

Optimization Automates Emergency Department Nurse Scheduling at Hartford Hospital

Liangyuan Na

Operations Research Center, Massachusetts Institute of Technology, Cambridge, MA, 02139, USA, lyna@mit.edu

Jean Pauphilet

Management Science and Operations, London Business School, London NW1 4SA, UK, jpauphilet@london.edu

Ali Haddad-Sisakht, Louis Raison

Dynamic Ideas LLC, Waltham, MA, 02452, USA, {ali, lraison}@dynideas.com

Audrey Silver, Patricia Veronneau, Nicole Vogt

Hartford HealthCare, Hartford, CT, 06103, USA, {Audrey.Silver, Patricia.Veronneau, Nicole.Vogt}@hhchealth.org

Dimitris Bertsimas

Sloan School of Management, Massachusetts Institute of Technology, Cambridge, MA, 02139, USA, dbertsim@mit.edu

To optimize nurse staffing in the Emergency Department (ED), Hartford Hospital has been collaborating with academics and consultants to schedule nurse-shifts over each 6-week staffing cycle. We develop and implement two-phase optimization models: a robust optimization model to find optimal staffing levels given the uncertainty in patient demands, followed by a pair of mixed-integer problems to generate individual schedules including work, trainee, and preceptor shifts for each nurse. Our approach leads to less costly (5–8%) staffing with better coverage of patient care (8–25%) and higher nurse satisfaction (5%). Moreover, nurses can work fewer shifts on week-ends (17%), holidays (14%), and overtime (85%) as well as be assigned to more diverse positions (3.6) and more daily training opportunities (0.95). We implement our framework into an automated end-to-end scheduling optimization software, deployed for use at Hartford Hospital since March 2023. The software collects preferences from over 200 ED nurses and enables managers to optimize schedules with guided dynamic adjustments. This transformative implementation streamlines a previous labor-expensive staffing process (currently taking over 88 manual hours per cycle) and delivers schedules that are more suitable for patients and nurses together, with an annual projected cost saving of around \$720,000.

Key words: robust optimization; mixed integer optimization; software automation; nurse scheduling; emergency department

Introduction

Emergency Departments (ED) are a crucial component of hospitals, providing urgent care to patients in need. However, ED overcrowding can lead to increased waiting times, compromise care, and result in adverse outcomes for patients, such as dissatisfaction and

increased mortality rates (Bernstein et al. 2009). Additionally, ED congestion can affect patient flows into the hospital, further impacting care delivery (Elder et al. 2015).

Nurse staffing is vital for patients, especially ED patients in critical condition, having a direct impact on their health outcomes. Among other evidence, more hours of care provided by registered nurses, a lower nurse workload, and a higher nurse skillset have been associated with reduced mortality rate and shorter hospital stays (Needleman et al. 2002, Aiken et al. 2014, Twigg et al. 2019). For the ED as a whole, shortage of nurses also negatively impacts throughput metrics, such as length of stay or abandonment (i.e., patient leaving without being seen; Ramsey et al. 2018). However, healthcare systems are facing increasing difficulties to meet appropriate nurse-to-patients ratios, due to nurse shortages, burnout, and dissatisfaction (Aiken et al. 2002), as exacerbated since the COVID-19 pandemic (Peters 2023). For example, systemic overtime and dissatisfaction with schedule flexibility are common reasons for nurses to leave the profession (Leineweber et al. 2016).

One challenge is to schedule the “right” number of nurses, i.e., enough to accommodate future patient demand without overstaffing and wasting precious and limited nursing resources. Another challenge is to increase the satisfaction of each nurse regarding their individual schedule, by taking into account their individual preferences and requirements. To improve both patients’ and nurses’ experiences, we use optimization for nurse staffing and scheduling in the ED and address the aforementioned challenges. Our final objective is to positively impact nurses, patients, as well as the overall performance of the ED and the hospital, which we achieve through a tight and long-standing collaboration between ED practitioners and academics.

ED Nurse Staffing Problem at Hartford Hospital

Hartford HealthCare is Connecticut’s largest healthcare network, operating in over 400 locations and providing a broad range of services. Hartford HealthCare comprises 7 hospitals, with over 2,400 inpatient beds. Hartford HealthCare’s flagship facility is Hartford Hospital (HH), a teaching hospital and tertiary care center with 867 beds, in partnership with the University of Connecticut School of Medicine. With over 160 years of experience, HH is renowned for its exceptional performance in various procedures, conditions, and specialties. The hospital serves over 40,000 discharges and 100,000 ED visits annually.

Nurses working in the ED are typically assigned to one of 13 ‘positions’. Five positions cover managerial and logistical duties: Clinical Leader (CNL), first nurse, resource nurse,


triage nurse, and Front End Provider (FEP). The remaining eight positions, also referred as ‘pods’, correspond to treating different categories of patients. There are four main pods (blue, green, orange pods, and their hallways) where most of the patients are treated, one red pod for resuscitation/emergent patients, one purple pod for behavioral patients, an iTrack pod for less urgent patients, and one Emergency Department Observation Unit (EDOU) for observation patients. This organization was disrupted during the early surge of COVID-19: HH opened up a new position, a ‘trailer pod’, to treat COVID-19 patients specifically, and stopped using the hallways to reduce patient cross-infection.

Every nurse cannot work at any position. At HH, nurses are classified into nine different tiers based on their experience and qualifications, which dictate the positions they can fill. In addition, based on their years of experience, nurses can receive preferential treatment regarding week-end schedules (see Anderson et al. 2022, for another example of ‘Armstrong’ seniority requirements in nurse scheduling).

Each day is divided into three 12-hour shifts (7 am–7 pm, 11 am–11 pm, and 7 pm–7 am). Accordingly, ED nurse scheduling, i.e., the problem of assigning nurses to pods and shifts, can be divided into two phases. First, one needs to estimate the ‘right’ number of nurses needed at each pod for each future shift, i.e., staffing targets. Second, one needs to find a reasonable assignment of individual nurses to shifts in order to achieve these targets. In addition, at this phase, the ED leadership might want to schedule additional training shifts, where a trainee seconds a preceptor.

Regarding the first phase, prior to our effort, HH was using predetermined staffing targets. These targets were the same for all days (among others, they did not differentiate between weekdays and week-ends). For most positions (i.e., all positions except triage and iTrack), the staffing targets were also constant across shifts. However, patient arrivals, hence the need for ED nurses, vary across days and within each day. Figure 1 demonstrates such inconsistency between fluctuating demand levels and fixed staffing levels (see more in Appendix Section Demand Patterns and Standard Scheduling). This contrast can lead to overstaffing at some times, increasing hospital staffing costs, and understaffing at other times, compromising patient care quality. Furthermore, COVID-19 has further complicated the situation by disrupting patient demand patterns and increasing employee attrition.

Regarding the second phase, ED nurse schedules were planned for 6-week periods according to the following timeline: Before each staffing cycle, nurses enter their preferences and



(a) Day-of-week Demand. (b) Hour-of-day Demand. (c) Daily Schedule.

Figure 1 Patterns of Average Patient Demand vs Standard Staffing Levels.

availability. Schedulers and nurse managers manually generate a schedule that aims to balance staffing levels and preferences between nurses. In addition, the ED leadership manually adds training sessions to the schedule whenever possible. After the schedule is announced, nurses can ask for amendments to better satisfy individual preferences and increase fairness among nurses. In turn, managers and schedulers need to constantly re-compute schedules to accommodate as many requests as possible. To announce the schedule, scheduling managers take an additional step to manually convert the schedule of approximately 3,600 nurse-shifts (on average, three weekly shifts per nurse for 200 nurses over 6 weeks) into a “team sheet” of a specified format. HH estimates that the entire scheduling process takes managers and schedulers 88 hours of manual work per cycle and can be prone to errors.

Since 2020, we, a collaboration of nursing providers, academics, and data consultants, have been working to improve the ED nurse scheduling process at HH. We develop an integrated approach consisting of a two-phase optimization methodology and a software implementation. First, we use a data-driven optimization approach that allocates limited staffing resources cost-effectively and determines sufficient staffing targets to reach appropriate nurse-to-patient ratios, taking into account variability in patient arrivals and needs. Second, we generate an optimized nurse-to-shift allocation that aims to match their individual requests and preferences, increase nurse satisfaction, as well as provide opportunities to schedule training shifts. Finally, we develop scheduling software that relieves manual labors and operations burdens for nurse managers and schedulers. Our tool was partially used at HH ED in March 2023 (in preparation for the staffing cycle of April 9 to May 20) and started to be fully used for future cycles. Overall, our decision-support tool leverages a combination of data, optimization, and software to achieve a holistic improvement for ED staffing.

Related Work

The operations research literature contains numerous methods for ED operations, which we summarize here. We refer to Saghafian et al. (2015) for a comprehensive review.

Simulation-based approaches use simulations of the ED environment to evaluate alternative scenarios and, combined with optimization, to better allocate resources (Chen and Wang 2016), staff (e.g., doctors, lab technicians, and nurses; see Ahmed and Alkhamis 2009), and particularly nurses (Draeger 1992). Alternatively, patient arrivals can be well approximated by Poisson processes with non-stationary (Kim and Whitt 2014) or uncertain (Maman 2009) arrival rates. Hence, queuing approaches have been used to model the ED and address staffing problems such as, recently, dynamic shift assignment (Chan et al. 2021) or surge staffing (Hu et al. 2021). There is a rich literature on formulating personnel scheduling problems in general (Brucker et al. 2011), and nurse scheduling problems in particular (Svirsko et al. 2019), as formal optimization problems with constraints and often multiple objectives (e.g., workload, staffing costs, individual preferences; see Rerkjirattikal et al. 2022, Hamid et al. 2020, Ang et al. 2018, Mohammadian et al. 2019), solved via iteratively solving for each objective or using goal programming techniques.

Optimization methods can account for uncertainty in their input parameters, here in patient arrivals/needs and in nurse availability. Lim and Mobasher (2011), Van Hulst et al. (2017), for example, use robust optimization to generate workforce plannings that are robust against uncertain per-patient workload. In addition to demand uncertainty, frequent change in nurse availability due to sickness or absenteeism is another bottleneck for the implementation of optimized schedules in practice. Clark et al. (2015) highlights the importance and challenges of shift rescheduling and advocates for mathematical models to support the rescheduling process. For example, Wickert et al. (2019) propose an integer optimization formulation for the nurse-shift rescheduling problem. In this work, we directly account for variability in patient arrivals and needs in the first phase of our optimization approach, by proposing a robust optimization problem to find appropriate staffing targets. While we do not directly account for uncertainty in nurse availability, we aim to maximize nurse engagement and adherence to the schedule by developing the second phase of our approach, which takes into account their individual preferences. In addition, the software interface we develop allows nurses to update their availability and preferences directly in the system and enables ED managers to dynamically re-compute schedules with the latest

input from nurses, hence materializing the rescheduling process advocated by Clark et al. (2015).

Most related to our work is Ang et al. (2018), who develop a mixed-integer sequential goal programming model to optimize nurse-to-patient ratio and shift preferences among others. They integrate their model into an online decision support system to ease practitioners' adoption. Their paper provoked a vibrant discussion in the nursing community with critics calling such solutions "unimplementable" (Park et al. 2022). Despite the abundance of academic work in the past decade on nurse scheduling, there is still a gap between research and practice, with supposedly only 30% of nurse scheduling models from research ever being implemented, let alone still being used today (Kellogg and Walczak 2007).

Another obstacle to implementation is the lack of mathematical modeling and optimization skills among nurses (Park et al. 2022). To bridge this gap and advance nursing science, we collaborate closely with nurses and nurse managers, to collectively reach an optimization formulation that can be realistically implemented. We also develop a user-friendly end-to-end software interface for nurses to use and we train nurse leadership to run the scheduling optimization on their own. Our model-software integration, jointly with trust built with the medical team, results in the successful deployment at HH.

Main Contributions

We summarize our contributions as follows.

1. We develop a two-phase optimization approach for ED nurse staffing: In the first phase, we leverage robust optimization and historical data to optimize aggregate nurse staffing levels. Then, taking these aggregate optimized staffing levels as inputs, we develop mixed-integer optimization models to generate an individual-level schedule that prioritizes individual nurse preferences and training opportunities.
2. On computational experiments, we demonstrate the benefit of our two-phase optimization approach: Our first-phase optimization model reduces costs by 5–8% during overstaffing periods and reduces insufficient nurse-to-patient-ratio coverage by 8–25% during understaffing periods. In addition, we conduct various experiments that compare different modeling variants, illustrate the trade-offs between the schedule's different objectives, and provide strategic insights for ED leadership. The second-stage model for individual-level schedules improves individual nurse satisfaction by 5%, while increasing position diversity, training opportunities, and fairness.

