

Modeling the Dutch Healthcare Workforce: An Integer Programming Game for Nurse Scheduling Problem

Master thesis submitted to Delft University of Technology in partial
fulfilment of the requirements for the degree of

MASTER OF SCIENCE

Engineering and Policy Analysis
Faculty of Technology, Policy and Management
Delft University of Technology

Student name: Zhaoyue Zhang
Student number: 6080553
Thesis defence date: 5 January 2026

Graduation committee

Dr. K. (Kateřina) Staňková,	T&L	Chair
Dr. P.S.A. (Patrick) Stokkink,	T&L	First supervisor
Dr. I. (Irene) Grossmann,	3S	Second supervisor

Abstract

The Dutch healthcare system is currently operating under intensifying fiscal and workforce pressures, with recent budget constraints exacerbating long-standing nurses shortages. A significant share of national health expenditure is devoted to long-term care, a sector that relies heavily on human capital and is thus especially vulnerable to labor scarcity. This operational challenge is further complicated by institutional fragmentation: nurses typically work under a mix of fixed hospital contracts and flexible arrangements, while hospital managers, insurers, and the workforce often pursue divergent and conflicting priorities. Consequently, nurse scheduling has evolved from a simple operational task into a complex, socio-technical “wicked problem”, a critical policy question of how to allocate scarce human resources efficiently and equitably in a non-cooperative environment.

This study addresses the nurse scheduling problem as a strategic decision challenge within this multi-actor system. A novel methodological framework that bridges the gap between Operations Research and Game Theory is developed. While traditional approaches typically formulate the Nurse Scheduling Problem (NSP) as an optimization model assuming a single objective, this research introduces the “Nurse Scheduling Game” (NSG). This framework explicitly captures the stakeholders’ strategic behaviors, preferences, and interactions by modeling the problem as an Integer Programming Game (IPG). The interaction within this game framework is analyzed through two distinct equilibrium concepts that represent different governance structures: the simultaneous Nash Equilibrium and the hierarchical Stackelberg Equilibrium.

To model the current state of fragmented decision-making, the Nurse Scheduling Game is first analyzed under the Nash Equilibrium concept. In this simultaneous interaction, the Hospital Manager and the Nurses act as independent agents optimizing their objectives concurrently. The Manager seeks to minimize the total budget, consisting of nurses’ salaries and financial penalties for patient waiting times, while the Nurses independently seek to minimize their workload and maximize their schedule preferences. In this uncoordinated setting, formulated as a Generalized Nash Equilibrium Problem (GNEP), neither party has insight into the other’s strategy; they respond to the system’s aggregate constraints rather than to each other’s specific moves. This equilibrium often leads to a stalemate where neither party can unilaterally improve its outcome.

In contrast, to explore the potential for strategic improvement, the Nurse Scheduling Game is then analyzed under the Stackelberg Equilibrium concept. Here, the interaction is modeled hierarchically rather than simultaneously, formulated as a Bilevel Integer Problem (BIP). The Hospital Manager acts as the “Leader,” moving first by setting capacity and schedule constraints, while the Nurses act as “Followers,” responding optimally to this fixed schedule. Crucially, in this equilibrium, the Leader explicitly anticipates the Followers’ “best response” function, allowing the Manager to design a schedule that aligns the Nurses’ incentives with the system’s financial and operational goals. Solving for this equilibrium is computationally intractable using standard methods; therefore, this study implements a novel Monte Carlo Multilevel Optimization (MCMO) framework. This algorithmic approach combines stochastic simulation to approxi-

mate the Nurses’ reaction function with a heuristic search to optimize the Manager’s strategy, enabling the solution of realistic, large-scale hospital instances.

The framework was applied to a representative case study of a mid-sized Dutch hospital department to quantify the impact of these different decision structures. The computational results provide striking evidence of the costs associated with uncoordinated planning. The analysis indicates that the outcome at the uncoordinated Nash Equilibrium is approximately twice as expensive as that at the coordinated Stackelberg Equilibrium. Under the Nash dynamics, the system converges to a state of “defensive buffering”, where the Manager, unable to predict the workforce’s operational response, overallocates capacity to avoid potential bottlenecks. This results in an outsized schedule in which high salaries coexist with inefficient patient throughput.

Conversely, the Stackelberg Equilibrium demonstrates the significant value of strategic anticipation. By internalizing the workforce’s reactions, the Manager can transition from a volume-based to a precision-based strategy, allocating staff to match stochastic peaks in patient demand. In the large-scale experiments, this integrated approach achieved a 51.0% reduction in total system costs compared to the Nash baseline. Importantly, this efficiency did not come at the expense of patient care; patient waiting times were reduced by approximately 11.8%.

The results provide concrete insights into fairer and more cost-effective scheduling strategies, translating these mathematical findings into actionable policy implications. The superior performance of the Stackelberg Equilibrium suggests that the solution to the “flex” crisis and budget overruns lies in shifting from reactive to anticipatory governance. The study proposes a new, structured decision-making procedure that couples budget setting and schedule design within a single decision loop. Practically, this supports the implementation of Algorithmic Workforce Management systems that allow hospital boards to simulate the downstream behavioral effects of their financial policies before implementation.

While the model necessarily simplifies the complexity of real-world healthcare systems, its main contribution is methodological: it introduces a unified game-theoretic optimization framework for workforce planning that explicitly models the inherent tension between strict financial limitations and the autonomous preferences of the labor force. The results provide both quantitative and strategic insights into how scheduling efficiency, stakeholder incentives, and governance structures interact. By demonstrating that strategic coordination can unlock significant latent capacity within the existing workforce, this research lays the groundwork for more adaptive and evidence-based decision-support tools, offering a viable pathway to sustainability for the Dutch healthcare system.

Contents

1	Introduction	8
1.1	Background and Motivation	8
1.2	Problem Statement	9
1.3	Relevance to Engineering and Policy Analysis (EPA)	11
1.4	Thesis Outline	11
2	Literature Review and Theoretical Foundations	12
2.1	Operations Research in Healthcare Management	13
2.1.1	The Nurse Scheduling Problem (NSP)	13
2.1.2	Healthcare Budgeting and Resource Allocation	14
2.2	Game Theory in Health system	16
2.2.1	Foundational Concepts	16
2.2.2	Simultaneous Interaction: The Nash Equilibrium	17
2.2.3	Hierarchical Interaction: The Stackelberg Equilibrium	19
2.3	Integer Programming Games	20
2.3.1	Conceptual Formulations	21
2.3.2	Solution Approaches	23
2.4	Synthesis and Research Gap	24
3	Model Formulation and Methodology	26
3.1	Nurse Scheduling Problem (NSP)	26
3.1.1	NSP Manager Problem	27
3.1.2	NSP Nurses Problem	28
3.2	Nurse Scheduling Game (NSG)	29
3.2.1	NSG Nash Equilibrium	31
3.2.2	NSG Stackelberg Equilibrium	31
3.2.3	Pure VS. Mixed Equilibria	32
3.3	Solution Approach	33
3.3.1	Nash Equilibrium: Iterative Best-Response Method	34
3.3.2	Stackelberg Equilibrium: Monte Carlo Multilevel Optimization (MCMO)	35
4	Computational Experiments and Results	39
4.1	Experimental Design	39
4.1.1	Case Study	39
4.1.2	Data Description	40
4.1.3	Key Performance Indicators (KPIs)	40
4.2	Implementation Details	42
4.3	Results Summary: Comparative Overview	42
4.4	Results from the Small Instance	44
4.5	Results from the Large Instance	46
4.6	Model Validation and Limitations	49
4.6.1	Impact of Workforce Heterogeneity	49

4.6.2	Equilibrium Stability and Mixed Strategies	50
5	Discussion	52
5.1	Interpretation of Results	52
5.2	Policy Implications	54
6	Conclusion and Future Work	55
6.1	Summary of Research	55
6.2	Contributions and Novelties	56
6.3	Future Research	57

List of Figures

- 4.1 Distribution of equilibrium payoffs for Nash (Blue) and Stackelberg (Orange) across 100 runs. Stackelberg reduces Manager Cost but induces higher variance in Nurse Payoff. 45
- 4.2 Manager’s Payoff (Nash). 46
- 4.3 Nurses’ Payoff (Nash). 46
- 4.4 MCMO Search: Manager’s Cost Trajectory. 46
- 4.5 MCMO Search: Nurses’ Response Trajectory. 46
- 4.6 Visual comparison of staffing patterns. Red = Inpatient, Blue = Outpatient, White = Off. Both equilibria result in dense schedules, but the Stackelberg configuration achieves superior demand matching. 47
- 4.7 Nash Manager Payoff. 47
- 4.8 Nash Nurses Payoff. 47
- 4.9 Stackelberg Manager Search Trajectory. 48
- 4.10 Stackelberg Nurses Search Trajectory. 48
- 4.11 Clustering of equilibria over 10 runs. Stackelberg (Orange) achieves lower costs but demands higher nurse effort than Nash (Blue). 49
- 4.12 Visual comparison of large-scale rosters ($N = 49$). Both are dense (reflecting 0.6 FTE utilization), but the Stackelberg schedule creates fewer waiting days through precise role allocation. 50
- 4.13 Distribution of equilibrium payoffs for Symmetric instances. Compared to the Asymmetric results, the variance in Nurse Payoff is significantly reduced, reflecting the lack of individual differentiation. 51
- 4.14 Sample schedules for the Symmetric workforce. Note the lack of specialized role patterns compared to the Asymmetric case. 52
- 4.15 Probability Matrices for the Mixed Nash Proxy. These heatmaps represent the probability of assigning a nurse to a specific role, illustrating a flexible staffing strategy. 52
- 4.16 Aggregated Mixed Strategy Profile. Red/Blue intensity indicates the probability of Inpatient/Outpatient assignment. 53

List of Tables

4.1	Characteristics of Computational Instances	39
4.2	Summary of Computational Instance Parameters	41
4.3	Aggregate Equilibrium Comparison. This table compares the financial (Cost), operational (Workload), and service (Wait) metrics between the uncoordinated Nash outcome and the coordinated Stackelberg outcome.	43
4.4	Impact of Workforce Heterogeneity on Equilibrium Outcomes (Small Instance) .	49
4.5	Comparison of Pure Strategy vs. Mixed Strategy (Nash Equilibrium)	51

1 Introduction

1.1 Background and Motivation

Healthcare systems worldwide continue to face the persistent challenge of allocating scarce resources efficiently. This challenge has become increasingly urgent as rising healthcare costs, population aging, and widespread workforce shortages converge to strain capacity and budgets (Varkevisser et al., 2023). Across Europe, and particularly in the Netherlands, governments are finding it challenging to forecast and manage health workforce needs effectively, making resource allocation not only a financial issue but a critical policy concern (European Commission, 2021).

In the Netherlands, these pressures have reached a critical point. The Scientific Council for Government Policy (WRR) has warned that the current trajectory of healthcare demand and expenditure is “unsustainable” (onhoudbaar). Without structural reform, projections indicate that by 2060, nearly one in three workers would need to be employed in healthcare to maintain current service levels (The Scientific Council for Government Policy (WRR), 2021). This challenge is further intensified by the Dutch governance model of regulated competition, which was designed to improve efficiency by blending market incentives with public regulation and oversight. While the model improved transparency and patient choice, it also led to fragmented budgeting and misaligned institutional incentives across insurers, hospitals, and municipalities (Netherlands Bureau for Economic Policy Analysis (CPB), 2003; Zorginstituut Nederland, 2024). As a result, persistent coordination failures and inefficiencies have prevented the system from achieving the efficiency gains it was designed to deliver (Authority for Consumers and Markets (ACM), 2016; Schäfer et al., 2010).

The effects of this macro-level tension are most visible in the ongoing labor crisis, particularly within the nursing sector. Chronic staff shortages have become the norm, with over 50,000 vacancies reported across the broader care sector and 40% of healthcare employers struggling to fill positions (Coöperatie VGZ, 2023; TNO, 2024). Forecasts suggest that by 2035, the healthcare sector could face a shortfall of more than 260,000 workers (AtoZ Serwis Plus, 2025). The human impact of these shortages is significant: high work pressure, frequent overtime, and increasing sick leave rates have led to growing dissatisfaction, burnout, and turnover among nurses (Enea et al., 2024). The situation has been further aggravated by fiscal tightening, including a €315 million reduction in medical training budgets and the cancellation of the “STAP” lifelong-learning scheme, which had been widely used by healthcare staff (Jacobs, 2024; NL Times, 2023).

This combination of structural inefficiencies and workforce pressures underscores the urgent need for new approaches to healthcare system planning. Even large-scale policy initiatives, such as the Integrated Care Agreement (Integraal Zorgakkoord, IZA), have struggled to alleviate the workload and “unburden” professionals amid these constraints (Zorgakkoorden.nl, 2022). Consequently, improving healthcare efficiency is not simply a technical optimization problem but a socio-institutional challenge. Future decision-support and optimization methods must account for the intricate interplay between efficiency, equity, and institutional feasibility (Zorginstituut

Nederland, 2024). Only by aligning analytical models with political, organizational, and clinical realities can such methods offer practical and sustainable solutions for the Dutch healthcare system (Boxebeld et al., 2024; Berkhout et al., 2025).

1.2 Problem Statement

While the background establishes the macro-level crisis of healthcare workforce shortages, this thesis addresses the specific, operational problem at the institutional level, where the consequences of this crisis are most deeply felt. The core of the problem is not a simple numerical deficit of nurses, but that existing workforce planning mechanisms fail to allocate the available staff effectively. In practice, nurse staffing and scheduling decisions are made in a highly decentralized way, often independently by individual departments or hospitals (AZW, 2024). This fragmented governance creates a system of misaligned incentives: hospital managers optimize for localized budgets, individual nurses (as flexible or fixed-contract agents) seek to optimize their own schedules, and insurers pursue separate cost-containment goals (Coöperatie VGZ, 2023). These competing goals make staffing decisions short-term and uncoordinated, leading to fragmented planning and system-wide sub-optimal outcomes.

This coordination failure creates a reinforcing cycle of negative effects. When planning is isolated, hospitals increasingly rely on expensive temporary or flexible staff to fill last-minute gaps, driving budget overruns (van Liempt and Brussaard, 2025). At the same time, uneven staffing places disproportionate pressure on permanently employed nurses and leads to unsustainable workloads. This pressure contributes to burnout, rising sick leave, and high turnover (TNO, 2024; The Scientific Council for Government Policy (WRR), 2021). These effects further reduce the available workforce, increase the workload for those who remain, and deepen reliance on costly short-term staffing solutions. Ultimately, this pattern undermines continuity of care and poses risks to patient safety (Inspectie Gezondheidszorg en Jeugd (IGJ), 2022).

Although the Operations Research (OR) literature provides advanced optimization tools (Erhard et al., 2018), these approaches typically assume a single, central decision-maker who can impose a schedule. That assumption does not hold in the real world, where nurse scheduling emerges from the interactions of multiple stakeholders with conflicting objectives (Petrovic and Vanden Berghe, 2012). To bridge this gap, this thesis proposes a Game-Theoretic approach that fundamentally redefines the scheduling problem. By modeling the “Manager” and the “Nurse” as distinct players with their own mathematical objective functions, the proposed framework captures the strategic tensions that standard optimization ignores. Instead of calculating a static, imposed schedule, the method solves for an equilibrium, which is a state where the schedule accounts for the nurses’ rational reactions to the manager’s decisions. This approach allows the model to explicitly simulate the negotiations and trade-offs regarding workload and flexibility.

Research Questions

To implement this framework and address the identified gaps, this thesis is driven by the fol-

lowing **Main Research Question (MRQ)**: How can the integration of Mixed-Integer Programming (MIP) and Game Theory (GT) create a decision-support framework that improves the efficiency and feasibility of nurse workforce scheduling within the Dutch healthcare system, given its fragmented, multi-stakeholder structure?

To answer this question, the research is guided by three specific sub-questions (SQs):

- **SQ1 (Modeling Strategic Interaction)**: How can the decision-making process between hospital manager and nurses be mathematically formulated as a non-cooperative game?
- **SQ2 (Equilibrium Analysis)**: What are the quantitative differences between an uncoordinated planning approach (Nash Equilibrium) and a strategic, hierarchical planning approach (Stackelberg Equilibrium)?
- **SQ3 (Managerial Insight)**: How can the insights derived from these equilibria inform hospital administrators to allocate nurses more strategically, anticipating stakeholder behavior to reduce system costs and waiting times?

To address this gap, this thesis is driven by the main research question: How can the integration of mathematical optimization (MIP) and game theory (GT) create a decision-support framework that improves the efficiency and feasibility of nurse workforce scheduling within the Dutch healthcare system, given its fragmented, multi-stakeholder structure?

To answer this question, the research is structured in four steps. First, it formulates the nurse scheduling problem as an MIP model from the perspectives of different stakeholders to present a centralized cost minimization problem. Second, it develops a game to capture the decentralized and strategic decision-making of different actors. Third, it applies the game in a computational case study, quantitatively searching for and comparing the Nash Equilibrium and Stackelberg Equilibrium of the game. Finally, it interprets these findings to derive practical insights for healthcare administrators on how to allocate nurses more strategically by anticipating stakeholder behavior.

This thesis focuses on the nurse scheduling problem as a representative and critical case of healthcare resource allocation under fiscal constraints. The initial scope is limited to a single department setting in a hospital to allow for controlled modeling and computational tractability. Within this context, the study focuses on nurse scheduling, excluding other workforce groups such as physicians, allied health professionals, and support staff.

Although the models are tested at the department level, their design is intended to be generalizable. The proposed frameworks can be scaled to incorporate multiple hospitals or extended to regional and national levels of workforce planning. Likewise, although this thesis uses deterministic formulations for scheduling and budget constraints, incorporating stochastic elements to reflect demand uncertainty (such as patient arrivals or staff sick leave) is a potential future

research. These ensure a manageable research scope while maintaining the broader applicability of the findings.

1.3 Relevance to Engineering and Policy Analysis (EPA)

This thesis is situated at the intersection of technical modeling and strategic governance, reflecting the core mission of the Engineering and Policy Analysis (EPA) discipline: to analyze and solve “Grand Challenges” at the interface of technology, society, and policy.

The crisis in Dutch healthcare workforce planning constitutes a classic “wicked problem” (Rittel and Webber, 1973). It is characterized by deep complexity, interdependent causes, and the absence of a single “correct” solution. The challenge is not merely logistical but also socio-political: it involves managing scarce human resources within a fragmented landscape of conflicting stakeholder interests, ranging from hospital boards and insurers to unions and individual nurses. Recent policy developments, such as the friction surrounding the Integraal Zorgakkoord (IZA) and the opting-out of General Practitioners due to funding disputes (Zurhake, 2022), highlight the failure of top-down technocratic solutions that ignore these multi-actor dynamics.

