BMJ Innovations Publish Ahead of Print, published on March 29, 2017 as doi:10.1136/bmjinnov-2016-000162

ORIGINAL ARTICLE

Rostering general medicine physicians to balance workload across inpatient wards: a case study

Thomas Adams,[1] Michael O'Sullivan,[1] Jonathan Christiansen,[2] Paul Muir,[2] Cameron Walker[1]

[1]Department of Engineering Science, University of Auckland, Auckland, New Zealand
[2]Division of Medicine and Health of Older People, Waitemata District Health Board, Auckland, New Zealand

Correspondence to Thomas Adams, Department of Engineering Science, University of Auckland, Room 215, Level 2, 70 Symonds Street, Auckland 1142, New Zealand; tada955@aucklanduni.ac.nz

Received 30 September 2016
Revised 22 November 2016
Accepted 7 March 2017

To cite: Adams T, O'Sullivan M, Christiansen J, *et al*. *BMJ Innov* Published Online First: [*please include* Day Month Year] doi:10.1136/bmjinnov-2016-000162

ABSTRACT

During an expansion of the General Medicine (GM) department at Waitakere Hospital, a new roster was required for its teams. Patients being transferred to another hospital due to uneven patient workloads between the teams were identified as an issue that could be addressed by an improved roster. A more even distribution of patients among the teams is desirable because it is fairer for the staff and it also improves the continuity of care for patients. Continuity of care is improved by balanced workloads because fewer patients need to be transferred from teams with high workloads. A novel rostering technique, using a mixed integer programme (MIP), which uses past data on patient admissions to track team workloads, was implemented. This allowed multiple rosters for different configurations of admitting teams to be created and evaluated against past data traces, in terms of the difference in workloads of the teams and estimated ward occupancies. The best performing of the constructed rosters reduced the median difference in workload between the team with the most patients and the team with the fewest from 14.5 to 12.5 (13.8% reduction) for the data traces considered, when compared with a roster from an internal rostering group. Waitakere Hospital has put this roster in place and has observed a reduction in the variation in workload between the teams, with fewer patients being transferred to other hospitals due to high team workloads. These improvements cannot be solely attributed to rostering improvements, as other factors such as an increase in the number of inpatient teams have also contributed. However, the generation and evaluation of multiple different rosters via MIP was central to the process that determined the final configuration of the GM department.

BACKGROUND

Rostering personnel is often a complex, time-consuming and expensive task. This is especially true in the healthcare sector where many facilities are required to be staffed 24/7. One such facility, the General Medicine (GM) department of Waitakere Hospital in Auckland, New Zealand, was planning an expansion of the medical teams in 2014. This process involved a restructuring of the teams within GM, which meant that they would not be able to continue using the current roster.

The expansion project was seen by the management team at Waitemata District Health Board (WDHB) as an opportunity to experiment with a number of models of care. The primary goals of the project were to ensure that acutely presenting patients would be seen and assessed rapidly and would be able to have their medical care provided in Waitakere Hospital. A secondary goal was to improve the working and training environment for the medical staff in GM at Waitakere Hospital.

The intention was to increase the number of teams working in GM and introduce dedicated short-stay teams in addition to the teams that look after the long-stay wards. One of the main reasons that patients were unable to be cared for at Waitakere was because the team that was responsible for their care had too many patients and had to transfer some of their patients to North Shore Hospital.

If the number of patients that each team at Waitakere Hospital was caring for could be better balanced, then the need for transferring patients to North Shore Hospital would be reduced. With this in mind, a number of potential staffing configurations were determined, each requiring rosters to be generated with an aim to balance workload across the GM teams.

The rostering problem is well studied in Operational Research.[1–4] In the modelling of healthcare services, rostering models have mainly been applied to nurses, rather than physicians. A problem similar to the one presented here has been considered, although the physicians' shifts are not linked to the number of patients they are caring for.[5] Instead, they aim to minimise the amount by which a roster breaks some of the rostering rules.