3. We implement our optimization models into an end-to-end scheduling software that covers all staffing steps, from configuration preparation and input collection, to schedule generation and output production. For every scheduling cycle, nearly 200 ED nurses enter their availability and preferences into the system. Through an interactive interface, nurse managers can run the optimization models on their own to generate, edit, and announce schedules in a few clicks. In addition to the benefit of optimization, this unifying interface and software, has accelerated our collaboration, revolutionized ED nurse scheduling, and reduced manual burdens to a minimum.

ED Aggregate Staffing Optimization

In the first phase, we develop a robust optimization model to determine the right number of nurses to staff in the ED, at each position and each shift, for the next six weeks.

Optimization Model Overview

Our key decision variables are the number of nurses to staff for each position j , each shift i , and each day d of the next 42-day period. We differentiate nurses based on their tier q and their week-end group g (some nurses can work every other week-end while others work every third week-end only). So, we denote z_{qgjid} the number of nurses of tier q and week-end group g that are staffed at position j on the shift i of day d .

We then evaluate our decision based on three concurrent objectives. First, based on historical data, we can estimate the number of nurses required to treat patients present in the ED, for each position and each hour. We call *insufficiency* the number of nurse-shifts missing, i.e., the gap between the number of nurse-shifts needed and the number staffed. Since the number of nurse-shifts needed is uncertain, we adopt a robust perspective and try to minimize the worst-case insufficiency, where the worst-case is taken over a range of possible scenarios we construct from data, as described in the following paragraph. Second, we want to minimize staffing costs which can be approximated by the sum of scheduled nurse-shifts, $\sum_{q,g,j,i,d} z_{qgjid}$ (one could easily add weights based on tiers for example). Finally, to increase the adoption of our recommendation, we want to minimize the number of changes recommended compared with the current schedule. We incorporate these three objectives into a single objective function with weight parameters, which we tune based on the desirable trade-offs.

Our model can incorporate many operational constraints. For example, we incorporate constraints accounting for the maximum number of weekly working hours for each

nurse, their week-end schedule patterns, or for imposing some week-over-week regularity in staffing levels. We detail our optimization formulation, its objective, and constraints in Appendix Section Aggregate Optimization Formulation.

Estimating and Modeling Demand

Using historical data, we estimate how many nurses are needed for each position and each hour.

Data Collection and Processing. HH’s IT system daily transfers records of ED patients to our data repository on Amazon Web Services using a secure file transfer protocol server. The data records detailed information about each patient such as Emergency Severity Index (i.e., acuity level), service, ICD-10 code, and discharge disposition (e.g., admitted to ICU, surgery, interventional radiology, or discharge). Based on current ED practice, we use this information to match every patient to an appropriate pod and a target nurse-to-patient ratio (1:1 or 1:2 for those in the red pod, 1:7 in the purple pod, 1:12 in iTrack, and 1:5 in the remaining ones). Hence, at any point in time (in practice, we aggregate the data at an hourly level), we compute the number of patients present in each pod and the number of nurses needed.

Modeling Uncertainty. Let us denote \tilde{h}_{jws} the number of nurses needed at position j on week $w \in \{1, \dots, 6\}$ at hour $s \in \{1, \dots, 24 \times 7 = 168\}$. The schedule for the next 6-week cycle is announced three weeks in advance. Hence, we start formulating, calibrating, and solving our optimization problem five weeks before the beginning of the cycle, as described in Figure 2. At this time, we estimate the number of nurses needed in the past six weeks and use it as our baseline estimate for demand in the coming six weeks, \bar{h}_{jws} . The actual demand for nurses \tilde{h}_{jws} , however, might defer from this baseline estimate, in part because it uses data that will be 10-week old by the time the schedule is implemented. Accordingly, we consider all demand vectors contained in an uncertainty set \mathcal{U} and calibrate the size of the uncertainty set based on the distance between demand vectors that are ten weeks apart, using data from $w_0 - 20$. The full uncertainty set construction is described in Appendix Section Aggregate Model Data.

Executing the Model

Solving the Robust Optimization Model. Due to the fact that both the decision variables and uncertainties are integers, duality theory is not applicable and thus a closed-form robust




Figure 2 Timeline of Uncertain Demand Modeling for Every Staffing Cycle.

counterpart cannot be obtained. Hence, we employ a cutting plane approach to identify a subset of the most restrictive constraints and approximate the worst-case objective value among the uncertainty set. All models in this work are implemented using the JuMP package (Dunning et al. 2017) in Julia programming language (Bezanson et al. 2017) and are solved by the Gurobi solver (Gurobi Optimization, LLC 2023). Details of solving the model, such as constraint linearization, algorithmic and programming implementation, and specifications are included in Appendix Section Solving the Robust Optimization Model.

Parameter Tuning. In preparation for each upcoming staffing cycle starting on week w^0 , we can manually tune a combination of parameters, including two weights in the objective terms, the size of the uncertainty set, and relative penalty for each week’s insufficiency (penalize more on later week’s demand from the uncertainty set). To validate our choices, we examine the performance of different combinations of parameters over a validation period of $w^0 - 10$ to $w^0 - 5$. By comparing the values of the objective terms, we determine the optimal combination of parameters. For comparative purposes, we also consider a deterministic model where we only use the nominal demand instead of incorporating the uncertainty set. We demonstrate some of the process in Section Results.

Quantifying Strategic Decisions.

Our model not only optimizes staffing levels for each cycle but can also help nurse leadership inform some strategic decisions for the ED. For example, we can evaluate the impact of having staffing targets that are the same for all days or the same from one week to another. We can simulate the benefit of creating a fourth shift (e.g., 2 pm–2 am). Details on incorporating these variations into the model are presented in Appendix Section Model Variants for Informing Strategic Decisions.

ED Individual Scheduling Optimization

For the second phase, we develop another optimization model to schedule individual nurses based on the recommended aggregate staffing levels.

Optimization Model Overview

Our core decision variables are individual nurse assignments, i.e., when and where each nurse works (and trains). We use binary variables $b_{\ell jid}$, $r_{\ell jid}$, and $p_{\ell jid}$ to denote whether nurse ℓ works, trains, or mentors, respectively, at position j during shift i on day d .

Our primary objective is to maximize nurse satisfaction across various metrics, such as individual preference on dates, times, and shift patterns, diversity in pod assignments, and reduction in week-end, holiday, and overtime shifts. Meanwhile, we penalize shortages and surpluses in aggregate staffing levels while rewarding scheduled training shifts. We jointly optimize for a combination of four objective terms with adjustable trade-offs between the components.

In addition, we include a wide range of constraints. Some are to keep track of the different objective values, such as nurse satisfaction scores and deviations from aggregate staffing levels. Other constraints enable a feasible assignment, subjective to labor regulations, ED requirements, and eligibility, among others. We present a detailed formulation of the optimization model in Appendix Section Individual Shift Scheduling Model.

Individual Model Data

One main input of the second phase model is the output from the first phase, namely the number of nurses of each tier/group to schedule for each position during each shift on each day. In addition, the model takes into account a range of individual preferences and the availability of each nurse:

- Feasibility and preference to work at each shift type.
- Total number of weekly shifts, Paid Time Off (PTO), and education time of each week.
- Week-end group.
- Preference and availability to work on each day.
- Preference for different work patterns. Typically, nurses can express a preference for one of three patterns:
 1. All shifts in a row (e.g., working on Monday, Tuesday, and Wednesday consecutively),
 2. Every other day (e.g., Monday on, Tuesday off, Wednesday on, ...),
 3. Two on–two off (e.g., Monday and Tuesday on, Wednesday and Thursday off, ...).
- Eligibility to be a trainee or a preceptor (training the trainees) at each position.

Integer Optimization Model on Preceptor Scheduling

After obtaining the optimal solution for scheduled work and trainee shifts, we develop and solve another mixed integer linear optimization model to schedule preceptor shifts to train the corresponding trainees. The overall goal is to distribute trainees more evenly among eligible preceptors. Thus we set the objective function to minimize the maximum number of preceptor shifts assigned to each nurse. We formulate the preceptor scheduling optimization model in Appendix Section Individual Preceptor Scheduling Model.

Flexible Options for Nurse Leadership

We provide various options for ED nurse leadership to generate alternative schedules based on their preferences. These options include:

- The option to forbid shift type change from each nurse's pre-assigned shift type.
- The ability to exclude trainee scheduling.
- The option to decompose the model into ED and EDOU and solve for each of the two partitions separately.
- Control over the priority importance of different terms in the objective function by adjusting the weighting factors. Currently, the order of importance is shortage, satisfaction, unassigned/surplus, and then training. If training is not intended to be scheduled, it can be excluded by setting its weight to zero.
- The option to penalize shortages differently for each position j through weights w_j^{shortage} . Currently, positions that require more senior nurse tiers (such as CNL and triage) have higher penalization weight for the shortage.
- Fine-tuning of the relative importance of satisfaction score metrics using λ 's. To reduce model solving time, the preferred pattern and/or unpreferred pattern term(s) can be turned off by assigning a weight of zero if desired.
- Inclusion of fairness among nurses as a hard constraint controlling worst-case penalty score, or as a soft constraint by adding the worst-case score term in the objective function with a corresponding weight factor. Alternatively, fairness can be excluded altogether, which could also reduce solving time.

Results

In this section, we illustrate the benefits of our two-phase optimization approach. We decide to report these results chronologically, by following the timeline and progress of the

collaboration, for three reasons. First, we believe other researchers might find it insightful to know the successive iterations and milestones our collaboration has gone through. Second, because of the COVID-19 pandemic, the number of ED visits significantly changed between the first (end of 2020) and the latest (end of 2022) evaluations of our model. Accordingly, the value of optimization switched from reducing staffing costs to reducing insufficiency risk. Finally, at this stage in the project, we had not developed our user-friendly interface yet. Hence, iterations on the models between the research and nursing teams were slow and primarily paced by the 6-week staffing cycles. These frictions in our collaboration were the prime motivation to develop and deploy software for nurses to use directly.

End of 2020: Reducing Staffing Costs during Low-Demand Periods

We first evaluated the benefit of our first-stage optimization problem (staffing levels) at the end of 2020, using data from October 26 to December 20, 2020. In particular, we generated three instances of 6-week staffing cycles by considering three overlapping 6-week periods (beginning one week apart) from October 26 –December 6 to November 9–December 20.

We measure the quality of our staffing levels based on two metrics: *Cost*, i.e., the total number of nurses staffed, on average per day. Note that this cost is not subject to any uncertainty. *Insufficiency*, i.e., the number of extra nurse-shifts needed to satisfy the demand, on average per day. We obtain this number by computing the number of extra nurses needed to satisfy the target nurse-to-patient ratios on the realized demand at each hour of the day, and sum across all hours in a shift and all shifts in a day. We then average over all days in the 6-week cycle. For example, an average insufficiency of 0.17, as is the case for the historical schedule (see Table 1), means that we need to add on average 0.17 nurse-shifts per day to satisfy the target nurse-to-patient ratios. We aim to reduce both costs and insufficiency.

Table 1 reports the daily staffing cost and insufficiency for different policies, as well as their relative reductions compared with the current staffing levels. The current staffing levels require 56 nurse-shifts per day and lead to an average insufficiency of 0.17 nurse-shifts. As a comparison, we compute the optimal staffing level obtained by an oracle that knows the realized demand for the coming 6-week cycle. It reduces daily costs to 49.81 (11.05% reduction) and insufficiency to 0.02 (86.96% reduction), suggesting significant potential for improvement from using analytics, on both quality measures. Although the potential to reduce insufficiency is significant in relative terms, we should emphasize that

Table 1 Results of Average Outcomes (\pm Standard Deviation) from 6 Approaches.

Schedule approach	Shift types	Daily cost	Cost reduction [%]	Daily insufficiency	Insufficiency reduction [%]
Current	7-7, 11a-11p	56	-	0.17 (\pm 0.06)	-
Oracle	7-7, 11a-11p	49.81 (\pm 0.46)	11.05 (\pm 0.82)	0.02 (\pm 0.00)	86.96 (\pm 0.33)
Non-robust	7-7, 11a-11p	50.42 (\pm 0.38)	9.95 (\pm 0.67)	0.21 (\pm 0.04)	-22.55 (\pm 79.57)
Robust (a)	7-7, 11a-11p	51.33 (\pm 0.50)	8.33 (\pm 0.90)	0.12 (\pm 0.03)	27.70 (\pm 81.31)
Robust (b)	7-7, 12p-12a	50.57 (\pm 0.49)	9.69 (\pm 0.88)	0.15 (\pm 0.07)	13.18 (\pm 72.64)
Robust (c)	7-7, 11a-11p, 2p-2a	50.57 (\pm 0.38)	9.69 (\pm 0.67)	0.15 (\pm 0.07)	11.45 (\pm 74.95)

the current schedule experiences low insufficiency levels in absolute terms so we are primarily interested in reducing costs at this stage. While these results are reported under a certain selection of hyper-parameter values, we present metric trade-offs under different parameters in Appendix Section Supplementary Results during Low Demand Periods.