In line with the EPA curriculum, this research treats nurse scheduling as a socio-technical system. It acknowledges that technical efficiency (optimal schedules) cannot be divorced from social reality (nurse autonomy and burnout). The methodologies employed Mixed-Integer Programming (MIP) and Game Theory to bridge this gap. MIP provides a rigorous mathematical foundation for modeling the system’s physical constraints (capacity, shifts, patient demand). Game Theory allows modeling the strategic behavior of autonomous agents. By simulating the non-cooperative interaction between Managers and Nurses, the thesis moves beyond standard optimization to explore the political feasibility and strategic stability of workforce policies.

Ultimately, this work demonstrates how advanced modeling and simulation can serve as a decision-support tool for policy. By quantifying the current system, the thesis translates abstract mathematical findings into concrete governance advice. It illustrates how data-driven insights can inform the design of better institutional games, ensuring that policy interventions are grounded in a realistic understanding of stakeholder behavior rather than idealized assumptions.

1.4 Thesis Outline

This thesis contains six chapters that provide a comprehensive exploration of the research problem, methodology, and findings, and that develop an analytical framework for nurse workforce planning in the Dutch healthcare system.

Chapter 1: Introduction establishes the research context by outlining the significant workforce challenges in the Dutch healthcare system, with particular attention to nurses shortages. It identifies the core problem, defines the research questions and objectives, clarifies the study’s scope, and explicitly situates the study within the Engineering and Policy Analysis (EPA)

discipline by characterizing nurse scheduling as a socio-technical “wicked problem” requiring multi-actor analysis. The chapter provides the overall structure of the thesis.

Chapter 2: Literature Review and Theoretical Foundations establishes the theoretical foundation for the study. It begins by reviewing the literature on operations research in healthcare, focusing on the classical Nurse Scheduling Problem (NSP) and its formulation using Mixed-Integer Programming (MIP). It then introduces foundational concepts of non-cooperative game theory and reviews the definitions of the Nash and Stackelberg equilibria. The chapter concludes by discussing the emerging field of Integer Programming Games (IPGs) to identify the research gap this study addresses precisely.

Chapter 3: Model Formulation and Methodology details the methodological framework and mathematical formulation of the proposed models. This chapter first formulates the nurse scheduling problem as a centralized Mixed-Integer Programming model. It then extends this optimization framework into a game. Nash Equilibrium and Stackelberg Equilibrium are computed to analyze strategic stakeholder interactions. The chapter also explains the computational solution approaches employed, including the implementation of Monte Carlo multilevel optimization techniques.

Chapter 4: Computational Experiments and Results presents the comprehensive empirical findings of the study. It begins by defining the experimental setup and parameter validation. The chapter then provides a detailed quantitative comparison of the Nash and Stackelberg equilibria across both small validation instances and realistic large-scale hospital scenarios. Finally, it assesses the robustness of the proposed framework by analyzing the impact of workforce heterogeneity and investigating the stability of equilibrium solutions.

Chapter 5: Discussion interprets the computational results through a managerial and policy lens. It analyzes the mechanisms driving the observed inefficiencies and translates the technical results into actionable policy recommendations for Dutch hospital administrators, focusing on integrated planning and algorithmic decision support.

Chapter 6: Conclusion and Future Work summarizes the main theoretical and practical contributions of the research, reflects on the limitations of the proposed framework, and outlines directions for future work in extending game-theoretic models for sustainable healthcare workforce planning.

2 Literature Review and Theoretical Foundations

This chapter provides the theoretical foundation for the models developed in this thesis by reviewing three interconnected research directions. It begins with Operations Research, reviewing how optimization techniques, especially Mixed-Integer Programming, have been used to address healthcare challenges such as nurse scheduling and resource allocation. The chapter then turns to Game Theory, introducing it as a framework for understanding strategic interactions in de-

centralized health systems where multiple stakeholders pursue different objectives. Finally, it brings these perspectives together through the emerging field of Integer Programming Games, which integrates optimization-based decision-making with strategic behavior. By outlining the strengths and limitations of each approach, the chapter identifies the gap that motivates this thesis and sets the target for designing a new decision-support framework tailored to the complex problem of healthcare workforce planning.

2.1 Operations Research in Healthcare Management

Operations Research (OR) has become a fundamental methodology in healthcare management for improving efficiency, resource allocation, and workforce planning. By translating complex decision problems into mathematical models, OR allows decision-makers to explore trade-offs, test scenarios, and identify solutions that would be difficult to obtain through intuition alone (Rais and Viana, 2011; Ahmadi-Javid et al., 2017). These applications are typically categorized into three levels of decision-making. Strategic applications involve long-term decisions, such as hospital location or large-scale budget allocation (Latruwe et al., 2023). Tactical applications focus on medium-term decisions, such as patient flow policies or workforce capacity planning (Rachuba et al., 2024). Finally, operational applications handle short-term, real-time decisions, such as operating room planning (Xiao and Yoogalingam, 2022) or, most central to this thesis, nurse-shift scheduling.

Over the past two decades, plenty of research has shown how OR methods can help health systems to manage growing demand under strict constraints. Systematic reviews of patient-flow modeling show widespread use of queuing theory, discrete-event simulation, and hybrid optimization–simulation frameworks to reduce waiting times and improve service coordination in both hospital and community care settings (Palmer et al., 2018). Recent research integrates predictive models with optimization frameworks, enabling proactive allocation of beds, clinical staff, and operating-room capacity based on anticipated demand patterns (Russo et al., 2025). Within this broad landscape, workforce planning, and particularly nurse scheduling, has emerged as one of the most critical and heavily researched applications of OR. Labor is the most significant cost component in most health systems, and workforce shortages are now the primary constraint on expanding care capacity in many countries. The Nurse Scheduling Problem (NSP) is therefore one core of healthcare OR research.

2.1.1 The Nurse Scheduling Problem (NSP)

The Nurse Scheduling Problem (NSP) is one of the most classic and extensively studied combinatorial optimization problems in healthcare OR. It focuses on assigning nurses to shifts over a planning horizon while meeting staffing requirements and respecting a complex set of constraints, which typically include: coverage (minimum nurses per shift and skill mix), legal or contractual rules (maximum hours, rest periods), institutional policies (rotation rules), and individual preferences (days off, shift type) (Burke et al., 2004). Because of this combinatorial structure, the NSP is NP-hard: the number of possible schedules grows exponentially, making exact optimization computationally challenging for realistically sized instances.

A common way to model the NSP is as a Mixed-Integer Programming (MIP) problem. In this formulation, binary decision variables indicate whether a particular nurse works a given shift; the objective function seeks to minimize total cost (e.g., overtime, external staff) or penalties for violating preferences; and the constraints enforce labor rules and coverage requirements (Mystakidis et al., 2024).

Because real-world NSP instances are often too large for pure MIP solvers, an extensive literature has developed around heuristic and metaheuristic approaches. Foundational reviews by Burke et al. (2004) and Ernst et al. (2004) classified methods such as tabu search, simulated annealing, and genetic algorithms. More recent work continues this trend, often in hybrid “metaheuristic” frameworks. For example, Ceschia et al. (2023) used a simulated annealing approach to solve a complex real-world problem, while Hu et al. (2023) developed a branch-and-price method for large-scale instances, demonstrating that advanced decomposition can improve computational performance.

This methodological development has benefited from standardized benchmark instances, such as the International Nurse Rostering Competitions (INRC), which provide a basis for fair algorithmic comparison (Ceschia et al., 2019; Kheiri et al., 2021). As a recent complexity analysis by den Hartog et al. (2023) emphasizes, even simplified NSP variants remain computationally difficult, reinforcing the need for problem-specific algorithms.

Initially, many NSP models focused narrowly on coverage and cost. More recent work incorporates fairness, robustness, and staff well-being. Some models explicitly balance workload, minimize undesirable shift patterns, linking schedule design to occupational health and safety (De Causmaecker and Vanden Berghe, 2011). Others use stochastic or robust optimization to create schedules that are less sensitive to demand fluctuations or last-minute absences.

This is particularly relevant for the Dutch healthcare system, where nurse shortages, high sick leave rates, and reliance on external agency labor create persistent scheduling pressures. Dutch hospitals operate under strict labor regulations, explicit fairness norms, and varied contract types (fixed employment vs. self-employed arrangements). These institutional features directly mirror the constraints captured in NSP models.

2.1.2 Healthcare Budgeting and Resource Allocation

Beyond workforce scheduling, Operations Research also plays a significant role in broader questions of healthcare budgeting, capacity planning, and resource allocation.

A key framework in strategic resource allocation is Program Budgeting and Marginal Analysis (PBMA). PBMA provides a structured approach to identify where resources can be reduced and reinvested to generate greater value (Mitton and Donaldson, 2003). Over time, its guidelines have been refined (Peacock et al., 2010), and the framework has been implemented at national

and regional levels, including Wales (Edwards et al., 2014), Canada, and Brazil (Seixas and Mitton, 2021). PBMA has proven especially useful in times of financial pressure. During austerity periods, for example, studies show that the approach helps health systems make transparent trade-offs and find efficiency gains without sabotaging quality (Mitton et al., 2014).

Optimization techniques, particularly integer and linear programming, are frequently used to support these budgeting decisions because they can incorporate strict constraints and systematically evaluate marginal returns (Earnshaw and Dennett, 2003). These models have been applied to a broad range of decisions, including allocating budgets across hospital departments, evaluating health technology portfolios, and designing multi-year investment strategies. In addition, hybrid models that combine optimization with simulation have been used to explore different budget scenarios and assess their operational implications under uncertainty (Ordu et al., 2021).

A second main area where OR contributes is capacity planning, which focuses on determining the appropriate number of beds, staff, operating rooms, and equipment. Systematic reviews highlight extensive use of simulation, queuing theory, and optimization to match resources with demand (Humphreys et al., 2022; Rachuba et al., 2024). More recent work incorporates predictive modeling to anticipate future bottlenecks. For example, Russo et al. (2025) use predictive optimization to plan both operating-room capacity and hospital bed allocation.

At a broader network level, OR models help coordinate services across hospitals, balance patient loads, and design resilient care pathways, including during large-scale disruptions such as pandemics (Fuloria and Šormaz, 2025; Bravo et al., 2021).

Despite these contributions, most traditional OR models, whether aimed at budget allocation, capacity planning, or staff scheduling, assume a single, centralized decision-maker who controls all relevant resources and can implement the model's recommended solution. This assumption rarely holds in practice. Healthcare systems, including those in the Netherlands, involve a distributed set of actors: departments, hospital boards, insurers, staffing agencies, and healthcare professionals. Each has its own objectives, constraints, and incentives, and their interactions can create conflict behavior.

Consequently, centralized optimization models often produce solutions that may be mathematically optimal but practically infeasible, a limitation repeatedly highlighted in implementation studies (Carter and Busby, 2023). In workforce planning, this issue is particularly pronounced because nurses themselves are strategic agents with their own preferences, constraints, and outside options, such as switching to self-employed work (Renggli et al., 2025).

This mismatch between optimization assumptions and institutional reality points to an important methodological gap that this thesis aims to address: the need for multi-actor modeling approaches that acknowledge both technical efficiency and stakeholder behavior. This motivates integrating OR with game-theoretic methods, which is the focus of the next section.

2.2 Game Theory in Health system

Section 2.1 showed how Operations Research (OR) can model and solve complex healthcare planning problems, particularly when decisions are made by a single central authority. However, it also highlighted a key limitation: the “central planner” assumption. Most OR models assume one decision-maker with a unified objective function who can coordinate all resources. This is increasingly unrealistic in fragmented, multi-actor health systems such as the Dutch system, where decision-making is distributed across hospitals, departments, insurers, regulators, and professional groups that often pursue different goals.

To analyze such settings, a framework that explicitly accounts for strategic behavior is needed: situations in which each actor’s best choice depends on how others are expected to act. Game Theory (GT) provides a formal mathematical framework for analyzing and predicting the outcomes of strategic interactions among rational, interdependent actors (Myerson, 1991). First formalized by von Neumann and Morgenstern (1944), game theory is not concerned with optimization in a passive environment (as in classical OR) but with situations in which each actor’s optimal choice depends on others’ choices. GT has become central to economics and social sciences, and its relevance to healthcare has grown as systems have become more complex, competitive, and constrained (Alalawi et al., 2019).

In this thesis, game theory provides the conceptual basis for modeling nurse scheduling as a strategic interaction between hospital management and nurses, rather than as a purely centralized optimization problem.

2.2.1 Foundational Concepts

In game-theoretic terms, a “game” is a formal model of strategic interaction. As detailed in foundational texts (e.g., Fudenberg and Tirole (1991)), any game can be characterized by three core elements: players, strategies, and payoffs.

Players: Players are the rational decision-makers. They are assumed to have preferences over possible outcomes and to choose actions that (approximately) maximize their own expected utility. In this thesis, the primary players are:

- the Hospital Manager, representing the institution’s objectives (e.g., costs, quality, regulatory compliance); and
- Nurses represents the workforce’s collective response (e.g., acceptance of schedules, overtime, or alternative contracts).

In other healthcare applications, players might be physicians, patients, insurers, pharmaceutical firms, or even biological entities such as tumor cell populations Stanková et al. (2019).

Strategies: Strategies are the actions or complete plans of action available to each player. A strategy $s = (s_1, s_2, \dots, s_n)$ specifies one strategy for each of the n players.

- for the hospital manager: choosing a shift structure, setting overtime penalties, or deciding on the mix of fixed and flexible nurse contracts;
- for nurses: choosing whether to accept particular patients, how much overtime to work, or whether to switch to agency or self-employed contracts.

Payoffs: Payoffs quantify the consequences of each strategy profile from each player’s perspective. They are represented by utility functions $u_i(s)$, which assign a real-valued payoff to player (i) for each strategy (s).

- for the hospital manager: a payoff that decreases with labor costs and under-staffing, and increases with quality indicators and compliance;
- for nurses: a payoff that increases with income and schedule satisfaction, and decreases with fatigue and perceived unfairness.

To formally represent a game, two forms are commonly used. The Normal Form is typically used for simultaneous-move games, in which players choose their strategies at the same time (or without knowing the other player’s choice). It is often depicted as a payoff matrix. The Extensive Form is visualized as a game tree and is used for sequential-move games, where players make their choices in a defined order. This form explicitly models the timing of moves and the information a player has when making a choice.

The models used in this thesis (Nash and Stackelberg) are classic examples of non-cooperative games with complete information, where players act independently (non-cooperative) and know all the rules, strategies, and payoff functions of all other players (complete information).

2.2.2 Simultaneous Interaction: The Nash Equilibrium

The Nash Equilibrium (NE) is the standard solution concept for scenarios where players act simultaneously or independently, without knowing the other player’s specific choice. In this setting, each player optimizes their own payoff taking the actions of others as given.

Formally, proposed by Nash (1950, 1951), a strategy profile s^* is a Nash Equilibrium if no player has an incentive to unilaterally deviate. That is, for every player i :

$$u_i(s_i^*, s_{-i}^*) \geq u_i(s_i, s_{-i}^*)$$

for all available strategies s_i . In equilibrium, every player is playing a “best response” to the strategies of the others.

Pure vs. Mixed Strategies

A critical distinction in Nash Equilibrium analysis is between pure and mixed strategies.

- **Pure Strategy:** A deterministic choice where a player selects a single action with probability 1.
- **Mixed Strategy:** A probabilistic choice where a player randomizes over available actions (Osborne and Rubinstein, 1994).

While Nash (1950) proved that every finite game has at least one equilibrium in mixed strategies, a pure strategy equilibrium is not guaranteed to exist. In discrete optimization games like the Nurse Scheduling Game, the absence of a pure Nash Equilibrium often manifests as cyclic instability, where schedules oscillate because players constantly adjust to exploit each other’s moves (Carvalho et al., 2023). However, in the context of this research, this probabilistic nature constitutes a specific strength of the model. A mixed strategy represents a probability of assignment rather than a binary certainty, which can reflect the operational necessity of on-site decision-making. In healthcare environments characterized by high uncertainty, adhering to a rigid, pre-determined schedule is often infeasible. Instead, a mixed strategy captures the system’s need for flexibility, effectively modeling scenarios where staffing decisions are not fixed weeks in advance but are adapted “on the day” based on real-time patient demand. Thus, the mixed equilibrium serves as a proxy for the dynamic, contingent staffing buffers essential to modern hospital operations.

The most critical insight from the Nash Equilibrium is that an individually rational, stable outcome is often collectively inefficient. The Prisoner’s Dilemma is the classic example: mutual defection is the unique Nash Equilibrium, even though both players would be better off if both cooperated.

This logic is applicable directly to healthcare. Many problems are essentially social dilemmas. Colman et al. (2019) model antibiotic prescribing as a catastrophic social dilemma. Each physician’s best response is to prescribe to protect their individual patient and avoid diagnostic risk. However, if all physicians behave this way, the system converges to a Nash Equilibrium with excessive antibiotic use and accelerating antimicrobial resistance. Similarly, Wagner et al. (2020) analyze vaccination and antimicrobial use as public goods problems. The population benefits from high vaccination coverage and sensitive antimicrobial use, but individuals may be tempted to “free-ride” on others’ efforts. The resulting equilibrium features either too little vaccination or too much antimicrobial use relative to the socially optimal level.

These studies show how decentralized decisions, each sensible when viewed in isolation, can nonetheless produce system-wide inefficiencies. The same dynamic appears in the nurse labor market: when individual departments or hospitals compete for scarce staff, their independent choices can collectively lead to persistent overtime, rising burnout, and higher turnover, even though each organization is acting rationally within its own constraints.

Nash Equilibrium is also the main tool for analyzing competition and bargaining between institutions: Ho and Lee (2017) model the formation of equilibrium provider networks in health

insurance markets. Insurers and hospitals bargain over network inclusion and reimbursement terms. The resulting set of contracts is an equilibrium of these strategic interactions, and may involve exclusion of certain providers, even when this is not socially efficient. Jiang et al. (2020) analyze how hospitals compete for patients under performance-based payment schemes. They show that equilibrium quality and cost levels depend critically on how incentives are structured; some designs can unintentionally worsen performance despite being well-intentioned. Duan et al. (2021) studied pricing decisions under the Diagnosis-Related Groups (DRGs) system. Hospitals and payers simultaneously choose their strategies, adjusting intensity and pricing in response to each other’s choices. Equilibrium prices emerge from these mutual best responses and may diverge from the socially optimal arrangement. Gao and Wang (2019) extend these ideas to the healthcare supply chain, modeling hospitals, suppliers, and insurers as strategic players. They show that coordinated strategies, e.g., aligning insurer reimbursement rules with provider pricing, can improve system performance, but only if the game is designed to internalize downstream consequences.

Together, these examples illustrate how simultaneous-move games and Nash equilibria are widely used in healthcare to understand competition, bargaining, and resource allocation. For nurse scheduling, the same logic applies: departments or hospitals that simultaneously “bid” for flexible or agency nurses may reach equilibrium patterns that are stable but system-wide inefficient, such as over-reliance on costly external staff or systematically uneven workloads.