Improving patients' continuity of care by considering the number of patients that a physician hands over at the end of their shift has been attempted.[6 7] Both attempt to minimise the handovers or a representative metric. Neither study models the number of patients each physician is caring for, however.

CASE DESCRIPTION

WDHB provides health services to the North Shore, Rodney and West Auckland areas. It serves the largest population of any District Health Board in New Zealand, with a population of more than 575 000 people, and is expected to grow by ~110 000 people (20%) in the next 10 years. WDHB operates 2 hospitals, Waitakere Hospital and North Shore Hospital, and 14 other clinics.

In each of the hospitals, there is a GM department. These departments deal with the diagnosis and treatment of patients that do not require surgery. The staff in GM is split into teams. Each team consists of a consultant (a specialist general physician), a registrar (a doctor training to be a specialist) and a house officer (a more junior doctor also in training).

This case study focuses on Waitakere Hospital where the GM department is responsible for two areas: the Assessment, Diagnostic and Cardiology Unit (ADCU) and the inpatient wards. Patients are admitted to GM either from the emergency department (ED) or via referral from an external source, for example, their general practitioner. If the patient is expected to stay in hospital for <48 hours they are classified as short stay, otherwise they are classified as long stay.

Before the redesign, some of the patients that were classified as short stay would move from ED to the ADCU, or remain in the ADCU if they were already there, but they were cared for and discharged by inpatient ward teams. Under the new model, dedicated ADCU teams would care for and discharge the short-stay patients in the ADCU. Two systems were proposed for limiting the number of short-stay patients that are cared for in the ADCU; either a limit on the number of patients an ADCU team can care for, with any additional short-stay patients moving to an inpatient ward, or a proportion (eg, 50%) of all short-stay patients remaining in the ADCU to be cared for by an ADCU team with the remaining fraction moving to inpatient wards.

When creating a roster for GM, the goal was to minimise the difference in patient workloads between the teams, as this would reduce the number of patients that had to be transferred to another hospital. The roster directly influences the workloads of the teams because the number of patients being cared for by a team is affected by the number of patients admitted to and discharged from that team, and the roster determines when the team will be admitting new patients.

Rosters for GM are generated in a three-step process. First, the shifts for registrars that either involve admissions or have restrictions are scheduled. Second, all other registrar shifts are scheduled. Finally, schedules are generated for all the other personnel. The registrars are rostered first because they perform the admissions, and therefore, it is their roster that influences the workloads of the teams. This process is also attractive because only the first step is difficult—once the admission and restricted shifts are determined for the registrars, their remaining shifts and the rosters for other personnel are easily determined. We therefore focus on the first step. The relevant shifts are given in table 1.

We decided to generate cyclic rosters for the registrars in GM. This means that, for example, a 6-week roster is constructed where each of six registrars works one of the weeks, then at the end of the week, everyone moves on to the next week in the cyclic roster. This approach was chosen because all the registrars have the same skill set, and it results in an equitable roster, that is, the registrars work the same number of hours and the same amount of each of the different types of shifts in the long run.

Each day, one registrar must work the night shift. This is accomplished by having a night 'tour' which means one registrar works 7 days of consecutive night shifts, then gets 3 days off. During the night tour, the registrar's other shifts are covered by a relief registrar.

Table 1 Registrars' shifts

Shift name	Description
Admission (A)	The registrar is on call and all new patients will be admitted to their ward
Postacute (P)	The registrar performs this shift the day after their A shift, to follow-up on patients
Night (N)	The registrar is on call at night, but admissions still go to the ward of the registrar who was working the A shift during the day
Sleep (Z)	After working night shifts, registrars are given time to catch up on sleep
No shift (X)	Registrars are given some days off
Other (O)	All other shifts are grouped together as they do not affect admissions

This means that during the night tour, the registrar can be assigned two shifts; the night shift that they perform and the one during the day that the relief registrar performs.

