

Risk-Aware Parallel Simulation–LBBD Framework for Integrated Multi-Skill Nursing-Home Staffing and Rostering

Omar Salim Abdullah 

Institute of Fine Art, Ministry of Education, Diyala, 32001, Iraq
omarhasibat@gmail.com

Abstract

Healthcare worker scheduling is an optimization problem that is difficult because of the interaction between skills and workers. Uncertain service demand, regulatory constraints, requirements, and workload balance. This study proposes a Risk-Aware Parallel Simulation–Logic Based Benders Decomposition (LBBD) framework for integrated Stochastic demand nursing-home staffing and rostering with multi-skill workers. The proposed approach uses a risk-aware mathematical framework, along with Conditional Value-at-Risk (CVaR) and parallel scenarios. The techniques used are simulation, adaptive cut generation, and a hybrid Bandit-CMSA metaheuristic embedded in the LBBD optimization process. The framework separates strategic rostering from "scenario"- related service evaluations, allowing efficient exploration of uncertain operating conditions while maintaining valid evaluations. Optimality bounds. We ran computational experiments on a nursing-home benchmark instance with identical scenarios, computational resources and a time limit of 300 seconds. The 13 people were six CNAs, four nurses, and two pharmacists. The dataset includes three LPNs, three RNs, three competence levels, sixteen 1-hour intervals, twelve lawful shifts, four stochastic demand scenarios, and one thousand out-of-sample validation scenarios. The proposed framework achieved an optimal objective value of 11,169.019, and the lower bound of 11,169.019 was achieved and is valid. The lower bound was 10,754.446, leading to a final optimality gap of 3.712%. This schedule improved the objective value by 19.36% over the first feasible solution. The results also show the effectiveness of the integrated learning-based search mechanisms, with the bandit strategy identifying the most useful ones: an improved neighborhood operator and more exploration of the primal solution space. The proposed framework is a scalable and powerful optimization method that can create high-quality nurse schedules while ensuring the rostering process runs smoothly. It manages operating costs, risk, and employee balance under fluctuating demand.

Keywords: Nurse rostering, multi-skill workforce scheduling, Stochastic optimization, Logic-Based Benders Decomposition (LBBD), Conditional Value-at-Risk (CVaR).

Article history: Received: 13 Aug 2026, accepted: 8 Sep, published: 15 Sep 2026.

This article is open-access under the CC BY 4.0 license (<http://creativecommons.org/licenses/by/4.0/>).

1. Introduction

Nurse scheduling is an essential healthcare optimization challenge that allocates nurses to shifts while meeting staffing needs, skill competencies, contractual stipulations, rest regulations, and personal preferences. Roster quality strongly affects continuity of care, labor efficiency, employee satisfaction, and hospital operational performance. Nevertheless, roster design is computationally challenging because it involves interrelated

constraints across many days, shifts, nursing categories, and qualification levels. The problem becomes more challenging when staffing needs are uncertain, as a schedule feasible for the expected scenario can be severely understaffed in the worst case [1] and [3].

Most traditional nurse scheduling approaches assume that staffing needs are fixed and known in advance during the planning period. Although this assumption simplifies the optimization problem, it

does not accurately reflect healthcare settings subject to fluctuations in patient arrivals, care intensity, absenteeism, and emergency demand. Scenario-based stochastic programming provides a systematic way to model uncertainty, but optimizing expected costs may lead to solutions that are insufficient in rare but high-impact scenarios. Conditional Value-at-Risk (CVaR) is a suitable risk measure to control expected losses in worst-case scenarios. It is relevant to nurse staffing and scheduling under uncertain patient demand [4] and [5].

The computational complexity of stochastic nurse scheduling escalates rapidly as the number of nurses, shifts, competencies, and demand scenarios grows. Consequently, monolithic mixed-integer programming models can be challenging to solve within feasible time constraints. Logic-Based Benders Decomposition (LBBD) offers a strong alternative by separating high-level assignment decisions from detailed scheduling or scenario-assessment subproblems. Unlike classical Benders decomposition, LBBD generates cuts using logical inference and can use optimization techniques that exploit the natural structure of the subproblems to solve them [14]. Researchers have previously addressed complex planning and scheduling problems and shown LBBD's usefulness in these settings, such as two-stage stochastic problems.

Additionally, once the main decisions are made, the subproblems are conditionally independent. They can be solved in parallel for each scenario. Deterministic nurse scheduling has been well researched, and uncertainty-aware workforce planning is garnering more attention. However, uncertainty-aware nurse scheduling with risk aversion, scenario decomposition, and parallel solutions has not been adequately researched in a reproducible nurse rostering framework. Coverages, skill sets, workload, and order limits are feasible combinations taken from existing nurse roster benchmarks [6]. Stochastic methods generally focus on mathematical risk modelling but do not fully exploit scenario-level computational parallelism. This leaves a big gap between pragmatic risk modelling and scalable solutions. This research proposes a Risk-Aware Parallel Logic-Based Benders Decomposition architecture to solve the

multiskilled nurse scheduling problem under demand uncertainty [8-12].

The main contributions are four in number:

1. A scalable Parallel Simulation-LBBD architecture for effective scenario-based decomposition and evaluation;
2. An integrated risk-aware staffing and rostering formulation that explicitly addresses stochastic healthcare demand using CVaR;
3. A Bandit-CMSA hybrid approach that dynamically recognizes and utilizes advantageous solution elements
4. A secure MIP-probing approach that enhances existing solutions without compromising the approved lower limit or the optimality gap.

Current healthcare workforce models offer no defense against severe demand-related shortfalls and frequently divide staffing from rostering. A disconnect remains between computationally scalable solution techniques and theoretical risk-aware formulations. Current LBBD techniques have not yet combined parallel scenario evaluation, CVaR-based risk control, workload fairness, and learning-driven adaptive search. By combining CVaR and Bandit-CMSA adaptive search into a single Risk-Aware Parallel Simulation-LBBD framework, this study fills these gaps.

The proposed approach aims to develop rosters that are viable and cost-efficient under expected conditions, while remaining robust under poor demand. The rest of the paper is structured as follows. Section 2 offers a detailed review of related studies on multi-skill nurse rostering, stochastic workforce scheduling, risk-aware optimization, CVaR-based approaches, and Logic-Based Benders Decomposition. Section 3 describes the proposed Risk-Aware Parallel Simulation-LBBD framework, including the mathematical formulation, scenario-based simulation mechanism, adaptive cut generation strategy, and the integrated Bandit-CMSA metaheuristic with safe MIP-bound probing. Section 4 introduces the computational experimental design, including the benchmark instance, stochastic demand generation, algorithmic configurations, comparative

evaluation, and performance metrics. Section 5 presents the computational results and discussion, including convergence analysis, incremental performance evaluation, ablation analysis, and risk-based solution assessment. Section 6 presents the study's conclusion and possible future research directions.

