Aalto University
MS-E2177 - Seminar on Case Studies in Operations Research

Final report: Optimisation models for long-term workforce planning

Group 6: SOK

Adrian Alfatoni
Anna-Maija Kangaslahti
Iina Manninen (Project Manager)
Salla Nicholls

May 31, 2026

Contents

Notice on AI usage

Grammarly was used for checking grammatical errors. In addition, ChatGPT and Claude Sonnet 4.6 were used to rephrase some text.

Alphabetical list of symbols

$\mathrm{a}_{w,c_\mathrm{f}}$	Number of absences of fixed-term employees with contract hours c_f in week w
$\mathrm{a}_{w,c_\mathrm{p}}$	Number of absences of permanent employees with contract hours c_p in week w
ca_{c_1,c_2}	Permission to increase hours from c_1 to c_2
$\mathrm{cc}_{\mathrm{lower}}$	Minimum of average contract hour coverage
$\mathrm{cc}_{\mathrm{upper}}$	Maximum of average contract hour coverage
$\mathrm{ch}_{\mathrm{w},c_1,c_2}$	Number of contract hour changes from contract type c_1 to c_2 in week w
cont_w	Reserve of contract hours in week w
$\mathrm{cost}_{\mathrm{change}}$	Cost of an hour increase
$\mathrm{cost}_{\mathrm{fixed\text{-}term}}$	Weekly cost of a fixed-term employee
$\mathrm{cost}_{\mathrm{permanent}}$	Weekly cost of a permanent employee
$\mathrm{cost}_{\mathrm{recruit,fixed\text{-}term}}$	Cost of recruiting a fixed-term employee
$\mathrm{cost}_{\mathrm{recruit,permanent}}$	Cost of recruiting a permanent employee
$\mathrm{efe}_{w,c_\mathrm{f},d}$	Number of fixed-term employees with contract type c_f and duration d in week w
$\mathrm{eh}_{\mathrm{max}}$	Maximum amount of extra hours the employees are allowed to work per balancing period
$\mathrm{emplo}_{\mathrm{min},w}$	Required number of workers at week w
$\mathrm{epe}_{w,c_\mathrm{p}}$	Number of permanent employees with contract type c_p in week w
$\mathrm{fh}_{\mathrm{max},w}$	Maximum number of fixed-term hires at week w
flex_w	Reserve of flexibility in week w
fu_w	Flexibility used in week w
$\mathrm{f}_{c_\mathrm{f}}$	Weekly flexibility of a fixed-term employee with contract hours c_f in week w
$\mathrm{f}_{c_\mathrm{p}}$	Weekly flexibility of a permanent employee with contract hours c_p in week w
$\mathrm{hours}_{\mathrm{r,max}}$	Maximum rental work hours per person at week w
$\mathrm{hours}_{\mathrm{r,min}}$	Minimum rental work hours per person at week w
$\mathrm{h}_{c_\mathrm{f}}$	Weekly hours of a fixed-term employee with contract hours c_f in week w
$\mathrm{h}_{c_\mathrm{p}}$	Weekly hours of a permanent employee with contract hours c_p in week w
max_{c_f}	Maximum number of fixed-term employees with contract type c_f on any given week
max_{c_p}	Maximum number of permanent employees with contract type c_p on any given week

min_{c_p}	Minimum number of permanent employees with contract type c_p on any given week
$\text{nhf}_{w,c_\text{f},d}$	Number of recruited fixed-term employees with contract type c_f and duration d in week w
$\text{nhp}_{w,c_\text{p}}$	Number of recruited permanent employees with contract type c_p in week w
n_min	Minimum number of full-time employees that must be available
$\text{qe}_{\text{min},c_p}$	Global minimum number of the predetermined employees of contract type c_p
qe_{w,c_f}	Number of predetermined fixed-term employees with contract type c_f and duration d in week w
qe_{w,c_p}	Number of predetermined permanent employees with contract type c_p in week w
rh_w	Rental work hours at week w
rw_w	Rental workers at week w
r_max	Maximum amount of rental work hours allowed per balancing period
share_max	Maximum share of particular contracts on any given week
share_min	Minimum share of particular contracts on any given week
$\text{share}_{\text{p},w}$	Required share of permanent workers at week w
td_w	Total demand in week w
wage_ren	Hourly cost of rental work
wd_w	Work demand in week w
$\tilde{C}_\text{f}$	Set of fixed-term contract hours that are available for new recruitments
$\tilde{C}_\text{p}$	Set of permanent contract hours that are available for new recruitments
B	Set of balancing periods
C_f	Set of fixed-term contract types
C_p	Set of permanent contract types
D	Set of fixed-term contract durations
W	Set of weeks

1 Introduction

1.1 Background

Suomen Osuuskauppojen Keskuskunta (SOK) is the central organisation of S Group, whose member cooperatives are among the leaders in the Finnish retail market, particularly in grocery retail and in the hospitality sector through hotels and restaurants. In the context of SOK's business operations, human resource planning is crucial to ensuring that the available workforce meets the business's demand [1]. A sufficient workforce is particularly relevant where demand varies seasonally, and the workforce supply needs to be adjusted accordingly to meet that variation. Retail is a human-resource-intensive industry in which strategic long-term planning enables a shift from reactive to proactive operations. However, the longer the planning horizon, the greater the uncertainty in demand forecasting and personnel scheduling.

Decisions on recruitment planning are useful not only for meeting the business's seasonal demands but also for reducing the additional costs associated with overstaffing and understaffing. Overstaffing occurs when a business unit has a surplus of employees relative to demand, resulting in unnecessary costs that are slow to eliminate. On the other hand, understaffing occurs when a business has too few employees to meet demand, leading to resource constraints, overburdening employees, undermining productivity and well-being, and making it difficult to run the business as expected. Thus, it is in the best interests of SOK and in line with the objective of this project to optimise recruitment planning to reduce personnel costs, ensure employee well-being, and operate business units effectively.

To optimise recruitment planning, it is helpful to employ a mathematical optimisation model, such as mixed-integer linear programming (MILP), based on demand forecasts for each business unit. An MILP seems appropriate due to its ability to optimise hiring decisions under a finite budget and capacity constraints, while handling integer decisions such as employee recruitment and changing their contract types. It is crucial to account that the adjustments and decisions made with the model adhere to current employment regulations on working time and wages, such as the commerce sector collective agreement [2].

1.2 Objectives

The overall objective of this project is to develop an optimisation tool for long-term workforce planning over an 18-24 month time horizon, updated weekly. This tool aims to reduce both over- and understaffing. The optimisation tool must account for known future workloads, absences, and fixed personnel changes. Initially, the project also aimed to build an additional model on top of the first one to account for uncertainty related to personnel changes, etc., by incorporating randomness. However, as the project progressed, we redefined the scope together with SOK representatives to focus on identifying literature-based methods for including stochasticity into the model.

The optimisation tool is required to provide recommendations on which recruitments or increases in part-time employees' contract hours to make, and when. The aim is to minimise expected personnel costs resulting from existing employees, new recruitments and rental work. The constraints taken into account include future workloads, absences, fixed personnel changes, part-time employees' flexibility for extra work, balancing periods, and business units' constraints on the mix of contract hours. In addition, the optimisation solutions must comply with the requirements set by law and collective bargaining.

The optimisation tool should be usable by business units without optimisation expertise. In addition, the optimisation model should be flexible enough to account for minor differences in needs and practices across business units.

2 Literature review

2.1 Human resource management and planning

Human resource management encompasses a broad set of organisational functions, such as recruitment, training, development, and employee relations [1]. Managing human resources is critical for an organisation's productivity, performance, reducing staffing costs and aligning employee efforts with the company's strategy. In addition to efficiency objectives, human resource management also addresses employees' welfare and interests [3]. Moreover, it is essential for maintaining occupational health and safety standards, ensuring compliance with labour legislation, and facilitating organisational change.

Human resource planning supports human resource management by determining the right number of people for the right job at the right time and place [1]. It is the process of forecasting an organisation's future demand and workforce supply. The planning horizon ranges from short-term operational planning that addresses immediate, day-to-day staffing requirements to long-term strategic planning that aligns workforce capabilities with future business objectives. Especially in labour-intensive organisations, recruitment and effective employee selection practices are necessary [1]. Industries that rely heavily on human resources include retail, transportation, healthcare, manufacturing, etc.[4].

Human resource planning has to account for labour legislation and collective agreements to ensure legal compliance and fulfil employer obligations. Furthermore, modern human resource planning increasingly integrates technology, such as ergonomic shift-scheduling tools examined by Karhula et al. [5], to transcend basic compliance by embedding evidence-based well-being recommendations into daily workforce management. Beyond constraints imposed by regulations and employee availability, human resource planning must align with organisational strategy and policies.

2.2 Optimisation models for human resource planning

Human resource planning can be time-consuming, and quantitative methods are increasingly utilised to support staffing and scheduling decisions. Ernst et al. [4] provide an extensive bibliography of personnel scheduling reviewing articles across diverse sectors. Models vary considerably across industries due to differing operational requirements [4]. All models represent simplifications of real-world systems and therefore cannot account for all relevant factors. However, they remain highly valuable decision-support tools. The requisite level of accuracy depends on several factors, including the planning horizon, industry-specific characteristics, and demand variability. Different stages of the scheduling process include demand modelling, days-off scheduling, shift scheduling, task assignment, and staff assignment [4].