The rosters must also abide by the following rules, some of which are set by WDHB and others by the national collective agreement between the District Health Boards and medical professionals:

1. One A shift must be performed each day;
2. The A shift on Saturday is followed by the A shift on Sunday, for all other days, the A shift is followed by the P shift on the subsequent day;
3. A night 'tour' consists of 7 days of N shifts starting on Friday and then 3 Z shifts;
4. Registrars can only be assigned A, P, N, Z, X or O shifts;
5. Registrars must be assigned one shift each day, or two during relief week (N and one other);
6. Registrars must have every second weekend off.

To generate rosters that obey all the rules, a mixed integer programme (MIP) was formulated. Its objective function was to minimise the largest difference in workloads between the team with the most and the team with the fewest patients over the planning horizon of the roster. One of the advantages of using a mathematical modelling framework is the ability to easily change parameters and generate new rosters if the rules change, for example, after new collective agreements are negotiated.

The model uses binary variables to determine whether the cyclic roster includes a particular shift on a particular day on a given week. To translate these assignments into the workloads of the teams, past data on admissions and discharges to and from GM for a year (31 December 2013–31 December 2014) were used.

Each of the days for which there were data were split into three segments: day, evening and night. The arrivals were also classified according to their length of stay, which allowed the use of the ADCU for short-stay patients to be modelled. The number of each type of patient that arrived to, and was discharged from, GM for each segment of each day could then be calculated from the data.

The number of patients, of each type, each team is caring for in each segment of each day can be calculated, given: an assignment of shifts by the binary variables; the admission numbers and discharge rates; and the patient flow structure.

Figure 1 shows how the workloads of the teams are calculated for one segment of 1 day. The actual number of admissions, taken from the data, is fed into the model. The admissions follow the appropriate pathway chosen based on the proportion of short-stay and long-stay patients, as well as the limits in place for the ADCU. The admissions that are allocated to the inpatient wards are then divided equally among the teams of the on-call registrars (working the admission shift) at that time; note that this may result in fractional numbers of patients for each team.

We assume patients are split evenly by all admitting teams and the number and structure of the teams may be different from when the data were collected, hence the workloads of the teams recorded by the model will not match the workloads that the teams actually experienced. This discrepancy between the estimated and past workloads means that the actual discharge numbers cannot be subtracted from the team workloads because the teams may not have the same number of patients in the model as they did when the discharges were performed.

Instead, the proportion of patients discharged in each segment of each day is calculated from the data and each team discharges that proportion of their patients. This means that the total number of patients that are discharged (and therefore that remain in hospital) is the same as in the data. This process is repeated for each segment of each day, with the number of patients (after admissions and discharges) each team is caring for being carried over to the next segment. Calculating the number of patients each team is caring for enables the largest difference in workloads to be calculated for every segment of every day.

The MIP was formulated in Python using the package 'PULP'. It was solved using Gurobi on a

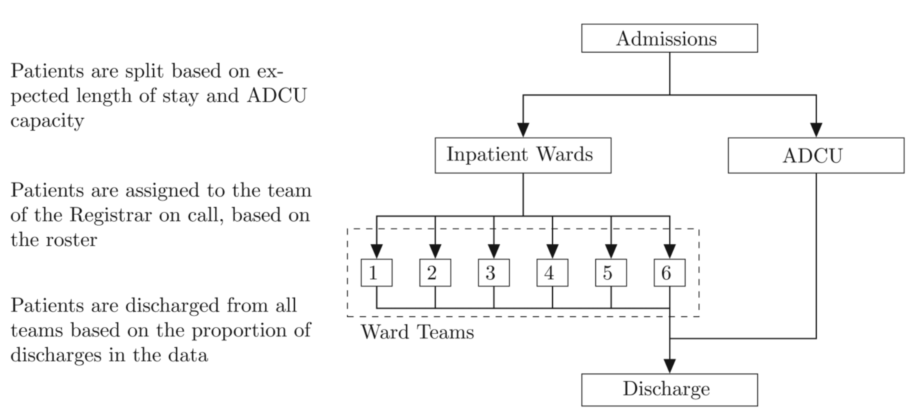

Figure 1 Determining team workloads.

