AMSI SUMMER RESEARCH SCHOLARSHIPS 2025–26

Get a taste for Research this Summer

A Mathematical Investigation of Palliative Care Workforce Planning

Thomas Caldecott

Supervised by Professor Andreas Ernst

Monash University

Jane Street®

AMSI AUSTRALIAN MATHEMATICAL SCIENCES INSTITUTE

Abstract

Community palliative care services face the ongoing challenge of scheduling client visits across a workforce with limited capacity, while adhering to clinical guidelines that vary with each client's condition. This report, undertaken in collaboration with Palliative Care South East (PCSE), develops a perfect information integer programming model to explore what an ideal visit schedule looks like under the clinical guidelines, and what this reveals about the capacity requirements of delivering it. The model assigns visits to days for 322 clients across five staffing scenarios, minimising total days late subject to nursing capacity and frequency constraints. The analysis finds that the existing schedule does not closely follow the clinical guidelines, with stable clients overserviced and deteriorating clients frequently seen later than recommended. Optimisation alone, holding resources fixed to what was used in reality, reduces average days late per client from 9.03 to 6.36, suggesting a meaningful share of the lateness in the real schedule is attributable to resource allocation rather than staffing levels alone. Current staffing levels are insufficient to support the active client base under the guidelines, with 11 to 12 nurses per weekday required to deliver schedules that adhere closely to them. New client first visits are identified as a significant capacity bottleneck, and spare weekend capacity, which is routinely underutilised, is shown to provide meaningful relief across a range of staffing levels.

1 Introduction

Delivering community palliative care requires balancing a limited nursing workforce against the needs of clients whose conditions vary considerably and change over time. Working with Palliative Care South East (PCSE), this report develops a perfect information integer programming model to determine what an ideal visit schedule aligned with the clinical guidelines would look like, and what this implies for the nursing capacity required to achieve it. We find that the existing schedule departs from the guidelines, that scheduling more closely aligned to the clinical guidelines can achieve meaningful improvements in schedule quality without additional resources, and that current staffing levels fall short of what is needed to fully support the active client base under the guidelines.

2 Statement on Authorship and Acknowledgements

The model formulation and analysis contained within this report are the work of Thomas Caldecott, completed under the supervision of Professor Andreas Ernst. The code used to implement the model was developed with the assistance of Claude. The author would like to thank the team at PCSE for their support throughout the project, and in particular Jodi Lynch for their assistance in understanding the data and the operational context that motivated this work.

3 Background and Motivation

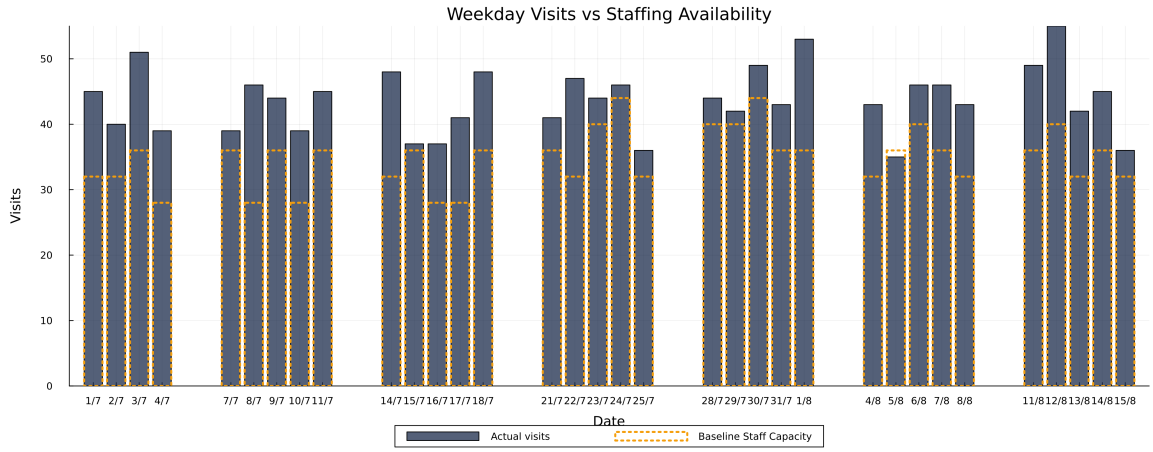

Figure 1: Baseline Staffing Capacity On Weekdays.

Figure 1 shows the total daily visit count across the study period, where new client admissions are assumed to be two visit slots, plotted alongside the nursing capacity of the baseline staff (nurses who actually showed up each day). The baseline capacity assumes each nurse can conduct four in-person visits per shift; both assumptions were derived from consultation with the PCSE staffing team. On most weekdays, the visit count exceeds this baseline capacity, placing enormous pressure on staff and the timeliness of care, with shortfalls routinely having to be backfilled by additional staff. Understanding whether this pressure could be reduced through better allocation of existing resources, and what staffing levels would be required to deliver visits in line with the clinical guidelines, were therefore key questions for PCSE.

The data used in this study was provided by PCSE and consists of de-identified client visit records; each record contains the de-identified client, date and palliative care phase of that visit. The study includes 487 clients and spans 1st July 2025 to 17th August 2025, with a prestudy window extending back to the 4th April 2025 to capture the visit history of clients already receiving care at the study start. Clients are classified as either pre-existing, those already receiving care at the commencement of the study period, or new, those first admitted during the study period. Further detail on how each client type is handled is given in Section 6.