Our robust optimization approach (Robust (a)) recovers part of the benefit from the oracle by both costs and insufficiency, to 51.33 (8.33% reduction) and 0.12 (27.70% reduction) respectively. Compared with the oracle, this optimization model does not use the realized demand for each 6-week cycle, but estimates it using historical data and constructs an uncertainty set around this estimate. Indeed, we observe that the non-robust version of this model, which takes demand estimates at their face value, reduces costs by a comparable amount (to 50.42) but significantly increases insufficiency compared with the current schedule because it fails to account for variability in nurse needs.

Finally, we leverage our optimization model to inform strategic decisions for the ED nurse leadership. For example, we consider moving the 11 am–11 pm shift one hour later (Robust (b)) or introducing a fourth shift, 2 pm–2 am (Robust (c)). Compared with Robust (a), we observe that these changes can further reduce costs, with comparable (yet slightly worse) insufficiency. After reviewing these changes with the nurse leadership, we decided to keep the existing shift structure and recommended the robust schedule (a).

We provide details in Appendix Section Supplementary Results during Low Demand Periods that illustrate how the recommended schedule patterns better match demand patterns.

2021: Accounting for Frictions to Implementation

Given the positive results obtained at the end of 2020, our optimization approach gained support from the hospital executives. Hence, in 2021, we worked closely with the nurse leadership to refine the model and anticipate any friction to implementation. In particular, it was during this phase that we started penalizing the number of changes compared with

the current schedule in our model for the staffing levels and that we started building the individual nurse-to-shift assignment model.

Thanks to the flexibility provided by our mathematical model, we were able to incorporate many requirements from the nurse leadership to better align the solution of our model with the needs of the nursing staff. For instance, we fixed the staffing level in the red pod (the one for the most critical patients) to the current level. We introduced the possibility for nurses to ‘float’ (or reassign) between pods, i.e., between main pods or from the red pod to the main pods when red pod demand is low. These improvements led to the first staffing level recommendation, displayed in Table 2a, for the May 30–July 10, 2021 cycle. The optimization-based solution schedules more 11 am–11 pm shifts, but fewer 7 am–7 pm, 7 pm–7 am, and week-end shifts, resulting in an 8% reduction in the total number of shifts and a better match between demand and staffing levels.

However, the nurse leadership was reluctant to implement such a significant reduction in staffing levels, citing various concerns. First, the current ED staffing levels had been in place for decades, thus providing stability and visibility to nurses. To address this concern, we added the third component in our objective function that penalizes the distance between the current and the proposed schedule. Second, despite long and extenuating working hours, nurses tend to prioritize quality of care (i.e., insufficiency) over cost reduction. Based on this observation, we refined our calibration of the trade-off parameters in the objective and of the size of the uncertainty set to reflect their preferences and worked on a software interface that allows them to explore different solutions for different configurations and chose the one they deem is best. Third, aware of the disruptions and non-stationarities due to the COVID-19 situation, they raised concerns about the relevance of past data to predict the next 6-week period demand. In response to this concern, we added weights to the definition of insufficiency to allow us to put more emphasis on the most recent data, which should be more predictive of future demand.

After taking into account all these changes, we obtained a second proposal shown in Table 2b, with less drastic changes, which made the nurse leadership more comfortable with our tool. These iterative adjustments have undeniably helped build a closer relationship and trust between the research and nursing teams.

Table 2 Recommended changes to overall staffing levels.

(a) Iteration 1

Pod Shift	Blue			Green			Orange			Purple		
	7a	11a	7p	7a	11a	7p	7a	11a	7p	7a	11a	7p
Sun	-1	0	-1	-1	0	-1	-2	0	-2	-1	0	-1
Mon	-1	+1	-1	0	0	0	-1	+1	-1	0	0	0
Tue	0	0	0	0	0	0	-1	+1	-1	0	0	0
Wed	0	0	0	0	0	0	0	0	0	0	0	0
Thur	-1	+1	-1	0	0	0	-1	+1	-1	0	0	0
Fri	0	0	0	0	0	0	-1	0	0	0	0	0
Sat	-1	+1	-1	-1	0	0	-2	+1	-2	-1	0	0

(b) Iteration 2

Pod Shift	Blue			Green			Orange			Purple		
	7a	11a	7p	7a	11a	7p	7a	11a	7p	7a	11a	7p
Sun	-1	0	-1	-1	0	-1	-2	0	-2	-1	0	-1
Mon	0	0	0	0	0	0	0	0	0	0	0	0
Tue	0	0	0	0	0	0	-1	+1	-1	0	0	0
Wed	0	0	0	0	0	0	0	0	0	0	0	0
Thur	0	0	0	0	0	0	0	0	0	0	0	0
Fri	0	0	0	0	0	0	-1	0	0	0	0	0
Sat	-1	+1	-1	-1	0	-1	-1	0	-1	-1	0	0

2021: Building Individual Schedules

The development of optimized staffing levels justifies the need to automate the design of individual schedules as well, because optimization might introduce more changes from one staffing cycle to the next than the current practice (which uses the same staffing levels for all cycles).

Based on staffing levels from Iteration 2, we solve individual schedules for the May 30–July 10, 2021 cycle on a roster of 110 ED nurses (excluding per diem nurses) and 26 EDOU nurses based on their preferences. Table 3 compares the standard schedule with the one returned by our optimization model (see Appendix Section Supplementary Results during Iteration Periods for additional results with alternative model variations). We observe that the optimization model reduces the number of daily nurse-shifts by 5%, with a higher reduction during week-ends (17%) and holidays (14%). Additionally, the optimization leads to a remarkable 85% reduction in overtime shifts. Such savings in staffing cost and nurse workload is achieved while keeping the insufficiency level below 0.1%. In addition to cost savings, the optimized schedule has mainly additional practical benefits: It creates more opportunities for training new nurses by scheduling nearly one training shift per day (compared with 0 in the current schedule). It provides more diversity to nurses by assigning them to 4.6 different positions across the staffing cycle compared with

Table 3 Comparison of Standard vs Optimized Schedule.


Metric	Standard	Recommended	Reduction
Work shifts per day	48.3	45.7	5%
Per week-end day	48.0	40.0	17%
Per holiday	48.5	41.5	14%
Overtime per day	1.3	0.2	85%
Insufficiency per day	< 0.1	<0.1	-
Training shifts per day	0	0.95	-
Different positions per nurse	1	4.6	-
Weighted dissatisfaction score	101	96	5%

1 currently. Finally, it decreases the individual nurses' dissatisfaction penalty score by 5%, while also ensuring fairness among their scores. Overall, our optimization significantly improves operational efficiency and nurse satisfaction in the ED.

2022: Reducing Insufficiency during High Demand Periods

In 2022, the ED patient volumes recovered to (and exceeded) their pre-COVID-19 levels compared with our initial evaluation in 2020. Consequently, HH changed the structure of the ED and decided to increase the number of staffed nurses. We now illustrate how the same model as the one we developed in 2020 (after some adjustments to account for new positions and tiers) continues to provide substantial benefits, yet primarily by reducing insufficiency instead of costs.

Figure 3a shows the insufficiency vs. costs trade-off for different solutions obtained from our optimization models with different weight parameters—i.e., varying the relative penalty applied to each objective component and varying the size of the uncertainty set. For simplicity, we do not impose a penalty on the changes compared with the current schedule for these simulations. First, let us observe that, due to higher demand, the insufficiency levels are much higher than previously experienced (around 5 nurse-shifts per day vs. 0.17). Qualitatively, this curve informs the ED leadership on how additional staffing translates into insufficiency reduction. Based on these simulations, we selected a combination of parameter values (yellow star) that was providing a strict improvement over the current schedule, both in terms of costs and insufficiency. We then fine-tuned the other parameters (penalty on deviation from the current schedule in the objective and weight on the most recent data in the definition of insufficiency) as displayed in Figure 3b. Our final parameter combination reduces insufficiency by 8% (from 5.08 to 4.66 nurse-shifts) with 2.38% lower staffing costs (from 66 to 64.43 daily shifts). We provide details of the parameter tuning process in Appendix Section Supplementary Results during High Demand Periods.



(a) Tuning insufficiency parameters (weights and uncertainty set size). (b) Fine-tuning other parameters (penalization on changes and recency).

Figure 3 Objective Trade-offs during Staffing Cycle December 4, 2022 - January 14, 2023.

To further facilitate automating the scheduling in adjustment to demand variation, we investigate the benefit of resolving the aggregate model and adjusting the schedule every week during the 6-week staffing cycle in Appendix Section Supplementary Results during High Demand Periods.

Implementation

In collaboration with a data consultancy company and a development team, we integrate our models into end-to-end software interface.

The process, described in Figure 4, outlines the four phases of the scheduling process: preparation, input collection, solution generation, and schedule production. The flowchart depicts the interactions between the software (blue rounded rectangle), ED nurse managers (red rectangle), and individual nurses (green oval) at HH. Our software revolutionizes the ED nurse scheduling process as follows: The process starts with nurse managers setting up general configurations and nurse information in the software, which then enables individual nurses to log into their accounts and enter general preferences. This phase is completed prior to the first use and can be updated whenever necessary, whereas the remaining three phases are performed for every 6-week staffing cycle. Approximately five weeks before the start of each cycle, the software collects nurses' and managers' requests specific to the upcoming cycle and then passes them as inputs to the optimization model. Four weeks before the cycle, the software automatically solves the first-stage aggregate model, allowing managers to edit the output aggregate levels and run the program to generate schedules

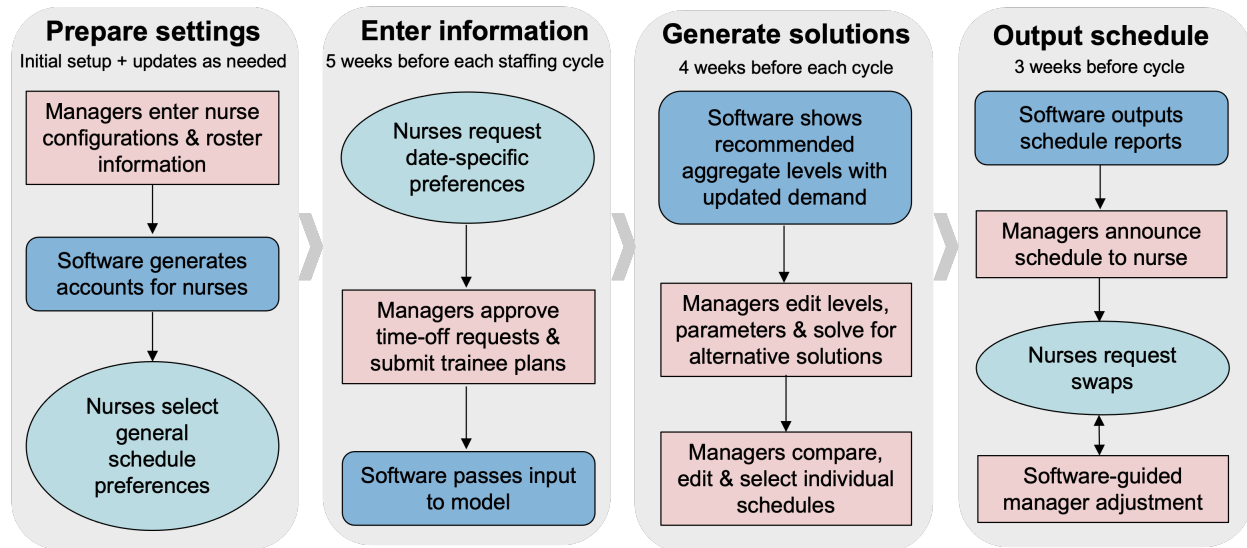


Figure 4 Process of Integrating Decision-support Software into ED Nurse Staffing.

with parameters of their choice. After managers edit, compare, and select the schedules, the software outputs reports based on templates for managers to announce to nurses three weeks ahead of the staffing cycle. If nurses request to swap or change shifts from the announced schedule, the software guides managers to accept or reject the changes based on staffing shortages and surpluses at each shift. The end-to-end implementation was partially used in March 2023 and fully deployed since April 2023.

Software Illustration

In this section, we demonstrate various components of the software from input collection to output generation.