2. Literature Review

This section reviews the main research directions related to the proposed framework. It begins with nurse rostering and multi-skill workforce scheduling, then stochastic approaches for tackling demand uncertainty. It then discusses risk-aware optimization using Conditional Value-at-Risk (CVaR), Logic-Based Benders Decomposition (LBBD), and parallel scenario evaluation. This review highlights the methodological limitations of existing studies. It identifies the research gap covered by the proposed Risk-Aware Parallel Simulation–LBBD framework.

2.1 Multi-Skill Workforce Scheduling and Nurse Rostering

Nurse rostering is the assignment of nurses to shifts in such a way that coverage, skill, workload, rest, and legal requirements are satisfied. It also considers soft objectives such as staff preferences, fairness, and workload balance. This combination of requirements creates a highly constrained combinatorial problem, especially as the workforce, planning horizon, and number of shifts increase. Therefore, researchers have widely used exact, heuristic, and hybrid optimization methods to produce feasible, high-quality rosters [1], [8], and [9].

Multi-skill scheduling further differentiates nurses by qualifications and allows eligible staff to meet different care needs. Flexibility may improve resource use but also increases dependencies, as adequate overall staffing does not ensure coverage in each skill category. Previous nurse scheduling models have accounted for skill-specific staffing, qualifications, workload constraints, and preferences [4], [5], [19], [20]. Some models. However, they decouple staffing and rostering or assume deterministic demand, reducing their applicability in nursing homes where skill-specific care needs vary over time.

2.2 Demand Uncertainty in Stochastic Nurse Rostering

Deterministic rostering models assume staffing demand is known in advance. At the same time, nursing requirements often vary due to changes in patient volume, care intensity, and unplanned absences. Stochastic programming overcomes this limitation by modeling uncertain demand through scenarios. In two-stage formulations, managers make staffing and rostering decisions before observing demand, and recourse decisions adjust staffing after demand is realized. Such models can reduce staffing costs and improve demand coverage compared to deterministic models [14], [2] and [9].

Recent work also included short-term schedule adjustments to account for dynamic patient demand (Aydas et al., 2023). Increasing the number of nurses, skills, shifts, and situations greatly increases the resulting stochastic model and its computational burden. Also, models that minimize predicted cost may perform well on average but are still susceptible to high-impact, low-probability events. These limitations call for scalable decomposition techniques and explicit risk measures to reduce the likelihood of undesirable outcomes.

2.3 CVaR and Risk-Aware Optimization

Risk-neutral stochastic models generally minimize expected cost or delay over a range of demand possibilities. This works well on average, but it may over- or underestimate the importance of rare scenarios with critical understaffing or long service delays. Risk-aware optimization addresses this shortcoming by explicitly accounting for the distribution of negative outcomes. A commonly used measure of coherence risk is Conditional Value-at-Risk (CVaR), which measures the expected loss in the worst percent of cases at a given confidence level [10]. Value-at-Risk considers only the loss threshold, but CVaR also accounts for the magnitude of values above that threshold.

In nurse scheduling, He et al. introduced CVaR in a two-stage stochastic model to manage understaffing under uncertain patient demand, showing its ability to balance cost and coverage reliability. However, risk-aware applications in multi-skill nurse rostering remain limited,

particularly when detailed simulation and numerous demand scenarios are involved. Furthermore, Conventional optimization approaches often oversimplify real-world system complexity, but integrating discrete-event simulation with risk-constrained mathematical programming enables a thorough examination of operational options and tail-risk exposure across high-dimensional scenario frameworks (Juan et al., 2026).

This motivates integrating expected delay and CVaR within a scalable optimization framework that produces cost-efficient schedules while reducing exposure to severe service outcomes.

2.4 Benders Decomposition Based on Logic and Parallel Scenario Assessment

High-level allocation choices are separated from specific scheduling or scenario-evaluation subproblems using Logic-Based Benders Decomposition (LBBD). Unlike standard Benders decomposition, which obtains its cuts from linear programming duality, LBBD builds problem-specific cuts through logical inference and allows each subproblem to be solved using any optimization or simulation technique appropriate to the subproblem. LBBD has been used with substantial computational gains over the integer L-shaped method for two-stage stochastic scheduling [9].

Another study combined simulation with the master solutions and the resulting performance estimates to generate validation [9]. Their system solved up to 100 nursing home shift-scheduling scenarios. Because scenario assessments for a fixed roster are conditionally independent, parallel processing can accelerate computation. A cohesive framework integrating individual multi-skill rostering, simulation-based delay evaluation, CVaR, workload fairness, parallel scenario processing, and adaptive cut generation has not been fully explored. This gap drives the proposed Risk-Aware Parallel Simulation-LBBD paradigm. Recent studies have advanced stochastic healthcare staffing, decomposition-based optimization, and learning-guided metaheuristics.

Overall, previous studies have addressed multi-skill nurse rostering, stochastic demand, risk-aware optimization, and simulation-based LBBD

separately. However, their joint integration within a scalable framework remains limited. In particular, individual nurse assignments, skill-dependent coverage, service-delay simulation, CVaR, workload fairness, parallel scenario evaluation, and adaptive cut generation have not been sufficiently combined. The proposed Risk-Aware Parallel Simulation-LBBD framework addresses this gap by integrating these components into a unified approach for nursing-home rostering under demand uncertainty.

3. Methodology

3.1 Problem Formulation

The proposed framework addresses a risk-aware multi-skill nurse rostering problem under uncertain care demand. Let N , S , K , T , and Ω denote the sets of nurses, legal shifts, skills, planning periods, and demand scenarios, respectively. Binary variable $x(n,s)$ equals one if nurse n is assigned to legal shift s . The resulting coverage of skill k in period t is denoted by $y(k,t)$. Feasible rosters satisfy skill coverage, shift compatibility, workload, shift-count, non-overlap, minimum-rest, and consecutive-workday constraints. The total deviation between assigned and target hours represents workload fairness:

$$F(x) = \sum_{n \in N} |H_n(x) - H_n^{target}| \quad (1)$$

In (1), H with subscript n is the workload assigned to nurse n ; its superscript 'target' denotes the contractual target. For a coverage profile y , a discrete-event simulation returns the total care delay $Q(\omega, y)$ for each scenario ω . The risk-aware objective jointly minimizes staffing cost, expected delay, tail risk, and workload unfairness:

$$\min Z = \lambda_c C(x) + \lambda_m \sum_{\omega \in \Omega} p_\omega Q_\omega(y) + \lambda_r \text{CVaR}_\alpha(Q) + \lambda_f F(x) \quad (2)$$

$$\text{CVaR}_\alpha(Q) = \eta + \left[\frac{1}{1-\alpha} \right] \sum_{\omega \in \Omega} p_\omega \xi_\omega, \quad (3)$$

$$\xi_\omega \geq Q_\omega(y) - \eta, \quad \xi_\omega \geq 0 \quad (4)$$

In (2)-(3), $C(x)$ denotes staffing cost, p with subscript ω is the scenario probability, α is the CVaR confidence level, and the four λ weights control the relative importance of cost, expected delay, risk, and fairness. This objective discourages both persistent average delay and extreme outcomes in high-demand scenarios.