In the literature, the objective of optimisation models for human resource planning is often to minimise costs, though alternative formulations may prioritise, for example, workforce composition. Llort et al. [6] examine optimisation models in workforce planning within consultancies, where the composition of highly educated professionals substantially influences organisational competitiveness. Li et al. [7] apply multiple objective linear programming (MOLP) to minimise the shortfall of professional development and overtime. Different industries are also characterised by varying hard and soft constraints. In safety-critical industries, meeting demand requirements is more important than in the service sector. Planning decisions encompass, for example, hiring, firing, training, holiday allocation, and promotions [6]. Other factors that influence workforce capacity include sick leave, part-time contracts, and retirements.

The planning horizon, workforce size, differing contract types, and skills affect whether the planning focus is on aggregate staffing levels or individual scheduling. In strategic long-term planning, category-level planning is sufficient, especially when individual preferences need not be taken into account. For instance, Llort et al. [6] model the proportions of consultants working in certain categories, business lines, and industries. Absences are also modelled as average proportions of capacity loss, as individual absences are unpredictable at the aggregate level [6]. However, when specific shifts must be assigned to specific individuals, employee-level optimisation becomes necessary [8]. Even though Talarico & Maya-Duque [8] model full-time and part-time employees at the individual level when determining the ideal weekly mix of different contract types in a chain of supermarkets, the contract-hour categories could alternatively be modelled as groups, with all employees within each group assumed homogeneous.

Li et al. [7] state that, despite differences in time horizons, both aggregate-level staff sizing and shift scheduling should be coordinated because they are interdependent. They argue that significant overstaffing often occurs in service organisations as a result of separate, independent decisions regarding long-term workforce planning and short-term staff assignment [7]. Therefore, they introduce an approach that recursively improves workforce planning and shift scheduling, so that the solution at one stage can be fed back to revise the solution at the other stage [7].

To respond to seasonal demand, the retail sector requires different contract types, lengths, and staff flexibility [9]. In addition, retail stores have extended opening hours that far exceed standard employee working hours, and therefore, there is considerable use of flexible types of employment contracts [8]. Several models also capture planned days off [7] and absences, such as sick leave [6]. Corominas et al. [9] highlight that in the service sector, the product is not storable, necessitating precise workforce scheduling to match demand fluctuations. Some studies utilise annualised working hours, which allow the distribution of the total number of working hours irregularly throughout the year to achieve flexibility [9]. Employees may accept overtime, but prolonged working hours can lead to fatigue and reduced service quality [7].

In turn, optimisation models can have constraints that capture employee preferences to address work-life balance. The preferences can concern vacations, days off, preferred tasks, or coworkers [10]. İşeri et al. also incorporate fairness considerations to ensure that employee preferences are given equal priority across the workforce. Considering employee preferences enhances organisational loyalty and commitment [10] and can improve productivity [11].

Several models consider employee skills, training and development. Li et al. [7] state that matching the labour skills with job requirements is crucial, especially when personnel are highly specialised. Depending on the role's characteristics, learning periods can be lengthy, and considerable training costs may arise [6]. Li et al. [7] incorporate a constraint that some employment time must be dedicated to professional development. They also seek to minimise the unachieved professional development time [7]. Moreover, they account for the different costs associated with employees at various skill levels. In some approaches, higher-qualified workers can substitute for lower-qualified ones [10]. The ability to apply one's skills is associated with greater job satisfaction [10]. However, Karimi et al. [11] emphasise that there are trade-offs between immediate costs and the long-term benefits of integrating training decisions into workforce planning.

2.3 Modelling uncertainty in planning

Personnel scheduling in the retail sector is highly dynamic due to sales fluctuations [8]. Additionally, there are other unpredictable phenomena such as employee absenteeism [12]. Annualised hours, multiskilled employees, and overtime are among the main labour flexibility strategies [12]. In mathematical modelling, strategies for incorporating uncertainty into the formulation include robust optimisation, two-stage stochastic optimisation, and simulation-based approaches [12].

Porto et al. [12] incorporate demand uncertainty through a two-stage stochastic optimisation (TSSO) model, where staffing and training decisions are made before demand realisation, while working hour allocation adapts to specific demand scenarios. Rather than optimising for a single demand instance, the model minimises the expected average cost across multiple scenarios generated via Monte Carlo simulation from a truncated normal distribution. The stochastic approach yields a single robust workforce configuration that does not necessarily perform optimally for all scenarios but remains feasible for a variety of demand realisations [12].

In his thesis, Wilhelmsson [13] evaluates contract structures under uncertainty using stochastic agent-based simulations, in which employees probabilistically become ill and decide whether to accept additional shifts. The model uses parameters including the probability that an employee falls sick, the distribution of sick leave durations, and the probability that an employee accepts a vacant shift offer. Cost and flexibility metrics are used to assess the results.

Karimi et al. [11] introduce a mixed-integer linear programming resource selection model that accounts for uncertainty in technician availability. They aim to minimise costs while simultaneously maximising fairness in training assignment under uncertain conditions. They apply robust optimisation methods to ensure feasibility across different worker-absenteeism scenarios. Additionally, they incorporate various fairness scenarios where either individual fairness or collective fairness, balancing fairness across workers, is maximised. They utilise a scenario tree structure that captures multiple possible combinations of technician availability. Lliort et al. [6] use scenarios to analyse varying demand realisations and promotion policies that influence workforce planning.

3 Methods & model

In this section, we shall discuss the foundations and assumptions of our resulting model, as well as its mathematical formulation.

3.1 Mixed-integer linear programming

In this project, the objective is to determine the optimal recruitment plan while satisfying multiple constraints. Mixed-Integer Linear Programming (MILP) was selected as the primary modelling approach as it is a common approach in the related literature and also suggested by the client. MILP is a mathematical optimisation method in which the objective function and constraints are represented in linear form, while some decision variables are restricted to integer values. The method is widely used in operations research and workforce planning because it can effectively model decision-making problems involving both continuous and discrete variables under multiple constraints. The general form of a MILP model can be expressed as

$$\begin{aligned} \min \quad & c^T x \\ \text{subject to} \quad & Ax \leq b \\ & x_i \in \mathbb{Z}, \quad i \in I, \end{aligned} \tag{1}$$

where x represents the decision variables, $c^T x$ denotes the objective function, and $Ax \leq b$ represents the set of constraints. The set I contains the variables that are restricted to integer values.

3.2 Description and assumptions of the problem

Strategic workforce planning aims to align workforce supply with fluctuating, season-dependent demand, enabling timely recruitment decisions while minimising labour costs and maintaining operational efficiency. The objective is to minimise expected labour costs by scheduling employer decisions to increase total weekly contract hours and the number of weekly recruitments, ensuring sufficient resources to fulfil the forecasted demand in all weeks. Additionally, there is a possibility of using rental work to cover peak demand.

The planning horizon is 18–24 months, but the user must be able to choose the time period under study. Due to the long timeframe, the modelling precision is one week. For instance, it is assumed that if an employee is absent, they are not present during the whole week. Balancing periods are approximately 18-week periods during which employers must compensate an employee if they have not provided the working hours specified in the employee's contract. If an employee is recruited in the middle of a balancing period, the hours are calculated for the period they have been employed. The contract hour supply must be between lower and upper limits, expressed as percentages, during a balancing period.

Workforce units range from approximately 30 to 60 employees, with a maximum of 200. They can have a permanent or fixed-term contracts. The user can choose which contract hours are allowed. There are at most ten different possibilities for contract hours. Full-time contacts (37.5 hours per week) are always included. In reality, there could be more than ten different contract hours in use. Fixed-term contract lengths vary from 3 to 20 weeks. In addition to contract hours, there is flexibility associated with each contract type. Flexibility refers to the employee's willingness to work extra hours during one week. It is assumed that each employee with a fixed-term contract with the same contract hours has the same flexibility, and the same holds for permanent employees. Full-time employees do not have flexibility. It is possible to constrain the minimum number of employees, the number and share of different contract types, or the number of recruitments in a given week.

The number of employees with different weekly contract hours fluctuates due to known personnel changes such as retirements, the start and end of long absences, and known changes in weekly contract hours, as well as the employer's decisions regarding recruitment, increases in existing contract hours, or hiring for rental work. It is assumed that there is available additional workforce whenever the employer decides to recruit new employees, and that employees are willing to increase their contract hours. Employees cannot be fired, and their hours cannot be decreased. Fixed-term employees cannot be changed to permanent employees, or the changes are modelled as new recruitments. In reality, there are also random changes from an employer's perspective. These fluctuations include resignations that unexpectedly decrease the number of staff. However, this model does not account for uncertainty but assumes future workloads, absences and personnel changes as known parameters.

There are different costs associated with the decisions. Recruitment costs differ for permanent and fixed-term employees. Contract hour increase has a small cost. Rental work is considered the last option, and it has a higher cost. In addition, there are contract-hour-independent fixed costs due to, for instance, occupational healthcare and work clothing. This cost is slightly lower for fixed-term employees than for permanent employees. Apart from differences in the costs of permanent and fixed-term employees, it is assumed that there are no other cost differences, even though, in a more realistic model, wage differences could arise from factors such as work experience, job roles, and responsibilities.