Collecting input. The software has two sections to gather input from nurse managers and individual nurses. The Config tab enables nurse managers to set general staffing inputs: the RN Tier page specifies eligibilities for nurses of each tier to work at each position; the Groups page assigns nurses into different cohorts with corresponding unavailability dates for each cohort (e.g., holiday); the Schedule Date page defines the start and end date of the staffing cycle to be scheduled. These settings define the structure of the Employee tab shown in Figure 5. Managers (ED leadership, head nurses, and scheduling assistants) can modify the nurse roster by adding/removing employees or importing an Excel file. They can also edit each nurse's account settings and basic information. Once the forms are released, nurses can log into their accounts and enter their shift availability, shift pattern preferences, dates they prefer to work on, plus any request for paid or unpaid time off.

Figure 5 Software Section for Employee Information and Preference Collection.

After nurses enter their preferences, managers approve/deny their request for days off, hence validating the inputs for the subsequent optimization models.

Generating solutions. After collecting nurse input, nurse managers can use the software’s output sections. In the Solutions tab, managers can set some parameters (e.g., penalties for the different objective components), solve the model, and obtain a solution (identified using a user-specified name). They can review and compare the recommended aggregate staffing levels in a table, and select some modeling options such as whether or not to allow nurses to change shift types, to include training schedules, and to incorporate fairness considerations. The software offers additional flexibility to solve for a subset of nurses or tiers, such as solving for ED only or EDOU nurse only. The managers can experiment with different parameters and compare alternative aggregate/individual solutions, as shown in Figure 6. In the schedule table in Figure 6, each nurse is represented by a row containing the assigned shifts, positions (in different colors), and dates for that nurse.

Adjusting schedules. Due to dynamic changes in staffing needs and nurse availabilities, the support for schedule adjustments is deemed essential. As seen at the bottom of Figure 6, aggregate staffing levels from the individual schedule above (first number) are presented in comparison to target levels (second number) for each shift and any subset of positions. For example, Friday, April 14 has a surplus of one nurse during the 7p–7a shift (highlighted in red), whereas Saturday, April 15, has a shortage of one nurse during the 7a–7p shift (highlighted in orange). Every week, the software also resolves the model with updated demand and alerts managers of any expected changes. This component raises managers’ attention on potential shortages or surpluses across the pods, suggesting adjustments to be made. Also demonstrated in Figure 6, managers can execute any schedule changes by editing shifts, dates, positions, or nurses with simple clicks. The combination of these functionalities supports a variety of schedule adjustments: It suggests how each shortage can be filled if specific nurses can work overtime, advising managers of specific nurses that can help reduce the shortages best. Since some nurses sometimes request shift swaps due to changes in their availability, the aggregate summary helps their managers approve these requests by visualizing whether the affected shifts are currently experiencing a surplus or shortage of nurses. Similarly, if two nurses request to swap with each other, the software automatically checks the feasibility of the switch. To cover additional shortages, managers can add shifts on per diem nurses (who do not work full-time but pick up shifts on demand anytime) to fill in the shortage shifts.

Announcing Outputs. Finally, the Schedule tab shows the determined schedule, where nurse managers have a holistic view and each nurse can view their own schedule. In addition, the Report tab enables managers to download and print reports of the schedule based on the ED template, including a summary of staffing for each cycle and a daily team sheet for each date. Figure 7 includes an example team sheet on April 9, 2023, which assigns nurse names at each position during each time block for both day and night shifts of the day. Overall, the software provides decision support as well as aligns with ED operations.

Financial Benefit Estimation

The schedule optimization with its implementation translates into substantial financial benefits projected by HH. The reduction of nursing hours from Table 3 and Table 4 converts into estimated cost savings in Table 5. By saving 2.6 shifts (12-hour long) daily, ED will save an average of 31.2 nurse-hours per day, 11,388 hours per year. With overtime salary

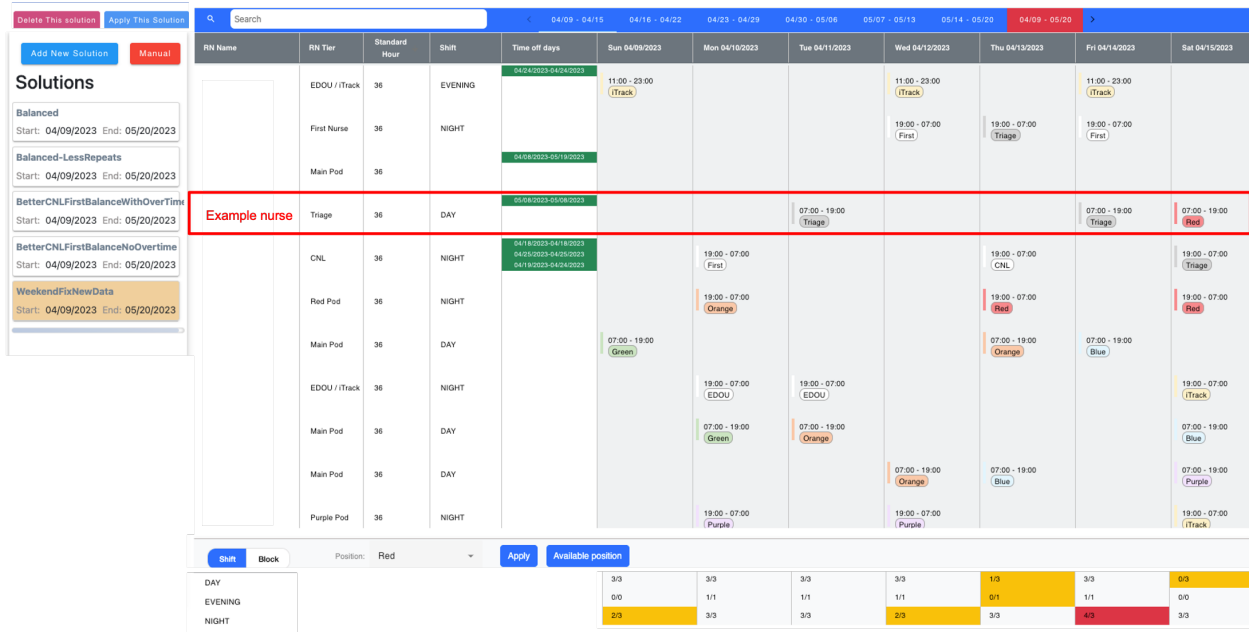


Figure 6 Software Section for Generating and Adjusting Alternative Individual Schedules.

		ED ASSIGNMENT RECORD: 2023-04-09 DAY SHIFT			ED ASSIGNMENT RECORD: 2023-04-09 NIGHT SHIFT	
		07:00:00-11:00:00	11:00:00-15:00:00	15:00:00-19:00:00	19:00:00-23:00:00	23:00:00-07:00:00
CNL	RN name	RN name	RN name	RN name	RN name	RN name
Triage	RN name	RN name	RN name	RN name	RN name	RN name
Triage	RN name	RN name	RN name	RN name	RN name	RN name
Triage	RN name	RN name	RN name	RN name	RN name	RN name
Triage	RN name	RN name	RN name	RN name	RN name	RN name
Blue						
	12, 13, 14, 15, 16 LEAD	RN name	RN name	RN name	RN name	RN name
	7, 8, 9, 10, 11	RN name	RN name	RN name	RN name	RN name
	17, 18, 19, 20, 21	RN name	RN name	RN name	RN name	RN name
	22, 23, 24, 25, 26	RN name	RN name	RN name	RN name	RN name
Green						
	38, 39, 40, 41 LEAD	RN name	RN name	RN name	RN name	RN name
	27, 28, 29, 32, 32H	RN name	RN name	RN name	RN name	RN name
	30, 31, 33, 34, 34H	RN name	RN name	RN name	RN name	RN name
	35, 35H, 36, 36H, 37	RN name	RN name	RN name	RN name	RN name
Orange						
	73, 74, 75, 83, 84 LEAD	RN name	RN name	RN name	RN name	RN name
	62, 63, 64, 65, 66, 67	RN name	RN name	RN name	RN name	RN name
	68, 69, 70, 71, 72	RN name	RN name	RN name	RN name	RN name
	75, 77, 78, 80a, 80b	RN name	RN name	RN name	RN name	RN name
	79, 81a, 81b, 82a, 82b					
Red						
	Blue Functional Float (BFF)	RN name	RN name	RN name	RN name	RN name
	Green Functional Float (GFF)	RN name	RN name	RN name	RN name	RN name
	Orange Functional Float (OFF)	RN name	RN name	RN name	RN name	RN name
Purple						
	West RN (LEAD)	RN name	RN name	RN name	RN name	RN name
	West RN	RN name	RN name	RN name	RN name	RN name
	South RN	RN name	RN name	RN name	RN name	RN name
	ITrack	RN name	RN name	RN name	RN name	RN name
	ITrack	RN name	RN name	RN name	RN name	RN name
	ITrack	RN name	RN name	RN name	RN name	RN name
	EDOU	RN name	RN name	RN name	RN name	RN name
	EDOU	RN name	RN name	RN name	RN name	RN name
	EDOU	RN name	RN name	RN name	RN name	RN name
	EDOU	RN name	RN name	RN name	RN name	RN name
	First			RN name	RN name	

Figure 7 Software Section for Outputting Final Schedules.

being 50% above the base salary, the reduction of 1.1 overtime shifts per day cuts the cost by an additional \$25 per hour. Combined with the 88-hour manual work per 6 weeks

Table 4 Estimated Time Spent by ED Leaders on Manual Scheduling.

Responsibility component	Responsible staff	Total time
Schedule build/balancing	Scheduler	12 hours
Schedule approval	Scheduler/Manager	1 hour
Management of shift switches	Scheduler/Manager/Assistants	4 hours weekly
Orientation scheduling	Educators/Managers/Scheduler	3 hours
Daily team sheet building	Manager/Assistants	8 hours weekly
Total hours per 6-week staffing cycle		88 hours

Table 5 Projected Savings of Nursing Hours and Cost for HH ED.

Component	Daily hours	Annual hours	Hourly cost	Annual cost
Base shifts	31.2	11,388.00	\$50	\$569,400.00
Overtime (extra)	13.2	4,818.00	\$25	\$120,450.00
Manual scheduling	2.1	764.76	\$50	\$38,238.10
Total				\$728,088.10

that could be saved from Table 4, the hospital financial department projects that the optimization could save \$728,088.10 annually in nursing costs.

Overcoming Challenges in Deployment

Deploying an optimization tool for ED nurse staffing, which impacts the lives of 200 nurses, nurse schedulers, and managers, carries numerous challenges. The deployment of the tool also involved transitioning from an offline to an online scheduling process. In this section, we discuss how we overcame these challenges, progressed with the deployment, and achieved practical impact.

Offline strategic adoption. The development started with offline iterations in late 2020. Researchers generated schedules offline for ED leadership to review, and then incorporated feedback and constraints into the models. As described in Section 2021: Accounting for Frictions to Implementation, researchers adapted the optimization models and parameters to overcome the nurse leadership’s conservatism in reducing staffing levels. With refined models, the team communicated with hospital leadership to adopt our optimized schedules. The model suggested fewer 7-7 shifts and more 11-11 shifts in response to demand

patterns. However, as each nurse was pre-assigned to a fixed shift type upon being hired, it was difficult to swap shifts without disrupting their lives. After strategic discussions, hospital leadership decided to hire more 11 am–11 pm shift nurses as opposed to 7-7 shift nurses to be able to realize the recommendation. The ED reports having tripled the number of 11 am–11pm positions (from 4.9% to 14.7% of total nurses). During this overstaffing period, the optimized schedule needed fewer nurse-shifts than previously as shown in Section End of 2020: Reducing Staffing Costs during Low-Demand Periods. This presented another challenge as the hospital was still required to utilize all nursing hours as per the contract. After evaluating several options, the ED decided to turn the saved work shifts into additional training shifts, which are counted as working hours but do not serve patient demand. Besides, scheduling fewer week-end shifts also changed the ED’s conventional policy to move nurses from working every other week-end to every third week-end after working for two years as a seniority incentive. Since January 2021, the ED was able to move a total of 80 nurses from every other week-end to every third week-end early (between two and seven months prior to the 2 year mark), which enhances retention rewards to nurses.

Motivation from offline to online. Following the successful strategic adoptions at the ED, our next goal was to implement the generated individual nurse schedule for the ED to use. However, experimentation for the next staffing cycle showed that offline iterations between researchers solving models and managers providing feedback were not feasible to be adopted. Managers needed to propose changes to the schedule frequently for reasons such as nurse dissatisfaction, availability changes, and manager preferences. The frequent back-and-forth communications with the research teams caused delays and were not sustainable. On the other hand, manually reshuffling the machine-generated schedule was infeasible for nurse managers without optimization training. To address these issues, we decided to develop our software tool, which automates the optimization part and allows nurse managers to regenerate and edit schedules in an online fashion.