Windows 64-bit computer with an Intel Core i7-4790 CPU at 3.60 GHz and 8 GB of RAM.

By linking the roster to the workloads of the teams within the MIP, a roster that minimises the difference in workloads between the teams can be found. The MIP framework also enables different numbers of teams and interactions between the ADCU and inpatient wards to be experimented with easily, as well as the straightforward examination of the effects that these factors have on the roster, workload balance and continuity of care.

RESULTS

A number of different ward team configurations were tested, using past admission and discharge data from 31 December 2013 to 31 December 2014. The number of teams working in the inpatient wards varied from 5 to 10 (excluding 7 and 9), and either 1 or 2 ADCU teams were considered. For the rosters with 8 and 10 inpatient ward teams, 2 teams would perform admissions each day, splitting the arrivals evenly. Several models of the interaction between the ADCU and the inpatient wards were examined and are described in more detail in table 2, along with all of the rosters. In all cases, it is assumed that any team working in the ADCU shares the patients assigned to the ADCU evenly among themselves. 'Cur' refers to the current roster, 'Opt' to rosters generated by finding the optimal solution to a MIP and 'Prop' to rosters proposed by the Northern Regional Training Hub (NoRTH), which provides several services to District Health Boards.

The notation (a, b, c, d, e) is used to describe the rosters where a is the number of inpatient ward teams, b is the number of ADCU teams, c is the number of inpatient ward teams working admissions, d is the percentage of short-stay patients that are treated in the ADCU and e is the limit on the number of new patients an ADCU team can admit in a day.

The aim, when generating the rosters for GM, is to balance the workloads of the teams. Therefore, when comparing the rosters, the difference between the teams with the most and fewest patients, over the planning horizon, is considered. Figure 2 shows this measure for two representative rosters and the current roster; the rosters themselves are given in figure 3. The proposed and optimised rosters shown have variability and trends similar to the other proposed and optimised rosters, respectively.

Figure 4 shows a box and whiskers plot summarising the difference between the team with the most and the team with the fewest patients for each of the rosters considered. The roster Opt (8, 2, 2, 50%, –) had the lowest median difference over the planning period. The rosters Opt (8, 2, 2, 50%, –), Opt (8, 2, 2, 100%, –), Opt (10, 2, 2, 100%, 8) and Opt (10, 2, 2, 100%, 6) had very similar IQRs which were smaller than all of the other rosters. The current roster, Cur (6, –, 1, –, –), had the highest median and the largest IQR.

The results presented in figure 4 show that the roster with the lowest median, Opt (8, 2, 2, 50%, –), had a median difference between the number of patients being cared for by the teams with the most and fewest patients of 9.8 compared with 21.1 for the current roster, a reduction of 53.5%.

For the configurations for which there was a proposed as well as an optimised roster, the optimised rosters performed 1.9 patients (13.5%) better on average, in terms of the median difference between the team with the most and fewest patients. The IQRs for the optimised rosters were also smaller, by an average of 2.2 patients (36.4%).