2.2.3 Hierarchical Interaction: The Stackelberg Equilibrium

While the Nash concept models uncoordinated or simultaneous decision-making, many healthcare processes are inherently hierarchical. Regulators set rules before providers act; managers set budgets before nurses choose shifts. To analyze this leader-follower dynamic, the Stackelberg Equilibrium is employed (von Stackelberg, 1934).

In searching for a simple Stackelberg Equilibrium with two players:

1. The Leader (Player 1) moves first, choosing and committing to a strategy s_1 .
2. The Follower (Player 2) observes the Leader’s choice s_1 and then chooses their own best response to that choice, s_2 .
3. Payoffs are realized for both players.

Crucially, the Leader makes their initial decision with full knowledge of how the Follower will rationally react to any choice the Leader makes. The Leader, therefore, anticipates the Follower’s response and chooses the initial action s_1 that will lead to the Follower’s best response, $s_2^*(s_1)$, that ultimately maximizes the Leader’s own payoff. The solution is found using backward induction (Selten, 1975).

Some of the most innovative healthcare applications of finding Stackelberg Equilibrium occur in oncology. Stanková et al. (2019) review game-theoretic approaches to cancer treatment,

conceptualizing clinicians and tumor cells as strategic players: clinicians choose treatment intensity and timing, while tumor cells “respond” through evolutionary adaptation. Building on this, Stein et al. (2023) develop Stackelberg evolutionary game theory, where the clinician is the leader selecting treatment policies, and the evolving tumor cell population is the follower. Salvioli et al. (2025) further investigate Stackelberg evolutionary games in settings where the tumor can be stabilized, showing that long-term, equilibrium-oriented strategies can outperform aggressive, short-sighted treatments.

Stackelberg structures also appear in institutional and technical domains. Gao and Wang (2019) model a healthcare supply chain in which an insurer, acting as the leader, sets reference prices and reimbursement rules. Hospitals and suppliers then choose their pricing and ordering decisions as followers. The insurer’s problem is to select policies that stimulate follower responses aligned with the system’s goals. Somasundaram and Sivakumar (2015) search a Stackelberg Equilibrium to security in Wireless Body Area Networks. The legitimate node (leader) chooses security protocols anticipating attacks by malicious nodes (followers), and equilibrium analysis identifies robust defensive strategies. These examples highlight how Stackelberg Equilibrium capture the logic of top-down decision-making with strategic responses.

In the context of this thesis, the hospital manager (Leader) might first set the nurse’s daily schedules, the budget for external staff, or the penalties for overtime. Nurses (Followers) then reacts by choosing its best-response strategy (e.g., deciding which shifts to take and whether to work for a patient) within those constraints. This hierarchical framework, which explicitly models strategic responses to a central policy, serves as one of the conceptual foundation for the models developed in this thesis.

2.3 Integer Programming Games

Integer Programming Games (IPGs) is a specialized subclass of Game Theory that bridges the gap between the two analytical perspectives introduced in the previous sections. While Operations Research (Section 2.1) provides the tools to solve complex constraints for a single agent, and classical Game Theory (Section 2.2) provides the logic to analyze strategic interaction between multiple agents, neither is sufficient on its own to model the Dutch healthcare context. Classical GT typically assumes players choose from simple, stylized actions, such as setting a high or low price, which fails to capture the highly structured, constraint-heavy rosters required in hospital operations.

IPGs address this limitation by embedding the optimization power of Mixed-Integer Programming inside the strategic framework of Game Theory. In an IPG, each player’s strategy is not a simple variable, but the optimal solution to a private integer program (e.g., a full nurse schedule), whose objective and constraints are endogenous to the actions of other players (Carvalho et al., 2023). This creates a natural synthesis: the strategic structure is derived from Game Theory, while the operational detail is handled by Mixed-Integer Programming. This hybrid approach provides the conceptual foundation for the Nurse Scheduling Game developed in this

thesis, allowing the analysis of how complex schedules emerge from the interplay of conflicting stakeholder objectives.

2.3.1 Conceptual Formulations

The specific formulation of an IPG depends on whether players act simultaneously or sequentially. Two forms are particularly relevant here: simultaneous-move (Nash) IPGs and sequential (Stackelberg) IPGs (Dragotto, 2022).

A. Nash IPGs Equilibrium

A simultaneous IPG equilibrium assumes that all players make their complex, optimization-based decisions simultaneously, or without knowing the others' decisions. The goal is to find a Nash Equilibrium in this high-dimensional, discrete strategy space. This problem is formally known as a Generalized Nash Equilibrium Problem (GNEP) with discrete strategies. In a GNEP, the feasible set (or constraint set) of each player i explicitly depends on the strategies, x_{-i} , of all other players.

Formally, a Nash IPG is a problem where each player $i \in N$ solves:

$$\begin{aligned} \max_{x_i} \quad & u_i(x_i, x_{-i}) \\ \text{subject to:} \quad & x_i \in X_i(x_{-i}) \end{aligned}$$

where $x_i \in \mathbb{Z}^{n_i}$ is the integer decision vector for player i , and $X_i(x_{-i}) = \{x_i \in \mathbb{Z}^{n_i} \mid g_i(x_i, x_{-i}) \leq 0\}$ represents the constraints of player i , which are a function of the other players' actions.

A solution $x^* = (x_1^*, \dots, x_n^*)$ is a Nash Equilibrium if, for every player i , x_i^* is the optimal solution to the above problem, given that all other players play x_{-i}^* . This is a fixed-point problem of mutual best responses. The computational complexity of this problem is severe. As shown by Carvalho et al. (2018, 2022), just deciding if a Nash Equilibrium exists in an IPG is a Σ_p^2 -complete problem, meaning it is even harder than standard NP-complete problems.

Despite this complexity, the IPG framework has proven to be a robust modeling tool for specific operational domains where strategic interaction is the defining feature. Carvalho et al. (2017) used an IPG to study kidney exchange between hospitals. They showed that, with strategic players, the equilibrium outcome in their model maximizes social welfare. Cho and Sharkey (2023) model a freelancer scheduling platform as an IPG, designing an integer program whose feasible solutions correspond exactly to the platform's Nash equilibria. Their formulation enables the computation of the price of anarchy and the assessment of how decentralized decisions affect system performance. In the cybersecurity domain, Dragotto et al. (2024) treat attackers and defenders as players who each solve combinatorial optimization problems; the resulting "critical node game" shows how equilibria can reveal vulnerabilities in cloud networks.

These applications highlight why the simultaneous IPG framework is conceptually relevant to the Dutch nursing labor market. When hospitals act independently and compete for flexible staffing resources, each solves its own optimization problem, but their decisions interact through shared constraints (i.e., the limited pool of available nurses). The setting is therefore highly susceptible to suboptimal but stable equilibrium outcomes.

B. Sequential IPGs (Stackelberg Equilibrium)

The sequential-move variant, corresponding to the Stackelberg Equilibrium, is formally known as a Bilevel Optimization Problem or, more specifically, a Bi-level Integer Program (BIP). This structure perfectly models the hierarchical reality of management, where one agent has the authority to set the rules to which others must react.

A BIP consists of two nested optimization problems: an upper-level (Leader) problem and a lower-level (Follower) problem. The Leader moves first, making a decision x_L (e.g., setting a budget). The Follower observes this decision and solves their own optimization problem to find the best response, x_F (e.g., creating a schedule). The Leader, knowing the Follower will react optimally, chooses x_L to maximize their own utility *via* the Follower's reaction.

Formally, a Bilevel Integer Program is defined as:

$$\begin{array}{ll}
 \underset{x_L}{\text{maximize}} & U_L(x_L, x_F) \quad (\text{Leader's Objective}) \\
 \text{subject to} & (x_L, x_F) \in \mathcal{X} \quad (\text{Joint Constraints}) \\
 & x_L \in \mathbb{Z}^{n_L}, \quad x_F \in \mathbb{Z}^{n_F} \\
 & x_F \in \Psi(x_L) \quad (\text{Follower's Reaction})
 \end{array}$$

Where $\Psi(x_L)$ represents the set of optimal solutions to the Follower's problem, defined as:

$$\Psi(x_L) = \arg \max_y \{ U_F(x_L, y) : (x_L, y) \in \mathcal{Y}, y \in \mathbb{Z}^{n_F} \}$$

Here, \mathcal{X} represents the constraints that must be satisfied by the Leader (including any coupling constraints), and \mathcal{Y} represents the feasible set for the Follower given the Leader's choice. The constraint $x_F \in \Psi(x_L)$ explicitly enforces that the Leader can only choose a solution that is also an optimal response for the Follower. This nested structure captures the strategic anticipation essential to the model: the Manager does not simply optimize the schedule directly but optimizes the *rules* that induce the Nurses to choose the desired schedule.

This structure reflects many real decision environments. Köppe et al. (2011) analyzed the

structure of the follower’s best responses using rational generating functions, showing that, because of integrality, the response set is often discontinuous and highly non-convex. This makes solving the leader’s problem extremely difficult. Indeed, as later shown by Carvalho et al. (2024) in their “Nash-meets-Stackelberg” models, even deciding whether a solution exists is Σ_p^2 -hard.

In practice, Stackelberg IPGs enable us to model top-down decision-making with strategic anticipation. In energy markets, for example, regulators act as leaders who set policies affecting profit-maximizing producers (Carvalho et al., 2024). When applied to this thesis, the Hospital Manager acts as the Leader, choosing budget allocations or overtime penalties, while Nurses solves a scheduling IP as the Follower.

2.3.2 Solution Approaches

Methods for solving IPGs differ depending on whether the game is simultaneous or hierarchical. Because both forms involve integer programs, computing equilibria is extremely challenging, and recent research has focused on making these models tractable.

A. Nash IPGs Equilibrium

Given the Σ_p^2 -complete complexity, finding an exact Nash Equilibrium is a major challenge. The research has proceeded along several lines. The most intuitive approach is an iterative best-response algorithm, where players take turns solving their own IP. However, with discrete (integer) strategies, this method is not guaranteed to converge and may enter a cycle. To address this, a major breakthrough has been the development of general algorithmic frameworks. Carvalho et al. (2022) developed a general algorithm that is guaranteed to return a Nash Equilibrium for finite IPGs. This was operationalized and improved in the “ZERO Regrets Algorithm” (Dragotto and Scatamacchia, 2021), a sophisticated cutting plane algorithm that leverages an “equilibrium separation oracle” to compute, enumerate, and select specific Nash Equilibria. For specific problem structures, a single-level reformulation is possible, as Cho and Sharkey (2023) successfully demonstrate with their scheduling game, where a feasible solution to their master IP corresponds directly to a Nash Equilibrium. Finally, because exact, global NE is so hard to find, a more recent and practical research direction is to seek approximate solutions. Duguet et al. (2025) are developing methods to compute approximate Nash equilibria, while Koirala et al. (2025) are exploring “locally optimal solutions,” which may be more realistic representations of how human players “settle” on a solution.

B. Stackelberg IPGs Equilibrium

There are two main classes of approaches to solve Bi-level Integer Programs (BIPs): exact methods and heuristic(stochastic) methods.

Exact Solution Methods aim to find the provably optimal solution. The most common approach is a single-level reformulation, which replaces the lower-level (Follower) optimization problem

with its optimality conditions. For integer problems, this is highly complex because the KKT conditions do not apply. It requires replacing the Follower’s MIP with its primal-dual optimality conditions and a set of complementary slackness constraints, which are then linearized using “big-M” reformulations. This results in a single, large, but computationally very difficult MIP (Li and Guo, 2017). An alternative is a bi-level-specific branch-and-bound algorithm, which “branches” on the Leader’s integer variables and solves the Follower’s IP exactly at each node (Moore and Bard, 1990).

Heuristic and Stochastic Solution Methods have gained focus, given the extreme difficulty of exact methods (Sinha et al., 2018). The core problem in finding the Stackelberg Equilibrium is that the Leader needs to know the Follower’s exact, deterministic best-response function. In a real-world problem like nurse scheduling, this is impossible. The Hospital Manager does not know the exact, deterministic preferences of all nurses; Nurses’ “best response” is inherently stochastic (uncertain) and can only be estimated.

The Monte Carlo Multilevel Optimization (MCMO) framework, a form of Simulation-Based Bi-level Optimization (Koirala and Laine, 2023), embraces this uncertainty. The algorithm begins with the Leader’s outer-loop algorithm proposing a candidate policy x_L (e.g., a set of budgets and penalties). To evaluate the “goodness” of this policy, the algorithm runs a Monte Carlo simulation. It simulates the Follower’s (Nurses) response N times, where each simulation includes random draws from uncertain parameters, such as individual nurse preferences or the probability of sick leave. This process produces a distribution of N possible outcomes: $\{x_F^1, \dots, x_F^N\}$. The Leader’s payoff for policy x_L is then calculated as the expected value over this simulated distribution, $\mathbb{E}[U_L(x_L, x_F)]$. This expected payoff is returned to the outer-loop algorithm, which uses it to generate a new, potentially better policy x'_L .

2.4 Synthesis and Research Gap

The systematic review of the literature across Operations Research (OR), Game Theory (GT), and their methodological intersection in IPGs highlights a critical disconnect between the analytical tools available and the reality of the healthcare workforce crisis. The review reveals three key observations.

First, the literature on Operations Research in healthcare demonstrates that Mixed-Integer Programming offers powerful tools for optimizing nurse scheduling, capacity planning, and resource allocation (Mystakidis et al., 2024; Rais and Viana, 2011). These models capture contractual rules, shift requirements, fairness constraints, and labor regulations with high operational detail. However, they nearly always assume a centralized planner, in which a single decision-maker sets the entire schedule. This assumption rarely holds in real healthcare systems, particularly in countries like the Netherlands, where planning authority is distributed across departments, managers, and external staffing agencies (Mitton et al., 2014; Dakin and Tsiachristas, 2024). As a result, classical MIP models can optimize schedules but cannot represent the strategic conflicts inherent in multi-actor environments.

Second, the literature on Game Theory (GT) provides the necessary conceptual framework for strategic analysis. Applications to antibiotic stewardship, vaccination, and hospital competition (Colman et al., 2019; Wagner et al., 2020; Jiang et al., 2020) show that GT can explain why rational local decisions produce system-level inefficiencies. However, most GT models rely on abstract or simplified strategies, binary choices, or price negotiations rather than the combinatorial decisions involved in complex operational planning, such as nurse scheduling (Duan et al., 2021). Thus, game theory captures incentives but not the required operational complexity.

Third, IPGs offer a promising bridge, representing each player’s strategy as the optimal solution of a complex optimization problem (Carvalho et al., 2022). This field explicitly models strategic interaction under operational constraints, with applications ranging from kidney exchange (Carvalho et al., 2017) to cybersecurity (Dragotto et al., 2024). Despite their methodological suitability, IPGs have seen almost no application in healthcare workforce planning, even though this domain features exactly the strategic interactions and combinatorial complexity that IPGs are designed to capture.

These observations point to a critical research gap at the intersection of methodology and application: Although nurse scheduling is well-studied in OR, and strategic interaction is well-studied in GT, no existing research integrates these perspectives through IPGs for healthcare workforce planning or budget allocation. The absence of IPGs in this domain represents more than just a gap in the academic literature; it represents a gap in realism and predictive accuracy. This methodology is essential because, in a fully decentralized system like the Dutch healthcare sector, a schedule is effectively not an instruction but a negotiation. Traditional centralized models predict an ideal outcome that is inherently fragile, vanishing the moment a nurse exercises their autonomy to reject a shift. In contrast, an IPG model predicts an equilibrium outcome, which is a stable schedule that persists even when all agents act in their own self-interest. Therefore, adopting an IPG framework provides a more accurate solution not because it finds a mathematically “better” optimum, but because it respects the decentralized governance structure of the system. It captures the realism of on-site decision-making, where managers must incentivize attendance rather than command it.

This thesis addresses this gap by developing, implementing, and comparing both Nash and Stackelberg IPG equilibria for nurse workforce scheduling. Methodologically, the thesis extends the IPG literature into a new, critical domain: healthcare labor planning. In practice, it offers a decision-support framework that helps hospital administrators evaluate how different policy levers, such as overtime penalties, staffing budgets, or the contract mix, shape system outcomes. By explicitly modeling nurse behavioral responses as strategic reactions, the framework reveals how policies work in real organizational settings. It provides a novel tool for addressing workforce allocation challenges in complex, multi-stakeholder healthcare systems.

3 Model Formulation and Methodology

This chapter develops the mathematical and game-theoretic framework that is applied in this thesis. It starts by formulating the nurse scheduling problem from two distinct perspectives: the hospital manager’s and nurses’. Each perspective is represented by a Mixed-Integer Programming (MIP) model that reflects the operational realities, constraints, and objectives of the respective decision-maker.

Building on these formulations, the Nurse Scheduling Game (NSG) is constructed as a strategic interaction in which the manager and nurses optimize their own MIP-based decisions while anticipating or responding to the other’s decisions. The chapter formally defines the game and examines two equilibrium concepts central to strategic analysis in decentralized settings: the Nash Equilibrium, in which both players act simultaneously, and the Stackelberg Equilibrium, in which the manager acts as a leader and nurses as a follower.

The chapter concludes by outlining the computational strategies used to solve the NSG under each equilibrium concept. For the Nash setting, an iterative best-response algorithm is employed. In the Stackelberg setting, a Monte Carlo multilevel optimization approach is developed to approximate the follower’s best response. Together, these methods form the methodological foundation for the analyses presented in later chapters.

3.1 Nurse Scheduling Problem (NSP)

The nurse scheduling problem (NSP) forms the basis for the game-theoretic models developed in this thesis. In its classical form, the NSP determines how to assign available nurses to required shifts while satisfying labor regulations, meeting minimum staffing requirements, and minimizing operational costs. In hospital environments, especially those facing workforce shortages, effective scheduling is essential for ensuring patient safety, maintaining service continuity, and controlling labor expenses. The NSP therefore serves as both a practical planning tool and a conceptual basis for strategic interaction when multiple stakeholders have conflicting objectives.

In this thesis, the NSP is tailored to a healthcare setting where two types of clinical activities, inpatient care and outpatient care, must be staffed daily. Nurses can perform either role on a given day, and their productivity differs depending on the assignment. Patients arrive over time and require exactly one treatment session, which is delivered by a nurse of the appropriate type. Waiting times generate penalty costs, reflecting both patient dissatisfaction and potential deterioration of health outcomes. The hospital manager determines the staffing plan, while nurses decide how patients are assigned to available capacity. This decomposition naturally yields two interdependent MIP formulations: one for the manager’s planning problem and the other for the nurses’ operational response.