3.2 Simulation-Based LBBD and Adaptive Cut Generation

Because scenario delay is evaluated by simulation and has no explicit algebraic form, the model is solved using Simulation-based Logic-Based Benders Decomposition (LBBD). The unconstrained master includes nurse assignments, nurse size, pre-defined roster restrictions, fairness metrics, and lower limits on scenario delays. The total scenario delay is measured by the variable $\theta(\omega)$, and the skill period delay is measured by the variable $\theta(\omega, k, t)$. The binary variable $z(k, t, \ell)$ indicates the coverage level ℓ , where for each skill-period combination, exactly one coverage level is specified. During each iteration, the master creates a potential roster $\bar{x}$ along with a coverage profile $\bar{y}$. The simulation engine considers all the possible scenarios. If the simulated delay exceeds the main estimate, both components and the overall estimate are reduced, with adaptive separation. The delay is the same or shorter in the skill pools where you have focused your efforts, and you can get the right cut of the component:

$$\theta_{\omega kt} \geq Q_{\omega kt}(\bar{y}) [1 - D_k(\bar{y})] \quad (4)$$

$$D_k(\bar{y}) = \sum_{t \in T} \sum_{\ell > \bar{y}_{kt}} z_{kt\ell} \quad (5)$$

The simulated latency in (4) is represented by Q with subscripts ω , k , and t . When the number of staff exceeds the value of the assessed profile, then the term deactivation D_k in (5) becomes positive, thus inactivating a value of the delay, which is not relevant.

Further staffing changes to the delay framework cancel the cut for profiles that remain within the coverage. Violations of cuts are ranked by severity, or by the number of scenarios they violate. If the current threshold or minimum threshold is not changed, their allocation for selection also increases. At the same time, complete separation is required near convergence and during the conclusive-proof stage. Authorized reductions are stored in a common global store.

3.3 Proposed Bandit-CMSA Metaheuristic

Version 0.8.1 combines three free methods into the LBBD loop: persistent bandit-informed ALNS, internal CMSA-LBBD, and safe MIP-bound probing. These strategies raise the upper threshold, provide

useful coverage profiles, and reduce the search space without compromising the lower threshold.

Bandit-guided ALNS

ALNS has five neighborhood operators: shift switch, nurse reassignment, shift mutation, block shifting, and coverage-focused destroy-and-repair. Before simulation, we pre-mine the candidates according to the set of criteria. A continually discounted multi-armed bandit controls the operator choice:

$$S_o = R_o + c \sqrt{\{\ln(P + |O| + 1) / [p_o + 1]\}} \\ - \min(0.75, 0.02f_o) b I(o) \quad (6) \\ = \text{Mutation}$$

In (6), $S(o)$ is the score of the operator o , $\bar{R}(o)$ is the discounted reward of the operator o , $p(o)$ is the frequency of the operator being selected, and $f(o)$ is the frequency of the operator failing. P is the total number of picks. Parameter c is supposed to control exploration, while parameter b gives the measured inclination in advance of Mutation. Scores are mapped to softmax probabilities with a small probability threshold, so that no operator is "marginalized forever. Continuous lack of progress intensifies the failure penalty, while improvement increases the incentive. The acquired state remains until the end of ALNS iterations. Simulated annealing encourages diversification, and local branching restricts the Hamming distance from the current solution during intensification.

Internal CMSA-LBBD

CMSA organizes nurse-shift components by considering both selection frequency and component age. Each restricted core includes current assignments, frequently observed elements, and a randomized selection of the remaining variables. Variables outside the core are temporarily set to zero, and a novelty constraint ensures the incumbent is not replicated. Unlike a single sub-MIP approach, each CMSA iteration involves an internal LBBD cycle: solving the restricted master, simulating the schedule, identifying underestimated delay components, introducing valid cuts, and resolving. The core expands after a failure and contracts following an improvement. Restricted fixings are kept local, while globally valid simulation cuts are added to the shared pool.

Safe MIP-bound probing

Targeted probing replaces reduced-cost fixing. For an unselected assignment variable $x(j)$, an independent master containing all global cuts is solved after forcing $x(j)=1$. The variable can be fixed to zero only when that restricted model is infeasible or when

$$LB(M \mid x_j = 1) \geq UB^{inc} - \varepsilon \quad (7)$$

Condition (7) proves that $x(j)=1$ cannot generate a solution better than the incumbent. Because the incumbent already satisfies $x(j)=0$, the fixing preserves a valid solution. The restricted bound is used only for this certificate and is never reported as the global lower bound.




Fig. 1 Overall structure of the proposed Bandit-CMSA Simulation-LBBBD framework.

Fig. 1 illustrates the full procedure of the proposed system. The procedure starts with analyzing the nursing-home dataset and generating stochastic demand scenarios for uncertain care requirements. These conditions are input into the Logic-Based Benders Decomposition (LBBBD) master problem, which produces an initial nurse staffing and scheduling solution under deterministic scheduling constraints. The candidate lineup is then run through

a parallel simulation engine to evaluate service interruptions and staffing efficiency by running all the demand scenarios in parallel. The simulation outputs enable estimation of expected delay and Conditional Value-at-Risk (CVaR) to assess average- and worst-case operational performance. If the simulated performance breaches the current master approximation, adaptive Benders cuts are formulated and added to the master problem. Next, Bandit-guided ALNS, limited CMSA, and safe MIP-bound probing refine the current solution and accelerate convergence. This cyclic process continues until the stopping condition is met, yielding better risk-aware nurse allocation under a reasonable optimality constraint.

In contrast to traditional LBBBD approaches, the proposed framework integrates parallel scenario simulation, CVaR risk analysis, adaptive cut generation, and hybrid metaheuristic search into a cohesive optimization process. This integration significantly improves computational efficiency and solution robustness under demand uncertainty.