Skills, skill levels or training are not considered in this model. It is assumed that employee training entails negligible cost and time requirements. Therefore, different departments in the store are not taken into account. However, there are constraints on the minimum number of employees per week, as a certain number is required to perform all relevant tasks.

Since the planning horizon is long, which increases uncertainty, the model operates on aggregated contract types rather than individual-level data. Hence, no assumptions are made regarding an individual employee's willingness to transition to higher contract hours. However, the current approach only accounts for the absences of employees already working in the unit, and missing hours are estimated without considering potential increases in hours.

Input data includes the predicted week-level work demand, the contract hours and contract types of employees already working in the unit, flexibilities and absences, vacation plans, and known changes to contracts. As an output, the model provides the number of recruitments and contract hours increases per contract hour and type for each week during the planning period.

3.3 Parameters

The set of weeks consists of the planning period's weeks. For instance, for a 24-month period, the set of weeks is $W = \{1, ..., 104\}$. The set of balancing periods B contains approximately 18-week periods during which employees' actual hours should be close to their contract hours for the same period. C_{p}denotes permanent contract hour options. $\tilde{C}_{\mathrm{p}}$ represents the set of contract-hour options available for new recruitments and hour increases. These sets are separate because the original set of contract hours may contain options that are not available for change. Similarly, for the fixed-term contract hours options, the notation C_{f} is used, and $\tilde{C}_{\mathrm{f}}$ represents the fixed-term contract hour options for new recruitments. D represents the set of allowed fixed-term contract durations. When 3–20 week durations are allowed, the set of durations is $D = \{3, ..., 20\}$. All the sets are defined in Table 1.

Table 1: List of sets.

Sets	Definition
$W \subseteq \mathbb{N}$	Set of weeks
B	Set of balancing periods
C_{p}	Set of permanent contract hours
$\tilde{C}_{\mathrm{p}} \subseteq C_{\mathrm{p}}$	Set of permanent contract hours that are available for new recruitments
C_{f}	Set of fixed-term contract hours
$\tilde{C}_{\mathrm{f}} \subseteq C_{\mathrm{f}}$	Set of fixed-term contract hours that are available for new recruitments
$D \subseteq \mathbb{N}$	Set of fixed-term contract durations

The model's parameters relate to contracts, work demand, costs, and constraints that, for example, limit the number of recruits in a given week. Variables $\mathrm{qe}_{w,c_{\mathrm{p}}}$ and $\mathrm{qe}_{w,c_{\mathrm{f}}}$ denote the predetermined number of employees for each contract hour and length category in the input data. The parameter $\mathrm{qe}_{w,c_{\mathrm{p}}}$ is constructed directly from the input data by first counting, for each contract type c_p, the number of permanent employees contracted to start on or before each week w. Then we count, from another table, how many predetermined hour changes, retirements or terminations of each contract type happen in each week w and adjust the previous result accordingly. For example, if the data shows that the hours for one employee with a 10-hour contract have been increased to 25 hours a week but only for weeks 2-4, we add one to all of the base values of $\mathrm{qe}_{w,25}$ and reduce one from the base values of $\mathrm{qe}_{w,10}$ on the corresponding weeks $w \in \{2, 3, 4\}$. This system allows us to record known changes, such as decreases in the hours or retirements, even though our model must not consider similar changes. Similarly, we construct the parameters $\mathrm{qe}_{w,c_{\mathrm{f}}}$ by counting, from the input data, the number of fixed-term employees contacted to work on week w, separately for each contract type. Other parameters are explained in detail in Table 2.

Table 2: List of parameters.

Parameters	Definition
Given contracts	
$\mathrm{qe}_{w,c_{\mathrm{p}}}$	Number of predetermined permanent employees with contract type c_p in week w
$\mathrm{qe}_{w,c_{\mathrm{f}}}$	Number of predetermined fixed-term employees with contract type c_f and duration d in week w
Contract hours	
$\mathrm{h}_{c_{\mathrm{p}}}$	Weekly hours of a permanent employee with contract hours c_p in week w
$\mathrm{h}_{c_{\mathrm{f}}}$	Weekly hours of a fixed-term employee with contract hours c_f in week w
Flexibility	

(continued on next page)

(continued)

Parameters	**Definition**
f_{c_p}	Weekly flexibility of a permanent employee with contract hours c_p in week w
f_{c_f}	Weekly flexibility of a fixed-term employee with contract hours c_f in week w
Allowed increases	
ca_{c_1,c_2}	Permission to increase hours from c_1 to c_2
Absences	
a_{w,c_p}	Number of absences of fixed-term employees with contract hours c_p in week w
a_{w,c_f}	Number of absences of fixed-term employees with contract hours c_f in week w
Work demand	
wd_w	Work demand estimate on week w
Costs	
$cost_{permanent}$	Weekly cost of a permanent employee
$cost_{fixed\text{-}term}$	Weekly cost of a fixed-term employee
$cost_{recruit,permanent}$	Cost of recruiting a permanent employee
$cost_{recruit,fixed\text{-}term}$	Cost of recruiting a fixed-term employee
$cost_{change}$	Cost of an hour increase
$wage_{ren}$	Hourly cost of rental work
Contract coverage	
cc_{lower}	Minimum of average contract hour coverage
cc_{upper}	Maximum of average contract hour coverage
Min/Max etc.	
$fh_{max,w}$	Maximum number of fixed-term hires at week w
$share_{p,w}$	Required share of permanent workers at week w
n_{min}	Minimum number of full-time employees that must be available
eh_{max}	Maximum amount of extra hours the employees are allowed to work per balancing period
$hours_{r,max}$	Maximum rental work hours per person at week w
$hours_{r,min}$	Minimum rental work hours per person at week w
qe_{min,c_p}	Global minimum number of the predetermined employees of contract type c_p
max_{c_p}	Maximum number of permanent employees with contract type c_p on any given week
min_{c_p}	Minimum number of permanent employees with contract type c_p on any given week
max_{c_f}	Maximum number of fixed-term employees with contract type c_f on any given week
$share_{max}$	Maximum share of particular contracts on any given week

(continued on next page)

(continued)

Parameters	Definition
$\mathrm{share}_{\mathrm{min}}$	Minimum share of particular contracts on any given week

3.4 Decision variables

The decisions regarding new recruitments are represented by variable $\mathrm{nhp}_{w,c_\mathrm{p}}$, denoting the number of permanent employees with contract type c_p recruited in week w, and variable $\mathrm{nhf}_{w,c_\mathrm{f},d}$, denoting the number of fixed-term employees with contract type c_f and duration d recruited in week w. Contract hour increases are captured by variable $\mathrm{ch}_{\mathrm{w},c_1,c_2}$, representing the number of contract hour changes from contract type c_1 to c_2 in week w. Contract hour changes are only allowed for permanent contracts $c_1, c_2 \in C_\mathrm{p}$. The employer also makes decisions on both hours covered by rental work in week w, denoted by rh_w, and the number of rental workers in week w, denoted by rw_w. The number of rental workers must be modelled as a separate variable, as it affects the total number of employees in the unit for a given week. Even if fewer workers cover the hours, there might be tasks that require a certain number of workers to be present. The decision variables are summarised in Table 3.

Table 3: List of decision variables.

Decision variable	Definition
$\mathrm{nhp}_{w,c_\mathrm{p}} \in \mathbb{N}$	Number of recruited permanent employees with contract type c_p in week w
$\mathrm{nhf}_{w,c_\mathrm{f},d} \in \mathbb{N}$	Number of recruited fixed-term employees with contract type c_f and duration d in week w
$\mathrm{ch}_{\mathrm{w},c_1,c_2} \in \mathbb{N}$	Number of contract hour changes from contract type c_1 to c_2 in week w
$\mathrm{rh}_w \in \mathbb{R}_+$	Rental work hours in week w
$\mathrm{rw}_w \in \mathbb{N}$	Rental workers hired in week w

3.5 Helper quantities

To simplify the notation, we define several helper quantities, which are described in this section.

One of the most fundamental of these helper quantities is $\mathrm{epe}_{w,c_\mathrm{p}}$ defined in (2) below. This helper quantity represents the number of existing permanent employees with c_p-hour contract, employed on week w: we derive $\mathrm{epe}_{w,c_\mathrm{p}}$ by summing the number of predetermined employees $\mathrm{qe}_{w,c_\mathrm{p}}$ with c_p contract on week w with the number of new recruitments done for that contract type and the differences between the changes into the contract type and out of it, for all weeks preceding the week w.

$$\mathrm{epe}_{w,c_\mathrm{p}} = \mathrm{qe}_{w,c_\mathrm{p}} + \sum_{i=0}^{w} \Big(\mathrm{nhp}_{i,c_\mathrm{p}} + \sum_{c_1 \in C_\mathrm{p}} \sum_{c_2 \in C_\mathrm{p}} (\mathrm{ch}_{\mathrm{i},c_1,c_\mathrm{p}} - \mathrm{ch}_{\mathrm{i},c_\mathrm{p},c_2}) \Big) \tag{2}$$

For example, if our number of predetermined permanent employees on week 5 is 2 for 10 hour contracts and 3 for 20 hour contracts, and we have chosen to hire a 2 new employees with 10 hour contracts on week 2, as well as, change one 10 hour employee into a 20 hour employee on week 3, then for week 5 for 10 hour contract $\mathrm{epe}_{5,10} = 2 + (2 + (0 - 1)) = 3$ and for 20 hour contract $\mathrm{epe}_{5,20} = 3 + (0 + (1 - 0)) = 4$.