Table 2 Description of rosters

Roster name	Description
Cur (6, –, 1, –, –)	The current roster has 6 inpatient ward teams and no ADCU teams. One of the ward teams is responsible for all of the patients admitted to GM each day
Opt (5, 2, 1, 100%, –)	Five ward teams with one admitting new patients each day. Two ADCU teams responsible for all of the short-stay patients. Generated by finding an optimal solution to the MIP
Opt (6, 2, 1, 100%, –)	Six ward teams with one admitting new patients each day. Two ADCU teams responsible for all of the short-stay patients
Opt (8, 2, 2, 100%, –)	Eight ward teams with two admitting new patients each day. Two ADCU teams responsible for all of the short-stay patients
Opt (8, 2, 2, 50%, –)	Eight ward teams with two admitting new patients each day. Two ADCU teams responsible for 50% of the short-stay patients
Opt (10, 1, 2, 100%, 6)	Ten ward teams with two admitting new patients each day. One ADCU team responsible for 100% of the short-stay patients, up to a limit of 6 new patients each day
Opt (10, 1, 2, 100%, 8)	Ten ward teams with two admitting new patients each day. One ADCU team responsible for 100% of the short-stay patients, up to a limit of 8 new patients each day
Prop (10, 1, 2, 100%, 6)	Ten ward teams with two admitting new patients each day. One ADCU team responsible for 100% of the short-stay patients, up to a limit of 6 new patients each day. A roster proposed by an internal rostering group for WDHB
Prop (10, 1, 2, 100%, 8)	Ten ward teams with two admitting new patients each day. One ADCU team responsible for 100% of the short-stay patients, up to a limit of 8 new patients each day

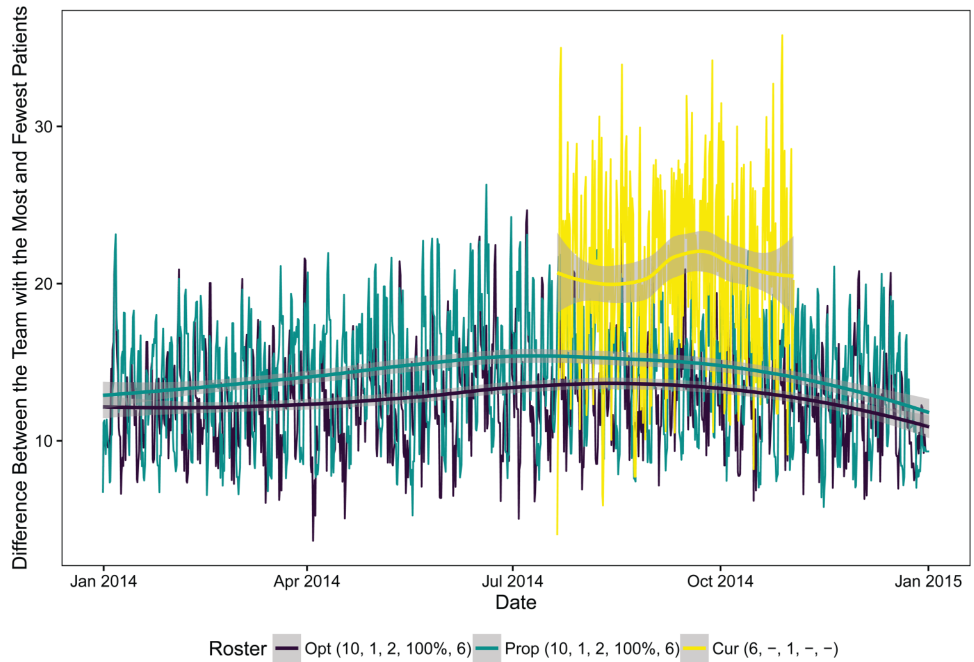

Figure 2 Time series of differences in patient workloads.

Week	Mon	Tue	Wed	Thu	Fri	Sat	Sun
1	O	O	A	P	N	N	N
2	O,N	O,N	A,N	P,N	O,Z	X,Z	X,Z
3	O	A	P	O	O	A	A
4	P	O	O	O	O	X	X
5	A	P	O	O	O	A	A
6	P	O	O	O	O	X	X
7	A	P	O	O	A	P	X
8	O	O	O	A	P	X	X
9	O	A	P	O	A	P	X
10	O	O	O	A	P	X	X

Opt (10, 1, 2, 100%, 6)

Week	Mon	Tue	Wed	Thu	Fri	Sat	Sun
1	A	P	O	O	O	A	A
2	P	O	A	P	O	X	X
3	A	P	O	O	O	X	X
4	O	A	P	O	O	A	A
5	P	O	O	A	P	X	X
6	O	O	O	O	A	P	X
7	O,N	O,N	O,N	A,N	P,Z	X,Z	X,Z
8	O	O	A	P	O	X	X
9	O	O	O	O	A	P	X
10	O,N	A,N	P,N	O,N	O,Z	X,Z	X,Z

Prop (10, 1, 2, 100%, 6)

Figure 3 Example rosters.