3.4 Integrated Solution Procedure

The global master is solved through bounded time slices. Each solution is simulation-validated to update the upper bound and generate valid cuts. Bandit-ALNS is activated after incumbent stagnation, CMSA after prolonged search stagnation, and probing after lower-bound stagnation. During the reserved proof phase, heuristic search is suspended, and full separation focuses the remaining budget on improving the global bound. The hyperparameters for Bandit-ALNS, CMSA, and MIP probing were fine-tuned by initial experiments and sensitivity assessments on a typical subset of occurrences. In cross-skilling, the compatibility of nursing skills is illustrated using an eligibility matrix. Precise LBBBD cuts retain their validity when obtained from subproblem certificates, while dedicated-skill cuts are either elevated using total eligible capacity or disregarded if their validity criteria are unmet. Consequently, heuristic information does not influence the certified lower bound or the optimality gap. Algorithm 1 summarizes the procedure.

Algorithm 1. Proposed v0.8.1 Bandit-CMSA
 Simulation-LBBD

N, S, K, T, Ω , and time limit T_{max}

- 1 Construct unrestricted risk-aware LBBD master M .
- 2 Initialize cut pool P , bandit state B , $UB = +$, and $LB = -\infty$.
- 3 Simulate maximum staffing and add initial valid cuts.
- 4 while elapsed time $< T_{max}$ do
- 5 Solve M within the current time slice.
- 6 Extract $\bar{x}$ and $\bar{y}$; simulate $\bar{y}$ over all scenarios.
- 7 Update the validated incumbent and UB .
- 8 Generate, select, and store violated LBBD cuts.
- 9 Update global LB only from unrestricted M .
- 10 If the incumbent stagnates: run bandit-ALNS/local branching.
- 11 If search stagnates: run an internal CMSA-LBBD cycle.
- 12 If LB stagnates: perform safe MIP-bound probing.
- 13 Inject improved incumbents, valid cuts, and certified fixings.
- 14 In proof phase, suspend heuristics and fully separate cuts.
- 15 if M is optimal and no simulation cut is violated then
- 16 return incumbent, UB , LB , Gap, OPTIMAL.
- 17 end if
- 18 end while
- 19 return incumbent, UB , LB , Gap, TIME_LIMIT.

Validity of the Bounds and Termination

The unrestricted LBBD master is solely responsible for the reported global lower bound. ALNS and CMSA are primal heuristics; their objective values may update UB but cannot update LB . Similarly, probing bounds certify only individual fixings. As a result, the optimality gap continues to hold mathematical validity:

$$\text{Gap} = (UB - LB) / \max(1, |UB|) \quad (8)$$

In this context, UB signifies the optimal simulation-validated roster objective. At the same time, LB represents the most precise bound obtained

from the unrestricted master problem. The procedure achieves an optimal status only when the unrestricted master problem is solved to optimality, and no simulation-induced cuts are violated. Alternatively, it provides the most validated roster and global boundary within the specified time frame.

Execution Principle

Scenario simulations are evaluated concurrently using deterministic seeds and coverage-profile caching. The reciprocal interaction is essential to the approach: the global master supplies reference incumbents to ALNS and CMSA, while newly simulated profiles generate cuts that refine the global master. This process simultaneously facilitates primal improvement and dual-bound advancement without conflating heuristic and precise data.

4. Computational Experiments

This section outlines the computational configuration, utilizing consistent instances, scenarios, objective weights, and time constraints across all algorithmic variants.

4.1 Demand and Test Case Scenarios

The reference scenario, which uses a one-day planning period split into sixteen hourly intervals, is derived from DataNursingHome.csv. The nursing staff includes thirteen individuals, divided into three specific skill categories: three registered nurses (RNs), four licensed practical nurses ($LPNs$), and six certified nursing assistants ($CNAs$). The candidate rotations allow for twelve possible shifts, each starting every two intervals and lasting either four or eight hours. Minimum simultaneous coverage requires three $CNAs$, one LPN , and one RN . Four stochastic demand scenarios are produced from a shared seed for each optimization cycle. The mean interarrival time for unscheduled requests is 20 minutes; 90% of requests are brief, 10% are urgent, and the skill demand probabilities for CNA , LPN , and RN are 0.70, 0.20, and 0.10, respectively. The desired service level is 10 minutes. To ensure competing plans face the same demand realizations, we perform out-of-sample validation using 1,000 new scenarios and common random values.

Table 1: Nursing-home dataset characteristics.

Characteristic	Value
Planning horizon	One day with 16 one-hour periods
Workforce	13 nurses: 6 CNAs, 4 LPNs, and 3 RNs
Legal shifts	12 shifts of 4 or 8 hours
Minimum coverage	3 CNAs, 1 LPN, and 1 RN
Hourly costs	CNA: 22; LPN: 30; RN: 42
Maximum workload	8 hours and 2 shifts per nurse
Minimum rest	8 periods
Skill assignment	Dedicated-skill assignment

The planning horizon, workforce composition, legal shifts, skill-specific coverage requirements, staffing costs, workload restrictions, rest requirements, and skill-assignment rules are all summarized in Table 1.

4.2 Algorithmic Configuration

The suggested system integrates adaptive LBBD cut separation, bandit-guided ALNS, local branching, constrained CMSA–LBBD optimization, and MIP-bound probing. To avoid outcome-dependent tuning, we set all parameters before running the tests.

Adaptive separation: The candidate-cut budget starts at 32, doubles after search stagnation, and is capped at 256. Full separation is triggered after extended stagnation or when the optimality gap drops below 5%, and aggregate cuts are enabled.

Bandit-led ALNS: After three iterations without improvement, the process activates and runs for up to eight rounds. Each round runs for four seconds, or 120 neighborhood moves. The reward-decay factor, temperature, probability floor, exploration rate, and Mutation prior are set to 0.35, 0.03, 0.20, 0.80, and 0.25, respectively.

Local branching: The neighborhood radius is 10. After six iterations without improvement, the process is triggered. It can run for up to three rounds of eight seconds each.

Restricted CMSA-LBBD: The process runs for a maximum of three rounds and triggers after four iterations without improvement. The initial,

incremental, and maximum core fractions are 0.45, 0.15, and 0.90, respectively. Each round allows three internal iterations and six seconds. **MIP-bound probing:** Probing runs for a maximum of two rounds and is triggered after six iterations without bound improvement. Each round looks at four candidate variables, giving each candidate 0.75 seconds and a fixing tolerance of 10–6.

4.3 Ablation and Comparative Design

We measure each algorithmic component's contribution via a nested ablation study. The nurse instance, demand scenarios, goal function, solver parameters, computational resources, and 300-second time restriction are all the same across all variations. To ensure a fair comparison, all tests begin without externally provided warm-start solutions.

Table 2: Algorithmic variants and their enabled components used in the ablation study.