We also defined a similar helper quantity $\mathrm{efe}_{w,c_\mathrm{f}}$ below in (3), to represent the number of fixed employees with c_f-hour contracts employed on week w. Similar to $\mathrm{epe}_{w,c_\mathrm{p}}$, we sum the number of predetermined fixed employees $\mathrm{qe}_{w,c_\mathrm{f}}$ with the total of new fixed employee recruitments for that contract type for all possible fixed contract lengths d. Since fixed-term employees are not employed after their contract ends, we only account for each duration length d for recruitments made at the latest in the current week w and at the earliest in the week $w - d + 1$, with w being their last working week.

$$\mathrm{efe}_{w,c_\mathrm{f}} = \mathrm{qe}_{w,c_\mathrm{f}} + \sum_{d \in D} \Big(\sum_{i=w-d+1}^{w} \mathrm{nhf}_{i,c_f,d} \Big) \qquad \forall w \in W, c_f \in C_f \tag{3}$$

For example, if we have no predetermined employees for fixed 10-hour contracts but choose to hire two new fixed employees for 10-hour contracts lasting 3 weeks, one starting in week 2 and one starting in week 3. The number of existing fixed employees for a 10-hour contract would be 0 in week 1, 1 in week 2, 2 in weeks 3 and 4, again 1 in week 5 and 0 from week 6 onwards.

Using the helper quantities above, we define another helper quantity cont_w below in (4), to denote the total of person-hours our existing employees w are contracted to work. We get this by multiplying the number of existing employees in week w by the hours associated with each contract, and summing the products over all possible contracts:

$$\mathrm{cont}_w = \sum_{c_p \in C_\mathrm{p}} (\mathrm{epe}_{w,c_\mathrm{p}} \cdot \mathrm{h}_{c_\mathrm{p}}) + \sum_{c_f \in C_\mathrm{f}} (\mathrm{efe}_{w,c_\mathrm{f}} \cdot \mathrm{h}_{c_\mathrm{f}}) \tag{4}$$

Additionally, we define a separate helper quantity in (5) to represent the total number of extra hours available through employees' flexibility. When calculating this by deducting from the number of existing employees for a given contract type in a given week, the number of employees absent for that contract type in that week is obtained. Then the deductions are each multiplied by the flexibility associated with the contact type, and summed over all possible contracts.

The number of absences is deducted from the number of existing employees, since only employees present in a given week are eligible to work extra hours that week. For permanent contracts, the number of existing employees is denoted by $\mathrm{epe}_{w,c_\mathrm{p}}$, absences by $\mathrm{a}_{w,c_\mathrm{f}}$, and flexibility associated with the contract is denoted by $\mathrm{f}_{c_\mathrm{p}}$. Thus, for fixed-term contracts, the corresponding quantities are denoted by $\mathrm{efe}_{w,c_\mathrm{f}}$, $\mathrm{a}_{w,c_\mathrm{f}}$, and $\mathrm{f}_{c_\mathrm{f}}$.

$$\mathrm{flex}_w = \sum_{c_p \in C_\mathrm{p}} (\mathrm{epe}_{w,c_\mathrm{p}} - \mathrm{a}_{w,c_\mathrm{p}}) \cdot \mathrm{f}_{c_\mathrm{p}} + \sum_{c_f \in C_\mathrm{f}} (\mathrm{efe}_{w,c_\mathrm{f}} - \mathrm{a}_{w,c_\mathrm{f}}) \cdot \mathrm{f}_{c_\mathrm{f}} \tag{5}$$

In addition to the helper quantities describing the available workforce, we wish to define a helper quantity for the total workforce demand. In our approach, we interpret employee absences as increasing workforce demand, since additional work hours are needed to fulfil the work the absentee is unable to do. This approach is most in line with reality, as absences such as paid and sick leave are still counted towards the contracted hours the employee is mandated to work and do not have to be made up afterwards.

Therefore, to simplify the representation of the total work demand, we define a helper quantity in (6) to represent the sum of predicted work demand and the additional work needed created by absences in a given week w. We derive this quantity by multiplying the number of absences for permanent contracts $\mathrm{a}_{w,c_\mathrm{p}}$ and fixed-term employees $\mathrm{a}_{w,c_\mathrm{f}}$ by the working hours of these contracts, c_p and c_f, summing the products over all contract types, and adding them to the estimated work demand $\mathrm{wd_w}$.

$$\mathrm{td}_w = \mathrm{wd}_w + \sum_{c_p \in C_\mathrm{p}} (\mathrm{a}_{w,c_\mathrm{p}} \cdot \mathrm{h}_{c_\mathrm{p}}) + \sum_{c_f \in C_\mathrm{f}} (\mathrm{a}_{w,c_\mathrm{f}} \cdot \mathrm{h}_{c_\mathrm{f}}). \tag{6}$$

For example, if we estimate that we require 300 hours of work on week 2, and we know that one permanent employee with a 10-hour contract and one with a 37.5-hour contract are absent that week, the total work demand for week 2 is $300+(1\cdot10+1\cdot37.5)+(0) = 347.5$.

Using the previous helper quantities, we also define a helper quantity to represent the number of flexibility hours actually utilised on a given week w in (7). This quantity is derived by deducting from the total work demand on the week w, td_w, the total number of contract hours available cont_w, and the number of rental work hours hired rh_w for week w.

$$\mathrm{fu}_w = \mathrm{td}_w - (\mathrm{cont}_w + \mathrm{rh}_w). \tag{7}$$

3.6 Constraints

3.6.1 Functional constraints

In this section, we describe the constraints of our model necessary to prevent results that are impossible in reality. The first set of these constraints is the constraints (8), (9), and (10), which ensure that the model does not recommend actions that the user has forbidden.

We have defined our parameter ca_{c_1,c_2} to be 1 if a change from a permanent contract c_1 is allowed, and 0 if it is not. Consequently, constraint (8) sets the change decision variables $\mathrm{ch}_{\mathrm{w},c_1,c_2}$ as zero for all weeks that have been forbidden by setting $\mathrm{ca}_{c_1,c_2} = 0$. This constraint, for example, is used to prevent changes from higher contract hours to lower ones in the model:

$$\mathrm{ch}_{\mathrm{w},c_1,c_2} = 0 \qquad \forall w \in W, (c_1, c_2) \in C_\mathrm{p} : \mathrm{ca}_{c_1,c_2} = 0. \tag{8}$$

Similarly, as set $\tilde{C}_\mathrm{p}$ and $\tilde{C}_\mathrm{f}$ contain only those permanent and fixed-term contract types that the model is allowed to recruit, by constraining the hiring decision variables $\mathrm{nhp}_{w,c_\mathrm{p}}$ and $\mathrm{nhf}_{w,c_\mathrm{f},d}$ to zero for all weeks $w \in W$ and for all contract types not included in these sets, the constraints (9) and (10) prevent the model from considering hiring decisions forbidden by the user.

$$\mathrm{nhp}_{w,c_\mathrm{p}} = 0 \qquad \forall w \in W, c_p \in C_\mathrm{p} \setminus \tilde{C}_\mathrm{p} \tag{9}$$

$$\mathrm{nhf}_{w,c_\mathrm{f},d} = 0 \qquad \forall w \in W, c_f \in C_\mathrm{f} \setminus \tilde{C}_\mathrm{f}, d \in D \tag{10}$$

Additionally, we define the constraint (11) below to ensure that we only recommend changes to hours only on contracts that actually exist that week: for each week w and permanent contract type c_p, we sum number of changes out of the contract type c_p across all permanent contract types and set it to be at most the global minimum number of the predetermined employees of contract type c_p, $\mathrm{qe}_{\mathrm{min,c_p}}$ summed with number of new hires, and the net flow between hour changes into an out of the contract type for all the weeks before the given week w. We use the global minimum number of predetermined employees for contract type c_p instead of the number of predetermined employees for week w to prevent edge cases where hour changes would need to be assigned in postprocessing to employees who are known to leave the company due to retirement, quitting, etc.

$$\sum_{c_2 \in C_\mathrm{p}} \mathrm{ch}_{w,c_p,c_2} \leq \mathrm{qe}_{\mathrm{min,c_p}} + \sum_{i=0}^{w-1} \Big\{ \mathrm{nhp}_{i,c_p} + \sum_{c_1 \in C_\mathrm{p}} \mathrm{ch}_{\mathrm{i},c_1,c_\mathrm{p}} - \sum_{c_2 \in C_\mathrm{p}} \mathrm{ch}_{\mathrm{i},c_\mathrm{p},c_2} \Big\} \; \forall w \in W, \; \forall \; c_p \in C_\mathrm{p} \tag{11}$$

where $\mathrm{qe}_{\mathrm{min,c_p}} = \min_{w \in W} \mathrm{qe}_{\mathrm{w,c_p}}$

For example, if the number of predetermined employees working 25 hours on week on week 4 is 10, with one of them quitting on week 6, and model has chosen to hire one employee of this contract type on week 3, and increase the hours of one of these employees to full time on week 2, left side of the constraint (11) is $(10-1)+(1+0-1) = 9$, meaning that the total number of hour changes made for 25-hour contracts on week 4 can at most be 9, as that is the number of people we have employed on that contract type for the foreseeable future.

