yield a quantitative improvement, it runs the risk of widening the range of gains among the preference physicians to the point of being perceived as unfair. Nonetheless, it should be explored for its potential.

Our work provides insights on iteratively engaging with physician groups to maximize the acceptance of optimization methods. The schedule visualization tool helped the group identify and converge on a rational set of work-practice constraints and reach agreement on processes for incorporating vacation requests. In addition, the opt-in or opt-out aspect of the HPSM allowed physicians to adopt the preference-based approach at their own pace.

To the physician group, however, this method demonstrated the value of an objective, computer-based approach that resulted in a less physically demanding schedule, covered all required shifts, and provided improved continuity of care. These schedule characteristics should promote physician longevity and help in meeting the increased demand for healthcare services.

Appendix

The Hybrid Preference Scheduling Model

Model parameters include the following:

P = set of preference physicians.
n = number of preference physicians.
I = total number of physicians.
J = number of shifts to be covered in one week.
K = number of weeks in the scheduling horizon.
$p_{i,j}$ = physician i's preference value for shift j, $i \in P$.
E = set of equality physicians.
N = maximum number of night shifts a physician may work in a given week.
D = maximum number of shifts a physician may work in a given week.
W = number of weekend shifts an equality physician must work in each half of the scheduling horizon.
Q = minimum number of shifts a physician must work in a nonvacation week.
H = minimum number of weekday shifts a physician must work in each half of the scheduling horizon.
M = number of different shift types.

T_m = set of shifts in the mth shift type set.
t_m = number of members in set T_m.
$v_j = \sum_{i\in P} p_{i,j}/n$, average preference value for shift j among preference physicians.
$V = \sum_{j=1}^{J} v_j$, the preference physicians' total preference score, a surrogate for the total K-week physician workload for the preference physicians.
$R(i) = \sum_{j=1}^{J} p_{i,j}$, doctor i preference for an equality schedule (schedule where each doctor works an equal number of shifts for each shift type), where, $i \in P$.

The decision variables are as follows:

w = minimum proportional improvement in preference schedules among all preference physicians.

$$x_{i,j,k} = \begin{cases} 1, & \text{if doctor } i \text{ is assigned to shift } j \text{ in week } k; \\ 0, & \text{otherwise.} \end{cases}$$

The MIP is presented as follows:

$$\text{Maximize } w \tag{1}$$

subject to:

Workload Constraints for the Preference Physicians

$$R(i) - \sum_{j=1}^{J}\sum_{k=1}^{K} p_{i,j} x_{i,j,k} \ge w \cdot R(i), \quad \forall i \in P, \tag{2}$$

$$\sum_{j=1}^{J}\sum_{k=1}^{\lfloor K/2\rfloor} v_j x_{i,j,k} \ge V \cdot \frac{1}{K}\left(\left\lfloor \frac{K}{2} \right\rfloor - 1\right), \quad \forall i \in P, \tag{3}$$

$$\sum_{j=1}^{J}\sum_{k=\lceil K/2\rceil}^{K} v_j x_{i,j,k} \ge V \cdot \frac{1}{K}\left(\left\lceil \frac{K}{2} \right\rceil - 1\right), \quad \forall i \in P. \tag{4}$$

Workload Constraints for Equality Physicians

$$\sum_{k=1}^{K}\sum_{j\in T_m} x_{i,j,k} = t_m, \quad \forall i \in E, \forall m \in \{1,2,\ldots,M\}. \tag{5}$$

The HPSM can be customized to accommodate any type of shift structure and any number and type of practice-specific work rules, work patterns, lifestyle, quality-of-care, or related constraints. Table 1 gives the descriptions of the seven shift types adopted by the neonatal practice at UTMC, and Figure 1(a) gives the indices of the 19 weekly shifts. The model can be easily adapted to generate either a pure preference schedule in which all physicians are scheduled based solely on shift preference or a pure equality schedule in which all physicians work an equal number of each shift type.