ID	Method	Enabled components
A	Baseline hybrid	Adaptive LBBD, conventional ALNS, and local branching
B	Bandit-enhanced hybrid	Baseline hybrid with adaptive bandit operator selection
C	Bandit–CMSA hybrid	Bandit-enhanced hybrid with restricted CMSA–LBBD
D	Proposed framework	Bandit–CMSA hybrid with MIP-bound probing

Table 2 displays four hierarchical algorithmic variations intended to delineate the increasing impact of each element. Variant A establishes the foundation, whereas Variant B incorporates adaptive bandit operator selection. Variant C also integrates constrained CMSA–LBBD. In contrast, Variant D implements the full proposed framework by further enabling secure MIP-bound probing. This architecture enables a controlled evaluation of each improvement under uniform experimental conditions.

4.4 Metrics, Statistical Analysis, and Repetition

The studies are based on five paired scenario seeds, from 20260802 to 20260806. Each procedure runs once per seed under a 300-second time limit, yielding 20 controlled runs. The paired design guarantees that rival methodologies are evaluated on equivalent demand realizations. Given the small number of

replications, the analysis focuses on paired differences, effect sizes, and confidence intervals; inferential outcomes are carefully scrutinized. The assessment involves four categories of performance indicators:

- Optimality primal and gap integrals, time to best solution, final optimality gap, best feasible goal (UB), and best lower bound (LB).
- Solution quality indicators are staffing cost, mean total delay, 95% CVaR delay, fairness deviation and out-of-sample service performance.
- Decomposition efficiency metrics are LBBD iterations, generated and selected cuts, cut-selection ratio, rejected cuts, simulation runs, and cache hits.
- Matheuristic contributions include probing attempts and valid variable fixings, ALNS operator attempts and improvements, final operator probabilities, and CMSA rounds and improvements.

The relative optimality gap for a minimization issue is computed as:

$$\text{Gap (\%)} = 100 \times (UB - LB) / \max(|UB|, 10^{-10})$$

A zero gap shows that the developed mathematical model is globally optimal. We present the mean, median, standard deviation, minimum, maximum, and paired improvement with respect to the baseline for each performance metric. Paired Wilcoxon signed-rank tests provide additional evidence, with bootstrap confidence intervals for the observed effect sizes and a Holm adjustment for multiple comparisons. Computational flexibility is evaluated by examining convergence curves and performance outlines.

4.5 Reporting and Replicability

Comparisons are valid provided all methodologies use the same objective function, scenario seed, time limits, hardware, solver options and stopping criteria. When you conclude at the specified time, it means you have found a feasible solution. However, it is unknown whether it is optimal. This is not a math error. Thus, we present the optimality gap, the lower bound, and the upper bound together. We provide all experimental parameters, random seeds, quantitative

summaries, and convergence trajectories to enable independent validation and reproduction of the figures and tables.

5. Results and Discussion

We used the nursing-home scheduling scenario to test the proposed framework, with a computation time of 300 seconds. The study comprised 13 nurses, 3 proficiency categories, 12 regulatory shifts and 4 random demand situations. The study involved 13 nurses, 3 proficiency categories, 12 regulatory shifts and 4 random demand situations. The results are consistent with the planned framework that consists of local branching, BILS, ALNS, constrained CMSA-LBBD, adaptive LBBD, and MIP-bound exploration.

5.1 Computational Performance

This section compares the computational performance of the proposed Bandit-CMSA-Probing Logic Based Benders Decomposition (BCP-LBBD) framework on a 300-s constraint. Three interrelated aspects – quality of optimization, convergence properties, and computational efficiency – are analyzed. We measure optimization quality by the superior feasible objective (UB), the universal lower bound (LB), and the resulting optimality gap. Meanwhile, the number of LBBD iterations, cuts generated, simulation runs, and cache usage measure computational efficiency. Together, these measures offer a comprehensive view of the framework's performance in generating quality, risk-aware nursing shifts and ensuring a favorable performance/benefit balance between simulation-based and optimization-based assessments.

Table 3 shows the computation time for the proposed Bandit-CMSA-Probing Logic-Based Benders Decomposition (BCP-LBBD), with a runtime of 300s. In conclusion, the framework found a high-quality feasible solution, with an upper bound (UB) of 11,169.019 and a valid global lower bound (LB) of 10,754.446, resulting in an optimality gap of 3.712%. The optimality gap is very small, and because the global optimum could not be proven within the time limit, the proposed framework found a near-optimum solution while preserving the mathematical rigor of the established lower bound.

The optimization started with a feasible objective value of 13,849.635. It achieved an overall improvement of 19.36%, confirming the effectiveness of combining accurate decomposition with adaptive metaheuristic search. Furthermore, the framework ran 29 LBBD rounds and produced 1,568 adaptive Benders cuts to update the master problem. Scenario evaluation occurred over 364 simulation iterations, while 412 simulation cache hits avoided re-evaluating scenarios, reducing computational effort and improving optimization efficiency. Finally, the results show that the proposed framework performs well in solution quality, convergence speed, and computational resources, making it suitable for solving large-scale nurse scheduling problems under demand uncertainty.

Table 3: Summary of the computational performance.

Metric	Result
Termination status	Time limit
Best feasible objective, UB	11,169.019
Best lower bound, LB	10,754.446
Final optimality Gap	3.712%
Initial feasible objective	13,849.635
Improvement over initial solution	19.36%
Runtime	300.34 s
LBBD iterations	29
Incumbent evaluations	29
Cuts added	1,568
Simulation runs	364
Simulation cache hits	412

Fig. 2 shows the convergence trend of the proposed BCP-LBBD framework over the optimization process. In the first several iterations, the lower bound improved gradually, reflecting that the adaptive Benders cuts effectively strengthened the relaxation and improved the master problem at each iteration. The upper bound improved considerably at iterations 6, 12, and 24, indicating that the integrated heuristic mechanisms produced higher-quality nurse rosters that satisfied the constraints at various stages of the search. Improved primal and dual bounds gradually reduced the optimality gap, confirming the complementary interaction between exact decomposition and adaptive metaheuristic search. The convergence

pattern further suggests the framework could combine continuous lower-bound refinement with periodic incumbent improvements to produce steady optimization progress within the allowed computational budget, resulting in a near-optimal solution within the prescribed runtime.




Fig. 2 Convergence of the objective, lower bound, and optimality gap.

In the last four LBBD iterations, the optimization procedure entered the proof stage, during which no further improvement in the lower bound was detected. This behaviour indicates that the main relaxation was already significantly tightened, making it harder to reduce the residual optimality gap with the expended computational effort. However, the final optimality gap of 3.712% shows that the proposed decomposition-based framework obtained a near-optimal solution, with robust assurance of solution quality. These results demonstrate the proposed framework's ability to balance computational efficiency with high-quality, risk-aware optimization, even for complex stochastic nurse scheduling problems.