We also ensure that the model does not use more flexibility than the contract structure actually allows, by defining in (12) for each week w that the actual used flexibility, fu_w, is smaller than or equal to the total number of extra hours we have access to, flex_w. As exhibited in (12), if we substitute helper quantity fu_w by its definition in (7), we can show that this constraint also ensures that if the total demand cannot be covered by the total amount of contract hours and the maximum amount of extra hours available, we buy enough rental work hours to cover the difference. Therefore, this constraint also ensures that the total work demand is always met.

$$\begin{aligned} &\mathrm{fu}_w \leq \mathrm{flex}_w \qquad \forall w \in W \\ &\Leftrightarrow \\ &\mathrm{rh}_w \geq \mathrm{td}_w - \mathrm{cont}_w - \mathrm{flex}_w \qquad \forall w \in W. \end{aligned} \tag{12}$$

3.6.2 Additional constraints

This section describes all the model constraints that are used to customise the model to the specific situation at hand. These constraints all include user-defined parameters, through which the user can incorporate unit- and scenario-specific needs and limitations into the model.

Due to the negative material and immaterial costs of both consistent over- and understaffing, units often wish to bound the total number of contracted hours available over a specific balancing period b, cont_w, to some upper and lower relative to the sum of total demand across all weeks in the balancing period. For this purpose, we define the constraints (13) and (14): using these constraints, we can set that the sum of the contracted hours of the first balancing period must be at least 94% and at most 98% of the sum of the weekly total demands in that balancing period. The relative lower bound, $\mathrm{cc}_{\mathrm{lower}}$, and relative upper bound, $\mathrm{cc}_{\mathrm{upper}}$, used in these constraints are defined by the user as a decimal between 0 and 1, inclusive.

$$\sum_{w \in b} \text{cont}_w \geq \text{cc}_{\text{lower}} \cdot \sum_{w \in b} \text{td}_w \qquad \forall b \in B \tag{13}$$

$$\sum_{w \in b} \text{cont}_w \leq \text{cc}_{\text{upper}} \cdot \sum_{w \in b} \text{td}_w \qquad \forall b \in B. \tag{14}$$

In addition to the work done by in-house employees, the units often wish to bound the work done by each rental worker hired. We incorporate this into the model through constraints (15) and (16): The user defines parameters $\text{hours}_{\text{r,min}}$ and $\text{hours}_{\text{r,max}}$ that represent the minimum and maximum number of hours that individual rental workers hired in any given week will work. Then using these parameters, we constrain the amount of rental hours bought rh_w for week w to be at least the number of rental workers hired for the week w,rw_w, times the user defined lower hour limit in (15), and at most the number of rental workers hired for the week w,rw_w, times the user defined upper hour limit in (16) for all weeks.

$$\text{rh}_w \geq \text{rw}_w \cdot \text{hours}_{\text{r,min}} \qquad \forall w \in W \tag{15}$$

$$\text{rh}_w \leq \text{rw}_w \cdot \text{hours}_{\text{r,max}} \qquad \forall w \in W \tag{16}$$

Due to budget constraints, seasonal demand, or overall preferences, the units often want to limit the number of fixed-term employees for individual weeks. These limits are included in the model through the constraint (17) below: The user defines the maximum number of fixed-term hires allowed,$\text{fh}_{\text{max},w}$, for each week w separately. Then the number of fixed-term hires,$\text{nhf}_{w,c_\text{f},d}$, summed over all allowed contract types, $\tilde{C}_\text{f}$, and possible durations,D, for any week w is set to be at most the corresponding upper limit $\text{fh}_{\text{max},w}$

$$\sum_{c \in \tilde{C}_\text{f}} \sum_{d \in D} \text{nhf}_{w,c_\text{f},d} \leq \text{fh}_{\text{max},w} \qquad \forall w \in W \tag{17}$$

To ensure continuity in the workforce, the units often wish to ensure that the permanent employees make up at least a certain portion of the total existing workforce. Due to the seasonal nature of the workforce demand, this lower limit often varies from week to week. For example, during summer and winter holidays, the units often need additional seasonal employees to cover employee absences or temporarily meet increased demand. Therefore, we define the constraint (18) below for each week w: The user sets for each week w a relative lower limit for the existing permanent employees, $\text{share}_{\text{p},w}$, as a decimal between 0 and 1, inclusive. Consequently, a constraint can be used to ensure, for example, that the sum of the number of existing permanent workers on week w, $\text{epe}_{w,c_\text{p}}$, is at least 40% of the total workforce i.e. sum of the number of existing permanent workers across all contracts $c_p \in C_\text{p}$, the number of fixed term employees, $\text{efe}_{w,c_\text{f}}$, across all contracts $c_f \in C_\text{f}$, and the number of rental workers hired for week w, rw_w.

$$\sum_{c \in C_\text{p}} \text{epe}_{w,c_\text{p}} \geq (\sum_{c \in C_\text{p}} \text{epe}_{w,c_\text{p}} + \sum_{c \in C_\text{f}} \text{efe}_{w,c_\text{f}} + \text{rw}_w) \cdot \text{share}_{\text{p},w} \qquad \forall w \in W \tag{18}$$

Due to infrastructure limitations or role requirements across shifts, units often wish to set upper and lower limits on the number of permanent employees in any given week. These limits, however, may vary depending on the type of permanent contract. Therefore, the upper limit is incorporated into our model through constraint (19), and the lower limit is incorporated through constraint (18). In these constraints, $\max_{c_p}$ and $\min_{c_p}$ are the global upper and lower limits for the number of existing employees of contract type c_p, $\text{epe}_{w,c_\text{p}}$, and are positive integers defined by the user. We define the constraints (19) and (20) only for the set of permanent contracts for which we can make new hires, $\tilde{C}_\text{p}$, as the number of existing employees is only subject to change for these contract types.

$$\text{epe}_{w,c_\text{p}} \leq \max_{c_p} \qquad \forall w \in W, c_p \in \tilde{C}_\text{p} \tag{19}$$

$$\text{epe}_{w,c_\text{p}} \geq \min_{c_p} \qquad \forall w \in W, c_p \in \tilde{C}_\text{p} \tag{20}$$

Similarly to the number of permanent employees, the units often wish to set a global upper limit for the number of fixed-term employees, which can vary across fixed-term contract types. Therefore, the user defines a global upper bound for fixed-term employees, $\max_{c_f}$, as a positive integer, and we define the upper bound constraint (21), for all weeks, and all contract types $c_f \in \tilde{C}_\text{f}$ for which we can make new hires. As fixed-term employees are guaranteed to leave at the end of their contracts, they cannot be assigned responsibilities, so units tend not to need a lower bound on the number of existing fixed-term employees.

$$\text{efe}_{w,c_\text{f}} \leq \max_{c_f} \qquad \forall w \in W, c_f \in \tilde{C}_\text{f} \tag{21}$$

In addition to the constraints presented in this section, the unit may wish to make further customised constraints on the share of different contract types among the existing employees for any given week. For example, the unit may wish to ensure that at least 40% of existing employees work 30 hours or more per week, or that at most 20% work 15 hours or less per week. For these purposes, we defined the minimum share constraint (22) and the maximum share constraint (23) below, which are customised based on user input.

$$\sum_{c_p \in C^s_\text{p}} \text{epe}_{w,c_\text{p}} + \sum_{c_f \in C^s_\text{f}} \text{efe}_{w,c_\text{f}} \geq \text{share}_\text{min} \cdot (\sum_{c_p \in C_\text{p}} \text{epe}_{w,c_\text{p}} + \sum_{c_f \in C^*_\text{f}} \text{efe}_{w,c_\text{f}}) \qquad \forall w \in W \tag{22}$$

$$\sum_{c_p \in C^s_\text{p}} \text{epe}_{w,c_\text{p}} + \sum_{c_f \in C^s_\text{f}} \text{efe}_{w,c_\text{f}} \leq \text{share}_\text{max} \cdot (\sum_{c_p \in C_\text{p}} \text{epe}_{w,c_\text{p}} + \sum_{c_f \in C^*_\text{f}} \text{efe}_{w,c_\text{f}}) \qquad \forall w \in W \tag{23}$$

In our parameters, the user can define several rows of minimum and maximum share constraints in their specified columns. In the case of minimum constraints (22), the input is in the form $\text{share}_\text{min}\ ; R; h$ and for maximum constraints (23) the input is correspondingly of form $\text{share}_\text{max}\ ; R; h$, where share_min and share_max represent the minimum or maximum shares of the employees that need to satisfy the condition derived from the input. Additionally, in both cases, R is a relation operator $R \in \{(x, y) \mid x, y \in \mathrm{R}_+\}$ and h is a real constant $h \in \mathrm{R}_+$ representing the hour threshold of the constraint. The exhaustive list of the relation operators R and their syntax is provided in the Appendix 7.1.