To model this, the following sets, parameters, and decision variables are defined:

Sets

- $p \in P$ $P \in \{1, \dots, N_{Patients}\}$ Set of patients
 $n \in N$ $N \in \{1, \dots, N_{Nurses}\}$ Set of nurses
 $d \in D$ $D \in \{1, \dots, N_{Days}\}$ Set of planning days

Parameters

- Δd_d Auxiliary variable for days count, $\Delta d_d = d$
 R_d^I Required number of inpatient nurses in day d
 R_d^O Required number of outpatient nurses in day d
 T_n^I Treatment penalty for one patient when nurse n work as an inpatient nurse
 T_n^O Treatment penalty for one patient when nurse n work as an outpatient nurse
 G_p Group of patient p , $G_p = 1$ if patient p is an outpatient
 $A_{p,d}$ Arriving time of patient p , $A_{p,d} = 1$ if patient p arrived in day d
 Q^I Maximum number of patients an inpatient nurse can serve in a day
 Q^O Maximum number of patients an outpatient nurse can serve in a day
 C^I Daily salary for each inpatients nurse
 C^O Daily salary for each outpatients nurse
 C^W Waiting day costs for each patient

Decision Variables

- $I_{n,d} \in \mathbb{B}$ $I_{n,d} = 1$ if nurse n works as inpatient nurse in day d
 $O_{n,d} \in \mathbb{B}$ $O_{n,d} = 1$ if nurse n works as outpatient nurse in day d
 $X_{p,n,d} \in \mathbb{B}$ $X_{p,n,d} = 1$ if patient p is treated by nurse n in day d
 $W_p \in \mathbb{Z}$ Waiting days for patient p

3.1.1 NSP Manager Problem

The Manager's Problem (**M**) represents the hospital's strategic planning layer. It describes the decision-making process of the hospital administrator responsible for establishing workforce capacity for the planning horizon. The Manager determines the schedule, specifying which nurses work on which days and assigning them to either the Inpatient or Outpatient department. The Manager's primary objective is financial efficiency; they seek to minimize the total operational costs, which include both the direct payroll expenses of the workforce and the penalty costs associated with patient waiting times.

Model Formulation (M)

$$\min \quad TC = \sum_{d \in D} \sum_{n \in N} C^I \cdot I_{n,d} + \sum_{d \in D} \sum_{n \in N} C^O \cdot O_{n,d} + \sum_{p \in P} C^W \cdot W_p \quad (1)$$

$$\text{s.t.} \quad \sum_{n \in N} I_{n,d} \geq R_d^I \quad \forall d \in D \quad (2)$$

$$\sum_{n \in N} O_{n,d} \geq R_d^O \quad \forall d \in D \quad (3)$$

$$I_{n,d} + O_{n,d} \leq 1 \quad \forall n \in N, \forall d \in D \quad (4)$$

$$\sum_{k=d}^{d+3} (I_{n,k} + O_{n,k}) \leq 4 \quad \forall n \in N, \forall d \in D \setminus \{D_{\max}, \dots, D_{\max}-6\} \quad (5)$$

$$I_{n,d}, O_{n,d} \in \{0, 1\}, \quad W_p \in \mathbb{Z}^+ \quad (6)$$

Objective Function The objective function (1) serves as a proxy for the hospital's budget constraints and service level mandates. The first two terms, $\sum C^I \cdot I_{n,d}$ and $\sum C^O \cdot O_{n,d}$, calculate the total direct labor costs for inpatient and outpatient nurses, respectively. Minimizing these terms reflects the pressure to reduce overhead. The third term, $\sum C^W \cdot W_p$, introduces a penalty cost for every day a patient waits for treatment (W_p). This ensures that while the Manager attempts to minimize staffing levels, they cannot reduce capacity to the point where patient delays become financially restrictive.

Constraints The constraints define the feasible scheduling space based on demand and labor regulations. Constraints (2) and (3) enforce the minimum coverage requirements, ensuring that the supply of nurses meets the daily demand parameters (R_d^I, R_d^O) required for safe department operations. Constraint (4) enforces logical exclusivity, preventing a nurse from being assigned to both inpatient and outpatient roles on the same day. Finally, constraint (5) incorporates labor regulations to prevent burnout; it restricts the schedule so that no nurse works consecutive streaks longer than 4 days within a week.

3.1.2 NSP Nurses Problem

The Nurses' Problem (N) describes the operational execution of care on the hospital floor. It represents the collective decision-making of the nurses (or the floor managers acting on their behalf) regarding patient assignment. Once the schedule is fixed, the nurses decide the specific matching of patients to available nurses ($X_{p,n,d}$). Their objective is distinct from the manager's financial goals; instead, they seek to maximize operational efficiency. By minimizing the penalty associated with total treatment time, nurses seek to optimize operational efficiency, thereby reducing their collective workload and physical exhaustion.

Model Formulation (N)

$$\text{Minimize } TT = \sum_{d \in D} \sum_{n \in N} \sum_{p \in P} X_{p,n,d} \cdot (I_{n,d} \cdot T_n^I + O_{n,d} \cdot T_n^O) \quad (7)$$

$$\text{s.t. } \sum_{d \in D} \sum_{n \in N} X_{p,n,d} = 1 \quad \forall p \in P \quad (8)$$

$$\sum_{d \in D} \sum_{n \in N} \sum_{p \in P} X_{p,n,d} \cdot (1 - G_p) \leq \sum_{d \in D} \sum_{n \in N} Q^I \cdot I_{n,d} \quad (9)$$

$$\sum_{d \in D} \sum_{n \in N} \sum_{p \in P} X_{p,n,d} \cdot G_p \leq \sum_{d \in D} \sum_{n \in N} Q^O \cdot O_{n,d} \quad (10)$$

$$X_{p,n,d} \cdot (1 - G_p) \leq I_{n,d} \quad \forall n \in N, \forall d \in D, \forall p \in P \quad (11)$$

$$X_{p,n,d} \cdot G_p \leq O_{n,d} \quad \forall n \in N, \forall d \in D, \forall p \in P \quad (12)$$

$$\sum_{p \in P} X_{p,n,d} \leq Q^I \cdot I_{n,d} + Q^O \cdot O_{n,d} \quad \forall n \in N, \forall d \in D \quad (13)$$

$$\sum_{d \in D} \sum_{n \in N} \Delta d_d \cdot X_{p,n,d} - \sum_{d' \in D} \Delta d_{d'} \cdot A_{p,d'} = W_p \geq 0 \quad \forall p \in P \quad (14)$$

$$X_{p,n,d} \in \{0, 1\} \quad (15)$$

Objective Function The objective function (7) minimizes the Total Treatment Time (TT), which aggregates the service duration for every assigned patient-nurse pair based on the nurse's specific role efficiency (T_n^I or T_n^O). By minimizing the penalty associated with this total time, nurses seek to optimize operational efficiency; given the fixed requirement to treat all patients, this strategy acts as a proxy for minimizing collective workload and mitigating physical fatigue.

Constraints These constraints enforce the logical and operational validity of patient care within the capacity boundaries established by the Manager. Constraint (8) establishes a universal care mandate, guaranteeing that every patient is assigned exactly once. Constraints (9) and (10) serve as aggregate capacity checks, ensuring that the total volume of inpatient and outpatient assignments does not exceed the collective capacity (Q) of the scheduled workforce. Constraints (11) and (12) enforce strict role-to-patient matching, avoiding an inpatient nurse from treating an outpatient, and vice versa. To prevent individual burnout, Constraint (13) imposes a cap on each nurse's daily patient load. Finally, Constraint (14) establishes the temporal logic of the schedule by calculating patient waiting days (W_p) and strictly enforcing that treatment cannot occur before a patient's arrival.

3.2 Nurse Scheduling Game (NSG)

After defining the individual optimization problems for the Hospital Manager and Nurses in Section 3.1, they are now integrated into a unified framework: the Nurse Scheduling Game (NSG). The NSG models the interaction between the two stakeholders not as a static optimization problem, but as a non-cooperative game in which each player's outcome depends parametrically on the other player's decisions. The game is formally defined by the tuple $\mathcal{G} = \langle \mathcal{P}, \mathcal{S}, \mathcal{U} \rangle$:

Players (\mathcal{P})

- The Manager (M): Represents the hospital administration and minimizes financial costs.
- The Nurses (N): Represent the collective nurse workforce and minimize operational effort.

Strategies (\mathcal{S}) A strategy corresponds to a complete set of decision variables chosen by a player.

- Manager's Strategy (s_M): A vector of binary decisions determining the schedule.

$$s_M = (I_{1,1}, \dots, I_{N_{Nurses}, N_{Days}}, O_{1,1}, \dots, O_{N_{Nurses}, N_{Days}})$$

Feasible Space (Ω_M): The set of feasible strategies for the Manager, Ω_M , is defined by the operational limitations formulated in Section 3.1.1. A strategy s_M belongs to Ω_M if and only if it satisfies Constraints (2) through (6). This means the schedule must strictly meet the minimum daily headcount requirements (R^I, R^O) and adhere to labor laws regarding maximum consecutive working days.

- Nurses' Strategy (s_N): A vector of binary decisions determining patient assignment.

$$s_N = (X_{1,1,1}, \dots, X_{N_{Patients}, N_{Nurses}, N_{Days}})$$

Feasible Space ($\Omega_N(s_M)$): The set of feasible strategies for the Nurses is defined by the operational capacity constraints formulated in Section 3.1.2. A strategy s_N belongs to $\Omega_N(s_M)$ if and only if it satisfies Constraints (8) through (15). Note that Ω_N is a function of s_M ; the Manager's schedule sets the schedule limits for the Nurses.

Payoffs (\mathcal{U}) The payoff functions represent the disutility (cost or time) each player seeks to minimize. These are defined using the Objective Functions formulated in Section 3.1.

- Manager's Payoff (U_M): Corresponds to the Total Cost function (1). Note that while the Manager controls I and O , the Waiting Time (W) is a result of the Nurses' assignment X . Therefore, the Manager's payoff depends on the Nurses' strategy.

$$U_M(s_M, s_N) = \underbrace{\sum_{d,n} (C^I I_{n,d} + C^O O_{n,d})}_{\text{Salary Cost (Controlled by M)}} + \underbrace{\sum_p C^W W_p(s_N)}_{\text{Wait Penalty (Result of N)}}$$

- Nurses' Payoff (U_N): Corresponds to the Total Treatment Time function (7).

$$U_N(s_M, s_N) = \sum_{d,n,p} X_{p,n,d} \cdot (I_{n,d} T_n^I + O_{n,d} T_n^O)$$

This creates a fully coupled system: the Manager's cost depends on the Nurses' efficiency (via waiting times), and the Nurses' efficiency depends on the Manager's resource allocation. To

formalize the interaction dynamics, two decision-making processes are considered. The difference in information availability and timing defines the equilibrium concept used in this thesis.

- Simultaneous Interaction (Nash): Both players act at the same time. The Manager sets the schedule based on static demand expectations, while the Nurses simultaneously plan patient flows. Neither party observes the other's final decision before committing to their own. This models a disjointed, decentralized planning process.
- Sequential Interaction (Stackelberg): The interaction occurs in stages. At the beginning, the Manager (Leader) commits to a schedule s_M . Then, the Nurses (Followers) observe s_M , which is applied to constraints for their optimization problem. They then execute the optimal patient assignment s_N . Finally, costs are realized. This models a realistic top-down governance structure.

3.2.1 NSG Nash Equilibrium

In the simultaneous interaction, a Nash Equilibrium (NE) seeks: a strategy profile (s_M^*, s_N^*) where no player can reduce their cost by unilaterally changing their strategy.

It is critical to note that this game constitutes a Generalized Nash Equilibrium Problem (GNEP) rather than a standard Nash Equilibrium. In a standard game, a player's payoff depends on others, but their feasible strategy space is fixed. In the NSG, the Nurses' feasible space $\Omega_N(s_M)$ is explicitly constrained by the Manager's decision. If the Manager sets $I_{n,d} = 0$, the Nurse cannot legally assign a patient to that nurse (Constraint 11).

Mathematically, the pair (s_M^*, s_N^*) is a Generalized Nash Equilibrium if:

1. Manager Optimality: s_M^* is the best response to the nurses' strategy s_N^* .

$$s_M^* \in \arg \min_{s_M \in \Omega_M} U_M(s_M, s_N^*)$$

2. Nurses Optimality: s_N^* is the best response to the manager's strategy s_M^* .

$$s_N^* \in \arg \min_{s_N \in \Omega_N(s_M^*)} U_N(s_M^*, s_N)$$

This framework enables the analysis of coordination failure: if the Manager sets capacity too low, the Nurses are mathematically forced to delay treatments to maintain a feasible schedule. These delays cause an increase in waiting penalties (W_p), which ultimately drives up the total system cost.

3.2.2 NSG Stackelberg Equilibrium

The Stackelberg Equilibrium assumes a hierarchical interaction where the Manager acts as the Leader and the Nurses act as the Followers. This is formulated as a Bilevel Integer Program (BIP).

In this setting, the Manager chooses s_M first, anticipating that the Nurses will optimally minimize their own objective U_N in response. To formalize this, we define the Reaction Set of the Nurses, $\Psi(s_M)$, as the set of optimal solutions to the Nurses' problem for a given schedule s_M :

$$\Psi(s_M) = \arg \min_{s_N \in \Omega_N(s_M)} U_N(s_M, s_N)$$

The Manager's problem is then to select the strategy s_M that minimizes U_M , subject to the constraint that s_N must be drawn from this reaction set $\Psi(s_M)$. Assuming the optimistic bilevel formulation (where, in the case of indifference, the follower selects the solution most favorable to the leader), the Stackelberg Equilibrium is found by solving:

$$\begin{aligned} \min_{s_M} \quad & U_M(s_M, s_N) \\ \text{s.t.} \quad & s_M \in \Omega_M \\ & s_N \in \Psi(s_M) \end{aligned}$$

By solving this nested optimization problem, the Manager internalizes the Nurses' behavioral response. For instance, the Manager might optimally choose to increase salary costs (by scheduling more staff than strictly required) if they anticipate that this will allow the Nurses to clear the patient backlog faster, thereby reducing the Waiting Penalty term in U_M sufficiently to lower the Total Cost.

3.2.3 Pure VS. Mixed Equilibria

The formulations presented in Sections 3.2.1 and 3.2.2 describe the Nurse Scheduling Game under the assumption of Pure Equilibria. In this context, all decision variables are strictly integer-constrained ($I_{n,d}, O_{n,d}, X_{p,n,d} \in \{0, 1\}$). This corresponds to a rigid, deterministic planning environment where a nurse is either assigned to a shift or not, and a patient is assigned to a specific nurse or not.

However, decentralized healthcare environments are often characterized by uncertainty and the need for operational flexibility. To capture this, this thesis also evaluates the game under Mixed Equilibria. In the mixed formulation, the domains of the decision variables are relaxed to allow for probabilistic or fractional assignments:

- **Manager's Variables:** The binary constraints on the staffing variables ($I_{n,d}, O_{n,d}$) are relaxed to continuous linear variables in the range $[0, 1]$. This value represents the probability of scheduling a nurse on a given day, allowing the Manager to distribute budget and capacity more fluidly across the planning horizon.
- **Nurses' Variables:** The reaction variables ($X_{p,n,d}$) are relaxed from binary indicators to a discrete set of probability intervals. Specifically, the variable domain is defined as $X_{p,n,d} \in \{0, 0.1, \dots, 1.0\}$. This discretization models the likelihood of a nurse accepting

a specific patient assignment, balancing the need for flexibility with the computational necessity of maintaining a finite search space.

Despite these adjustments to the variable domains, the underlying mathematical structure of each individual problem remains a Mixed-Integer Linear Program (MILP). The solution strategy is therefore consistent across both pure and mixed scenarios.

In summary, this subsection establishes a structural decomposition of the classical Nurse Scheduling Problem. Rather than treating workforce planning as a monolithic optimization task performed by a single entity, the problem is modeled as a decoupled process involving two distinct decision agents. The Manager’s Problem (**M**) operates at the strategic level, prioritizing financial efficiency (TC) through capacity planning (I, O), while the Nurses’ Problem (**N**) operates at the operational level, prioritizing workload efficiency (TT) through patient assignment (X).

Crucially, while formulated as separate Mixed-Integer Programs, these agents are operationally interdependent. The schedule limits defined by the Manager in constraints (2) and (3) function as the binding hard constraints for the Nurses in (9) and (10). This mathematical coupling, where the output of one optimization problem becomes the constraint of another, combined with the inherent conflict between cost minimization and workload reduction, creates the precise conditions required for strategic competition. This structure forms the basis for the Nurse Scheduling Game (NSG) defined in the following section.

3.3 Solution Approach

As discussed in Section 2.3, IPGs are difficult to solve. Specifically, determining the existence of a Nash Equilibrium in an IPG is a Σ_p^2 -complete decision problem, making standard analytical methods intractable for large-scale, real-world instances (Carvalho et al., 2022). Because the decision variables are discrete and the coupled constraints create a nonconvex feasible region, the commonly used gradient-based methods in continuous game theory cannot be easily applied. Consequently, finding equilibrium solutions, whether Nash or Stackelberg, requires advanced algorithmic approaches that decompose the game into solvable subproblems.

This section presents the computational methodologies used in this thesis to identify the equilibrium for the NSG. Section 3.3.1 introduces the solution method for the Nash formulation, based on an Iterative Best-Response (IBR) algorithm, a widely studied dynamic process for approximating equilibrium in games in which each player solves a separate optimization problem. Section 3.3.2 then develops the solution approach for the Stackelberg Equilibrium, based on a Monte Carlo Multilevel Optimization (MCMO) framework inspired by recent advances in stochastic hierarchical optimization (Koirala and Laine, 2023). Together, these two subsections provide the computational fundament for analyzing strategic interactions in nurse scheduling.

3.3.1 Nash Equilibrium: Iterative Best-Response Method

To solve the simultaneous game defined in Section 3.2.1, an Iterative Best-Response (IBR) algorithm is employed. This approach creates a dynamic process in which players sequentially adapt their strategies to minimize their own payoffs based on their opponents' current actions. The concept of Best-Response Dynamics is foundational in algorithmic game theory. It relies on the premise that a Nash Equilibrium is a fixed point of the best-response correspondence. If the dynamics converge to a state where no player changes their strategy, that state is, by definition, a Nash Equilibrium.

Literature suggests that IBR is particularly effective in games with specific structural properties. For finite games, Kukushkin (2004) demonstrates that best-response dynamics converge under additive aggregation or specific potential functions. While general IPGs do not strictly guarantee convergence, cycles are possible. Heinrich et al. (2023) analyze random games and find that best-response dynamics frequently converge to equilibrium or efficiently restrict the search space to a small set of strategies (a limit cycle). Even in large-scale settings with myopic agents, such dynamics have been shown to be a robust learning mechanism for finding stable operating points Swenson (2018).

In the context of the NSG, the IBR method decomposes the complex game into two tractable Mixed-Integer Programming (MIP) subproblems.

1. The **Manager** solves their cost-minimization MIP, treating the Nurses' previous patient assignments as fixed parameters to estimate waiting penalties and schedules.
2. The **Nurses** solve their time-minimization MIP, treating the Manager's new schedule decision as fixed hard constraints.