5.2 Incremental Performance of the Proposed Framework

To assess the incremental contribution of the proposed optimization framework, we compared its computational performance with three progressively enhanced baseline approaches under identical experimental conditions. We evaluated the independent SCIP solver, the traditional Logic-Based Benders Decomposition (LBBD) model, and the Logic-Based Benders Decomposition (LBBD) model

with Adaptive Large Neighborhood Search (LBBD-ALNS), along with the proposed Bandit-CMSA-Probing Logic-Based Benders Decomposition (BCP-LBBD) framework.

This structured experimental framework enables us to assess the different roles of decomposition, adaptive neighborhood search, and the proposed Bandit-CMSA-Probing improvements methodically. We applied all techniques to the same benchmark instance, the same stochastic demand scenarios, the same solver parameter settings, and a fixed time limit of 300 seconds to ensure an equitable and impartial evaluation. To better understand the adaptive search mechanism, we examined the internals of the Bandit-guided ALNS during optimization. The heuristic evaluated 158 potential schedules over 7 search rounds (840 neighborhood iterations), accepting 91 schedules and yielding an acceptance ratio of 57.59%.

This acceptance rate further reflects that the search struck a suitable balance between exploration and exploitation, allowing the framework to explore other nursing schedules without losing solution viability while improving the solutions already obtained. The multi-armed bandit dynamically adjusted the search strategy based on each neighborhood operator's past effectiveness. As a result, the computing tools gradually shifted toward operators that consistently produced better solutions. Meanwhile, less skilled operators were given a limited probability of selection for exploration to reduce the risk of converging too early. The most interesting result is that the Mutation operator alone produced all seven improvements.

In contrast, the Swap, Reassignment, Block Shift, and Coverage-Oriented Destroy-and-Repair operators yielded no further improvements for this benchmark case. The persistent multi-armed bandit continued to increase the Mutation operator's ultimate selection probability to 84.85%. In comparison, the other operators remained stably around 3.8%, consistent with the exploration threshold. This behavior suggests that the adaptive learning system identified the right search strategy for the current optimization terrain rather than distributing computational resources evenly among

all neighborhood frameworks. Selecting adaptive operators is particularly beneficial for stochastic nurse scheduling, where the effectiveness of certain neighborhoods depends heavily on the current roster arrangement and changes during the search process.

Fig. 3 also displays this flexible distribution of search effort. Unlike fixed search behavior throughout the entire optimization process, the proposed framework constantly updates the probability distribution of neighborhood operators according to their previous effect on the improvement of the incumbent, thus introducing a self-regulating mechanism that allows the search to be increasingly specialized as the optimization process progresses, while still providing sufficient diversification to prevent search stagnation. Experimental results show that the main advantage of the bandit-guided ALNS is not only producing better workable schedules, but also intelligently allocating computational resources to the most promising areas. The proposed approach greatly improves the primal-search performance of the BCP-LBBD algorithm by combining adaptive learning with feasibility-preserving neighborhood search, producing high-quality incumbent solutions to guide the exact decomposition process.

5.3 Ablation Analysis and Contribution of Algorithmic Components

To further investigate the effectiveness of the proposed BCP-LBBD framework, we conducted an ablation-oriented analysis to assess the individual contribution of its integrated algorithmic components. While the overall performance results show the framework's capability, a detailed understanding of each mechanism's contribution is essential to explain the source of the improvements. Consequently, the primary elements, such as bandit-guided ALNS, local branching, restricted CMSA-LBBD, and MIP-bound probing, were examined for their impact on the augmentation of primal solutions, intensification of search, improvement of decomposition, and minimization of the computational search space. The examination highlights the synergistic relationship between heuristic exploration and precise optimization methods. It illustrates how each input contributes to the overall effectiveness of the proposed risk-aware

nurse scheduling system. The bandit-guided ALNS mechanism demonstrated effectiveness by flexibly allocating computational resources among available neighborhood operators. Unlike conventional ALNS methods that use static or manually tuned selection probabilities, the learning process adjusts operator preferences based on their historical effectiveness in improving the current solution.

The adaptive process for the solved scheduling instance showed that the Mutation operator was the most effective search strategy, with its selection probability increasing gradually, while ensuring a minimum exploration level for the other operators. This behavior demonstrates that the bandit mechanism successfully balanced exploration and exploitation, focusing the search on beneficial solution transformations while preventing the complete exclusion of alternative neighborhoods. Thus, the adaptive operator selection strategy enhanced the primal search capability of the BCP-LBBD framework by providing better feasible solutions to drive the subsequent decomposition procedure.




Fig. 3 Search allocation and final selection probability of bandit-guided ALNS operators.

We introduced a local branching mechanism to the adaptive neighborhood search as an effective intensification strategy to improve good regions in the neighborhood identified during the optimization process. Local branching restricted the neighborhood to high-quality solutions near the incumbent, allowing local search of nearby feasible solutions while retaining the overall framework's diversity capability. During the optimization process, local branching was turned on in three phases, and two

additional improvements to the incumbent solution were produced. The improved solutions were then rolled over to the master problem. These improved solutions were then passed to the master problem, and the LBBD process could use better primal information for decomposition. Thus, local branching had the dual effect of directly improving solution quality and making heuristic exploration and exact optimization more tightly coupled in the proposed BCP-LBBD framework.

We refined the proposed optimization framework step by step by adding more precise and heuristic elements, while preserving the LBBD framework structure. Each configuration adds to the previous version by including an extra optimization component, allowing evaluation of each mechanism's contribution in the same computational context. The ablation findings confirm that every element of the proposed BCP-LBBD system enhances overall efficacy. Removing some mechanisms considerably impairs solution quality, convergence rate, or computing performance, emphasizing the need for their unified integration.

The study shows that the cooperative dynamics of deconstruction, parallel scenario analysis, adaptive cut production, and learning-oriented optimization strategies strengthen the proposed framework. The results confirm the effectiveness of the BCP-LBBD approach for solving complex stochastic scheduling problems. Table 4 presents the amalgamation of all the algorithmic constituents of the proposed framework. First configuration: Adaptive LBBD to match the master problem to scenario specific evaluations. Sustained Bandit-influenced ALNS is used to improve the primal search, where the ALNS is optimized for bandits and the computational resources are allocated adaptively based on the neighborhood operators. Later the limited CMSA-LBBD is adopted in order to intensify search with precise optimization of potential constrained solution cores. In conclusion, reliable MIP-bound probing is integrated into the full BCP-LBBD setup to detect genuine variable fixes, lessen the effective master-search area, and potentially fortify bound progression. The setups provide a systematic basis for the measurement of the supplementary effect of

each component under uniform experimental conditions.