Online implementation. In 2022, we began the software development for staffing automation. The research team collaborated with the consultancy and IT teams and integrated the optimization models into the software. Given imported input data, users were able to solve the models and obtain the schedule under the Solution and Schedule tabs. However, the upcoming main challenge was to connect the software with the hospital to collect input from nurses and pass output for the ED to use. One complication arose as the entire nurse

staffing at Hartford HealthCare relied on commercial software that links shift assignments to the nurse payroll system. Despite the team’s attempts to integrate with the commercial software throughout 2022, the complexity of involving a third-party organization was beyond our control. After careful discussions between the team and executives at Hartford HealthCare, a decision was made to build a standalone software instead. In early 2023, we built the information collection component of the software. After frequent weekly meetings and iterations between researchers, the development team, nurse managers, and schedulers, the interface evolved to cover all ED nurse staffing functionalities. We further improved the schedule generation process, enabling dynamic schedule changes by nurse leadership to adapt to quick information changes. The ED executed more model recommendations after the long collaboration, such as swapping shift types for some nurses. The fine-tuned schedule announced on April 30, 2023 has no overtime shifts, satisfies 92% preferred date assignments, only 1 undesirable shift pattern among all nurses across 6 weeks, and assigns nurses with desirable diversity to average 2.43 different positions per week. The impact of the tool’s deployment can be evaluated more deeply as more data after the implementation is collected in the future. Following the pilot ED deployment at HH, the Hartford HealthCare executive team aims to extend the models and software to cover all nurse staffing in 7 hospitals of the network in the future.

Conclusions

Under a collaboration between Hartford HealthCare, MIT, and Dynamic Ideas, we develop and implement optimization models to automate the nurse staffing process in the Emergency Department at HH. Our methodology consists of two phases to optimize each 6-week staffing cycle. First, we learn an uncertainty set from patient demand data and develop a robust optimization model, solved via a cutting planes algorithm to compute aggregate staffing levels. Next, we develop two mixed-integer optimization models to assign individual nurses to work, trainee, and preceptor shifts across the coming 6-week period.

Experimental results demonstrate the versatile benefit of the first-phase aggregate model: reducing staffing costs by 5–8% in pre-COVID periods where demand was relatively low compared to nurse supply versus reducing insufficiency by 8–25% since the COVID-19 pandemic. In addition, we analyze how the outcomes change with different model variations and parameter trade-offs as well as illustrate schedule iterations and demand patterns.

The second optimization phase produces individual schedules of nurses to shifts and brings significant benefits to ED nurses by reducing 17% week-end, 14% holiday, and 85% overtime shifts while increasing 5% satisfaction score, 3.6 more diverse positions, and 0.95 training per day.

The models are integrated into end-to-end software that supports scheduling from staffing preparation, input collection, and solution generation, to schedule output. After overcoming numerous implementation challenges, the software was deployed starting March 2023 at HH, automating and transforming the nurse scheduling at the Emergency Department. The implementation relieved manual scheduling burdens (78 hours per 6-week period) and is projected to save \$728,088.10 in annual nursing costs. The integration brought more cost-effective, sufficient, and desirable staffing into practice, benefiting various stakeholders in the hospital system.

Author Contributions

L.N. led the efforts of model development, planning and performing experiments, writing code, analyzing results, and writing the manuscript, and contributed to software implementation. J.P. contributed to the design of research, modeling, and experiments as well as manuscript writing. A.H-S. and L.R. implemented the software and maintained the production. A.S., P.V., and N.V. led the nursing team, executed deployment at the hospital, provided feedback, and contributed to the model improvement. D.B. directed the overall project, from concept and research to implementation, and edited the manuscript.

Acknowledgments

The authors would like to thank the Hartford HealthCare team, especially Cheryl Ficara, Barry Stein, Yolanda Johnson, and Katherine Martinez-Taveras, for their support on the implementation.

References

- Ahmed MA, Alkhamis TM (2009) Simulation optimization for an emergency department healthcare unit in kuwait. *European Journal of Operational Research* 198(3):936–942.
- Aiken LH, Clarke SP, Sloane DM, Sochalski J, Silber JH (2002) Hospital nurse staffing and patient mortality, nurse burnout, and job dissatisfaction. *JAMA* 288(16):1987–1993.
- Aiken LH, Sloane DM, Bruyneel L, Van den Heede K, Griffiths P, Busse R, Diomidous M, Kinnunen J, Kózka M, Lesaffre E, et al. (2014) Nurse staffing and education and hospital mortality in nine european countries: A retrospective observational study. *The Lancet* 383(9931):1824–1830.

- Anderson M, Bodur M, Rathwell S, Sarhangian V (2022) Optimization helps scheduling nursing staff at the long-term care homes of the city of Toronto. *INFORMS Journal on Applied Analytics* 53(2):133–154.
- Ang BY, Lam SWS, Pasupathy Y, Ong MEH (2018) Nurse workforce scheduling in the emergency department: A sequential decision support system considering multiple objectives. *Journal of Nursing Management* 26(4):432–441.
- Bernstein SL, Aronsky D, Duseja R, Epstein S, Handel D, Hwang U, McCarthy M, John McConnell K, Pines JM, Rathlev N, et al. (2009) The effect of emergency department crowding on clinically oriented outcomes. *Academic Emergency Medicine* 16(1):1–10.
- Bertsimas D, den Hertog D, Pauphilet J (2021) Probabilistic guarantees in robust optimization. *SIAM Journal on Optimization* 31(4):2893–2920.
- Bezanson J, Edelman A, Karpinski S, Shah VB (2017) Julia: A fresh approach to numerical computing. *SIAM Review* 59(1):65–98.
- Brucker P, Qu R, Burke E (2011) Personnel scheduling: Models and complexity. *European Journal of Operational Research* 210(3):467–473.
- Chan CW, Huang M, Sarhangian V (2021) Dynamic server assignment in multiclass queues with shifts, with applications to nurse staffing in emergency departments. *Operations Research* 69(6):1936–1959.
- Chen TL, Wang CC (2016) Multi-objective simulation optimization for medical capacity allocation in emergency department. *Journal of Simulation* 10(1):50–68.
- Clark A, Moule P, Topping A, Serpell M (2015) Rescheduling nursing shifts: Scoping the challenge and examining the potential of mathematical model based tools. *Journal of Nursing Management* 23(4):411–420.
- Draeger MA (1992) An emergency department simulation model used to evaluate alternative nurse staffing and patient population scenarios. *Proceedings of the 24th Conference on Winter Simulation*, 1057–1064.
- Dunning I, Huchette J, Lubin M (2017) Jump: A modeling language for mathematical optimization. *SIAM Review* 59(2):295–320.
- Elder E, Johnston AN, Crilly J (2015) Systematic review of three key strategies designed to improve patient flow through the emergency department. *Emergency Medicine Australasia* 27(5):394–404.
- Gurobi Optimization, LLC (2023) Gurobi Optimizer Reference Manual. URL <https://www.gurobi.com>.
- Hamid M, Tavakkoli-Moghaddam R, Golpaygani F, Vahedi-Nouri B (2020) A multi-objective model for a nurse scheduling problem by emphasizing human factors. *Proceedings of the Institution of Mechanical Engineers, Part H: Journal of Engineering in Medicine* 234(2):179–199.
- Hu Y, Chan CW, Dong J (2021) Prediction-driven surge planning with application in the emergency department. *Submitted* .
- Kellogg DL, Walczak S (2007) Nurse scheduling: From academia to implementation or not? *INFORMS Journal on Applied Analytics* 37(4):309–399.

- Kim SH, Whitt W (2014) Are call center and hospital arrivals well modeled by nonhomogeneous poisson processes? *Manufacturing & Service Operations Management* 16(3):464–480.
- Leineweber C, Chungkham HS, Lindqvist R, Westerlund H, Runesdotter S, Alenius LS, Tishelman C, et al. (2016) Nurses’ practice environment and satisfaction with schedule flexibility is related to intention to leave due to dissatisfaction: A multi-country, multilevel study. *International Journal of Nursing Studies* 58:47–58.
- Lim G, Mobasher A (2011) Robust nurse scheduling problem. *Proceedings of the 2011 Industrial Engineering Research Conference*, 1 (Institute of Industrial and Systems Engineers (IISE)).
- Maman S (2009) *Uncertainty in the Demand for Service: The Case of Call Centers and Emergency Departments*. Ph.D. thesis, Technion–Israel Institute of Technology.
- Mohammadian M, Babaei M, Amin Jarrahi M, Anjomrouz E (2019) Scheduling nurse shifts using goal programming based on nurse preferences: A case study in an emergency department. *International Journal of Engineering* 32(7):954–963.
- Needleman J, Buerhaus P, Mattke S, Stewart M, Zelevinsky K (2002) Nurse-staffing levels and the quality of care in hospitals. *New England Journal of Medicine* 346:1715–1722.
- Park CSY, Kabak M, Kim H, Lee S, Cummings GG (2022) No more unimplementable nurse workforce planning. *Contemporary Nurse* 58(2-3):237–247.
- Peters M (2023) Time to solve persistent, pernicious and widespread nursing workforce shortages. *International Nursing Review* .
- Ramsey Z, Palter JS, Hardwick J, Moskoff J, Christian EL, Bailitz J (2018) Decreased nursing staffing adversely affects emergency department throughput metrics. *Western Journal of Emergency Medicine* 19(3):496–500.
- Rerkjirattikal P, Singhaphandu R, Huynh VN, Olapiriyakul S (2022) Job-satisfaction enhancement in nurse scheduling: A case of hospital emergency department in thailand. *Integrated Uncertainty in Knowledge Modelling and Decision Making: 9th International Symposium, IUKM 2022, Ishikawa, Japan, March 18–19, 2022, Proceedings*, 143–154 (Springer).
- Saghafian S, Austin G, Traub SJ (2015) Operations research/management contributions to emergency department patient flow optimization: Review and research prospects. *IIE Transactions on Healthcare Systems Engineering* 5(2):101–123.
- Svirsko AC, Norman BA, Rausch D, Woodring J (2019) Using mathematical modeling to improve the emergency department nurse-scheduling process. *Journal of Emergency Nursing* 45(4):425–432.
- Twigg DE, Kutzer Y, Jacob E, Seaman K (2019) A quantitative systematic review of the association between nurse skill mix and nursing-sensitive patient outcomes in the acute care setting. *Journal of Advanced Nursing* 75(12):3404–3423.

Van Hulst D, den Hertog D, Nuijten W (2017) Robust shift generation in workforce planning. *Computational Management Science* 14:115–134.

Wickert TI, Smet P, Berghe GV (2019) The nurse rostering problem: Strategies for reconstructing disrupted schedules. *Computers & Operations Research* 104:319–337.


Demand Patterns and Standard Scheduling

Demand trend over time. Figure A1a displays the number of ED patient arrivals for each day since late 2016. Before March 2020, the ED volume was about 250–350 patient arrivals per day with small fluctuations. At the beginning of the COVID-19 pandemic, the number of ED arrivals dropped significantly and then gradually increased throughout 2021, up to a level slightly lower than before, around 200–300 arrivals per day. As the pandemic progressed through different stages, people were recovering from the COVID-19 isolation period, which brought the ED volume back to 250–350 arrivals per day, while having larger fluctuations over time compared with pre-COVID-19. Given the significant disruption incurred during and since the COVID-19 outbreak, we use data starting from June 22, 2020 for our analysis, which we highlight in Figure A1b.

Demand patterns and baseline staffing levels. In Figure 1, we present the demand patterns from November 2 - December 13, 2020. In this period, the average number of patient stays over 6 weeks and the number of nurses recommended to staff each week in each pod type from Monday to Sunday (starting from 7 am to 7 am each day) are shown in Figure 1a and Figure A2a. For intra-day patterns, we show the average number of patients over 6 weeks in each pod type at each hour of the day (from hours starting at 0:00 to 23:00) in Figure 1b. For the main, red, and iTrack pods, we have higher demands in the afternoons and evenings, where the variability is highest for main pods. For the purple pod, demand is more constant throughout the day with slightly higher demand at nighttime. Conventionally, every day applies the schedule shown in Figure 1c for the number of nurses working in each position at each hour of the day. Only the staffing level in iTrack changes within the day, with more nurses in afternoons and evenings and fewer at night. However, in the red and main pods, where we also observe variability in demand (especially in the main pods), staffing levels remain constant throughout the day and throughout the week.

Robust Optimization Model for Aggregate Staffing

In this Appendix Section, we provide details on the robust optimization model for determining the aggregate staffing levels described in Section ED Aggregate Staffing Optimization.



(a) From November 7, 2016 to March 6, 2023.

(b) From June 22, 2020 to March 6, 2023.

Figure A1 ED Historical Volume Over Time.

Aggregate Model Data

Main Components. Allocating ED nurse staffing levels contains the following basic components.