Work-Rule Constraints

$$x_{i,j,k+1} + x_{i,j,k} \le 1, \quad \forall i \in P, \forall j \in T_7, \quad \forall k \in \{1,2,\ldots,K-1\} \tag{6}$$

$$x_{i,j,k} = x_{i,j',k}, \quad \forall i \in E, \forall j \in T_3, \forall j' \in T_6, \forall k, \tag{7}$$

$$\sum_{j\in T_4\cup T_5\cup T_6\cup T_7}\sum_{k=1}^{\lceil K/2\rceil} x_{i,j,k} \ge W, \quad \forall i \in E, \tag{8}$$

$$\sum_{j\in T_4\cup T_5\cup T_6\cup T_7}\sum_{k=\lceil K/2\rceil}^{K} x_{i,j,k} \ge W, \quad \forall i \in E, \tag{9}$$

$$\sum_{i=1}^{I} x_{i,j,k} = 1, \quad \forall j, \forall k, \tag{10}$$

$$x_{i,j,k} + x_{i,j+1,k} + x_{i,j+2,k} \le 1, \quad \forall i, \forall j \in T_{1A}, \forall k, \tag{11}$$

$$x_{i,j,k} + x_{i,j',k} \le 1, \quad \forall i, \forall j \in T_1\cup T_2\cup T_3\cup T_4\cup T_5\cup T_6, \quad \forall j' \in T_7, \forall k, \tag{12}$$

$$\sum_{j\in T_4\cup T_5\cup T_6} x_{i,j,k} \le 1, \quad \forall i, \forall k, \tag{13}$$

Work-Pattern Constraints

$$x_{i,j,k} + x_{i,j+1,k} + x_{i,j+2,k} \le 1, \quad \forall i, \forall j \in T_2, \forall k, \tag{14}$$

$$\sum_{j\in T_3\cup T_4\cup T_5} x_{i,j,k} \le 1, \quad \forall i, \forall k, \tag{15}$$

$$x_{i,j,k} + \sum_{j'\in T_{MAB}} x_{i,j',k+1} \le 1, \quad \forall i, j \in T_6, k \in \{1,2,\ldots,K-1\}, \tag{16}$$

$$x_{i,j,k} + x_{i,j',k+1} \le 1, \quad \forall i, j \in T_6, j' \in T_7, \quad \forall k \in \{1,2,\ldots,K-1\}, \tag{17}$$

$$x_{i,j,k} + \sum_{j'\in T_{MAB}} x_{i,j',k+1} \le 1, \quad \forall i, j \in T_7, \quad \forall k \in \{1,2,\ldots,K-1\}, \tag{18}$$

$$\sum_{j'\in T_{MABN}} x_{i,j',k+1} + x_{i,j,k} \le 1, \quad \forall i, j \in T_4, \quad \forall k \in \{1,2,\ldots,K-1\}, \tag{19}$$

$$\sum_{j\in T_2\cup T_3\cup T_5\cup T_6} x_{i,j,k} \le N, \quad \forall i, \forall k, \tag{20}$$

Lifestyle Constraints

$$\sum_{j\in T_{MN}\cup T_{TN}\cup T_{WN}} x_{i,j,k} \le 1, \quad \forall i, \forall k, \tag{21}$$

$$\sum_{j\in T_{TN}\cup T_{WN}\cup T_{HN}} x_{i,j,k} \le 1, \quad \forall i, \forall k, \tag{22}$$

$$\sum_{j\in T_{WN}\cup T_{HN}\cup T_3} x_{i,j,k} \le 1, \quad \forall i, \forall k, \tag{23}$$

$$\sum_{j\in T_{HN}\cup T_3\cup T_5} x_{i,j,k} \le 1, \quad \forall i, \forall k, \tag{24}$$