Table 4: Activity of the proposed algorithmic components.

Version	Newly introduced component	Purpose
Baseline	Adaptive LBBD	Exact decomposition
+Bandit	Persistent bandit-guided ALNS	Better primal search
+CMSA	Restricted CMSA-LBBD	Exact intensification
v0.8.1	Safe MIP-bound probing	Stronger lower bound

5.4 Solution Quality and Risk Analysis

The final schedule had a personnel cost of 2,856, an average cumulative delay of 5,875.624, a 95% CVaR of a cumulative delay of 10,467.749, and a fairness variance of 48. The objective weights were pre-determined and then used to consolidate these results. The unrefined quality metrics and their proportional impacts on the goal of the ultimate incumbent timetable are encapsulated within Table 5. The timetable achieved a staffing expenditure of 2,856.000, an average overall delay of 5,875.624, a 95% CVaR overall delay of 10,467.749, and a fairness variance of 48.000. Applying the respective weights of 0.02, 1.00, 0.50, and 0.05 resulted in weighted contributions of 57.120, 5,875.624, 5,233.875, and 2.400, respectively, producing a final upper-bound objective value of 11,169.019. The objective was composed of mean delay (52.61%) and CVaR delay (46.86%), together accounting to 99.47% of the objective.

Table 5: Quality indicators of the final schedule.

Indicator (weight)	Raw value	Contribution
Staffing cost (0.02)	2,856.000	57.120
Mean delay (1.00)	5,875.624	5,875.624
95% CVaR delay (0.50)	10,467.749	5,233.875
Fairness deviation (0.05)	48.000	2.400
Final objective (UB)	—	11,169.019




Fig. 4 Weighted composition of the final risk-aware objective.

The weighted composition of the final objective is shown in Fig. 4. Panel (a) presents the two dominant delay-related contributions: mean delay, with a weighted value of 5,875.624 (52.61%), and 95% CVaR delay, with a weighted value of 5,233.875 (46.86%). Panel (b) magnifies the comparatively smaller staffing and fairness contributions of 57.120 and 2.400, respectively, which would otherwise be difficult to distinguish in the full objective scale. Collectively, the mean-delay and CVaR components account for 99.47% of the final objective.

The results show that the target was affected by the expected delays and tail risk with nearly 99.47% of the target. This is in line with the risk-aware model, which focused on timely delivery and protection from extreme congestion scenarios. The expenses related to personnel and equity are distinctly outlined. However, they are now less significant than the bigger goal. The CVaR metric was significantly higher than the usual delay, highlighting the need to account for unusual demand scenarios. Focusing solely on minimizing expected delays can result in a schedule that performs well under normal loads, but breaks down when the load is increased.

The use of CVaR encourages strategies that are more robust with respect to large fluctuations and time lags in demand. The successful target composition further illustrates the success of the proposed risk-aware optimization strategy. The explicit incorporation of CVaR into the decision-making process results in the resulting timetable to be less sensitive to peak demand scenarios and also to have acceptable operational costs. This is the

confirmation that the proposed BCP-LBBD framework can yield solutions beyond improving expected performance that are efficient, resilient, and can include equality.

5.5 Overall Discussion

First, the results of the computation matched the key features of the proposed BCP-LBBD methodology as described below: The proposed BCP-LBBD methodology improved the original viable answer by 19.36%, which was a significant improvement. Secondly, it made the lower bound adequate to reduce the final optimality gap to 3.712% even under the computational limit of 300 seconds. Thirdly, the bandit mechanism was very successful in identifying the neighborhood operator that it was most advantageous to use, and in directing the search process toward more promising solution improvements.

The results also highlight ongoing computational challenges. The lower bound did not change in the last proof phase, indicating that to reduce the gap further, the bounds must be tightened by improving the master problem relaxations, adding valid inequalities, or using better bound-based decomposition methods. Augmenting heuristic investigation may improve the current solution quality, although it does not guarantee further progress in optimality unless the lower bound improves simultaneously. Finally, the experimental evaluation confirms the effectiveness and convergence aspects of the proposed framework. From a practical point of view, the results validate its ability to generate high-quality, robust, and fast schedules that minimize risk.

6. Conclusion

This research introduced an integrated multi-skill problem-solving approach for staffing and scheduling in the nursing home industry using risk-aware parallel simulation-LBBD. The proposed approach is scenario-based and accounts for demand fluctuations, unlike previous deterministic scheduling techniques. It integrates Conditional Value-at-Risk (CVaR) to mitigate the effects of difficult service-delay situations. The known method is a combination of Logic Based Benders

Decomposition, Concurrent Scenario Assessment, Adaptive Cut Formulation and hybrid learning-driven search (Bandit-guided ALNS, constrained CMSA-LBBD and secure MIP-bound probing). The proposed approach was computationally evaluated to test its effectiveness in implementing within the limited computational resources. The objective value of the framework was 11,169.019, the lower bound was 10,754.446 (legitimate lower bound globally), and the optimality gap was very small of 3.712%, in 300 seconds. Moreover, the proposed approach was found to be capable of reaching the initial feasible solution with 19.36% improvement, highlighting the viability of the proposed approach for integrating accurate decomposition and adaptive heuristic search. The ablation test identified that all the integrated elements contributed to the performance, and the bandit mechanism helped to focus the search on the most promising neighborhood methods. The framework can assist managers in making decisions to maximize efficiency in staffing cost, to manage the risk of staff shortage and to ensure the policy takes the account of risk appetite, care priority and financial consideration.

A final schedule evaluation confirmed the risk-aware optimization concept. The target function places a strong weight on CVaR, suggesting that it is important to incorporate extreme demand scenarios into the target function, and not just in terms of increased expected performance. The proposed approach consequently provides a pragmatic balance between efficiency, resilience and equity in planning healthcare workforce in uncertainty. Future research may shed more light on larger-scale nursing environments, more sophisticated uncertainty modelling, more flexible rescheduling strategies, and more sophisticated mathematical relaxations for further enhanced lower-bound convergence, and for reduced discrepancies in computation.

Acknowledgement

We sincerely thank the Department of Computer Science and Artificial Intelligence, Faculty of Pure Sciences, University of Diyala, for their contributions, services, resources, and supportive academic environment, which were essential to the completion of this study. The author is sincerely

grateful to all who made direct or indirect contributions to the successful implementation of this study.

Conflict of Interest

The authors declare no conflicts of interest related to the dissemination of this paper.