ED positions: Nurses work among $J = 13$ positions: CNL, first nurse, resource, triage, FEP, blue, green, orange, hallway, red, purple pod, iTrack, and EDOU. As some of the pods have the same functionalities of treatment, we define $N = 5$ pod types as main pods (blue, green, orange, and hallway), red pod, purple pod, iTrack, and EDOU. As nurses can ‘float’ between iTrack, red, and main pods within the same shift, we define $M = 3$ pod floating groups where each group is a set of pods among which nurses can float between.

Nurses: Based on each nurse’s years of experience and training qualifications, they are categorized into $Q = 9$ nurse tiers, and are also classified by the hospital into $G = 2$ groups that restrict the frequency of week-end shifts to every other week-end or every third week-end.

Time: There are $I = 3$ shift types: 7 am–7 pm, 7 pm–7 am, 11 am–11 pm. The optimization model schedules for an entire staffing cycle (6-week period) with some week-over-week regularity. We divide each cycle into $W = 6$ weeks or $D = 42$ days or $T = 1,008$ 1-hour time periods. Since there might be some week-over-week regularity, we also introduce indices to divide a typical week of a cycle into $E = 7$ days or $S = 168$ hours.

Other sets: For each component of size J , we use the corresponding lower case letter j to index the element, which takes values in set $[J]$ denoting $\{1, 2, \dots, J\}$. We define other discrete subsets needed: $J_n, J_m \subseteq [J]$: sets of pods j of type n and of floating group m , respectively; $D_n \subseteq [D]$: week-end days of the cycle; $D_e \subseteq [E]$: days of the first week of the cycle; $D_w, D_{n_w} \subseteq [D]$: days and week-end days in week w , respectively.

Demand uncertainty set. For a 6-week cycle starting on week w^0 , we construct an uncertainty set around the demand for each pod floating group m , on each week $w \in \{w^0, \dots, w^0 + 5\}$ and each hour of the week $s \in [S]$, \tilde{h}_{mws} in the following way: First, we compute a nominal demand \bar{h}_{jws} from the historical demand on weeks $w^0 - 10 - w^0 - 5$, i.e.,

$$\bar{h}_{mws} := h_{ms(w-10)}^{\text{hist}}, \quad \forall w \in \{w^0, \dots, w^0 + 5\}.$$

Second, we estimate the pairwise absolute differences of historical demands that are 10 weeks apart, for the weeks between $w^0 - 20$ and $w^0 - 5$:

$$\hat{\epsilon}_{msw} = |h_{msw}^{\text{hist}} - h_{msw-9}^{\text{hist}}|, \quad \forall m \in [M], s \in [S], w \in \{w^0 - 11, \dots, w^0 - 5\}.$$

Finally, we define the final uncertainty set as:

$$\mathcal{U} := \left\{ \begin{array}{l} \left| \bar{h}_{msw} - \tilde{h}_{msw} \right| \leq \epsilon_{ms}^1, \quad \forall m \in [M], s \in [S], w \in [W] \\ \tilde{h}_{msw} \geq 0 \left| \sum_{s \in [S]} \bar{h}_{msw} - \sum_{s \in [S]} \tilde{h}_{msw} \right| \leq \epsilon_m^2, \quad \forall m \in [M], w \in [W] \\ \left| \sum_{m \in [M], s \in [S]} \bar{h}_{msw} - \sum_{m \in [M], s \in [S]} \tilde{h}_{msw} \right| \leq \epsilon^3, \quad \forall w \in [W]. \end{array} \right\},$$

where we calibrate ϵ_{ms}^1 (i.e, the bound on pairwise inter-10-week demand fluctuations) as the 80th percentile of the $\hat{\epsilon}_{msw}$ we estimated on historical data. We scale , for each m, s , and scale ϵ_{1ms} accordingly to obtain ϵ_{2m} and ϵ_3 . From large deviation bounds (see, e.g., table 2 Bertsimas et al. 2021), we know that ϵ_m^2 and ϵ^3 should scale as \sqrt{S} and \sqrt{MS} respectively, so we use

$$\epsilon_m^2 = \beta_\epsilon \frac{1}{\sqrt{S}} \sum_{s \in [S]} \epsilon_{ms}^1, \quad \forall m \in [M], \quad \epsilon^3 = \beta_\epsilon \frac{1}{\sqrt{MS}} \sum_{m \in [M], s \in [S]} \epsilon_{ms}^1,$$

with an adjustable discount factor $\beta_\epsilon \in (0, 1]$.

Other data and parameters. Other than the demand data, we collect the following data from ED leadership: Z_q, Z_q^g : total number of nurses of tier q , and those of week-end group g . z_{jie}^{curr} : number of nurses in position j during shift i on day e of the week in standard schedule. $\underline{n}_j, \bar{n}_j$: minimum and maximum number of nurses required at position j , respectively. \overline{Tn}_q : maximum total number of shifts each nurse of tier q can work in a week. N_U : set of pairs (q, j) such that nurse of tier q cannot work in position j due to qualifications. E_{qt} : number of nurses of tier q available to work during the one-hour period t . We also introduce σ_{idt} (or equivalently σ_{idsw}) to indicate whether shift i on day d contains time period t .

Aggregate Optimization Formulation

Decision Variables. In the next 6-week cycle, on each day of the week, we decide:

- $z_{qjidg} \in \mathbb{Z}^+$: number of nurses of tier q , week-end group g working at position j on shift i , day d .

To capture some aspects of our objectives to minimize insufficiency as well as deviation from the current schedule, we introduce two auxiliary decision variables:

- npr : the weighted sum of insufficiency in the worst case (w.r.t. the demand uncertainty set), where insufficiency is defined as the number of nurse-shifts missing to satisfy the target nurse-to-patient ratios to treat the patients.
- Δz_{jie} : the absolute difference in the number of scheduled nurses working at position j during shift i on day e of the week compared with the standard schedule.

Objective. The objective of the model is to minimize the total number of nurse-shifts scheduled while penalizing shortage for demand and changes from the standard schedule, with parameters to control trade-offs between the three terms. Hence, our objective is to minimize a weighted combination of three terms:

$$\begin{aligned} \min \quad & \sum_{q \in [Q], j \in [J], i \in [I], d \in [D], g \in [G]} z_{qjidg} && \text{(Minimize number of scheduled work shifts)} \\ & + \mu_1 \cdot npr && \text{(Penalize when staffed below target ratios)} \\ & + \mu_2 \sum_{j \in [J], i \in [I], e \in [E]} \Delta z_{jie} && \text{(Penalize number of changed work shifts),} \end{aligned}$$

where the parameters $\mu_1, \mu_2 \geq 0$ control the relative importance of each term of the objective.

Constraints. We now present the constraints. Some of these constraints correspond to the definition of the auxiliary variables and the others capture various staffing requirements. We use the notation $f(x)_+$ to denote the positive part of the function $f(x)$, i.e., $f(x)_+ := \max\{f(x), 0\}$.

For defining the auxiliary variables, we impose the following constraints:

$$\begin{aligned} npr &\geq \sum_{m \in [M], s \in [S], w \in [W]} \omega_w \left(\tilde{h}_{msw} - \sum_{q \in [Q], j \in [J], i \in [I], d \in [D], g \in [G]} z_{qjidg} \sigma_{idsw} \right)_+, && \forall \tilde{h} \in \mathcal{U}, \quad (1) \\ \Delta z_{jie} &\geq \left| z_{jie}^{\text{curr}} - \sum_{q \in [Q], d \in [D], e \in [E]} z_{qjide} \right|, && \forall j \in [J], i \in [I], e \in [E]. \quad (2) \end{aligned}$$

Constraint (1) computes the difference between the nurse demand for each position/hour/week \tilde{h}_{msw} and the total number of nurse-shifts scheduled for this position/hour/week —the coefficients $\sigma_{idsw} \{0, 1\}$ map shifts (i) and day (d) to each hour s and week w . It is a robust constraint imposed for all demand realization $\tilde{h} \in \mathcal{U}$.

It also involves a weighting parameter ω_w , which allows us to put more weight on some weeks. For example, by increasing ω_w with w , we place higher penalties for the weeks estimated with the more recent data, hence emphasizing the most recent demand trend.

The schedule must adhere to a range of staffing constraints to be feasible. These include:

- Each position j has a minimum and maximum number of nurses to staff (staffing levels are typically fixed at logistical positions e.g., CNL and first nurse, to be the same as before, while some pods can be adjusted to demand): $\underline{n}_j \leq \sum_{q \in [Q], i \in [I], d \in [D], g \in [G]} z_{qj idg} \cdot \sigma_{idt} \leq \bar{n}_j, \quad \forall j \in [J], t \in [T]$.
- Certain units are not eligible based on nurse tiers: $z_{qj idg} = 0, \quad \forall (q, j) \in N_U, i \in [I], d \in [D], g \in [G]$.
- Staffing levels are kept to be consistent from week to week for a stable schedule:

$$\sum_{q \in [Q], g \in [G]} z_{q, j, i, k + E(w-1), g} = \sum_{q \in [Q], g \in [G]} z_{q, j, i, k + E(W-1), g}, \quad \forall k \in [E], w \in [W-1], i \in [I], j \in [J]. \quad (3)$$

ED leadership indicates a preference to consider only the constraints above when deciding on staffing levels, based on demand only. However, they have the option to also consider nurse availability and supply information in this stage of scheduling by including additional constraints in the model:

- Assignments are capped by the number of available nurses of each tier during each period: $\sum_{j \in [J], i \in [I], d \in [D], g \in [G]} z_{qj idg} \cdot \sigma_{idt} \leq E_{qt}, \quad \forall q \in [Q], t \in [T]$.
- Total number of weekly working hours: $\sum_{d \in D_w, j \in [J], i \in [I], g \in [G]} z_{qj idg} \leq \overline{Tn}_q Z_q, \quad \forall q \in [Q], w \in [W]$.
- Week-end shifts on consecutive weeks for nurses working every other week-end (and analogous constraints for every third week-end): $\sum_{d \in D_{n_w} \cup D_{n_{w+1}}, j \in [J], i \in [I], g \in [G]} z_{qj idg} \leq 2Z_q^1, \quad \forall q \in [Q], w \in [W-1]$.

Solving the Robust Optimization Model

We now describe our method and implementation of solving the robust optimization problem. We use a cutting plane approach shown in Algorithm 1. At each iteration κ , we solve a mixed-integer linear optimization problem (a master problem) similar to the one described earlier, except that the uncertainty set \mathcal{U} in Constraint (1) is replaced by a finite subset \mathcal{U}^κ . To solve the master problem, we linearize the positive part in Constraint (1) and the absolute values in Constraint (2) via additional auxiliary variables and linear constraints. At each iteration, we augment \mathcal{U}^κ with the worst-case demand scenario h^κ , and solve the master problem up to 0.01 optimality within a 20-minute time limit. We implement the cutting-plane algorithm using lazy callbacks to accelerate the model compilation and solving process. As termination criteria, we use a sub-optimality gap target $\eta = 0.1$ and a maximum number of iterations $\kappa^{\max} = 20$.

We fine-tune a combination of parameters, $\beta_\epsilon, \mu_1, \mu_2, \omega_w$, to select the best-performing values for the model. To solve the deterministic model version, we solve a single master problem using the optimization model defined in Section Optimization Model Overview, but with \mathcal{U} replaced by the set of historical demand only $\{h^{\text{hist}}\}$ in constraint (1), which is equivalent to setting $\epsilon^1 = \epsilon^2 = \epsilon^3 = 0$ in \mathcal{U} .

Algorithm 1 Cutting plane algorithm

-
- 1: Given the nominal demand $h^{\text{hist}} : \mathcal{U}^0 \leftarrow \{h^{\text{hist}}\}$, $\kappa \leftarrow 1$, and tolerance η , initiate the master problem
 - 2: **repeat**
 - 3: Solve the master problem and obtain a solution z^κ
 - 4: $npr_\kappa(\tilde{h}, z^\kappa) \leftarrow \sum_{m \in [M], s \in [S], w \in [W]} \omega_w (\tilde{h}_{msw} - \sum_{q \in [Q], j \in [J_m], i \in [I], d \in [D], g \in [G]} z_{qjdg}^\kappa \sigma_{idsw}))_+$
 - 5: $h^\kappa \leftarrow \arg \max_{\tilde{h} \in \mathcal{U}} npr_\kappa(\tilde{h}, z^\kappa)$ \triangleright Maximize insufficiency w.r.t. demand uncertainty set
 - 6: $\mathcal{U}^\kappa \leftarrow \{h^{\text{hist}}, h^1, \dots, h^\kappa\}$
 - 7: $\kappa \leftarrow \kappa + 1$
 - 8: **until** $(npr^\kappa - npr^{\kappa-1})/npr^{\kappa-1} < \eta$ or $\kappa = \kappa^{\text{max}}$ \triangleright Reach violation gap or maximum iterations
-

Model Variants for Informing Strategic Decisions

Stable staffing levels. Motivated by the within-week variability of demand, we allow the staffing levels to vary by day of the week—but remaining the same week over week as per Constraint (3). To evaluate the benefits of this additional flexibility, we compare it to two alternatives: *Daily stability*: Staffing levels can be forced to be the same every day, by making weekly constraints (3) daily. This option is easier for the leadership to manage but might result in unnecessary overstaffing and understaffing on some days. It corresponds to the historical practice at HH. *No stability*: Staffing levels can be allowed to vary every day, by excluding constraint (3). This option provides even more flexibility to fit demand patterns that may differ from one week to another but is more at risk of overfitting the data. We compare the operational costs and benefits of these options.