$$x_{i,j,k+1} + \sum_{j'\in T_5\cup T_6} x_{i,j',k} \le 1, \quad \forall i, j \in T_{MN}, \quad \forall k \in \{1,2,\ldots,K-1\}, \tag{25}$$

$$\sum_{j\in T_{MN}\cup T_{TN}} x_{i,j,k+1} + x_{i,j',k} \le 1, \quad \forall i, j' \in T_6, \quad \forall k \in \{1,2,\ldots,K-1\}, \tag{26}$$

$$x_{i,j,k} + x_{i,j+4,k} + x_{i,j+5,k} \le 1, \quad \forall i, \forall j \in T_{MN} \cup T_{TN} \cup T_{WN}, \forall k, \tag{27}$$

$$x_{i,j,k} + \sum_{j'\in T_{MAB}} x_{i,j',k+1} \le 1, \quad \forall i, j \in T_5, \quad \forall k \in \{1,2,\ldots,K-1\}, \tag{28}$$

$$x_{i,j,k} + \sum_{j'\in T_{TAB}} x_{i,j',k+1} \le 1, \quad \forall i, j \in T_6, \quad \forall k \in \{1,2,\ldots,K-1\}, \tag{29}$$

$$\sum_{j=1}^{J} x_{i,j,k} + x_{i,j',k} \le D, \quad \forall i, j' \in T_4, \forall k, \tag{30}$$

$$\sum_{k=1}^{K} x_{i,j,k} = 1, \quad \forall i, j \in T_7, \tag{31}$$

Quality-of-Care Constraints

$$\sum_{j=1}^{J} x_{i,j,k} \ge Q, \quad \forall i, \forall k, \tag{32}$$

$$\sum_{j\in T_1} \sum_{k=1}^{\lfloor K/2 \rfloor} x_{i,j,k} \ge H, \quad \forall i, \tag{33}$$

$$\sum_{j\in T_1} \sum_{k=\lceil K/2 \rceil}^{K} x_{i,j,k} \ge H, \quad \forall i, \tag{34}$$

$$x_{i,j',k} \le \sum_{j\in\{T_1 \mid j'\}} x_{i,j,k}, \quad \forall i, \forall j' \in T_1, \forall k, \tag{35}$$

$$\sum_{j\in T_{FAB}} x_{i,j,k} - x_{i,j',k} \ge 0, \quad \forall i, j' \in T_4, \forall k. \tag{36}$$

More specifically, the HPSM model objective function (1) seeks to maximize w, the minimum proportional improvement to individual physician preference relative to an equality schedule for the preference set of doctors. Constraint set (2) defines w. Constraint sets (3) and (4) require that at least a portion of each preference doctor's workload is scheduled in each of the first and second parts of the K-week scheduling period, where the total workload is represented as the group's preference score V for a K-week period. Constraint set (5) characterizes the workload constraints for the equality set of physicians and ensures that each doctor in the equality subset works an equal number of shifts for each shift type over the K-week period. Constraint sets (6)–(13) describe work-rule constraints. Constraint set (6) ensures that no preference doctor is assigned the 24-7 *WeekLong* call shift for two consecutive weeks. Constraint set (7) ensures that for doctors in the equality subset, the same doctor works both the *FriNight* shift and the *SunNight* shift in a given week. Constraint set (8), in conjunction with constraint sets (5) and (7), ensures that all physicians in the equality subset work a weekend shift during W different weekends during the first half of the scheduling time horizon; similarly, constraint set (9) together with constraint sets (5) and (7) ensures that all physicians in the equality subset work a weekend shift during W different weekends in the second half of the scheduling horizon. If the number of weeks in the scheduling horizon is odd, the middle week in the horizon is included in both halves of the scheduling horizon. Constraint set (10) ensures that one and only one doctor is assigned to each shift during the K-week period. Constraint set (11) ensures that a given doctor is assigned to no more than one shift in any 24-hour period Monday–Friday. Constraint set (12) ensures that a doctor assigned to work the 24-7 *WeekLong* call shift in any given week will not be assigned to work any other shift in the same week. Constraint set (13) ensures that for a given week, the *SatNight*, *SunNight*, and *SatSunDay* shift assignments are mutually exclusive.