DISCUSSION

A model linking patient pathways and the staff roster has not been proposed before, to the best of our knowledge. The ability of the model to track each team's workload means: that a roster that minimises the difference in workloads of the teams can be found; that it provides a quantitative measure for comparing different rosters, in terms of patient workloads; and that the demand for beds can be estimated, given current ward occupancy and a scheme for generating synthetic data traces (such as the bootstrap). The model could also be used to evaluate the effect of changes to scheduling rules and policies.

The optimised rosters performed better on average across the past data (in terms of workload inequality) than their proposed equivalents. Their lower median workload difference is preferable because it means the roster is fairer overall. They also exhibited a lower variance in the inequality of the workloads across the teams, which is important as the rosters' cyclic nature means periods of high workload cycles around the teams. The more stable the workload difference is across this cycle, the more equitable the distribution of workload is in the long run.

The roster that was chosen to be implemented by Waitakere Hospital was Opt (10, 1, 2, 100%, 6). This roster was chosen because the combination of having 2 teams admitting each day and 10 teams in total meant that more patients could be kept at Waitakere Hospital without the patient numbers for each team becoming excessive.

Owing to limitations in recruiting enough doctors, only one ADCU team was able to be used with this roster. The number of registrars working night shifts was also altered in the implemented roster because of a change in how the house officers supported the wards, as well as feedback from registrars requesting additional after hours support. It is important to note that the structure of the admission shifts, which

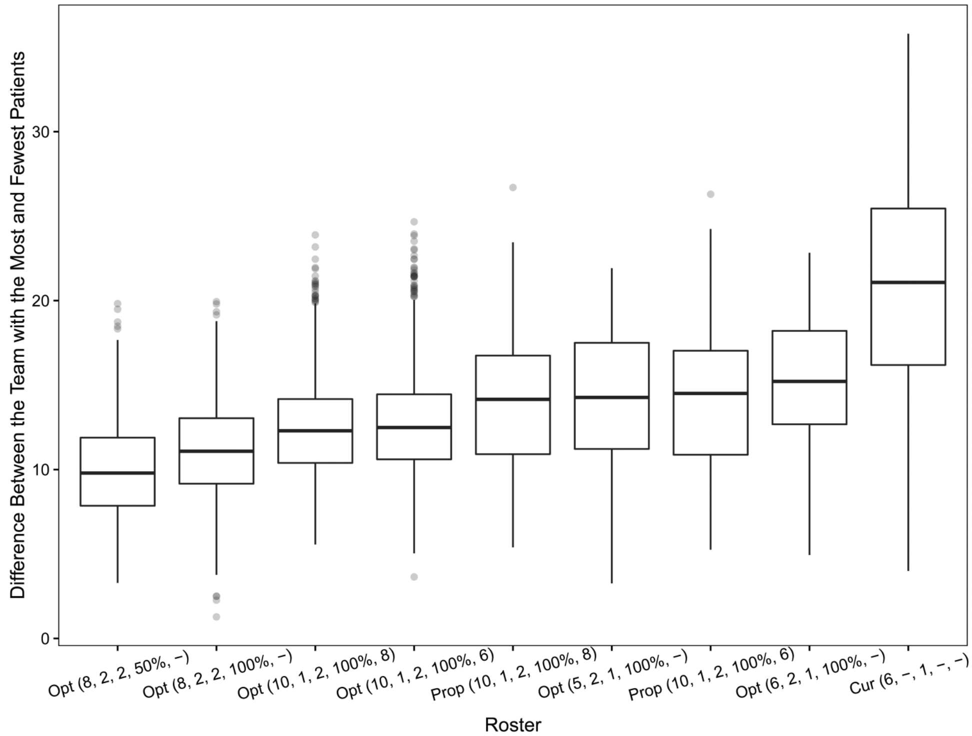

Figure 4 Box and whiskers plot of differences in patient workloads.

determine the workloads of the teams, was not altered in the implemented roster.