Change shift designs. ED leadership also considers changing or adding shift types. Shift changes would be disruptive and require nurses to change their lifestyles, to accommodate for the new working hours. However, new shifts could lead to more cost-effective staffing and a better matching of the demand patterns. With many potential shift type options available, it is essential to determine which ones could yield the greatest improvement. We support this investigation by introducing binary variables $y_i \in \{0, 1\}$ that indicate whether a shift type i is introduced. We have the set of existing shift types I_{exist} and consider a feasible set of potential shift types $I_{\text{feasible}} \supseteq I_{\text{exist}}$ with σ_{idt} for each $i \in I_{\text{feasible}}$. We add the following constraints:

- A shift type is generated if and only if used in any shifts: $\sum_{q \in [Q], j \in [J], d \in [D], g \in [G]} z_{qjdg} \leq M_y \cdot y_i, \quad \forall i \in [I]$
with a constant M_y .
- The number of shifts types is bounded with a parameter Y_{max} : $\sum_{i \in I_{\text{feasible}}} y_i \leq Y_{\text{max}}$.
- If considering only new shift types without changing the existing shift types, then all current shift types are kept: $y_i = 1, \quad \forall i \in I_{\text{exist}}$.

By solving different variations and comparing their objective values, we can identify the most valuable shift type candidates.

Mixed-Integer Optimization Models for Individual Scheduling

We provide details on the individual shift and preceptor scheduling optimization models from Section ED Individual Scheduling Optimization.

Individual Level Data

We define additional indices, sets, and data inputs to Section Individual Model Data.

Additional Components. We consider each individual nurse ℓ (with $L \approx 200$ total nurses), with L_q denoting the set of nurses of tier q . There are multiple ways to measure a nurse's dissatisfaction. We identify 6 possible criteria: the number of shifts on dates supposed to be off, of unassigned week-ends turned on, of overtime shifts, and the number of unpreferred shift types, dates, or patterns. We also have three criteria for satisfaction: the number of preferred shift dates, patterns ($U = 3$ on/off patterns for assigned shifts), and the number of unique positions a nurse is assigned to (as a measure of job diversity). We index these criteria via a subscript $o \in [O]$ ($O = 9$). For the ED, J^{slack} represents set of positions j where a shortage or surplus of at most one is allowed.

Data input. One of the inputs is the output aggregate levels z_{jid}^{agg} , representing the number of nurses to schedule for position j during shift i on day d from the first phase solution $z_{qj{id}g}^*$:

$$z_{jid}^{\text{agg}} = \sum_{q \in [Q], g \in [G]} z_{qj{id}g}^*, \quad \forall j \in [J], i \in [I], d \in [D].$$

While the assignment per nurse tier q and week-end group g in the first phase is included to ensure a feasible assignment, it is subject to change during the second phase. The other type of input is a range of individual preferences and availabilities for each nurse: $I_{\ell i}^{\text{pref}}, I_{\ell i}^{\text{unpref}}, I_{\ell i}^{\text{feas}}$: whether nurse ℓ prefers (typically their current shift type), does not prefer, and has the feasibility to work at shift type i , respectively. $F_{\ell w}$: maximum number of shifts for nurse ℓ on week w , obtained by subtracting Paid Time Off (PTO) and education time of each week from the total number of weekly shifts. $k_{\ell w}$: whether nurse ℓ is assigned to work on the week-end of week w . $P_{\ell d}^{\text{pref}}, P_{\ell d}^{\text{unpref}}, P_{\ell d}^{\text{off}}, P_{\ell d}^{\text{avail}}$: whether nurse ℓ prefers, does not prefer, is off, and is available to work on day d , respectively. $A_{\ell u}^{\text{pref}}, A_{\ell u}^{\text{unpref}}$: whether nurse ℓ prefers to have, and not to have the work pattern u , respectively. $r_{\ell j}^{\text{feas}}, p_{\ell j}^{\text{feas}}$: whether nurse ℓ is eligible to be a trainee, and a preceptor (training the trainees), respectively, at position j .

Individual Shift Scheduling Model

To generate an individual schedule that prioritizes various staff preferences as well as training opportunities given these inputs, we develop a mixed integer linear problem that comprises the following components.

Decision variables. The core decision variables track nurse assignments:

- $b_{\ell j{id}}, r_{\ell j{id}} \in \{0, 1\}$: whether nurse ℓ works, and resp. trains at position j during shift i on day d .
- $s_{\ell i} \in \{0, 1\}$: whether nurse ℓ is assigned to shift type i .

Objective. We introduce main auxiliary variables to track different terms of the objective:

- $c_{\ell o} \in \mathbb{R}$: computing nurse ℓ 's dissatisfaction penalty score according to the criterion o .
- $z_{jid}^-, z_{jid}^+ \in [0, 1]$: whether there is one shortage or surplus at position j during shift i on day d , resp.
- $f_{\ell}^- \in \mathbb{R}^+$: number of unassigned shifts for nurse ℓ .

Our objective function is to minimize a weighted combination of four terms:

$$\begin{aligned}
\min \quad & \mu_3 \sum_{\ell \in [L], o \in [O]} \lambda_o c_{\ell o} && \text{(Minimize total weighted dissatisfaction score)} \\
& + \mu_4 \sum_{j \in [J], i \in [I], d \in [D]} w_j^{\text{shortage}} z_{jid}^- && \text{(Penalize shortage to aggregate staffing levels)} \\
& + \mu_5 \left(\sum_{j \in [J], i \in [I], d \in [D]} z_{jid}^+ + \sum_{\ell \in [L]} f_{\ell}^- \right) && \text{(Penalize unassigned nurse-shifts and surpluses)} \\
& - \mu_6 \sum_{\ell \in [L], j \in [J], i \in [I], d \in [D]} r_{\ell jid} && \text{(Reward total training shifts assigned)}
\end{aligned}$$

with parameters:

- $\lambda_1, \dots, \lambda_5, \lambda_6 > 0, \lambda_7, \lambda_8, \lambda_9 < 0$ as weights for each penalty score metric,
- $\mu_3, \mu_4, \mu_5, \mu_6 > 0$ to control the trade-off between the objective terms,
- $w_j^{\text{shortage}} \geq 0$ as penalization weights for shortage at position j .

To optimize computational efficiency, we implement all variables as sparse matrices and tensors under conditions as follows. For instance, binary variables $b_{\ell jid}$ are only needed if the values are allowed to be one, and thus are only defined for indices ℓ, j, i, d such that nurse ℓ is eligible to work at position j and is available during shift i on day d . Such sparse indexing naturally incorporates some eligibility, availability, and feasibility constraints.

Constraints. To facilitate the model constraints, we define additional auxiliary binary variables on whether nurse ℓ : works an overtime shift on week w ($f_{\ell w}^+$); works on a week-end of week w supposed to be off ($k_{\ell w}^+$); has at least one shift at position j during the staffing cycle ($v_{\ell j}$); and starts work pattern u on day d ($a_{\ell ud}$). We then impose associated constraints to bound the objective function:

- Computes individual nurse penalty score for each $\ell \in [L]$ counting the number of:

- Shifts on dates supposed to be off: $c_{\ell 1} = \sum_{d \in [D], j \in [J], i \in [I]} P_{\ell d}^{\text{off}} (b_{\ell jid} + r_{\ell jid})$.
- Unassigned week-ends turned on: $c_{\ell 2} = \sum_{w \in [W]} k_{\ell w}^+$, where $\sum_{d \in D_{n_w}, j \in [J], i \in [I]} (b_{\ell jid} + r_{\ell jid}) \leq 2(k_{\ell w} + k_{\ell w}^+)$, $\forall \ell \in [L], w \in [W]$.
- Overtime shifts: $c_{\ell 3} = \sum_{w \in [W]} f_{\ell w}^+$, where $\sum_{d \in D_w, j \in [J], i \in [I]} (b_{\ell jid} + r_{\ell jid}) \leq Fw_{\ell w} + f_{\ell w}^+$, $\forall \ell \in [L], w \in [W]$.
- Unpreferred shift types: $c_{\ell 4} = \sum_{d \in [D], j \in [J], i \in [I]} I_{\ell i}^{\text{unpref}} b_{\ell jid}$.
- Unpreferred dates: $c_{\ell 5} = \sum_{d \in [D], j \in [J], i \in [I]} P_{\ell d}^{\text{unpref}} (b_{\ell jid} + r_{\ell jid})$.
- Preferred dates: $c_{\ell 6} = \sum_{d \in [D], j \in [J], i \in [I]} P_{\ell d}^{\text{pref}} (b_{\ell jid} + r_{\ell jid})$.
- Different positions assigned: $c_{\ell 7} = \sum_{j \in [J]} v_{\ell j}$, where $v_{\ell j} \leq \sum_{d \in [D], i \in [I]} b_{\ell jid}$, $\forall \ell \in [L], j \in [J]$.
- Unpreferred patterns: $c_{\ell 8} = \sum_{d \in [D-2], u \in [2]} A_{\ell u}^{\text{unpref}} a_{\ell ud} + \sum_{d \in [D-3]} A_{\ell 3}^{\text{unpref}} a_{\ell 3d}$.
- Preferred patterns: $c_{\ell 9} = \sum_{d \in [D-2], u \in [2]} A_{\ell u}^{\text{pref}} a_{\ell ud} + \sum_{d \in [D-3]} A_{\ell 3}^{\text{pref}} a_{\ell 3d}$, where all shifts in a row is tracked by $\sum_{j \in [J], i \in [I], d' \in \{d, d+1, d+2\}} (b_{\ell, j, i, d'} + r_{\ell, j, i, d'}) - 2 \leq 3a_{\ell 1d} \leq \sum_{j \in [J], i \in [I], d' \in \{d, d+1, d+2\}} (b_{\ell, j, i, d'} + r_{\ell, j, i, d'})$, $\forall \ell \in [L], d \in [D-2]$, and other patterns by analogous constraints.

- The individual nurse schedule matches aggregate staffing levels with shortages and surpluses: $z_{jid}^{\text{agg}} = z_{jid}^- - z_{jid}^+ + \sum_{\ell \in [L]} b_{\ell jid}$, $\forall j \in J^{\text{slack}}, i \in [I], d \in [D]$, where some positions do not allow shortages or surpluses: $z_{jid}^- = z_{jid}^+ = 0$, $j \in J \setminus J^{\text{slack}}, i \in [I], d \in [D]$.

- The number of unassigned shifts for each nurse is the number of maximum shifts minus assigned shifts: $f_{\ell}^- \geq \sum_{w \in [W]} F_{\ell w} - \sum_{j \in [J], i \in [I], d \in [D]} (b_{\ell jid} + r_{\ell jid})$.

Besides these constraints bounding the objective, the schedule is enforced to satisfy a variety of feasibility constraints, such as:

- Eligibility and availability:
 - Nurse tier eligibility in positions: $b_{\ell jid} = 0$, $\forall q \in [Q], (q, j) \in N_U, \ell \in L_q, i \in [I], d \in [D]$.
 - Shift only at one location on available dates: $\sum_{j \in [J]} (b_{\ell jid} + r_{\ell jid}) \leq P_{\ell d}^{\text{avail}}$, $\forall \ell \in [L], i \in [I], d \in [D]$.
- Shift type:
 - Assigned at most one shift type over the 6-week period: $\sum_{i \in [I]} s_{\ell i} \leq 1$, $\forall \ell \in [L]$.
 - Each nurse can be only assigned a feasible shift type: $s_{\ell i} \leq I_{\ell i}^{\text{feas}}$, $\forall \ell \in [L], i \in [I]$.
 - Work on assigned shift type: $\sum_{d \in [D], j \in [J]} (b_{\ell jid} + s_{\ell jid}) \leq M_s \cdot s_{\ell i}$, $\forall \ell \in [L], i \in [I]$ with $M_s = 2D$.
- Training:
 - Trainee eligibility: $r_{\ell jid} \leq r_{\ell j}^{\text{feas}}$, $\forall \ell \in [L], j \in [J], i \in [I], d \in [D]$.
 - At most three shifts per training: $\sum_{i \in [I], d \in [D]} r_{\ell jid} \leq 3$, $\forall \ell \in [L], j \in [J]$.
 - Capped by number of eligible preceptors: $\sum_{\ell \in [L]} r_{\ell jid} \leq \sum_{\ell \in [L]} b_{\ell jid} p_{\ell j}^{\text{feas}}$, $j \in [J], i \in [I], d \in [D]$.
- To ensure fairness among nurses, each nurse's penalty score cannot exceed a bound c^{max} :

$$\sum_{o \in [O]} c_{\ell o} \leq c^{\text{max}}, \quad \ell \in [L]. \quad (4)$$

Individual Preceptor Scheduling Model

The preceptor scheduling model from Section Integer Optimization Model on Preceptor Scheduling is as follows.