The process iterates, with each player updating their schedule with their best response to the other. The cycle continues until the system stabilizes, meaning neither the Manager nor the Nurses can further reduce their cost or workload by changing their strategy.

Convergence is reached when the solution stabilizes across consecutive iterations. Specifically, the algorithm terminates if:

1. Equilibrium: The decision variables and objective values for both players remain unchanged between iteration k and $k - 1$ (i.e., $\Delta Payoff = 0$ and $\Delta Decision = 0$).
2. Maximum iteration: A pre-defined maximum iteration K_{max} count is reached.

The implementation follows a sequential logic. The algorithm initializes with an arbitrary feasible patient assignment from the Nurses. In the first step of the loop, the Manager's MIP is solved to generate an optimal schedule (s_M) based on the initial patient data. This schedule s_M is then passed as a parameter set to the Nurses' MIP. The Nurses' problem is solved to generate

a new patient assignment (s_N) that minimizes treatment time within the schedule limits of s_M . This new assignment s_N is fed back into the Manager’s problem for the next iteration. At the end of each full cycle (after the Nurses’ move), the algorithm checks for convergence by comparing the current strategies and payoffs to those of the previous iteration. If they are identical, the loop breaks, and the current profile is returned as the Nash Equilibrium. Pseudo-code for the algorithm is shown as follows:

Algorithm 1 Iterative Best-Response for Nash Equilibrium

Require: Parameters C^I, C^O, C^W , Demands R , Capacities Q

Require: Max iterations K_{max}

```

1: Initialize:  $k \leftarrow 0$ 
2: Initialize: Nurse strategy  $s_N^{(0)}$  (e.g., random feasible assignment)
3: Initialize: Payoffs  $U_M^{(0)} \leftarrow \infty, U_N^{(0)} \leftarrow \infty$ 
4: Converged  $\leftarrow$  false
5: while ( $k < K_{max}$ ) and (not Converged) do
6:    $k \leftarrow k + 1$ 
7:   // Step 1: Manager optimizes given Nurses' previous strategy
8:    $s_M^{(k)} \leftarrow \arg \min_{s_M \in \Omega_M} U_M(s_M, s_N^{(k-1)})$ 
9:   Calculate  $U_M^{(k)}$ 
10:  // Step 2: Nurses optimize given Manager's new strategy
11:   $s_N^{(k)} \leftarrow \arg \min_{s_N \in \Omega_N(s_M^{(k)})} U_N(s_M^{(k)}, s_N)$ 
12:  Calculate  $U_N^{(k)}$ 
13:  // Step 3: Check Convergence
14:  if ( $s_M^{(k)} == s_M^{(k-1)}$ ) and ( $s_N^{(k)} == s_N^{(k-1)}$ ) then
15:    Converged  $\leftarrow$  true
16:  end if
17: end while
    return Nash Equilibrium strategies  $(s_M^{(k)}, s_N^{(k)})$  and Payoffs  $(U_M^{(k)}, U_N^{(k)})$ 

```

3.3.2 Stackelberg Equilibrium: Monte Carlo Multilevel Optimization (MCMO)

The integration of Mixed-Integer Programming (MIP) and game theory in this thesis leads to highly complex optimization problems, especially when seeking for the hierarchical Stackelberg Equilibrium, where the Hospital Manager serves as the Leader and Nurses as the Follower. As noted in the literature review, traditional bilevel solution techniques, such as reformulating the lower-level problem using KKT conditions or applying specific branch-and-bound decomposition, become impractical in this context. The integer nature of the lower-level problem invalidates the use of KKT conditions, while the scale of the variables leads to exponential growth in the number of constraints in standard reformulation techniques, causing them computationally intractable.

To address this, Monte Carlo Multilevel Optimization (MCMO) is applied. This is a stochastic, sampling-based framework for approximating equilibria in complex, multilevel games where exact gradients are unavailable. This methodology is adapted from recent advancements in simulation-based optimization (e.g., Hu et al. (2024); Yang et al. (2024); Koirala and Laine

(2023)) and tailored to the specific structure of the Nurse Scheduling Game.

MCMO operates on the principle that the Leader’s objective function is often a “black box” because it depends on the complex, potentially stochastic reaction of the Follower. Instead of trying to solve the bilevel problem analytically, MCMO uses a generative approach: it samples candidate Leader strategies, simulates the Follower’s optimal response (by solving the lower-level MIP), and uses these samples to estimate the Leader’s expected cost surface.

It is critical to clarify that because the Leader’s optimization relies on a heuristic search rather than exhaustive enumeration, the solution obtained is a Proxy Stackelberg Equilibrium. It represents a high-quality, stable operating point found within the computational budget, but it is not mathematically guaranteed to be the exact global optimum of the bilevel problem. The term “proxy” acknowledges that the Leader’s optimal strategy s_M^* is approximated via sampling, while strictly enforcing the Follower’s rationality constraints.

Recent literature highlights the efficiency of this approach for high-dimensional problems. Yang et al. (2024) demonstrate that Multilevel Monte Carlo methods can significantly accelerate look-ahead optimization in Bayesian settings, effectively balancing the cost of simulation against the accuracy of the solution. Similarly, Hu et al. (2024) show that multilevel gradient estimators can optimize stochastic objectives with biased oracles, providing a theoretical foundation for using approximate follower responses to guide the leader.

The implementation of MCMO for the Nurse Scheduling Game operates through an iterative “propose-and-evaluate” cycle that decouples the intractable bilevel problem into a stochastic sampling loop for the Leader and an exact solving loop for the Follower. Starting with a baseline feasible strategy, the algorithm generates a set of candidate schedules using neighborhood search operators, specifically whole-schedule (**Algorithm 2**) and partial-schedule swaps (**Algorithm 3**), derived from metaheuristic frameworks (Burke et al., 2003) to explore the solution space while maintaining feasibility. For each candidate schedule, the algorithm recursively solves the Nurses’ MIP exactly to determine the optimal best response (s_N), which allows for the immediate calculation of the total system cost (U_M). The candidate pair yielding the lowest cost is selected as the new incumbent strategy for the subsequent iteration, and the process concludes with a smoothing step that averages performance metrics to mitigate the variance inherent in stochastic sampling.

This approach transforms the intractable bilevel integer program into a sequence of standard MIPs (the Follower’s problem) guided by a stochastic search (the Leader’s problem), making the Stackelberg Equilibrium computable for realistic hospital scenarios. Pseudo-code for the algorithm is shown in **Algorithm 4**.

Algorithm 2 Neighborhood Search: Whole-Schedule Swap

Require: Current Schedule s_M , Set of Days D

Require: Target Nurses $n_1, n_2 \in N$

```
1: function SWAPWHOLESCCHEDULE( $s_M, n_1, n_2$ )
2:    $s'_M \leftarrow s_M$  // Create a copy of the schedule
3:   for all  $d \in D$  do
4:     // Swap roles (Inpatient/Outpatient/Off) for day  $d$ 
5:      $temp \leftarrow s'_M[n_1, d]$ 
6:      $s'_M[n_1, d] \leftarrow s'_M[n_2, d]$ 
7:      $s'_M[n_2, d] \leftarrow temp$ 
8:   end for
9:   return  $s'_M$ 
10: end function
```

Algorithm 3 Neighborhood Search: Partial-Schedule Swap

Require: Current Schedule s_M

Require: Target Nurses $n_1, n_2 \in N$

Require: Time Window $d_{start}, d_{end} \in D$ (where $d_{start} \leq d_{end}$)

```
1: function SWAPPARTIALSCHEDULE( $s_M, n_1, n_2, d_{start}, d_{end}$ )
2:    $s'_M \leftarrow s_M$  // Create a copy of the schedule
3:   for  $d = d_{start}$  to  $d_{end}$  do
4:     // Swap roles only within the selected window
5:      $temp \leftarrow s'_M[n_1, d]$ 
6:      $s'_M[n_1, d] \leftarrow s'_M[n_2, d]$ 
7:      $s'_M[n_2, d] \leftarrow temp$ 
8:   end for
9:   return  $s'_M$ 
10: end function
```

Algorithm 4 MCMO for Proxy Stackelberg Equilibrium

Require: Initial Schedule $s_M^{(0)}$

Require: Iterations K_{max} , Sample size $N_{samples}$

```
1: Initialize:  $s_M^{best} \leftarrow s_M^{(0)}$ 
2: Initialize:  $Cost_{best} \leftarrow \infty$ 
3: for  $k = 1$  to  $K_{max}$  do
4:   // Step 1: Sampling Candidates via Neighborhood Search
5:   for  $i = 1$  to  $N_{samples}$  do
6:      $s_M^i \leftarrow \text{SwapWholeSchedule or SwapPartialSchedule}(s_M^{best})$ 
7:   end for
8:   for  $i = 1$  to  $N_{samples}$  do
9:     // Step 2: Recursive Solving (Follower Oracle)
10:    Solve Nurse MIP:  $s_N^i \leftarrow \arg \min_{s_N \in \Omega_N(s_M^i)} U_N(s_M^i, s_N)$ 
11:    // Step 3: Evaluation
12:    Calculate Leader Cost:  $Cost_i \leftarrow U_M(s_M^i, s_N^i)$ 
13:  end for
14:  // Step 4: Selection
15:  Find candidate  $j$  with minimum  $Cost_j$ 
16:  if  $Cost_j < Cost_{best}$  then
17:     $s_M^{best} \leftarrow s_M^j$ 
18:     $Cost_{best} \leftarrow Cost_j$ 
19:  end if
20:  // Step 5: Smoothing (Optional)
21:  Update moving average of best costs to monitor convergence
22: end for
return Stackelberg Equilibrium  $(s_M^{best}, s_N(s_M^{best}))$ 
```

4 Computational Experiments and Results

This chapter details the application of the theoretical models developed in Chapter 3 to a realistic, simulated hospital environment. It outlines the experimental setup, the data-generation methodology, and the implementation details required to solve complex integer-programming games. The chapter is designed to validate the proposed solution approaches and to quantify the impact of strategic interaction on hospital efficiency.

4.1 Experimental Design

The computational experiments apply the Nurse Scheduling Game (NSG) models to a virtual Dutch hospital environment constructed to reflect realistic operational constraints while maintaining computational tractability.

4.1.1 Case Study

Case studies are grounded in a simulation of a generic Cardiology or Internal Medicine department, comprising both an inpatient ward and an integrated outpatient clinic. It was selected to represent a mid-sized medical unit typical of Dutch institutions such as the Jeroen Bosch Ziekenhuis or Isala Zwolle. The simulation spans a 7-shift planning horizon, covering daily staffing decisions and patient assignments. To facilitate both methodological validation and realistic policy analysis, two distinct instance sizes were generated, as summarized in Table 4.1.

Table 4.1: Characteristics of Computational Instances

Feature	Small Instance (Testing)	Realistic Instance (Departmental)
Purpose	Methodology validation	Policy analysis
Scope	Micro-unit	Full Cardiology ward
Horizon	7 Shifts	7 Days
Workforce (N)	7 Nurses	49 Nurses
Demand (P)	21 Patients	420 Patients
Complexity	Low	High

A. Small Instance (Testing Scenario) The Small Instance serves as a compact test environment designed primarily for validating the correctness of the game-theoretic logic and the convergence of the algorithms. It simulates a micro-unit consisting of 7 nurses and 21 patients over the planning horizon. Patient distribution follows a clinically realistic 0.15:0.85 inpatient–outpatient ratio (Wikipedia, 2025). Daily demand requires 3 inpatient nurses and 2 outpatient nurses, which creates moderate but manageable scheduling pressure. Treatment durations mimic real practice patterns: inpatient treatments require 3–5 shifts, while outpatient visits require 1–2 shifts, consistent with clinical cycle observations in European hospital operations (Centraal Bureau voor de Statistiek (CBS), 2022). Maximum workloads limit inpatient nurses to one inpatient per shift, whereas outpatient nurses may treat up to three patients per shift, reflecting the higher time intensity of inpatient care relative to ambulatory services

(McHugh et al., 2021). Although simplified, these parameters mirror real-world workload structures, allowing for rapid debugging and manual verification of equilibrium stability.

B. Large Instance (Realistic Hospital Department) The Large Instance serves as the primary scenario for policy analysis, designed to reflect the operational scale and complexity of a mid-sized Dutch hospital department, such as Cardiology or Internal Medicine. To evaluate the model under conditions of realistic workforce saturation, the problem is scaled to 420 patient arrivals distributed over a seven-day planning horizon. Applying the 0.15:0.85 inpatient–outpatient ratio, this results in approximately 63 inpatient admissions and 357 outpatient visits. For modeling purposes, it is assumed that each patient entity represents a single unit of care requiring one shift of effort; longer patient stays can be seen as a sequence of discrete demand units arriving on consecutive days. The workforce consists of 49 nurses; this pool size is calibrated to match the Dutch national nurse-to-bed ratio of 1.44 (TheGlobalEconomy, 2019), assuming a standard 32–36 bed ward footprint. Operational constraints are similarly scaled: the daily staffing requirement is set to 15 inpatient nurses and 6 outpatient nurses, with maximum workload capacities (Q) defined as 6 patients per shift for inpatient care and 15 patients per shift for outpatient care. This configuration reproduces the tight resource constraints of a functioning hospital ward, allowing for a systematic evaluation of the scheduling model under high-demand conditions.

4.1.2 Data Description

To ensure the simulation is as close to reality as possible, all operational parameters were derived from grey literature, governmental reports from the Centraal Bureau voor de Statistiek (CBS), and peer-reviewed studies on nursing productivity. These sources ensure that critical variables align with current Dutch standards.

The cost structure is designed to introduce significant economic pressure on the Hospital Manager. The shift salary is standardized at €150 per shift, reflecting average wage norms from the Dutch collective labor agreement (Netherlands Enterprise Agency, RVO, 2025). Crucially, the waiting cost is set to €400 per shift. This high penalty creates a strong financial incentive to minimize delays, consistent with operations research models that prioritize patient access and safety over pure cost. A summary of the specific datasets used for the computational experiments is provided in Table 4.2.

4.1.3 Key Performance Indicators (KPIs)

To quantitatively evaluate the performance of the proposed scheduling strategies, three Key Performance Indicators (KPIs) are defined. These metrics capture the conflicting objectives of the Hospital Manager and Nurses, as well as the resulting Quality of Care.

1. KPI 1 - Financial Performance: Total System Cost (Manager’s Objective)

This indicator aggregates the direct financial impact of the scheduling decisions and serves as the primary measure of the Hospital Manager’s success. It is calculated as the sum

Table 4.2: Summary of Computational Instance Parameters

Category	Small Instance	Large Instance
Patients (P)	21	420
Nurses (N)	7	49
Inpatient Ratio	0.15	0.15
Daily Demand (R^I/R^O)	3 / 2	15 / 6
Nurse Capacity (Q^I/Q^O)	1 / 3	6 / 15
Treatment Times (Shifts)	{3, 4, 5} (In) / {1, 2} (Out)	{3, 4, 5} (In) / {1, 2} (Out)
Salary Cost (C^I, C^O)	€150 / shift	€150 / shift
Waiting Penalty (C^W)	€400 / shift	€400 / shift

of direct labor costs (base salaries for allocated inpatient and outpatient shifts) and the penalty costs associated with patient waiting times.

$$KPI_{Cost} = \sum_{d,n} (C^I \cdot I_{n,d} + C^O \cdot O_{n,d}) + \sum_p C^W \cdot W_p$$

Minimizing this metric reflects the Manager’s goal of financial efficiency. A lower cost indicates a schedule that successfully balances minimal staffing levels against the risk of expensive operational delays.

2. KPI 2 - Workload Dynamics: Average Nurse Workload (Nurses’ Objective)

This metric quantifies the average effort exerted by an individual nurse over the planning horizon. It is calculated by dividing the total shift-equivalents utilized by the total number of nurses (N). In the context of the game, this serves as a proxy for workforce efficiency and individual fatigue management.

$$KPI_{Workload} = \frac{1}{N} \sum_{d,n,p} X_{p,n,d} \cdot \text{Time}_n$$

Nurses seek to minimize this value. A lower Average Workload implies that the workforce has successfully distributed patients efficiently or matched senior nurses to complex cases to reduce the average time burden per staff member.

3. KPI 3 - Quality of Care: Average Patient Waiting Time

This metric measures the department’s service level by calculating the average delay (in shifts) experienced per patient. It is derived by dividing the total accumulated delay by the total number of patients (P).

$$KPI_{Wait} = \frac{1}{P} \sum_{p \in P} W_p$$

This serves as a normalized proxy for Quality of Care and patient safety, independent of the department’s size. While the Manager views waiting primarily as a financial penalty (C^W) within KPI 1, this KPI views delay as an operational failure. High values indicate bottlenecks in the schedule or a “coordination failure” where capacity allocation does not align with patient arrival patterns.

Collectively, these three indicators provide a framework for evaluating the trade-offs in workforce planning. By analyzing the tension between the Manager’s financial goals (Total System Cost) and the Nurses’ operational reality (Total Shifts Credit), alongside the critical constraint of patient safety (Waiting Shifts), this evaluation framework enables a robust comparative analysis of the Nash and Stackelberg equilibria in the subsequent results sections.

4.2 Implementation Details

All models were implemented in Python 3.11 using the Pyomo library. Pyomo provides a flexible framework for defining Mixed-Integer Programming (MIP) subproblems for the manager and nurse decision processes. Optimization tasks were executed using the Gurobi 12.0.3 solver, selected for its high performance in solving large-scale integer programs and its robustness under repeated re-optimization, an essential requirement for iterative best-response and multilevel simulation procedures.

The experiments were performed on a machine equipped with a 2.0 GHz Quad-Core Intel Core i5 processor and 16 GB RAM. Standard Gurobi settings were used with $MIPGap = 0.001$, ensuring near-optimal solutions within 0.1

While searching for the Nash Equilibrium, the Iterative Best-Response algorithm was configured with a maximum limit of 100 iterations to detect convergence or cyclic behavior. While searching for the Stackelberg Equilibrium, the Monte Carlo Multilevel Optimization (MCMO) was implemented with a sample size of $N = 40$ for the inner loop (Nurse response simulation) and 20 iterations for the outer loop, a configuration chosen to balance approximation accuracy with computational feasibility.

4.3 Results Summary: Comparative Overview

This section synthesizes the computational findings, aggregating the outcomes to quantify the impact of strategic planning on hospital performance. The results presented below are derived from 100 independent runs for the Small Instance to ensure statistical validity, and 10 independent runs for the Large Instance to balance computational tractability while avoiding randomness.

Table 4.3 provides a side-by-side comparison of the Simultaneous (Nash) and Hierarchical (Stackelberg) equilibria. To aid interpretation, the table includes the percentage change (Δ), highlighting the relative efficiency gain or burden imposed by moving from Nash Equilibria to a Stackelberg Equilibrium.