Then, using this input, we define the subsets of permanent contracts C^s_p and fixed-term contracts C^s_f as defined below in (24) and (25). The input parameters also allow including only permanent contracts as part of the constraint, meaning that instead of the definition outlined in (25), we would define $C^s_\text{f} = \emptyset$.

$$C^s_\text{p} = \{c_p \in C_\text{p} \mid (c_p, h) \in R\} \tag{24}$$

$$C^s_\text{f} = \{c_f \in C_\text{f} \mid (c_f, h) \in R\} \tag{25}$$

For example, if the user would set that at most 20% work 15 hours or less per week through input $0.20; \leq; 15$, we would form the left side of the minimum share constraint (22) for a given week w by summing the number of existing permanent employees, $\text{epe}_{w,c_\text{p}}$, and the number of existing fixed-term employees, $\text{efe}_{w,c_\text{f}}$, over all possible contracts and multiplying it by the given minimum share 0.20. Additionally, we would define the sets C_p^s and C_f^s according to (24) and (24) by including only those permanent and fixed-term contract types with 15 hours or less, and then summing the number of existing permanent employees and the number of existing fixed-term employees for all the contracts in the sets C_p^s and C_f^s, with 15 hours or less.

3.7 Objective function

The general objective of this optimisation model is to find a long-term recruitment plan that meets the demand for work while minimising costs. As our constraints ensure that any feasible solution will fulfil the work demand, we shall include the total costs associated with the hiring plan in the objective function to minimise them. There are four types of costs associated with recruitment decisions: the nominal cost of having an employee, including, for example, employee benefits and administrative costs associated with employee records, the administrative costs of changing the hours of an employee, the costs related to hiring an employee, including, for example, recruitment and administrative costs, and the hourly cost of bought rental work. We do not need to account for individual employees' wages, since we assume wages are the same across employees. As all feasible solutions aim to cover the total demand exactly, this assumption implies that the total wage cost at all critical points will be the same.

We minimise the total of these costs in our objective function defined in (26) below. In this function we get the total costs of the plan by multiplying the cost of having a permanent employee, $\text{cost}_\text{permanent}$ with the number of existing permanent employees on week w, $\text{epe}_{w,c_\text{p}}$, across all permanent contracts C_p, the cost of having a fixed-term employee $\text{cost}_\text{fixed-term}$ with the number of existing fixed-term employees on week w, $\text{efe}_{w,c_\text{f}}$, across all fixed term contracts C_f, the administrative costs of changing a contract $\text{cost}_\text{change}$ with number of changes,$\text{ch}_{\text{w},c_1,c_2}$, across all possible options, the cost of recruiting a permanent employee $\text{cost}_\text{recruit,permanent}$ with the sum of hired permanent employees $\text{nhp}_{w,c_\text{p}}$, the cost of recruiting a fixed-term employee $\text{cost}_\text{recruit,fixed-term}$ with the sum of hired fixed-term employees $\text{nhp}_{w,c_\text{p}}$, and the hourly cost of rental work wage_ren with the hours of rental work bough on week w, and finally summing them over all considered weeks $w \in W$. We have defined separate cost parameters for having a permanent and a fixed-term employee, as well as for hiring a permanent and a fixed-term employee, because the costs associated with permanent employees tend to be higher than those for fixed-term employees in both cases.

$$\min \sum_{w \in W} \Big\{ \text{cost}_{\text{permanent}} \cdot \sum_{c_p \in C_{\text{p}}} \text{epe}_{w,c_{\text{p}}}$$
$$+ \text{cost}_{\text{fixed-term}} \cdot \sum_{d \in D} \sum_{c_f \in C_{\text{f}}} \text{efe}_{w,c_{\text{f}}}$$
$$+ \text{cost}_{\text{change}} \cdot \sum_{c_1 \in C_{\text{p}}} \sum_{c_2 \in C_{\text{p}}} \text{ch}_{\text{w},c_1,c_2}$$
$$+ \text{cost}_{\text{recruit,permanent}} \cdot \sum_{c_p \in \tilde{C}_{\text{p}}} \text{nhp}_{w,c_{\text{p}}}$$
$$+ \text{cost}_{\text{recruit,fixed-term}} \cdot \sum_{d \in D} \sum_{c_f \in \tilde{C}_{\text{f}}} \text{nhf}_{w,c_{\text{f}},d}$$
$$+ \text{wage}_{\text{ren}} \cdot \text{rh}_w \Big\} \tag{26}$$

4 Results

4.1 Model performance

The model's performance was tested with a sample dataset representing an average-sized retail store with approximately 50 employees. The model outputs indicate that the formulation behaves as expected and solves within five minutes. To validate the model's performance, the outputs are further analysed and visualized to evaluate whether the results are consistent, interpretable, and aligned with the intended objective of workforce planning.

As shown in Figures 1 and 2, the model satisfies work demand while minimizing cost through contract-hour flexibility, fixed-term hires, and contract hour changes. The timing of these decisions suggests that the model responds positively to variations in demand across the planning horizon. Table 4 describes individual components shown on Figures 1 and 2.

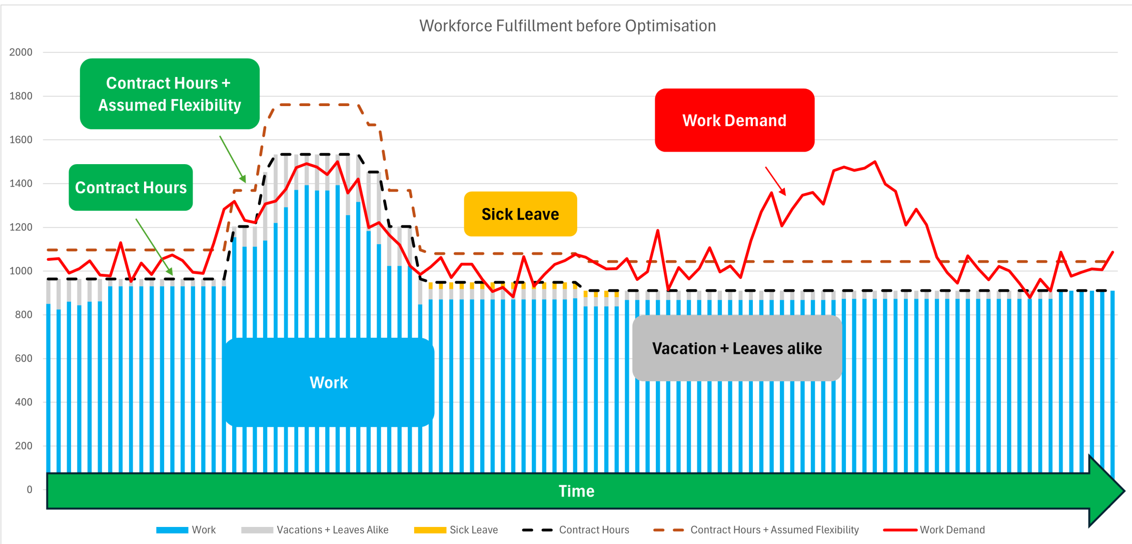

Figure 1: Workforce fulfilment for 2 years before optimisation.

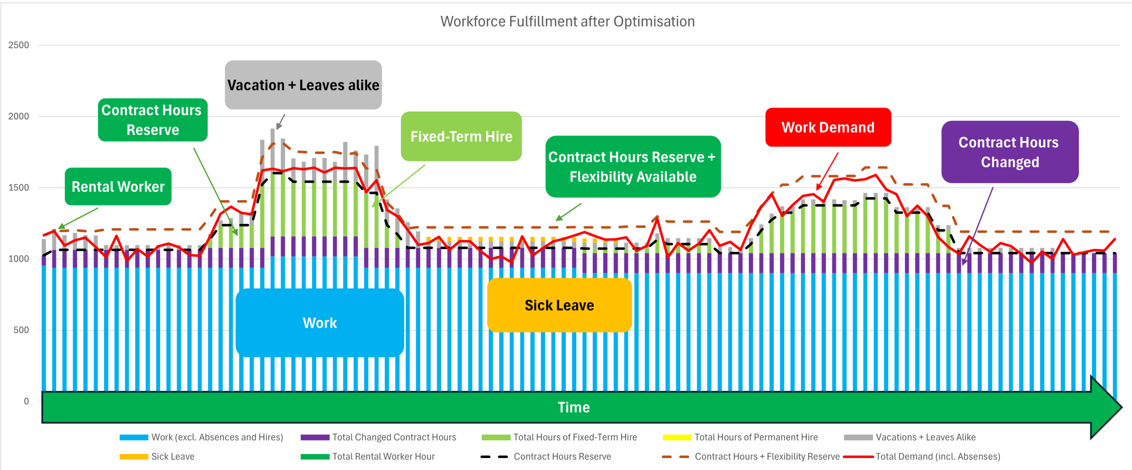

Figure 2: Workforce fulfilment for 2 years after optimisation.

Component	**Description**
Figure 1: Before optimisation	
Work	Available work hours from pre-determined employees.
Vacations + Leaves Alike	Contracted hours unavailable due to study leave, military leave, planned vacation, or maternity leave.
Sick Leave	Contract hours unavailable due to long-term sick leave.
Contract Hours	Total contracted hours, including hours unavailable due to absences.
Contract Hours + Assumed Flexibility	Total contracted hours plus assumed available flexibility.
Work Demand	Predicted work demand.
Figure 2: After optimisation	
Work	Available work hours from pre-determined employees, excluding new hires
Sick Leave	Contract hours unavailable due to long-term sick leave
Vacations + Leaves Alike	Contract hours unavailable due to study leave, military leave, planned vacations, or maternity leave
Contract Hours Changed	Increase in available work hours resulting from contract-hour changes
Fixed-Term Hire	Cumulative increase in available work hours resulting from fixed-term hires
Permanent Hire	Additional available work hours from permanent hires
Rental Workers	Available work hours from rental workers
Contract Hours Reserve	Total contracted hours available after excluding absences
Contract Hours Reserve + Flexibility Reserve	Contract hours reserve plus available flexibility reserve
Total Demand	Predicted work demand plus additional work required to cover absences

Table 4: Description of variables used in Figures 1 and 2 .