Decision variables. We introduce decision variables:

- $p_{\ell jid} \in \{0, 1\}$: whether nurse ℓ serves as a preceptor at position j during shift i on day d .
- $p^{\text{max}} \in \mathbb{R}$: tracking the maximum number of preceptor shifts assigned among all nurses.

Objective. We set the objective function to:

$$\min p^{\text{max}} \quad (\text{maximum number of preceptor shifts assigned to each nurse}).$$

Constraints. The model is subject to several constraints, where $b_{\ell jid}^*$ and $r_{\ell jid}^*$ denote the optimal solution from solving the model defined in Section Individual Shift Scheduling Model.

- p^{max} is at least the number of preceptor shifts for each nurse: $p^{\text{max}} \geq \sum_{j \in [J], i \in [I], d \in [D]} p_{\ell jid}$, $\forall \ell \in [L]$.
- Eligibility to train the trainees at each position: $p_{\ell jid} \leq p_{\ell j}^{\text{feas}}$, $\forall \ell \in [L], j \in [J], i \in [I], d \in [D]$.
- Preceptor and work shifts in tandem at the same position: $p_{\ell jid} \leq b_{\ell jid}^*$, $\forall \ell \in [L], j \in [J], i \in [I], d \in [D]$.
- Each trainee shift is covered by a preceptor: $\sum_{\ell \in [L]} r_{\ell jid}^* \leq \sum_{\ell \in [L]} p_{\ell jid}$, $\forall j \in [J], i \in [I], d \in [D]$.

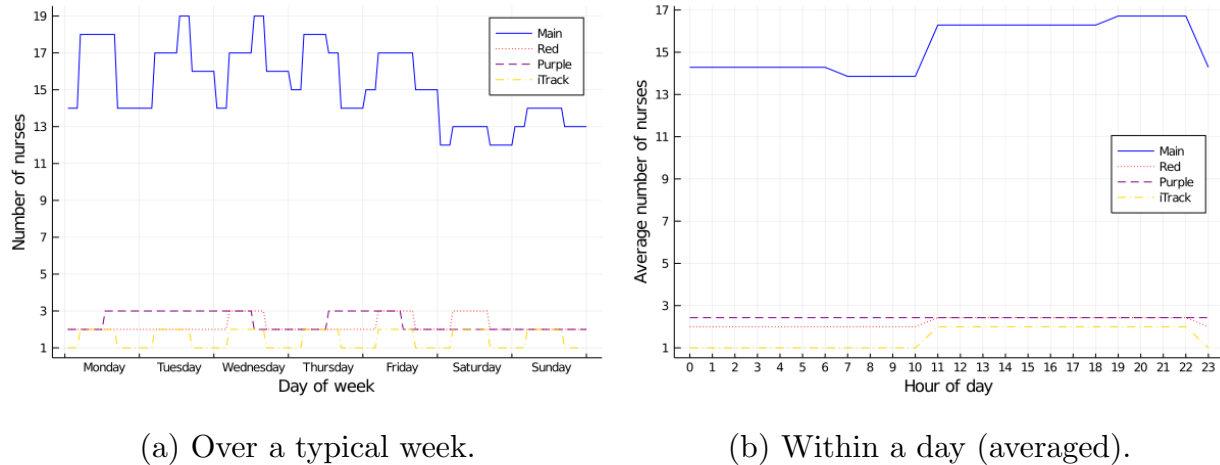


Figure A2 Recommended Schedules for the November 2 - December 13, 2020 cycle.

Supplementary Experimental Results

In this section, we provide additional experimental results to Section Results on each of the three periods.

Supplementary Results during Low Demand Periods

We augment Section End of 2020: Reducing Staffing Costs during Low-Demand Periods on data from October 26 to December 20, 2020, including the illustration of schedule patterns, additional information on different approaches, and elaboration on metric trade-offs.

Patterns of demand vs. staffing. Figure A2 provides a more detailed illustration of the recommended schedule for November 2 - December 13, 2020 obtained from the Robust (a) approach. A comparison with demand patterns from Figure 1 shows that our recommendation matches staffing with demand patterns and results in a more cost-effective schedule. For the main pods and the purple pod, as patient demands tend to be higher on weekdays than week-ends, we staff more nurses on weekdays than week-ends. For the red pod, we have a particularly higher demand on Friday when we also staff more nurses. Both demand and staffing for iTrack have less variability throughout days of the week. The recommended schedule shown in Figure A2b is consistent with the demand patterns by having smoother staffing in the purple pod and more staffing during 11 am–11 pm for other pods. Even though the schedule is generated with information only prior to the period, it can capture most of the day-of-week and hour-of-day demand patterns in the period prospectively, which justifies the robustness of our approach. By matching staffing with demand patterns, we reduce the staffing cost by 7.40%.

Trade-off between cost and insufficiency. There exists a trade-off between cost and insufficiency, as staffing more nurses facilitates more sufficiency with increased cost, and vice versa. We illustrate the flexibility of the optimization model to control such trade-offs by varying parameters μ_1 and β_ϵ in Figure A3, where scatter points from each approach represent the daily average insufficiency and cost of schedules generated with different parameters. This earlier version of the model did not include the objective term with μ_2 or the parameter ω_w . The current schedule has a daily average cost of 56 and an insufficiency of 0.17. Retrospectively, given perfect information, schedules optimized with different parameters range from having 50.43 cost and

0.01 insufficiency to having 48.57 cost and 0.06 insufficiency. Prospectively with the data-driven approach, the non-robust schedules lead to a reduction of cost to 49.43 - 50.90 with the trade-off of higher insufficiency between 0.26 and 0.20. The robust schedules (a) incorporate uncertain demand deviations to trade off some cost reduction for more sufficiency, ranging from giving 50.38 cost and 0.17 insufficiency to 53.24 cost and 0.08 insufficiency.




Figure A3 Cost-insufficiency Trade-offs.

Table A1 Comparison of Alternative Registered Nurse (RN)

Variation	Schedules.			
	Default	Less overtime	Shift change	EDOU float
Overtime shifts	72	52	8	48
Training shifts	72	60	40	60
Week-ends turned on	11	11	6	11
Holidays turned on	2	2	2	2
RN with shift type changed	0	0	4	0
EDOU RN relocated to ED	0	0	0	2

Supplementary Results during Iteration Periods

Alternative individual scheduling results. In addition to the results shown in Section 2021: Accounting for Frictions to Implementation, we solve and compare four schedules with alternative settings and parameters in Table A1. Compared with the solution with default parameters, a variation with increased weights on overtime shifts reduces overtime with a trade-off of less training. On top of the variation, allowing changes in nurses' shift types significantly reduces overtime shifts by changing four nurses' shift types. Alternatively, assigning two EDOU nurses to float to work at ED can bring several more overtime reductions. After reviewing with ED leadership, we decided to use the schedule with shift changes due to its best overall metrics.

Supplementary Results during High Demand Periods

We provide additional results to Section 2022: Reducing Insufficiency during High Demand Periods.

Parameter tuning details. We first tune the two main parameters, with μ_1 drawn from the list $\{0, 0.2, 0.3, 0.4, 0.6, 0.8, 1\}$ and β_ϵ from $\{0.1, 0.15, 0.2, 0.3, 0.4, 0.5, 0.6\}$ with fixed μ_2 and ω_w . With $\mu_1 = 0.2$ and $\beta_\epsilon = 0.3$ selected and fixed, we then vary μ_2 and ω_w . We use an auxiliary parameter $\omega \in [0, 1]$ and let $\omega_1, \dots, \omega_6$ to be uniformly distributed between $1 - \omega$ and $1 + \omega$. We vary μ_2 from $\{0, 0.1, 0.2, 0.3, 0.4, 0.5\}$ and vary ω from $\{0, 0.1, 0.2, 0.25, 0.3, 0.4, 0.5\}$. We finally select $\mu_2 = 0.2$ and $\omega = 0.3$ together with $\mu_1 = 0.2$ and $\beta_\epsilon = 0.3$.

Adjustment from resolving aggregate model weekly. We consider another variation to solve the aggregate model every week. We conduct experiments on three adjacent 6-week staffing cycles between July 31 and December 3, 2022. We compare three alternatives: (1) Current schedule. (2) The model is solved every 6 weeks and generates a schedule for each 6-week staffing cycle. This requires solving the model three times for the three staffing cycles. (3) For each staffing cycle, the model is resolved every week using the training

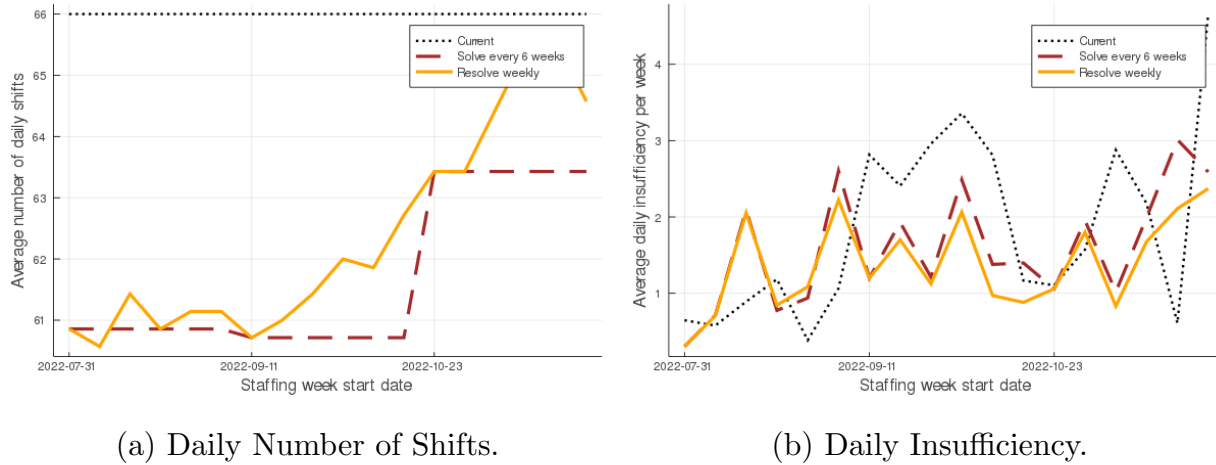


Figure A4 Change from Resolving Weekly during Three Staffing Cycles.

Table A2 Average Daily Results and Change (δ) from Current Schedules on Three Staffing Cycles during July 31 - December 3, 2022.

Schedule approach	Cost [δ]	Insufficiency [δ]	δz from current	δz by week
Current	66	1.85	0	0
Solve every 6 weeks	61.67 [-6.57%]	1.59 [-13.86%]	4.71 [7.14%]	0
Resolve weekly	62.37 [-5.51%]	1.39 [-24.83%]	4.26 [6.46%]	2.08 [3.33%]

data one week later and update the schedule for the remaining weeks of the cycle. This involves solving the model 18 times for the 18-week period.

For each week, we compute the average number of daily shifts and average daily insufficiency during that week's schedule, and plot the three variations during the time period in Figure A4. Resolving the model weekly changes in staffing levels from week to week and leads to improved insufficiency in most weeks. We summarize the average daily metrics among the 18 weeks for the three alternatives in Table A2. On average, resolving weekly changes the staffing levels by 2.08 shifts from week to week. This would require the nurse leadership to ask several nurses to request shift changes from their pre-arranged schedule. Solving the model every 6 weeks (resp. every week) on average changes the staffing levels from the current schedule by 4.71 (resp. 4.26) shifts. In return, resolving weekly can reduce the shortage of nurse-shifts to meet target nurse-patient-ratios by 24.83% from the current schedule, which is larger than 13.86% reduction by solving the model every 6 weeks. The benefit of reducing insufficiency is particularly useful for this period when the ED is extremely understaffed and is achieved with 5.51% (resp. 6.57%) less staffing cost from solving the model every 6 weeks (resp. every week).