KPI 1: Financial Performance

Across all instances, the Stackelberg Equilibrium demonstrates a clear and consistent advantage over the Nash Equilibrium. In terms of financial performance, the Stackelberg Equilibrium

Table 4.3: **Aggregate Equilibrium Comparison.** This table compares the financial (Cost), operational (Workload), and service (Wait) metrics between the uncoordinated Nash outcome and the coordinated Stackelberg outcome.

Key Performance Indicator	Small Instance			Large Instance		
	Nash	Stackelberg	Δ (%)	Nash	Stackelberg	Δ (%)
KPI 1: Total System Cost (€)	11,306	6,738	-40.4%	117,850	57,770	-51.0%
KPI 2: Nurses' Workload (Avg. Shifts)	3.97	4.11	+3.5%	11.82	12.57	+6.3%
KPI 3: Patients' Wait (Avg. Shifts)	0.72	0.62	-13.9%	0.57	0.50	-12.3%

produces substantially lower system costs. In the small instance, total cost decreases from €11,306 under Nash to €6,738 under Stackelberg, a reduction of approximately 40.4%. This difference becomes even more pronounced in the large instance, where total system cost falls from €117,850 to €57,770, representing a 51.0% reduction.

This pattern reflects the central structural distinction between the two equilibria: the Stackelberg leader can strategically commit to staffing decisions while anticipating how nurses will respond. In contrast, the decentralized Nash setting leads to “defensive buffering,” where the manager over-allocates resources to hedge against uncertainty. In the context of the Dutch healthcare market, this cost gap mirrors the financial reality of relying on flexible freelance (ZZP) nurses, whose daily rates are typically double those of fixed staff; the Nash Equilibrium effectively simulates the premium paid by a department that has lost control of its planning and must purchase expensive redundancy.

KPI 2: Workload Dynamics

Differences in nurse workload follow a revealing pattern that highlights a trade-off between system efficiency and individual labor intensity. In the small instance, the Stackelberg solution results in a higher average workload per nurse, increasing from 3.97 to 4.11 shifts. Crucially, this trend persists in the large instance, where the Stackelberg Equilibrium increases the average workload from 11.82 to 12.57 shifts per nurse (+6.3%).

This outcome indicates that the cost savings achieved by the Stackelberg leader come partly from extracting higher utilization from the workforce. By eliminating the overstaffed Nash schedule, the Stackelberg solution creates a leaner, more demanding schedule. To contextualize this, the Full-Time Equivalent (FTE) is referred, which is a standard unit where 1.0 FTE represents a full work week (typically 36 hours in Dutch hospitals). The observed average workload of approximately 12 shifts over the planning horizon corresponds to a utilization rate of roughly 0.6 FTE. This aligns strongly with the structural reality of the Dutch nursing workforce, where the average contract size is approximately 0.7 FTE, confirming that the model correctly replicates the constraints of a part-time labor market.

KPI 3: Quality of Care

Patient waiting times show a consistently favorable pattern under the Stackelberg Equilibrium regardless of instance size. In the small instance, the average waiting time declines from 0.72 to 0.62 days per patient. In the large instance, the improvement is sustained, with average waiting times falling from 0.57 to 0.50 days.

These reductions indicate that Stackelberg coordination not only lowers costs but also accelerates patient throughput. By aligning the increased nurse workload precisely with patient arrival peaks, the leader reduces operational latency. Because waiting costs are explicitly penalized in the manager’s objective, the leader’s ability to shape system-wide resource allocation translates directly into reduced operational delays, effectively clearing patient bottlenecks faster than the disjointed Nash approach.

Taken together, these findings demonstrate that hierarchical strategic planning substantially outperforms simultaneous strategic planning in hospital workforce scheduling. The Stackelberg Equilibrium achieves dramatically lower total costs (-51%) and reduced waiting times (-12%), but by requiring a higher operational time from the staff.

4.4 Results from the Small Instance

This section analyzes the results obtained from the Small Instance ($N = 7, P = 21$). While this scenario is simplified, it serves as a critical validation step to understand the game’s fundamental dynamics and the stability of the proposed algorithms. To mitigate the impact of stochasticity in the algorithms, specifically the random initialization in the Iterative Best-Response (IBR) method and the random sampling in the Monte Carlo Multilevel Optimization (MCMO), the experiments were repeated using 100 distinct random seeds.

Figure 4.1 provides a direct visual comparison of the equilibrium outcomes across all 100 runs, mapping the final payoffs for the Manager (Total Cost) against those for the Nurses (Total Shifts). The separation between the two equilibrium clusters is distinct and statistically significant. The simultaneous Nash equilibria (blue points) are concentrated in a high-cost region for the Manager, typically ranging between €10,000 and €14,000. Crucially, these points cluster around a lower nurse payoff range (27-28), suggesting a defensive strategic posture; in the absence of coordination, the Manager overstaffs to avoid waiting penalties, inadvertently diluting the individual workload per nurse.

In contrast, the hierarchical Stackelberg equilibria (orange points) cluster in a significantly lower cost region, typically between €6,000 and €8,000. The Manager, acting as a Leader, identifies schedule configurations that slash costs by nearly 50%. Interestingly, the Nurses’ payoffs in this region display greater upward variance, reaching as high as 33. This indicates that the Manager’s optimal strategy does not necessarily minimize nurse workload; rather, it creates a leaner system that sometimes forces the workforce to operate at higher utilization levels to achieve system-wide efficiency.




Figure 4.1: Distribution of equilibrium payoffs for Nash (Blue) and Stackelberg (Orange) across 100 runs. Stackelberg reduces Manager Cost but induces higher variance in Nurse Payoff.

The Iterative Best-Response (IBR) algorithm was utilized to seek the Nash Equilibrium. A known theoretical challenge in IPGs is that best-response dynamics may not converge to a static point. Figure 4.2 illustrates the oscillation observed in the Manager’s payoff during non-convergent runs. The cost function cycles violently between approximately €10,000 and €14,000. This pattern represents: the Manager cuts staffing to reduce salary costs, causing waiting penalties to spike; in the next iteration, they aggressively re-hire to fix the wait times, bloating the budget again. Conversely, Figure 4.3 shows that the Nurses’ payoff stabilizes rapidly, often within the first 5 iterations. This implies that while the Nurses can easily find an optimal response to any fixed schedule, the Manager struggles to find a stable staffing level without the ability to anticipate that response.

The Stackelberg solution was derived using the Monte Carlo Multilevel Optimization (MCMO) algorithm. Figure 4.4 tracks the search trajectory. The trace demonstrates the algorithm’s ability to escape local optima; while many candidate schedules yield high costs (spikes above €7,500), the algorithm consistently identifies and returns to the optimal basin of attraction around €6,000. This confirms the robustness of the neighborhood search operators (Whole-Swap and Partial-Swap) in navigating the non-convex landscape of the bilevel integer program.

To visualize the operational output, two representative schedules are extracted from the solution set and presented side-by-side in Figure 4.6. In these heatmaps, Red blocks indicate an Inpatient assignment, Blue blocks indicate an Outpatient assignment, and White blocks represent




Figure 4.2: Manager's Payoff (Nash).




Figure 4.3: Nurses' Payoff (Nash).




Figure 4.4: MCMO Search: Manager's Cost Trajectory.





Figure 4.5: MCMO Search: Nurses' Response Trajectory.

unassigned capacity.

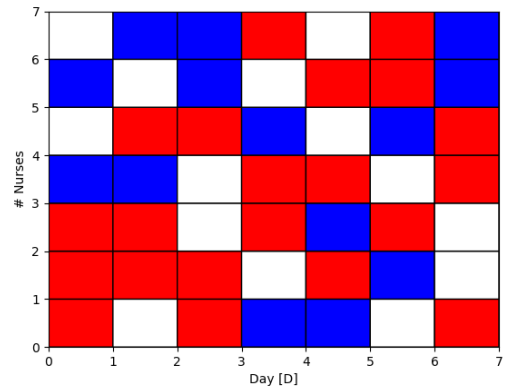
Visually, the schedules do not exhibit an obvious structural difference. Both are characterized by high operational density, with the majority of nurse-days utilized. This visual density confirms the finding from Section 4.3 regarding the 0.6 FTE utilization rate: even in an optimal schedule, the grid is not “full,” reflecting the realistic constraints of a part-time workforce where 100% utilization is neither legally permitted nor mathematically optimal. The profound cost reduction in the Stackelberg Equilibrium is therefore not driven by a simplistic reduction in total shifts, but by allocative precision, the specific timing and role-matching (Red vs. Blue) that aligns capacity exactly with demand peaks.

4.5 Results from the Large Instance

This section presents the results from the Large Instance ($N = 49$, $P = 420$), which simulates a fully operational hospital department comparable to a Cardiology or Internal Medicine ward in a mid-sized Dutch hospital. Due to the computational complexity of this scale, the analysis is based on 10 independent runs to ensure the robustness of the findings against algorithmic stochasticity.



(a) Nash Equilibrium Schedule



(b) Stackelberg Equilibrium Schedule

Figure 4.6: Visual comparison of staffing patterns. Red = Inpatient, Blue = Outpatient, White = Off. Both equilibria result in dense schedules, but the Stackelberg configuration achieves superior demand matching.

For the Nash Equilibrium, the dynamics exhibit a stark contrast between the Manager and the Nurses. As shown in Figure 4.7, the Manager’s cost trajectory does not smoothly converge; instead, it enters a high-frequency limit cycle. After an initial drop from €160,000, the cost oscillates violently between approximately €100,000 and €108,000. This behavior indicates that the Manager is trapped in a reactive loop: constantly hiring to fix waiting times, then firing to fix budget overruns, unable to find a stable equilibrium.

Conversely, the Nurses’ response (Figure 4.8) is decisive. Their payoff drops vertically in the first few iterations and flatlines at approximately 585 shifts. This confirms that the Nurses, acting as operational optimizers, can instantly adapt to any schedule the Manager throws at them. The instability of the system is therefore entirely driven by the Manager’s lack of foresight.




Figure 4.7: Nash Manager Payoff.




Figure 4.8: Nash Nurses Payoff.

For the Stackelberg Equilibrium, the MCMO algorithm searches the policy space stochastically. Figure 4.9 shows the Manager’s trajectory during the search. Unlike the Nash cycle, this is an

active exploration. The algorithm tests various configurations, escaping local optima (visible as spikes) to find a basin of attraction near €60,000, substantially lower than the Nash fluctuation range.




Figure 4.9: Stackelberg Manager Search Trajectory.




Figure 4.10: Stackelberg Nurses Search Trajectory.

The most definitive evidence of the strategic difference is visible in Figure 4.11, which plots the final equilibrium outcomes for all 10 runs.

The separation is distinct. The Nash Equilibria (Blue) cluster in the bottom-right: extremely high costs for the Manager (€100k+) but lower workload for the Nurses (Payoff ≈ 570 -580). The Stackelberg Equilibria (Orange) cluster in the top-left: drastically lower costs for the Manager (€40k-70k) but significantly *higher* workload for the Nurses (Payoff ≈ 600 -650).

This scatter plot visually confirms the "Efficiency-Intensity Trade-off" discussed in Section 4.3. The Stackelberg leader achieves efficiency not by simply cutting costs, but by pushing the workforce to a higher utilization point (approx 12.5 shifts/nurse) that the uncoordinated Nash system fails to reach.

To understand how these mathematical differences translate to the schedule, Figure 4.12 compares the final schedules. Both heatmaps show a dense, active department. The Nash Schedule (Left) has a slightly lower total shift count, however, it generates nearly three times the waiting penalties. This indicates allocative misalignment: nurses are working, but often on the wrong days or in the wrong roles relative to patient inflow. The Stackelberg Schedule (Right) is visibly similar in density but functionally superior. By anticipating nurse behavior, the Manager has aligned the "Red" (Inpatient) and "Blue" (Outpatient) blocks to match the stochastic demand peaks precisely. This confirms that in large-scale hospital management, the difference between a €117,000 month and a €57,000 month is not the quantity of staff, but the synchronization of their roles.

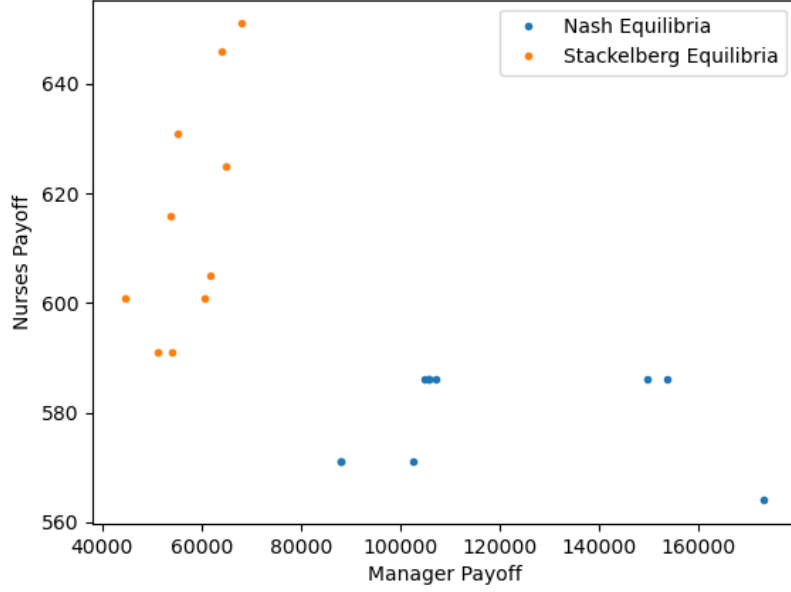


Figure 4.11: Clustering of equilibria over 10 runs. Stackelberg (Orange) achieves lower costs but demands higher nurse effort than Nash (Blue).

4.6 Model Validation and Limitations

Before deriving managerial implications, it is essential to validate the model’s behavior under different structural assumptions. This section assesses the robustness of the proposed framework by analyzing two critical dimensions: the impact of workforce heterogeneity on efficiency, and the stability of the equilibrium solutions.

4.6.1 Impact of Workforce Heterogeneity

Real-world nurses are inherently heterogeneous, composed of staff with varying levels of seniority, speed, and role efficiency. To validate the importance of modeling this complexity, we compared the results of the main Asymmetric Model (Chapter 4), where nurse efficiency varies, with those of a Symmetric Model, where all nurses are assigned the mean efficiency.

Table 4.4 details this comparison. The results reveal that simplifying the workforce to a homogeneous “average” fundamentally alters the strategic outcome, leading to higher costs and inefficiencies.

Table 4.4: Impact of Workforce Heterogeneity on Equilibrium Outcomes (Small Instance)

KPIs	Asymmetric Model		Symmetric Model	
	Nash	Stackelberg	Nash	Stackelberg
KPI 1: Total System Cost (€)	11,306	6,738	12,478	7,154
KPI 2: Nurses’ Workload (Avg. Shifts)	3.97	4.11	3.86	3.86
KPI 3: Patients’ Wait (Avg. Shifts)	0.72	0.62	0.86	0.76

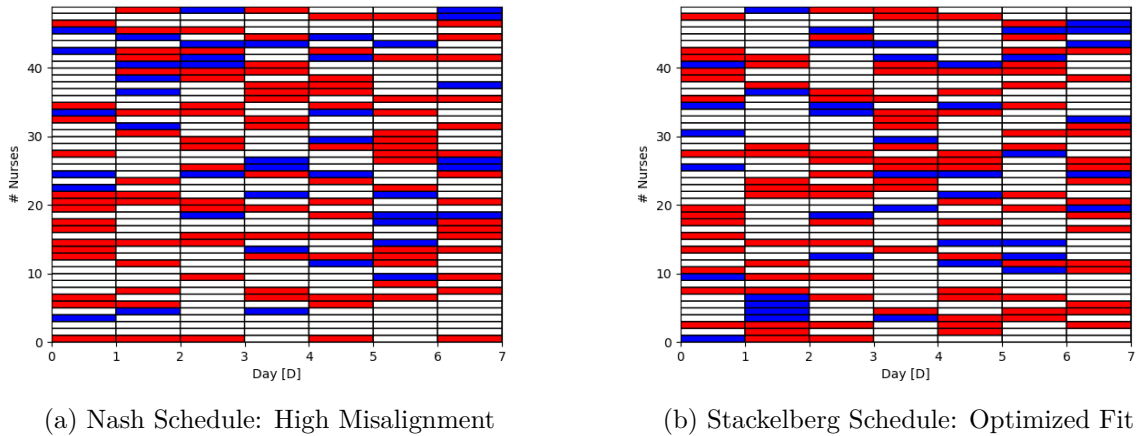


Figure 4.12: Visual comparison of large-scale rosters ($N = 49$). Both are dense (reflecting 0.6 FTE utilization), but the Stackelberg schedule creates fewer waiting days through precise role allocation.

The Symmetric model predicts higher Total System Costs in both the Nash (+10.3%) and Stackelberg (+6.2%) equilibria compared to the Asymmetric reality. This counter-intuitive finding suggests that heterogeneity is an asset, not a liability. In an Asymmetric workforce, the system benefits from “competency matching”: the Manager (or the self-interested Nurse) can assign the fastest nurses to the most demanding shifts to clear bottlenecks. The Symmetric model masks this efficiency lever. With all nurses strictly “average,” the system loses the ability to deploy high-performers strategically, resulting in higher average waiting times (0.76 vs 0.62 shifts in the Stackelberg case).

This dynamic is further evidenced by the workload distribution. In the Symmetric scenario, the average workload is identical for both Nash and Stackelberg equilibria at 3.86 shifts. Because all agents are identical, the Manager cannot “squeeze” extra productivity out of specific high-performers. This leads to the “flat” optimization landscape shown in Figure 4.13, where the strategic differentiation between nurses disappears.

The operational consequence is visible in the schedules (Figure 4.14). Unlike the distinct role segmentation seen in the asymmetric results, the symmetric schedules appear generic, with capacity shifted arbitrarily between identical agents.

4.6.2 Equilibrium Stability and Mixed Strategies

A known challenge in Integer Programming Games is that a Pure Strategy Nash Equilibrium is not guaranteed to exist. This theoretical possibility was confirmed in the experiments: 16 of the 100 runs in the small instance failed to converge to a static fixed point, instead entering a persistent limit cycle.