Model Configuration	Runtime	Integer Variables	Achieved Gap	Hires (P, F, HC, RW)
Baseline Model	299.5s	27560	0.44%	(0,34,8,1)
No Contract Changes	1.9s	27560	Infeasible	-
Less Fixed-Term	233.6s	23816	0.35%	(0,34,10,3)
Less Perm.	116.7s	25064	0.32%	(0,34,10,0)
Less Perm. and Fixed	101.8s	19448	0.41%	(0,43,10,2)
Cheaper Perm. Param.	337.2s	27560	0.58%	(0,32,10,1)

Table 5: Observation Analysis Summary. Runtime was averaged over three runs for each configuration with a MIP gap of 0.02. P = Permanent Hire, F = Fixed-Term Hire, HC = Contract Hour Change, RW = Rental Worker.

Experiments with different model configurations were conducted to evaluate the model's performance under varying user-defined parameter settings. The tested configuration, together with their performance metrics and resulting decision outputs are presented in Table 5. During testing, different MIP gaps were specified as solver parameters. However, varying the MIP gap did not noticeably affect the overall performance of the model. Therefore, a MIP gap of 0.02 was used for the experiments presented in this report.

According to Table 5, the baseline version of the model takes ≈5 minutes to obtain an optimal solution of 0.44% on the 1st run. This model accounts for 45 existing employees, 8 contract-hour options for permanent contracts, and 6 contract options for fixed-term contracts. On the 2nd and 3rd run, the model achieves the identical optimal solution of 0.44% using Feasibility Jump from memory in ≈44 seconds. To meet demand, the decision outputs are to hire 34 fixed-term contracts, mostly of 25 and 32 hour contracts during the summer time, increase 8 existing employee's contract hours from 5 and 6 hours to 20 and 28 hour contracts, and hire a rental worker on the 2nd week. The achieved solution of fixed-term hires and contract hour changes is due to the low cost of increasing the contract hours of an existing employee and the high cost of hiring a new employee entirely. In addition, the model is strongly against hiring permanent employees due to the less flexibility to meet seasonal demand compared to fixed-term employees.

Given that the achieved solution made extensive use of contract changes with predetermined employees in the start of the planning horizon, we implemented an instance in which contract changes are not allowed. However, the model became infeasible with a large number of constraints being violated and a runtime of ≈1.9 seconds. It was required that at least 25% of permanent employees had at least 30 hour contracts and a maximum of 30% of permanent employees had contracts with 15 hours. This requirement is a hard constraint which might be the reason for the model's infeasibility, if there are not enough contract options to meet these requirements.

To ensure that the model can handle different parameter settings, we implemented a configuration with varying contract options for fixed-term or permanent hires. In this instance, the model still found a feasible solution to meet the model's objectives. For a feasible solution to exist with less permanent and fixed term contract hours options, one must ensure that the maximum number of fixed term hires per each contract hour type each week is reasonable. During the summer, with absences due to summer vacations, fixed term hires are essential. If this condition is not adjusted accordingly, the model returns infeasible.

In addition, due to the achieved solution strongly discouraging permanent hires, we implemented a configuration in which the cost of keeping existing and hiring permanent employees is significantly cheaper than that of a fixed-term employee. However, testing suggests that costs relating to permanent employees need to be close to 0 for the model to promote hiring a permanent employee. One possible reason is that, compared to fixed-term employees, permanent hires do not provide enough flexibility to handle fluctuating demand. Another possible reason is the sample dataset having enough permanent employees for an optimal solution so additional permanent employees are not necessary.

4.2 Implementation

The recruitment optimisation model is implemented using the following framework to function as a usable tool. All packages used are licenced under MIT licence, making them usable in commercial contexts.

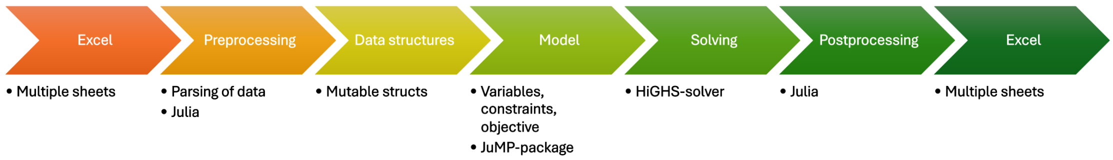

Figure 3: Implementation framework.

The implementation framework for the optimisation tool is described in Figure 3. The optimisation tool takes an Excel-file as input, containing multiple data sheets. The input is preprocessed using Julia [14] and transformed into mutable structs. Next, the model is built using these mutable structs and optimised using a mixed-integer linear programming solver. The model solution is postprocessed and an Excel-file is created as output.

The user interacts with the input Excel-file, a Jupyter notebook, and the output Excel-file. In the input Excel-file, the user defines the estimated weekly workload, predetermined employees, absences, changes in contract hours, and possible contract types. In addition, all of the adjustable parameters of the model are included in the Excel-file.

The optimisation tool is run from a Jupyter notebook. The functions for preprocessing, model building, and postprocessing are organised in separate files and called from the notebook. As requested, the Julia programming language [14] was used throughout the implementation. The input data from the Excel-file is parsed and preprocessed using the XLSX.jl and DataFrames.jl packages. Julia's built-in mutable structs allow the structuring of the input data for easier access during model building.

The mathematical modelling language JuMP [15] is used to build the optimisation model as described in Section 3. After a model is successfully built, the optimisation process is executed. The HiGHS-solver [16] is used to solve the model, utilising MILP techniques such as the Branch-and-Bound method. Solving time can be limited through the notebook.

After the optimisation model is implemented, the Jupyter notebook calls a separate Julia post-processing script to extract key results and decision-variable values. These outputs are then organized into structured tables to support clearer visualization and interpretation for analysis. After the results are converted into interpretable data structures, the Jupyter notebook produces a clean Excel output containing the workforce planning forecast solution. The output Excel-file allows the user to review the resulting recommendations for recruiting and contract-hour changes. In addition, the results are visualised. The Excel format allows easy import of the decisions into other systems.

5 Discussion

5.1 Limitations

Several aspects could be further developed in the created recruitment planning model. In addition, it is essential that the plan is continuously updated and also aligns with short-term planning objectives.

The model operates on aggregate numbers at the contract-hour category level rather than at the individual employee level. Because the input data includes individual-level absence information, but we do not track changes in contract hours, the missing hours may be underestimated. For the same reason, hour increases are not limited to employees whose hours have not yet been changed. One option would be to model the absences as a percentage of the workforce. In that way, we would not assume that the new employees will never be absent. Annual vacations are not accounted for in the current approach.

It is also assumed that there are no skill differences. Skills and skill levels, as well as work experience, different roles and responsibilities, could affect the wages that are currently assumed to be the same for all employees. Worker preferences are neither considered. Nevertheless, preferences are perhaps more related to short-term planning.

The flexibility is modelled at the contract-type level, and therefore, it is also assumed that employees will always accept extra work. Whenever the employer decides to recruit or increase hours, there is an available workforce.

The framework is highly dependent on the formatting of the input Excel-file. Validation of the input data is not included in the code, and thus, the framework assumes the input data does not contain errors. This limitation can be addressed by adding validation checks for the mutable structs. The model formulation could possibly be improved by defining the number of present employees iteratively using the information from the previous week.

The results also rely on the demand estimates. In addition to seasonal variability, surprising events such as a pandemic or global supply chain disruptions might affect the demand in unexpected ways. The variability within a week is not considered. The current model is deterministic, and considering uncertainty in both demand and workforce supply is also highly relevant for a more realistic recruitment plan.

Based on the decision output, the model heavily depends on increasing contract hours of pre-determined employees. In the current implementation, if contract hour changes are not allowed, the model may become infeasible, as seen with the example data used to validate the model. A possible reason for this are a large number of key constraints relating to minimum and maximum shares of employees being violated. Thus, in order to implement a model where contract hour changes are not allowed, several hard constraints need to be adjusted accordingly for the model to reach feasibility.

5.2 Modelling uncertainty

Long-term recruitment planning is inherently uncertain as in addition to demand fluctuations, personnel changes and unexpected absences cause uncertainty in workforce planning. The developed MILP model does not account for this uncertainty. Therefore, the model might underestimate the required recruitments. The personnel change uncertainty could be incorporated into the model by using a resignation rate that depends on contract hours, with part-time employees having a higher resignation rate to full-time employees.

Overall, uncertainty could be incorporated using robust optimisation, two-stage stochastic optimisation or using simulation-based approaches as mentioned in [12]. The optimisation could be evaluated over multiple scenarios. A scenario tree structure that captures multiple possible combinations [11] could be utilized. Example scenarios may include different resignation rates for different contract types and absences duet to seasonal sickness peaks. These scenarios could be based on historical data and similar estimates as in [13]. The optimal recruitment plan could be constructed either to perform well on average or to remain feasible under worst-case scenarios. Use of different scenarios also offers insights to the sensitivity of the model. The model should be return recruitments plans which are robust over different scenarios, such that worse-case workforce shortages are avoided.