Constraint sets (14)–(20) describe work-pattern constraints applied to both preference and equality physician schedules. Constraint set (14) ensures that if a doctor works a *WeekNight* nightshift, he (she) will not work the following *WeekDay* shift. Constraint set (15) ensures that for a given week, the *FriNight*, *SatNight*, and *SatSunDay* shift assignments are mutually exclusive. Constraint set (16) ensures that if a doctor works the *SunNight* shift, he or she will not work the following Monday *WeekDay* shift. Constraint set (17) ensures that if a given doctor works the *SunNight* shift in a given week, then he or she will not be assigned to work the *WeekLong* call shift the following week. Constraint set (18) ensures that if a doctor works the *WeekLong* call shift in a given week, he or she will not work the following Monday *WeekDay* shift. Constraint set (19) ensures that if a given doctor works a *SatSunDay* shift in a given week, he or she will not work the following Monday *WeekDay* shift nor the following Monday *WeekNight* shift. Constraint set (20) ensures that no physician works more than N night shifts in a given week.

Constraint sets (21)–(31) describe lifestyle constraints applied to both preference and equality physician schedules. Constraint sets (21)–(26) ensure that at most, one night shift in any given three consecutive nights is assigned to a given doctor with the exception of the *FriNight*–*SunNight* night shift pair assignment. Constraint sets (27)–(29) are required to ensure that if a given doctor works a night shift in a given week, then he or she will not be assigned to work the day shift two days after the respective night shift. For example, if a given doctor works a Monday night in a given week, then he or she will not be assigned to the following Wednesday day shift. Constraint set (30) ensures that no doctor works more than D shifts in any given week, where the *SatSunDay* shift is counted as two shifts. Physician-specific constraints are added to the lifestyle constraints to reflect requests for particular shifts or specific days off for vacation and professional leave during each scheduling horizon. An optional lifestyle preference, constraint set (31), ensures that the physicians work the 24-7 weeklong call shift only once in the scheduling horizon.

Constraint sets (32)–(36) describe quality-of-care constraints intended to increase quality and continuity in patient care and are applied to both preference and equality physician schedules. To maintain continuity in patient contact, constraint set (32) ensures that each physician works at least Q shifts in each week of the scheduling horizon excluding vacation weeks. The constraint set is customized to accommodate physician-specific vacation time. To further maintain patient contact during the more treatment-intensive *WeekDay* shifts, constraint sets (33) and (34) ensure that each physician works at least H *WeekDay* shifts during both the first and second portions of the scheduling horizon. Constraint set (35) imposes continuity of patient care because the constraints ensure that if a given doctor works a *WeekDay* shift in a given week, then he or she must work at least two *WeekDay* day shifts in that same week. Constraint set (36) ensures that the doctor assigned to work the weekend *SatSunDay* shift also works the preceding Friday *WeekDay* shift to promote consistency and continuity of care from the weekday to the weekend shifts.

References

American Board of Pediatrics (2014) Workforce data 2013–2014. Accessed October 29, 2015, http//:www.abp.org/sites/abp/files/pdf/workforcebook.pdf.

Baum R, Bertsimas D, Kallus N (2014) Scheduling, revenue management, and fairness in an academic-hospital radiology division. *Acad. Radiology* 21(10):1322–1330.

Beaulieu H, Ferland JA, Gendron B, Michelon P (2000) A mathematical programming approach for scheduling physicians in the emergency room. *Health Care Management Sci.* 3(3):193–200.

Bekkan C (2010) Quantitative methods of physician scheduling at hospitals. Master's thesis, University of Agder, Kristiansand, Norway.