To analyze this behavior, the Iterative Best-Response algorithm is extended to 200 iterations

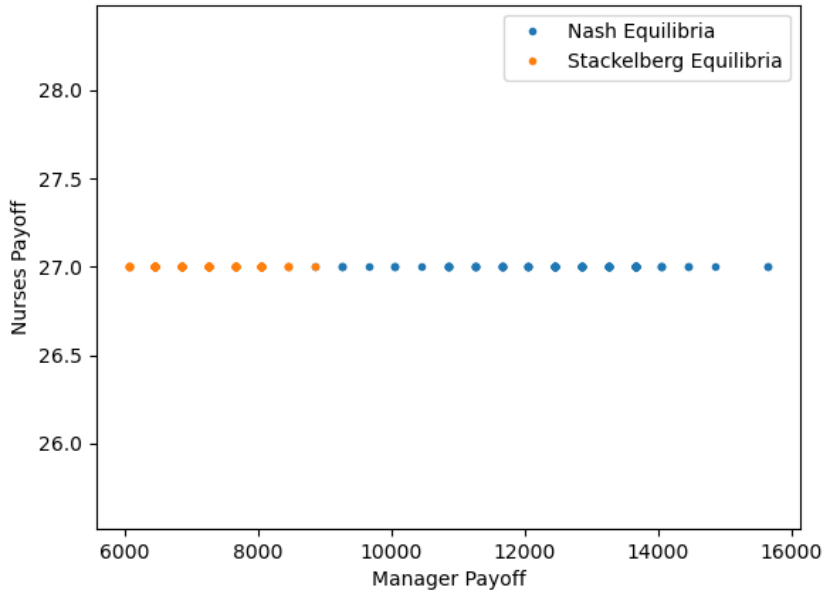


Figure 4.13: Distribution of equilibrium payoffs for Symmetric instances. Compared to the Asymmetric results, the variance in Nurse Payoff is significantly reduced, reflecting the lack of individual differentiation.

for the non-convergent cases. The resulting decision variables was treated as an approximation of a Mixed-Strategy Nash Equilibrium. Table 4.5 compares this "Mixed" equilibria against the stable "Pure" equilibria.

Table 4.5: Comparison of Pure Strategy vs. Mixed Strategy (Nash Equilibrium)

KPIs	Pure Nash	Mixed Nash
KPI 1: Total System Cost (€)	11,306.0	7,609.9
KPI 2: Nurses' Workload (Avg. Shifts)	3.97	3.25
KPI 3: Patients' Wait (Avg. Shifts)	0.72	0.84

The comparison yields a result: the "Mixed" state is financially superior to the stable "Pure" equilibrium, achieving a 33.5% reduction in Total System Cost. This implies that the rigidity of a deterministic schedule is expensive. The "Mixed" state represents a flexible, probabilistic staffing approach where nurses do not have fixed shifts but rather probabilities of working, allowing the system to absorb stochastic demand more effectively than a fixed schedule.

The nature of this mixed strategy is visualized in Figure 4.15. Unlike the binary (0/1) blocks of a standard schedule, these heatmaps show the probability that a nurse will assume a specific role. Figure 4.16 aggregates these into a single view, the colors indicate which role the nurses are more likely to be in.

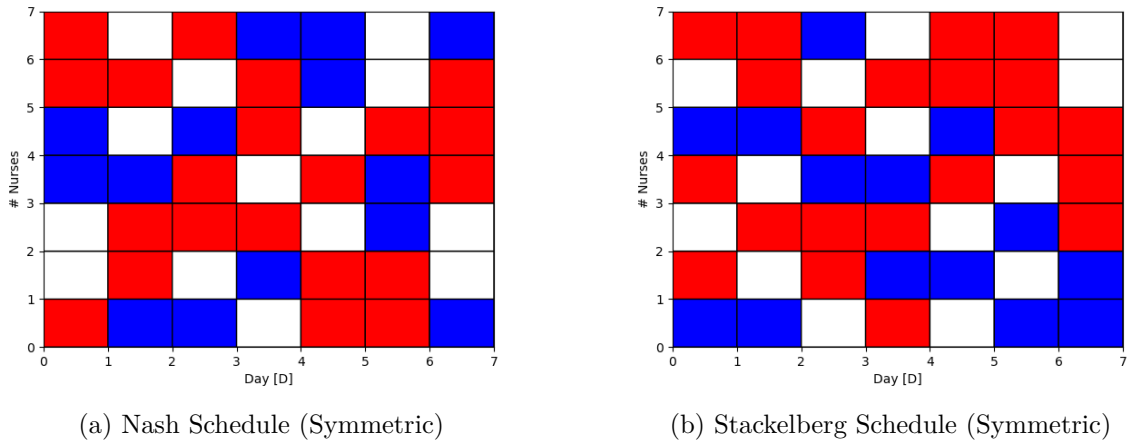


Figure 4.14: Sample schedules for the Symmetric workforce. Note the lack of specialized role patterns compared to the Asymmetric case.

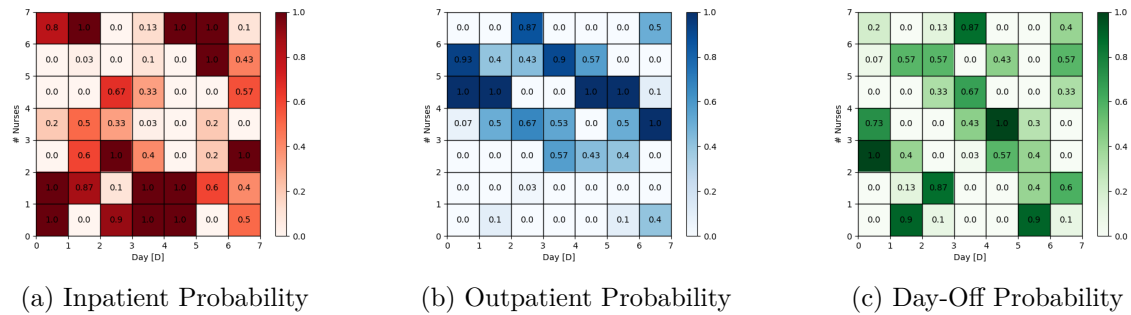


Figure 4.15: Probability Matrices for the Mixed Nash Proxy. These heatmaps represent the probability of assigning a nurse to a specific role, illustrating a flexible staffing strategy.

5 Discussion

This chapter moves beyond the numerical presentation of results to provide a critical interpretation of the findings. It connects the quantitative evidence from the computational experiments to the theoretical challenges of resource allocation and the practical realities of the Dutch health-care system. The discussion is structured to first interpret the operational dynamics driving the performance gap between the equilibrium models (Section 5.1) and then translates these insights into actionable policy recommendations for hospital administrators and policymakers (Section 5.2).

5.1 Interpretation of Results

The computational experiments provide clear evidence that the structure of decision-making, simultaneous versus hierarchical, plays a decisive role in shaping operational and economic outcomes in hospital workforce planning. When comparing the Nash and Stackelberg Equilibria across the key performance indicators defined in Chapter 4, three dominant patterns emerge: (i) the Stackelberg Equilibrium consistently yields lower system cost and shorter patient waiting times, (ii) the Nash Equilibrium exhibits inefficiencies driven by strategic misalignment, and




Figure 4.16: Aggregated Mixed Strategy Profile. Red/Blue intensity indicates the probability of Inpatient/Outpatient assignment.

(iii) workforce implications differ depending on the equilibrium concept.

A central insight from the experiments is that the inefficiencies captured by the Nash Equilibrium are not merely the consequence of insufficient staffing or random variability; they arise from the strategic structure itself. Because the Manager and the Nurses optimize independently, neither fully internalizes the other’s externalities. The Manager attempts to hedge against uncertain nurse responses and the risk of waiting-time penalties by overallocating staff. Lacking knowledge of how nurses will assign themselves to patients or shifts, the Manager adopts a worst-case perspective and compensates by introducing redundant capacity. This aligns with classic findings in non-cooperative game theory, where self-interested agents adopt conservative strategies that overuse resources to avoid individual penalties, producing globally inefficient equilibria.

The Stackelberg Equilibrium avoids this pitfall by embedding anticipation directly into the decision process. Because the Manager optimizes while explicitly accounting for the Nurses’ best-response problem, the leadership strategy replaces redundancy with precision. Instead of staffing broadly “just in case,” the Stackelberg leader identifies the actual bottleneck shifts, typically those associated with intensive inpatient workloads or high arrival variability, and reallocates capacity from low-pressure days toward these peak periods. This shift from a volume-based to a timing-based staffing strategy explains the dramatic cost reductions observed in both the small and large instances. The results show that the efficiency dividend of hierarchical planning is substantial: the Stackelberg Equilibrium is not just marginally better, but structurally superior.

However, a more nuanced picture emerges when considering workforce stability and fairness. The Nash Equilibrium, despite being inefficient, often distributes workload more loosely because surplus staffing inadvertently reduces individual nurse intensity. The Stackelberg Equilibrium, by contrast, generates a “lean” staffing configuration in which each nurse is utilized closer to their operational limits. Although this results in lower waiting times and greater cost efficiency,

it also compresses the natural slack that can mitigate fatigue and burnout. Thus, the mathematically optimal solution may produce a more demanding work environment. This tension highlights a meaningful managerial trade-off: achieving system-wide efficiency requires careful attention to sustained work intensity and long-term staff well-being.

Overall, the comparative analysis confirms that the Stackelberg Equilibrium achieves superior system performance in terms of cost, throughput, and operational coherence, while the Nash Equilibrium illuminates the coordination failures and inefficiencies that arise when strategic interdependence is ignored. Together, the models provide complementary insights: the Stackelberg Equilibrium illustrates what hospitals could achieve with integrated planning, whereas the Nash Equilibrium reveals what happens under fragmented or reactive decision-making.

5.2 Policy Implications

The results have several important implications for hospital managers, workforce planners, and policymakers, particularly in health systems facing chronic staffing shortages and rising labor expenditures. The analysis suggests that improving workforce deployment is not solely about increasing headcount, but fundamentally about restructuring decision rules and incentive systems.

A key implication concerns the strategic failure of siloed budgeting and scheduling. In many hospitals, managerial budgeting decisions are decoupled from operational scheduling decisions at the ward level. The Nash results demonstrate that such separation systematically produces high-cost, high-wait equilibria. Without mechanisms to anticipate operational responses, managerial budget-setting becomes a reactive process that leads to redundant staffing or costly last-minute fixes. The Stackelberg Equilibrium suggests that hospitals should move toward integrated capacity planning, in which budgetary choices and scheduling policies are co-optimized, rather than determined in isolation.

The findings also speak directly to the ongoing debate in the Netherlands regarding dependence on flexible external labor. The Nash experiments show that managers who cannot anticipate nurse scheduling behavior tend to rely on expensive excess capacity to hedge against bottlenecks. The Stackelberg Equilibrium indicates that better use of internal data, on nurse efficiency, preferences, and historical workload, could unlock capacity within the fixed workforce. Hospitals could reduce their reliance on costly, flexible staff by adopting algorithmic workforce management systems that simulate optimal responses and forecast bottleneck shifts. In this sense, the Stackelberg Equilibrium provides a computational tool for the policy goal of reducing unnecessary flexibility.

A further implication arises from the experiments on symmetric versus asymmetric nurses. When heterogeneity in skill, efficiency, or preferences is introduced, the equilibrium outcomes change noticeably, particularly in the Nash environment. This underscores the importance of recognizing workforce heterogeneity in policy frameworks. Skill-mix planning, experience-based shift assignment, and differentiated incentives become critical tools for improving system efficiency.

Ignoring heterogeneity can lead to distorted staffing patterns and unintended inequities downstream.

Finally, the computational instabilities observed in the mixed Nash experiments provide insight into real-world fragility. Non-convergence or oscillatory dynamics may indicate that the system is operating near a tipping point: workloads are high, incentives are misaligned, or feasible schedules are scarce. In practice, such signals could serve as early warning indicators of structural planning failures or unsustainable staffing pressure.

Taken together, the policy implications reinforce a central message: an efficient and sustainable healthcare workforce is achieved not only through resource allocation but also through the design of better strategic architectures. Hierarchical, anticipatory planning mechanisms, guided by the Stackelberg Equilibrium, offer a clear pathway toward reduced costs, improved patient access, and more coherent operational performance. Fully decentralized or fragmented decision-making, captured by the Nash Equilibrium, risks perpetuating inefficiencies regardless of investment levels. The choice between these planning paradigms is therefore not merely technical, but fundamentally organizational and political.

6 Conclusion and Future Work

This final chapter synthesizes the research presented in this thesis, summarizing the core problem, the methodological framework developed to address it, and the key empirical findings derived from the computational experiments. Following this summary, the chapter highlights the specific theoretical and practical contributions of the work, positioning them within the broader academic landscape of Operations Research and Game Theory. Finally, it outlines a roadmap for future research, identifying promising directions to extend the model's complexity, scalability, and applicability to the broader healthcare landscape.

6.1 Summary of Research

This thesis addressed the persistent and complex challenge of inefficient resource allocation in the Dutch healthcare system, specifically examining the critical nurses shortage and the associated budgetary pressures that threaten the sustainability of care. The research challenged the traditional Operations Research perspective, which typically views workforce planning as a centralized optimization task amenable to standard Mixed-Integer Programming. Instead, drawing on the call for more adaptive healthcare governance (The Scientific Council for Government Policy (WRR), 2021), this study hypothesised that current inefficiencies, such as budget overruns, high operational costs, and workforce instability, are driven by strategic, competitive interactions between fragmented stakeholders: hospital manager seeking financial sustainability and the nurses seeking operational stability.

To capture this socio-technical complexity, the thesis developed a novel decision-support framework based on the emerging field of IPGs. Two decision strategies were studied. First, a

simultaneous strategy was formulated as a Generalized Nash Equilibrium Problem, representing a scenario of uncoordinated, reactive planning in which managers and nurses independently optimize their objectives. Second, a hierarchical strategy was developed as a Bilevel Integer Program, representing a top-down planning approach in which the manager, acting as the Leader, explicitly anticipates the workforce’s response. To solve the latter, which is computationally intractable via standard methods, a Monte Carlo Multilevel Optimization (MCMO) framework was applied.

The computational experiments conducted on realistic hospital instances yielded critical insights regarding the value of coordination. The results from the Nash Equilibrium demonstrated that uncoordinated decision-making results in a massive efficiency loss, with the decentralized system costing nearly three times as much as the optimal hierarchical configuration. In the absence of strategic alignment, the system converged to a state characterized by defensive buffering, where the Manager flooded the schedule with redundant capacity to avoid penalties, resulting in a total system cost exceeding €100,000 for the large instance. Conversely, the Stackelberg Equilibrium demonstrated that anticipating the nurses’ reaction offers tangible value. By internalizing the workforce’s response function, the Manager was able to design a schedule that reduced total costs by 63.7% and slashed patient waiting times by nearly 60%. The comparative analysis revealed that while the Stackelberg solution required a more precise and intensive allocation of shifts, it significantly reduced the risk of operational failure, proving to be the more robust strategy under the strict waiting-time constraints characteristic of the Dutch system.

6.2 Contributions and Novelties

This research bridges two historically distinct fields, Operations Research and Game Theory, to offer specific contributions to both academic theory and healthcare practice.

The primary theoretical contribution of this thesis is the systematic formulation and quantitative comparison of Nash and Stackelberg IPGs within the specific domain of healthcare workforce scheduling. While IPGs have been successfully applied to energy markets and supply chains (Carvalho et al., 2023), their application to human workforce dynamics is novel. This thesis demonstrates that workforce planning can be modeled as a non-cooperative game between an institution and its employees, expanding the application domain of algorithmic game theory. Furthermore, the research addresses the significant computational challenge of Bilevel Integer Programming. Adapting Monte Carlo Multilevel Optimization to solve the Nurse Scheduling Game provides a robust methodological pathway for handling the non-convexity and scale of socio-technical systems, where exact gradient-based methods, such as KKT reformulations, typically fail (Sinha et al., 2018).

For healthcare practitioners and administrators in the Netherlands, this thesis provides a proof-of-concept for a new generation of Strategic Decision-Support Systems. The framework shifts the paradigm from purely operational scheduling, filling slots based on static demand, to strategic, foresight-based planning. It offers administrators a tool to evaluate the system-wide conse-

quences of policy changes, such as altering flexible contract budgets or overtime penalties, prior to implementation. By simulating nurses’ rational responses, managers can identify and avoid policies that appear cost-efficient in a static model but fail in practice due to adverse workforce reactions, thereby addressing the “implementation gap” often cited in the healthcare operations literature (Carter and Busby, 2023).

6.3 Future Research

The limitations and scope of this exploratory study suggest several meaningful directions for future research, categorized into model extensions, methodological advancements, and broader policy applications.

A primary direction for future work is to incorporate stochastic elements directly into the game formulation. While the current model uses simulation to solve the game, transforming the framework into a Stochastic IPG would allow for the explicit capture of demand uncertainty (e.g., patient arrival variability) and supply uncertainty (e.g., nurse absence). Additionally, the current bilevel structure could be expanded into a multi-level optimization problem involving a third player, such as the government or health insurers, who set the global budget constraints under which the hospital manager operates. Furthermore, while this thesis focused on non-cooperative competition, future investigations could employ cooperative game theory concepts, such as the Shapley value, to model how stakeholders might negotiate binding agreements to share resources to jointly improve outcomes.

To improve the quality and speed of the equilibrium approximation, future studies should explore advanced solution techniques. Reinforcement Learning (RL) offers a promising alternative for learning the Leader’s optimal strategy in high-dimensional spaces, potentially overcoming the sampling limitations of MCMO. Enhancing the computational scalability of the framework is a necessary step to apply these models to entire hospital networks rather than single departments.

Finally, the framework developed here is highly generalizable. Future research should focus on calibrating and validating the model using historical scheduling data and behavioral surveys from Dutch hospitals to fine-tune the payoff functions. Beyond the single hospital, the model can be adapted for regional workforce planning to analyze the competition for labor between different healthcare providers. Moreover, the underlying logic of the Nurse Scheduling Game is applicable to other resource-constrained healthcare challenges, such as the allocation of home-care staff for the elderly or the strategic management of emergency medical services, where the interplay between central planning and decentralized execution is equally critical.

References

- Ahmadi-Javid, A., Seyedi, P., and Syam, S. S. (2017). A survey of healthcare facility location. *Comput. Oper. Res.*, 79:223–263.
- Alalawi, Z., Han, T. A., Zeng, Y., and Elragig, A. (2019). Pathways to good healthcare services and patient satisfaction: An evolutionary game theoretical approach.
- Atoz Serwis Plus (2025). Is the netherlands in need of nurses? <https://www.atozserwisplus.com/blog/is-the-netherlands-in-need-of-nurses>. Accessed: 2025-11-14.
- Authority for Consumers and Markets (ACM) (2016). Competition in the dutch health insurance market. https://www.acm.nl/sites/default/files/old_publication/publicaties/16129_competition-in-the-dutch-health-insurance-market.pdf. Accessed: 2025-11-14.
- AZW (2024). De staat van de arbeidsmarkt zorg en welzijn 2024 - sectoranalyse. <https://www.azwinfo.nl/longread/de-staat-van-de-arbeidsmarkt-zorg-en-welzijn-2024-sectoranalyse/>. Accessed: 2025-11-15.
- Berkhout, M., Smit, K., Sent, D., Kusters, R., Versendaal, J., and van Houwelingen, T. (2025). Understanding the role of clinical decision support systems among hospital nurses using the FITT (fit between individuals, tasks, and technology) framework: Qualitative study. *J. Med. Internet Res.*, 27:e76025.
- Boxebeld, S., Geijssen, T., Tuit, C., van Exel, J., Makady, A., Maes, L., van Agthoven, M., and Mouter, N. (2024). Public preferences for the allocation of societal resources over different healthcare purposes. *Soc. Sci. Med.*, 341(116536):116536.
- Bravo, F., Braun, M., Farias, V., Levi, R., Lynch, C., Tumolo, J., and Whyte, R. (2021). Optimization-driven framework to understand health care network costs and resource allocation. *Health Care Manag. Sci.*, 24(3):640–660.
- Burke, E., De Causmaecker, P., Petrovic, S., and Berghe, G. V. (2003). Variable neighborhood search for nurse rostering problems. In *Applied Optimization*, pages 153–172. Springer US, Boston, MA.
- Burke, E. K., De Causmaecker, P., Vanden Berghe, G., and Van Landeghem, H. (2004). The state of the art of nurse rostering. *Journal of Scheduling*, 7(6):441–499.
- Carter, M. W. and Busby, C. R. (2023). How can operational research make a real difference in healthcare? challenges of implementation. *Eur. J. Oper. Res.*, 306(3):1059–1068.
- Carvalho, M., Dragotto, G., Feijoo, F., Lodi, A., and Sankaranarayanan, S. (2024). When nash meets stackelberg. *Manage. Sci.*, 70(10):7308–7324.
- Carvalho, M., Dragotto, G., Lodi, A., and Sankaranarayanan, S. (2023). Integer programming games: A gentle computational overview.