In two-stage stochastic optimisation, decisions are divided into first stage decisions before uncertainty is realised and second stage decisions which are adjusted after uncertainty is realized in different scenarios. The total objective is to minimize the total excepted objective value, which considers both first and second-stage decisions. In the context of recruitment planning, first stage decisions would include the recruitment decisions and contract hour change decisions. The second stage would capture the consequences of different uncertainty realizations, such as cost of rental work or overtime. Two-stage stochastic optimisation would lead to recruitment decisions which perform the best considering the consequences of different uncertainty realizations. However, this approach is typically computationally heavy and complex which might make it infeasible in this context.

6 Conclusions

This project developed a workforce recruitment and staffing optimisation model for a company using a mixed-integer linear programming (MILP) approach. The model aimed to minimise labour costs while satisfying multiple constraints. Using the provided sample dataset, the model was able to represent the workforce planning conditions of a medium-sized retail store with approximately 50 employees.

The results showed that the model could generate feasible staffing solutions within a reasonable computational time. The model also allowed existing employees, newly recruited workers, and rental workers to be considered together when making staffing decisions under varying operational requirements. The study also demonstrates the applicability of mathematical optimisation techniques in workforce management problems within the retail industry.

Nevertheless, several limitations need to be acknowledged. The model relied on several simplifying assumptions and therefore could not fully represent the complexity of real-world workforce management. Uncertain demand fluctuations, unexpected employee absences, and human behavioural factors were simplified or excluded from the optimisation process. Consequently, the results should be interpreted as decision-support guidance rather than as exact workforce plans.

Future research could extend the model by incorporating stochastic optimisation methods to represent uncertainty in workforce demand and staffing conditions. However, such extensions would increase both model complexity and computational requirements. Additional improvements may further enhance the model's applicability to more realistic retail workforce planning.

References

[1] K. Sarwary, F. Faizi, and M. R. Banayee. The role of human resource planning on the improvement of employees' recruitment process. *Journal of Corporate Finance Management and Banking System (JCFMBS)*, 2(05):29–41, 2022. `https://doi.org/10.55529/jcfmbs.25.29.41`.

[2] Service Union United PAM and the Finnish Commerce Federation. Commerce sector collective agreement. 2025. https://www.pam.fi/en/tes/commerce-sector-collective-agreement/.

[3] C. Cayrat and P. Boxall. The roles of the HR function: a systematic review of tensions, continuity and change. *Human Resource Management Review*, 33(4):100984, 2023. https://doi.org/10.1016/j.hrmr.2023.100984.

[4] A. T. Ernst, H. Jiang, M. Krishnamoorthy, B. Owens, and D. Sier. An annotated bibliography of personnel scheduling and rostering. *Annals of Operations Research*, 127(1):21–144, 2004. https://doi.org/10.1023/B:ANOR.0000019087.46656.e2.

[5] K. Karhula, R. Shiri, J. Ervasti, A. Koskinen, A. Ropponen, M. Sallinen, J. Turunen, and M. Härmä. Effects of the use of a shift schedule evaluation tool with ergonomic recommendations on employee wellbeing - a quasi-experiment in the Finnish healthcare sector. *Applied Ergonomics*, 130:104638, 2026. https://doi.org/10.1016/j.apergo.2025.104638.

[6] N. Llort, A. Lusa, C. Martínez-Costa, and M. Mateo. A decision support system and a mathematical model for strategic workforce planning in consultancies. *Flexible Services and Manufacturing Journal*, 31(2):497–523, 2019. https://doi.org/10.1007/s10696-018-9321-2.

[7] Y. Li, J. Chen, and X. Cai. An integrated staff-sizing approach considering feasibility of scheduling decision. *Annals of Operations Research*, 155(1):361–390, 2007. https://doi.org/10.1007/s10479-007-0215-z.

[8] L. Talarico and P. A. Maya-Duque. An optimization algorithm for the workforce management in a retail chain. *Computers & Industrial Engineering*, 82:65–77, 2015. https://doi.org/10.1016/j.cie.2015.01.014.

[9] A. Corominas, A. Lusa, and R. Pastor. Using MILP to plan annualised working hours. *Journal of the Operational Research Society*, 53(10):1101–1108, 2002. https://doi.org/10.1057/palgrave.jors.2601309.

[10] A. İşeri, H. Güner, and A. R. Güner. Pareto-optimal workforce scheduling with worker skills and preferences. *Operational Research*, 25(27):1–27, 2025. https://doi.org/10.1007/s12351-025-00903-7.

[11] T. Karimi, S. Thevenin, and H. H. Benderbal. Workforce management and resource selection with fairness. *IFAC-PapersOnLine*, 59(10):2165–2170, 2025. https://doi.org/10.1016/j.ifacol.2025.09.364.

[12] A. F. Porto, A. Lusa, S. A. Herazo, and C. A. Henao. Improving the robustness of retail workforce management with a labor flexibility strategy and consideration of demand uncertainty. *Operations Research Perspectives*, 15:100345, 2025. https://doi.org/10.1016/j.orp.2025.100345.

[13] A. Wilhelmsson. Computational models for evaluating contract structures under uncertainty. Master's thesis, 2025. Master's Thesis. Department of Mathematics and Systems Analysis, Aalto SCI. https://sal.aalto.fi/publications/pdf-files//theses/mas/twil25_public.pdf.

[14] J. Bezanson, A. Edelman, S. Karpinski, and V. B. Shah. Julia: A fresh approach to numerical computing. *SIAM review*, 59(1):65–98, 2017. https://doi.org/10.1137/141000671.

[15] M. Lubin, O. Dowson, J. Dias Garcia, J. Huchette, B. Legat, and P. Vielma, J. JuMP 1.0: Recent improvements to a modeling language for mathematical optimization. *Mathematical Programming Computation*, 15:581–589, 2023. 10.1007/s12532-023-00239-3.

[16] Q. Huangfu and J. A. J. Hall. Parallelizing the dual revised simplex method. *Mathematical Programming Computation*, 10(1):119–142, 2018. 10.1007/s12532-017-0130-5.

7 Appendix

7.1 Exhaustive list of the relation operators supported in the minimum and maximum share constraints and their corresponding inputs

Relation definition	Corresponding input
$R = \{(x, y) \in \mathrm{R}^2 \mid x = y\}$	`=`
$R = \{(x, y) \in \mathrm{R}^2 \mid x \geq y\}$	`>=`
$R = \{(x, y) \in \mathrm{R}^2 \mid x > y\}$	`>`
$R = \{(x, y) \in \mathrm{R}^2 \mid x \leq y\}$	`<=`
$R = \{(x, y) \in \mathrm{R}^2 \mid x < y\}$	`<`

8 Group assessment

8.1 How closely did the actual implementation of the project follow the initial project plan? Were there any major departures and, if so, what?

The initial objective of the project was to develop optimisation models for recruitment planning. The first phase of the project was a deterministic model, where future demand and personnel changes are assumed to be known. The first phase of the project has been implemented, but the model formulation phase took longer than we expected. Therefore, the second phase of the project, which would have modelled the demand and personnel changes as random variables, was left undone. Instead, we researched the literature on how it could have been implemented. Otherwise, the project implementation mainly followed the initial plan.

8.2 In what regard was the project successful?

Despite not completing the second phase, the project succeeded in creating a functional recruitment planning model that finds a solution in a reasonable time. In the implementation of the first phase, all the required constraints are taken into account. One aspect that made the success possible was the use of git and the coordination of version control. Our initial modelling approach was successful because we avoided an employee-level model formulation that probably would have been difficult to solve. Although we could have divided the project tasks earlier in the project, one successful aspect was that group members were able to use their strengths and work on tasks that aligned with their interests.

8.3 In what regard was the project unsuccessful?

In retrospect, the main issue was not having time to implement the second phase of the project, which would have incorporated stochasticity in workforce planning. As recruitment planning relies on demand forecasts and assumptions about the available workforce, considering uncertainties related to both demand and supply would be crucial for a more realistic model and plan that would be feasible in various scenarios. We used quite a lot of time to understand the problem scope and formulate the problem. We also focused on some details that later turned out to have a minor impact. Organising the project differently by dividing responsibilities earlier and testing a simpler model that would have been further developed could have enabled implementing the second phase.

8.4 What could have been done better, in hindsight? (you may analyze this question from the roles of the project team, the client, and the teacher(s))

As mentioned in the section above, we should have split responsibilities among the team in an earlier phase. Our initial approach was to work together for a couple of hours per week. Even though that worked for researching the literature, problem scoping and project planning phases, in the problem formulation phase we should have tried to divide the tasks more clearly. However, most of the project tasks required linear progression rather than parallel execution. A more effective approach would have been to allocate different amounts of work to different weeks based on the actual demands of each phase, rather than attempting to equalize workload distribution across all team members each week. To succeed better, we should have ensured all team members have the complete picture while simultaneously dividing tasks efficiently. From a project management perspective, maintaining detailed meeting notes for each meeting and consistently writing down assigned tasks would have provided clarity on responsibilities assigned.