Bourdais S, Galinier P, Pesant G (2003) HIBISCUS: A constraint programming application to staff scheduling in health care. Rossi F, ed. *Proc. Principles Practice Constraint Programming—CP* 2003 (Springer-Verlag, Berlin-Heidelberg), 153–167.

Burke EK, De Causmaecker P, Vanden Berghe G, Van Landeghem H (2004) The state of the art of nurse rostering. *J. Scheduling* 7(6):441–499.

Carter MW, Lapierre SD (2001) Scheduling emergency room physicians. *Health Care Management Sci.* 4(4):347–360.

Cheang B, Li H, Lim A, Rodrigues B (2003) Nurse rostering problems: A bibliographic survey. *Eur. J. Oper. Res.* 151(3):447–460.

Cohn A, Root S, Kymissis C, Esses J, Westmoreland N (2009) Scheduling medical residents at Boston University School of Medicine. *Interfaces* 39(3):186–195.

De Causmaecker P, Vanden Berghe G (2010) A categorisation of nurse rostering problems. *J. Scheduling* 14(1):3–16.

Erhard M, Schoenfelder J, Fügener A, Brunner JO (2015) State of the art in physician scheduling. Working paper, University of Augsburg, Augsburg, Germany.

Ernst A, Jiang H, Krishnamoorthy M, Sier D (2004b) Staff scheduling and rostering: A review of applications, methods and models. *Eur. J. Oper. Res.* 153(1):3–27.

Ernst A, Jiang H, Krishnamoorthy M, Owens B, Sier D (2004a) An annotated bibliography of personnel scheduling and rostering. *Ann. Oper. Res.* 127(1):21–144.

Ferrand Y, Magazine M, Rao US, Glass TF (2011) Building cyclic schedules for emergency department physicians. *Interfaces* 41(6):521–533.

Frintner MP (2013) Pediatrician life and career experience study. Accessed October 22, 2015, http://www.aap.org/en-us/professional-resources/Research/pediatrician-surveys/Pages/Pediatrician-Life-And-Career-Experience-Study-PLACES.aspx.

Fügener A, Brunner JO, Podschaske A (2015) Duty and workstation rostering considering preferences and fairness: A case study at a department of anesthesiology. *Internat. J. Production Res.* 53(24):7465–7487.

Gendreau M, Ferland J, Gendron B, Hail N, Jaumard B, Lapierre S, Pesant G, Soriano P (2007) Physician scheduling in emergency rooms. Burke E, Rudová H, eds. *Practice and Theory of Automated Timetabling VI; Lecture Notes in Computer Science* (Springer-Verlag, Berlin-Heidelberg), 53–66.

Gunawan A, Lau HC (2010) The bi-objective master physician scheduling problem. Accessed October 29, 2015, http://www.patatconference.org/patat2010/proceedings/2_15.pdf.

Gunawan A, Lau HC (2012) Master physician scheduling problem. *J. Oper. Res. Soc.* 64(3):410–425.

Rousseau L, Gendreau M, Pesant G (2002) A general approach to the physician rostering problems. *Ann. Oper. Res.* 115(1):103–205.

Smart DR (2013) Physician characteristics and distribution in the US. (American Medical Association, Washington, DC).

Stolletz R, Brunner JO (2012) Fair optimization of fortnightly physician schedules with flexible shifts. *Eur. J. Oper. Res.* 219(3):622–629.

Van den Bergh J, Beliën J, De Bruecker P, Demeulemeester E, De Boeck L (2013) Personnel scheduling: A literature review. *Eur. J. Oper. Res.* 226(3):367–385.

Verification Letter

Dr. James Schmid, Regional Neonatal Associates, 1930 Alcoa Highway, Suite 145, Knoxville, Tennessee 37920, writes:

"My name is Dr. James Schmid, a neonatal physician with Regional Neonatal Associates (RNA) in Knoxville, Tennessee. RNA is a group of nine neonatal physicians that serve neonatal intensive care units in two hospital systems in East Tennessee, the University of Tennessee Medical Center and Tennova Health Systems. Our group staffs the University of Tennessee Medical Center's Neonatal unit twenty-four hours each day, seven days each week and serves the Tennova system by providing one physician on call at all times.