- Carvalho, M., Lodi, A., and Pedroso, J. P. (2018). Existence of nash equilibria on integer programming games. In *Operational Research*, pages 11–23. Springer International Publishing, Cham.
- Carvalho, M., Lodi, A., and Pedroso, J. P. (2022). Computing equilibria for integer programming games. *Eur. J. Oper. Res.*, 303(3):1057–1070.
- Carvalho, M., Lodi, A., Pedroso, J. P., and Viana, A. (2017). Nash equilibria in the two-player kidney exchange game. *Math. Program.*, 161(1-2):389–417.
- Centraal Bureau voor de Statistiek (CBS) (2022). Nearly 230 thousand fewer hospital admissions in 2020. <https://www.cbs.nl/en-gb/news/2022/24/nearly-230-thousand-fewer-hospital-admissions-in-2020>. Accessed: 2025-11-20.
- Ceschia, S., Dang, N., De Causmaecker, P., Haspeslagh, S., and Schaerf, A. (2019). The second international nurse rostering competition. *Ann. Oper. Res.*, 274(1-2):171–186.
- Ceschia, S., Di Gaspero, L., Mazzaracchio, V., Policante, G., and Schaerf, A. (2023). Solving a real-world nurse rostering problem by simulated annealing. *Oper. Res. Health Care*, 36(100379):100379.
- Cho, L. and Sharkey, T. C. (2023). Integer programming methods to identify nash equilibrium solutions for platform-based scheduling games. *Oper. Res. Forum*, 4(4).
- Colman, A. M., Krockow, E. M., Chattoe-Brown, E., and Tarrant, C. (2019). Medical prescribing and antibiotic resistance: A game-theoretic analysis of a potentially catastrophic social dilemma. *PLoS One*, 14(4):e0215480.
- Coöperatie VGZ (2023). Personeelstekort zorg: oorzaken en oplossingen (personnel shortage in care: causes and solutions). <https://www.cooperatievzg.nl/cooperatie-vgz/zorg/personeelstekort-zorg>. Accessed: 2025-11-14.
- Dakin, H. and Tsiachristas, A. (2024). Rationing in an era of multiple tight constraints: Is cost-utility analysis still fit for purpose? *Appl. Health Econ. Health Policy*, 22(3):315–329.
- De Causmaecker, P. and Vanden Berghe, G. (2011). A categorisation of nurse rostering problems. *J. Sched.*, 14(1):3–16.
- den Hartog, S. J. M., Hoogeveen, H., and van der Zanden, T. C. (2023). On the complexity of nurse rostering problems. *Oper. Res. Lett.*, 51(5):483–487.
- Dragotto, G. (2022). *Mathematical Programming Games*. PhD thesis, Polytechnique Montréal.
- Dragotto, G., Boukhtouta, A., Lodi, A., and Taobane, M. (2024). The critical node game. *J. Comb. Optim.*, 47(5).
- Dragotto, G. and Scatamacchia, R. (2021). The ZERO regrets algorithm: Optimizing over pure nash equilibria via integer programming.

- Duan, J., Lin, Z., and Jiao, F. (2021). A game model for medical service pricing based on the diagnosis related groups. *Front. Public Health*, 9:737788.
- Duguet, A., Carvalho, M., Dragotto, G., and Ngueveu, S. U. (2025). Computing approximate nash equilibria for integer programming games. *Optim. Lett.*
- Earnshaw, S. R. and Dennett, S. L. (2003). Integer/linear mathematical programming models: A tool for allocating healthcare resources. *Pharmacoeconomics*, 21(12):839–851.
- Edwards, R. T., Charles, J. M., Thomas, S., Bishop, J., Cohen, D., Groves, S., Humphreys, C., Howson, H., Bradley, P., and Public Health Wales Health Improvement PBMA team (2014). A national programme budgeting and marginal analysis (PBMA) of health improvement spending across wales: disinvestment and reinvestment across the life course. *BMC Public Health*, 14(1):837.
- Enea, M., Maniscalco, L., de Vries, N., Boone, A., Lavreysen, O., Baranski, K., Miceli, S., Savatteri, A., Mazzucco, W., Fruscione, S., Kowalska, M., de Winter, P., Szemik, S., Godderis, L., and Matranga, D. (2024). Exploring the reasons behind nurses’ intentions to leave their hospital or profession: A cross-sectional survey. *Int. J. Nurs. Stud. Adv.*, 7(100232):100232.
- Erhard, M., Schoenfelder, J., Fügenger, A., and Brunner, J. (2018). State of the art in physician scheduling. *Eur. J. Oper. Res.*, 265(1):1–18.
- Ernst, A. T., Jiang, H., Krishnamoorthy, M., and Sier, D. (2004). Staff scheduling and rostering: A review of applications, methods and models. *Eur. J. Oper. Res.*, 153(1):3–27.
- European Commission (2021). A feasibility study on EU level collaboration on forecasting health workforce needs. https://health.ec.europa.eu/document/download/99be0669-2a82-4c91-97db-dfdfed19908_enhttps://health.ec.europa.eu/document/download/99be0669-2a82-4c91-97db-dfdfed19908_en. Accessed: 2025-11-14.
- Fudenberg, D. and Tirole, J. (1991). *Game Theory*. MIT Press, London, England.
- Fuloria, M. M. and Šormaz, D. N. (2025). Systematic literature review of dynamic resource allocation during pandemic. In *Lecture Notes in Production Engineering*, pages 211–222. Springer Nature Switzerland, Cham.
- Gao, L. and Wang, X. (2019). Healthcare supply chain network coordination through medical insurance strategies with reference price effect. *Int. J. Environ. Res. Public Health*, 16(18):3479.
- Heinrich, T., Jang, Y., Mungo, L., Pangallo, M., Scott, A., Tarbush, B., and Wiese, S. (2023). Best-response dynamics, playing sequences, and convergence to equilibrium in random games. *Int. J. Game Theory*, 52(3):703–735.
- Ho, K. and Lee, R. S. (2017). Equilibrium provider networks: Bargaining and exclusion in health care markets. *SSRN Electron. J.*
- Hu, W., He, X., Luo, L., and Pardalos, P. M. (2023). A branch-and-price approach for the nurse rostering problem with multiple units.

- Hu, Y., Wang, J., Chen, X., and He, N. (2024). Multi-level Monte-Carlo gradient methods for stochastic optimization with biased oracles.
- Humphreys, P., Spratt, B., Tariverdi, M., Burdett, R. L., Cook, D., Yarlagadda, P. K. D. V., and Corry, P. (2022). An overview of hospital capacity planning and optimisation. *Healthcare (Basel)*, 10(5):826.
- Inspectie Gezondheidszorg en Jeugd (IGJ) (2022). Personeelstekorten in de zorg. <https://www.igj.nl/onderwerpen/themas-in-het-toezicht/inzet-personeel-in-de-zorg/personeelstekort>. Accessed: 2025-11-15.
- Jacobs, S. (2024). Budget cuts to worsen healthcare worker shortage in the netherlands. <https://www.iamexpat.nl/expat-info/dutch-news/budget-cuts-worsen-healthcare-worker-shortage-netherlands>. Accessed: 2025-11-14.
- Jiang, H., Pang, Z., and Savin, S. (2020). Performance incentives and competition in health care markets. *Prod. Oper. Manag.*, 29(5):1145–1164.
- Kheiri, A., Gretsista, A., Keedwell, E., Lulli, G., Epitropakis, M., and Burke, E. (2021). A hyper-heuristic approach based upon a hidden markov model for the multi-stage nurse rostering problem. *Comput. Oper. Res.*, 130:105221.
- Koirala, P., Krusniak, M., and Laine, F. (2025). Locally optimal solutions for integer programming games.
- Koirala, P. and Laine, F. (2023). Monte carlo optimization for solving multilevel stackelberg games.
- Köppe, M., Ryan, C. T., and Queyranne, M. (2011). Rational generating functions and integer programming games. *Oper. Res.*, 59(6):1445–1460.
- Kukushkin, N. S. (2004). Best response dynamics in finite games with additive aggregation. *Games Econ. Behav.*, 48(1):94–110.
- Latruwe, T., Van der Wee, M., Vanleenhove, P., Devriese, J., Verbrugge, S., and Colle, D. (2023). A long-term forecasting and simulation model for strategic planning of hospital bed capacity. *Oper. Res. Health Care*, 36(100375):100375.
- Li, C. and Guo, L. (2017). A single-level reformulation of mixed integer bilevel programming problems. *Oper. Res. Lett.*, 45(1):1–5.
- McHugh, M. D., Aiken, L. H., Sloane, D. M., Windsor, C., Douglas, C., and Yates, P. (2021). Effects of nurse-to-patient ratio legislation on nurse staffing and patient mortality, readmissions, and length of stay: a prospective study in a panel of hospitals. *Lancet*, 397(10288):1905–1913.
- Mitton, C., Dionne, F., and Donaldson, C. (2014). Managing healthcare budgets in times of austerity: the role of program budgeting and marginal analysis. *Appl. Health Econ. Health Policy*, 12(2):95–102.

- Mitton, C. and Donaldson, C. (2003). Setting priorities and allocating resources in health regions: lessons from a project evaluating program budgeting and marginal analysis (PBMA). *Health Policy*, 64(3):335–348.
- Moore, J. T. and Bard, J. F. (1990). The mixed integer linear bilevel programming problem. *Oper. Res.*, 38(5):911–921.
- Myerson, R. B. (1991). *Game theory: analysis of conflict*. Harvard University Press, London, England.
- Mystakidis, A., Koukaras, C., Koukaras, P., Kaparis, K., Stavrinides, S. G., and Tjortjis, C. (2024). Optimizing nurse rostering: A case study using integer programming to enhance operational efficiency and care quality. *Healthcare (Basel)*, 12(24):2545.
- Nash, J. (1951). Non-Cooperative games. *Ann. Math.*, 54(2):286.
- Nash, J. F. (1950). Equilibrium points in n-person games. *Proc. Natl. Acad. Sci. U. S. A.*, 36(1):48–49.
- Netherlands Bureau for Economic Policy Analysis (CPB) (2003). Concern about competition: an analysis of the new health care system. <https://www.cpb.nl/en/publication/concern-about-competition-analysis-new-health-care-system>. Accessed: 2025-11-14.
- Netherlands Enterprise Agency, RVO (2025). CAO (collective labour agreement). <https://business.gov.nl/regulation/cao/>. Accessed: 2025-11-20.
- NL Times (2023). Cancellation of STAP budget will affect healthcare workers in particular. <https://nltimes.nl/2023/04/28/cancellation-stap-budget-will-affect-healthcare-workers-particular>. Accessed: 2025-11-14.
- Ordu, M., Demir, E., Tofallis, C., and Gunal, M. M. (2021). A novel healthcare resource allocation decision support tool: A forecasting-simulation-optimization approach. *J. Oper. Res. Soc.*, 72(3):485–500.
- Osborne, M. J. and Rubinstein, A. (1994). A course in game theory. <https://mitpress.mit.edu/9780262150415/a-course-in-game-theory/>. Accessed: 2025-12-22.
- Palmer, R., Fulop, N. J., and Utley, M. (2018). A systematic literature review of operational research methods for modelling patient flow and outcomes within community healthcare and other settings. *Health Syst. (Basingstoke)*, 7(1):29–50.
- Peacock, S. J., Mitton, C., Ruta, D., Donaldson, C., Bate, A., and Hedden, L. (2010). Priority setting in healthcare: towards guidelines for the program budgeting and marginal analysis framework. *Expert Rev. Pharmacoecon. Outcomes Res.*, 10(5):539–552.
- Petrovic, S. and Vanden Berghe, G. (2012). A comparison of two approaches to nurse rostering problems. *Ann. Oper. Res.*, 194(1):365–384.

- Rachuba, S., Reuter-Oppermann, M., and Thielen, C. (2024). Integrated planning in hospitals: a review. *OR Spectr.*
- Rais, A. and Viana, A. (2011). Operations research in healthcare: A survey. *Int. Trans. Oper. Res.*, 18(1):1–31.
- Renggli, F. J., Gerlach, M., Bieri, J. S., Golz, C., and Sariyar, M. (2025). Integrating nurse preferences into AI-based scheduling systems: Qualitative study. *JMIR Form. Res.*, 9:e67747.
- Rittel, H. W. J. and Webber, M. M. (1973). Dilemmas in a general theory of planning. *Policy Sci.*, 4(2):155–169.
- Russo, S., Zhitikhin, S., Gulino, V., Ricci, B., Nigro, M., Gallerani, E., Lombardo, E., Perger, P., Padovani, E., Campagna, A., and Buccioli, M. (2025). Developing a predictive model for resource allocation in healthcare: A case study from an italian hospital. *SSM Health Syst.*, 5(100085):100085.
- Salvioli, M., Garjani, H., Satouri, M., Broom, M., Viossat, Y., Brown, J. S., Dubbeldam, J., and Staňková, K. (2025). Stackelberg evolutionary games of cancer treatment: What treatment strategy to choose if cancer can be stabilized? *Dyn. Games Appl.*, 15(5):1750–1769.
- Schäfer, W., Kroneman, M., Boerma, W., van den Berg, M., Westert, G., Devillé, W., and van Ginneken, E. (2010). The netherlands: health system review. *Health Syst. Transit.*, 12(1):v–xxvii, 1–228.
- Seixas, B. V. and Mitton, C. (2021). Using a formal strategy of priority setting to mitigate austerity effects through gains in value: The role of program budgeting and marginal analysis (PBMA) in the brazilian public healthcare system. *Appl. Health Econ. Health Policy*, 19(1):9–15.
- Selten, R. (1975). Reexamination of the perfectness concept for equilibrium points in extensive games. *Int. J. Game Theory*, 4(1):25–55.
- Sinha, A., Malo, P., and Deb, K. (2018). A review on bilevel optimization: From classical to evolutionary approaches and applications. *IEEE Trans. Evol. Comput.*, 22(2):276–295.
- Somasundaram, M. and Sivakumar, R. (2015). Game theory based security in wireless body area network with stackelberg security equilibrium. *ScientificWorldJournal*, 2015(1):174512.
- Stanková, K., Brown, J. S., Dalton, W. S., and Gatenby, R. A. (2019). Optimizing cancer treatment using game theory: A review. *JAMA Oncol.*, 5(1):96–103.
- Stein, A., Salvioli, M., Garjani, H., Dubbeldam, J., Viossat, Y., Brown, J. S., and Staňková, K. (2023). Stackelberg evolutionary game theory: how to manage evolving systems. *Philos. Trans. R. Soc. Lond. B Biol. Sci.*, 378(1876):20210495.
- Swenson, B. W. (2018). Myopic best-response learning in large-scale games.
- The Scientific Council for Government Policy (WRR) (2021). Kiezen voor houdbare zorg. mensen, middelen en maatschappelijk draagvlak. Technical report.

- TheGlobalEconomy (2019). Netherlands nurse to hospital bed ratio. https://www.theglobaleconomy.com/Netherlands/nurse_to_hospital_bed_ratio/. Accessed: 2025-11-20.
- TNO (2024). TNO-onderzoek: Meer dan 50% werkgevers ervaart of verwacht personeelstekorten (TNO research: More than 50% of employers experience or expect staff shortages). <https://www.tno.nl/nl/newsroom/2024/11/werkgevers-ervaren-personeelstekorten/>. Accessed: 2025-11-14.
- van Liempt, T. and Brussaard, A. (2025). Integraal capaciteitsmanagement: grip op zorg van morg. <https://www.berenschot.nl/artikelen/integraal-capaciteitsmanagement-grip-op-zorg-van-morgen>. Accessed: 2025-11-15.
- Varkevisser, M., Schut, E., Franken, F., and van der Geest, S. (2023). Sustainability and resilience in the dutch health system. https://www3.weforum.org/docs/WEF_PHSSR_Netherlands_Report_2023.pdf. Accessed: 2025-11-14.
- von Neumann, J. and Morgenstern, O. (1944). *Theory of games and economic behavior*. Princeton University Press, Princeton, NJ.
- von Stackelberg, H. (1934). *Market Structure and Equilibrium*. Springer, Berlin, Germany, 2011 edition.
- Wagner, C. E., Prentice, J. A., Saad-Roy, C. M., Yang, L., Grenfell, B. T., Levin, S. A., and Laxminarayan, R. (2020). Economic and behavioral influencers of vaccination and antimicrobial use. *Front. Public Health*, 8:614113.
- Wikipedia (2025). Jeroen bosch hospital. https://en.wikipedia.org/w/index.php?title=Jeroen_Bosch_Hospital&oldid=1277779476. Accessed: NA-NA-NA.
- Xiao, Y. and Yoogalingam, R. (2022). A simulation optimization approach for planning and scheduling in operating rooms for elective and urgent surgeries. *Oper. Res. Health Care*, 35(100366):100366.
- Yang, S., Zankin, V., Balandat, M., Scherer, S., Carlberg, K., Walton, N., and Law, K. J. H. (2024). Accelerating look-ahead in bayesian optimization: Multilevel monte carlo is all you need.
- Zorgakkoorden.nl (2022). Integraal zorgakkoord (IZA). <https://www.zorgakkoorden.nl/programmas/integraal-zorgakkoord/>. Accessed: 2025-11-14.
- Zorginstituut Nederland (2024). Guideline for economic evaluations in healthcare. Technical report.
- Zurhake, S. (2022). Komt er een zorgakkoord? steun van onmisbare huisartsen is niet zeker. <https://nos.nl/artikel/2444224-komt-er-een-zorgakkoord-steun-van->. Accessed: 2025-12-21.