References

- [1] Aydas, O. T., Ross, A. D., Scanlon, M. C., & Aydas, B. (2023). Short-Term nurse schedule adjustments under dynamic patient demand. *Journal of the Operational Research Society*, 74(1), 310–329.
- [2] BATUN, S., & KARPUZ, E. (2020). Nurse Scheduling and Rescheduling Under Uncertainty. *Hacettepe Üniversitesi İktisadi ve İdari Bilimler Fakültesi Dergisi*, 38(1), 75–95. <https://doi.org/10.17065/huniibf.483986>
- [3] Beddoe, G., Petrovic, S., & Li, J. (2009). A hybrid metaheuristic case-based reasoning system for nurse rostering. *Journal of Scheduling*, 12(2), 99–119. <https://doi.org/10.1007/s10951-008-0082-8>
- [4] Böðvarsdóttir, E. B., Bagger, N. C. F., Høffner, L. E., & Stidsen, T. J. R. (2022). A flexible mixed integer programming-based system for real-world nurse rostering. *Journal of Scheduling*, 25(1), 59–88. <https://doi.org/10.1007/s10951-021-00705-7>
- [5] Ceschia, S., Dang, N., De Causmaecker, P., Haspeslagh, S., & Schaerf, A. (2019). The Second International Nurse Rostering Competition. *Annals of Operations Research*, 274(1–2), 171–186. <https://doi.org/10.1007/s10479-018-2816-0>
- [6] Ceschia, S., Di Gaspero, L., Mazzaracchio, V., Policante, G., & Schaerf, A. (2023). Solving a real-world nurse rostering problem by Simulated Annealing. *Operations Research for Health Care*, 36, 100379. <https://doi.org/10.1016/j.orhc.2023.100379>
- [7] Combrink, A., Do, S., Bengtsson, K., Roselli, S. F., & Fabian, M. (2025). A Comparative Study of SMT and MILP for the Nurse Rostering Problem. 11th 2025 International Conference on Control, Decision and Information Technologies, CoDIT 2025, 2105–2110. <https://doi.org/10.1109/CoDIT66093.2025.11321384>
- [8] De Causmaecker, P., & Vanden Berghe, G. (2011). A categorization of nurse rostering problems. *Journal of Scheduling*, 14(1), 3–16. <https://doi.org/10.1007/s10951-010-0211-z>
- [9] Elçi, Ö., & Hooker, J. (2022). Stochastic Planning and Scheduling with Logic-Based Benders Decomposition. *INFORMS Journal on Computing*, 34(5), 2428–2442. <https://doi.org/10.1287/ijoc.2022.1184>
- [10] Filippi, C., Guastaroba, G., & Speranza, M. G. (2020). Conditional value-at-risk beyond finance: a survey. *International Transactions in Operational Research*, 27(3), 1277–1319. <https://doi.org/10.1111/itor.12726>
- [11] Forbes, M. A., Harris, M. G., Jansen, H. M., van der Schoot, F. A., & Taimre, T. (2024). Combining optimization and simulation using logic-based Benders decomposition. *European Journal of Operational Research*, 312(3), 840–854. <https://doi.org/10.1016/j.ejor.2023.07.032>
- [12] He, F., Chausalet, T., & Qu, R. (2019). Controlling understaffing with a conditional Value-at-Risk constraint for an integrated nurse scheduling problem under patient demand uncertainty. *Operations Research Perspectives*, 6(February), 100119. <https://doi.org/10.1016/j.orp.2019.100119>
- [13] Juan, A. A., Uguina, A. R., Escoto, M., & Medina, V. (2026). Sim-Exact Methods for Stochastic Optimization: A Complementary Approach to Simheuristics. *Mathematics*, 14(9). <https://doi.org/10.3390/math14091518>
- [14] Kim, K., & Mehrotra, S. (2015). A two-stage stochastic integer programming approach to integrated staffing and scheduling with application to nurse management. *Operations Research*, 63(6), 1431–1451. <https://doi.org/10.1287/opre.2015.1421>
- [15] Landeghem, H. V. A. N. (2004). THE STATE OF THE ART OF NURSE ROSTERING. 441–499.
- [16] Legrain, A., Omer, J., & Rosat, S. (2020). An online stochastic algorithm for a dynamic nurse scheduling problem. *European Journal of Operational Research*, 285(1), 196–210. <https://doi.org/10.1016/j.ejor.2018.09.027>
- [17] Nasirian, A., Zhang, L., Costa, A. M., & Abbasi, B. (2025). Multiskilled workforce staffing and scheduling: A logic-based Benders' decomposition approach. *European Journal of Operational Research*, 323(1), 20–33. <https://doi.org/10.1016/j.ejor.2024.11.033>
- [18] Rachuba, S., Reuter-Oppermann, M., & Thielen, C. (2024). Integrated planning in hospitals: a review. In *OR Spectrum* (Vol. 2020, Number 0123456789). Springer Berlin Heidelberg. <https://doi.org/10.1007/s00291-024-00797-5>

- [19] Respicio, A., Moz, M., Pato, M. V., Somensi, R., & Dias Flores, C. (2018). A computational application for multi-skill nurse staffing in hospital units. *BMC Medical Informatics and Decision Making*, 18(1), 1–9. <https://doi.org/10.1186/s12911-018-0638-2>
- [20] Ryu, M., & Jiang, R. (2025). Nurse Staffing Under Absenteeism: A Distributionally Robust Optimization Approach. *Manufacturing and Service Operations Management*, 27(2), 624–639. <https://doi.org/10.1287/msom.2023.0398>
- [21] Tsang, M. Y., Shehadeh, K. S., Curtis, F. E., Hochman, B. R., & Brentjens, T. E. (2025). Stochastic Optimization Approaches for an Operating Room and Anesthesiologist Scheduling Problem. *Operations Research*, 73(3), 1430–1458. <https://doi.org/10.1287/opre.2022.0258>
- [22] Van Den Bergh, J., Beliën, J., De Bruecker, P., Demeulemeester, E., & De Boeck, L. (2013). Personnel scheduling: A literature review. *European Journal of Operational Research*, 226(3), 367–385. <https://doi.org/10.1016/j.ejor.2012.11.029>
- [23] Yao, X., Shehadeh, K. S., & Padman, R. (2024). Multi-resource allocation and care sequence assignment in patient management: a stochastic programming approach. *Health Care Management Science*, 27(3), 352–369. <https://doi.org/10.1007/s10729-024-09675-6>
- [24] Zuhanda, M. K., Hartono, Rahman, S., Gio, P. U., & Ongko, E. (2026). Metaheuristic nurse scheduling with hospital clustering using flower pollination algorithm. *Bulletin of Electrical Engineering and Informatics*, 15(3), 2406–2416. <https://doi.org/10.11591/eei.v15i3.11243>