The new roster was transitioned to smoothly on 7 December 2015 at the beginning of a registrar change-over, which occurs every 6 months. With the new roster in place and the corresponding changes to the GM department's configuration, the number of patients being seen in the ADCU within 120 min has improved from 50% to over 80%.

The number of patients transferred to North Shore Hospital for purely staff capacity reasons is now close to nil. Some patients are still transferred for specific levels of care that are not available at Waitakere.

Overall, the length of stay for Waitakere Hospital patients has reduced, and staff responses to an anonymous survey report that workload is more appropriate, and training and morale have improved.

Although these outcomes cannot be attributed entirely to the roster that was created, since the number of inpatient teams increased at the same time, the approach presented allowed more staff configurations to be examined and provided a quantitative means of comparing the configurations.

CONCLUSIONS

Complex rostering problems, such as that described here, are frequently encountered in hospitals and can be expensive and time-consuming to overcome. This paper studies the GM department at Waitakere Hospital and describes the patient flow process and rostering problem encountered there.

The department has been reconfigured, resulting in an improved quality of care. A MIP framework enabling the exploration of a number of potential configurations has been central to the redesign process. In particular, a mathematical model has been used, which allows rosters to be created easily, many different scenarios to be explored and, by incorporating historical admissions data into the model, the workloads of the personnel to be balanced through adjustments to the roster. The model is flexible enough to handle many different rules for staff and configurations of patient flow, and produces rosters that are useable in hospital departments. Rosters were developed in this manner for the departmental configurations under consideration by Waitakere Hospital with the goal of reducing the need to transfer patients to North Shore Hospital by balancing the workloads of the teams in GM.

Contributors TA performed modelling, analysis and prepared the manuscript. MOS and CW supervised the modelling and editing. JC and PM collected and prepared data.

Funding University of Auckland Doctoral Scholarship. Waitemata District Health Board.

Competing interests None declared.

Provenance and peer review Not commissioned; externally peer reviewed.

REFERENCES

1 Ernst AT, Jiang H, Krishnamoorthy M, *et al*. An annotated bibliography of personnel scheduling and rostering. *Ann Oper Res* 2004;127:21–144.
2 Bailey J. Integrated days off and shift personnel scheduling. *Comput Ind Eng* 1985;9:395–404.

3 Rocha M, Oliveira JF, Carravilla MA. Cyclic staff scheduling: optimization models for some real-life problems. *J Scheduling* 2013;16:231–42.
4 Stolletz R, Brunner JO. Fair optimization of fortnightly physician schedules with flexible shifts. *Eur J Oper Res* 2012;219:622–9.
5 Ferrand Y, Magazine M, Rao US, *et al*. Building cyclic schedules for emergency department physicians. *Interfaces* 2011;41:521–33.
6 Kazemian P, Dong Y, Rohleder TR, *et al*. An IP-based healthcare provider shift design approach to minimize patient handoffs. *Health Care Manag Sci* 2014;17:1–14.
7 Smalley HK, Keskinocak P, Vats A. Physician scheduling for continuity: an application in pediatric intensive care. *Interfaces* 2015;45:133–48.

Rostering general medicine physicians to balance workload across inpatient wards: a case study

Thomas Adams, Michael O'Sullivan, Jonathan Christiansen, Paul Muir and Cameron Walker

BMJ Innov published online March 29, 2017