"Our group has been working with Dr. Charles Noon and Dr. Melissa Bowers for the last three years to develop an automated physician scheduling model for our group. Our work with Drs. Noon and Bowers was initially motivated by the retirement of one of the physicians in the group that had traditionally constructed the schedule manually. At the same time, the group sought a methodology to generate a more equitable workload across physicians.

"Drs. Noon and Bowers have developed several models over the course of the last three years. They developed a strict equality model where each doctor works exactly the same number of each type shift over the scheduling horizon, as well as a scheduling model based on individual physician preference. A third model was designed to combine the two to allow a number of our physicians to receive equality schedules and a smaller subset of our physicians to receive their schedules based on shift preferences, all within the same scheduling period.

"The group has been using the schedules generated by Drs. Noon and Bowers since January 2013. Their models have been successfully implemented at RNA. The schedule currently in use is a hybrid schedule which combines both the preference and equality approaches into one schedule.

"Some physician comments follow. One of the hybrid physicians commented that she 'was pleased with how her preferences were granted. I liked hybrid much better.' Another of the hybrid physicians commented, 'I have several blocks of time off for 5–6 days in a row—that is sweet! The three Saturday nights are a bummer, but to have 5–6 days off before and after is pretty nice.' On the other hand, some physicians do not like blocks of time off and have voiced a preference for more evenly spaced call shifts to permit having the weekend off before and after a weekend call shift. The number of physicians in the group, along with the volume and specific combinations of vacation requests from all physicians collectively, have a significant impact on this aspect of the resulting schedule."

Melissa R. Bowers, PhD, is an associate professor and the Beaman Professor of Business in the Department of Business Analytics and Statistics at the University of Tennessee, Knoxville. Her teaching and research interests include production planning, scheduling, supply chain optimization, lean manufacturing, lean MRO, and discrete optimization models. Her publications have appeared in a variety of academic and trade journals. She is the director of the Master's in Business Analytics Program in the Haslam College of Business and is a co-author of the book *Lean Maintenance, Repair, and Overhaul*: *Changing the Way You Do Business.* She works with numerous companies in industry in the areas of scheduling, lean manufacturing, and lean maintenance, repair, and overhaul.

Charles E. Noon, PhD, is the Regal Entertainment Group Professor of Business in the Department of Business Analytics and Statistics at the University of Tennessee. He serves as department head and teaches in the college's Physician's EMBA program, undergraduate program, and in a number of executive education programs. His teaching interests include operational improvement, optimization, spreadsheet modeling, and decision analysis. His research concerns computer-based models for healthcare operations improvement and his published works have appeared in a diverse range of publications. He is co-author of a leading book on Emergency Department operations improvement.

Wei Wu is a PhD graduate of the Haslam College of Business at the University of Tennessee. He serves as a research specialist at the University of Illinois. His work involves applying the principles and techniques of operations research to real-world problems and providing managerial insights to improve business operations. His research interests include OR, healthcare operations, data mining/modeling, and decision analysis.

J. Kirk Bass, MD, is an assistant professor in the Department of Obstetrics and Gynecology at the University of Tennessee Medical Center, Knoxville and is board certified in general pediatrics and neonatal-perinatal medicine. He provides neonatology services in the Newborn Intensive Care Unit at the UTMC-K for babies that are born sick, premature, or in need of surgery. He is involved in teaching and training medical students and resident physicians from the Graduate School of Medicine at the University of Tennessee. Dr. Bass is also the medical director of the newborn services at Turkey Creek Medical Center in Knoxville. His research interests are parenteral nutrition-associated liver disease and potential therapies to treat this disorder. He has research articles published in *Nutrition, American Journal of Perinatology* and *Pediatric Neurology*.