The ideal visit frequency for a client is determined by their palliative care phase, a clinical classification assigned by the attending nurse at each visit. There are four phases: Stable, Deteriorating, Unstable, and Terminal. Clients in the Stable phase require a visit every three to four weeks; those in the Deteriorating phase require weekly visits; and clients who are Unstable or Terminal require daily follow up. It is also important to note that phases are only updated during a visit.

To address PCSE's questions, we develop an integer programming model that assigns client visits to days subject to nursing capacity and frequency constraints, with much further detail given in Sections 7. With the scheduling guidelines and client data in hand, we now turn to how schedule quality is measured.

4 Measuring Schedule Quality

Schedule quality is measured by counting days late, capturing the extent to which the clinical recommendations for visit intervals are exceeded and thereby representing potential impacts on client care. For each client c, each visit v establishes a due date for the next visit based on the client's recorded phase at that time. Let $f_{c,v}^{\max}$ denote the maximum recommended interval (in days) between visits associated with that phase. If the next visit in the study period occurs A_v days later, then the late days accrued for that visit pair are defined as $\max(0, A_v - f_{c,v}^{\max})$. If no subsequent visit occurs before the end of the study period, we assume that the client will be visited immediately after the study period ends and late days are calculated accordingly.

For pre-existing clients whose most recent prestudy visit occurred sufficiently long before the study start, a phase-specific buffer is applied to the maximum recommended interval for their study period visit: 7 additional days for clients last seen as Stable, 3 days for Deteriorating, and 0 for Unstable and Terminal. This mirrors the buffer zone discussed in Section 6.6. Before turning to the model itself, we will examine the degree to which the existing schedule adheres to the ideal clinical guidelines.

5 Adherence to Ideal Scheduling Guidelines

The following analysis is based on visit data provided by PCSE for modelled clients during the study period (modelled clients are discussed in Section 6.1; roughly those without significant clinical considerations), and serves to motivate the optimisation by illustrating where and how the real schedule deviates from the ideal guidelines.

We focus on two transition types: Stable-to-Stable and Deteriorating-to-Deteriorating. A transition type refers to the phase recorded at a visit and the phase recorded at the subsequent visit. These are two of the most common transition types and the ones for which the guidelines are most clearly defined. Other transitions, such as Stable-to-Deteriorating, may involve clinical considerations that bring visits forwards and so are harder to interpret in this context. It should also be noted that visits with no subsequent visit are not represented in the following graphs.

Figure 2 shows the distribution of inter-visit gaps for all Stable-to-Stable visit pairs during the study period. The guidelines specify a maximum gap of 28 days between visits for clients in the Stable phase. As the figure shows, the large majority of Stable-to-Stable visits occur well before the 21 to 28 day window set by the the guidelines, indicating that Stable clients may be being overserviced with respect to the ideal scheduling guidelines.

Figure 3 shows the equivalent distribution for Deteriorating-to-Deteriorating visit pairs. The guideline specify a maximum gap of 7 days for clients in the Deteriorating phase. In contrast to the above result, a considerable proportion of Deteriorating-to-Deteriorating occur beyond this threshold, suggesting that many Deteriorating clients are not being seen as frequently as their condition warrants. Additionally, we see that many Deteriorating clients are seen noticeably more frequent than weekly, though without further clinical context it

is difficult to determine whether this reflects overservicing or clinical needs. We also note that some of longer gaps may be attributable to hospitalisations, which are not recorded in the data available to us and therefore cannot be accounted for here.

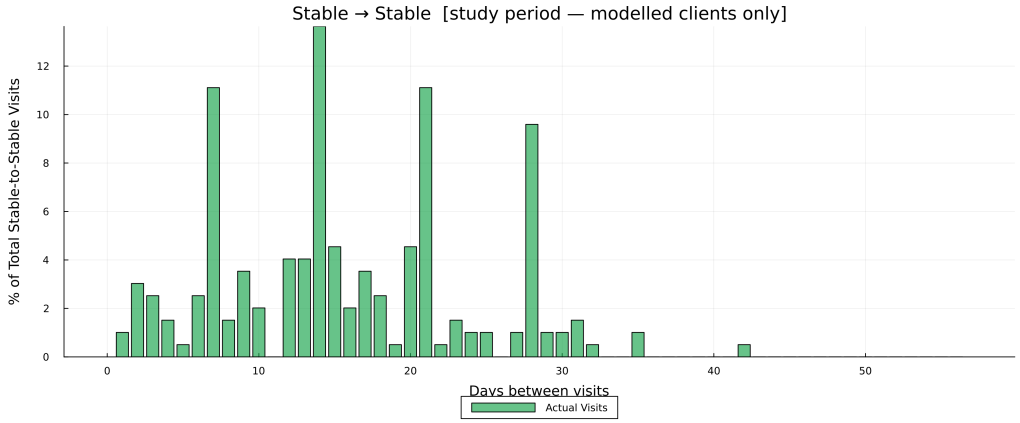

Figure 2: Stable-to-Stable Visit Histogram

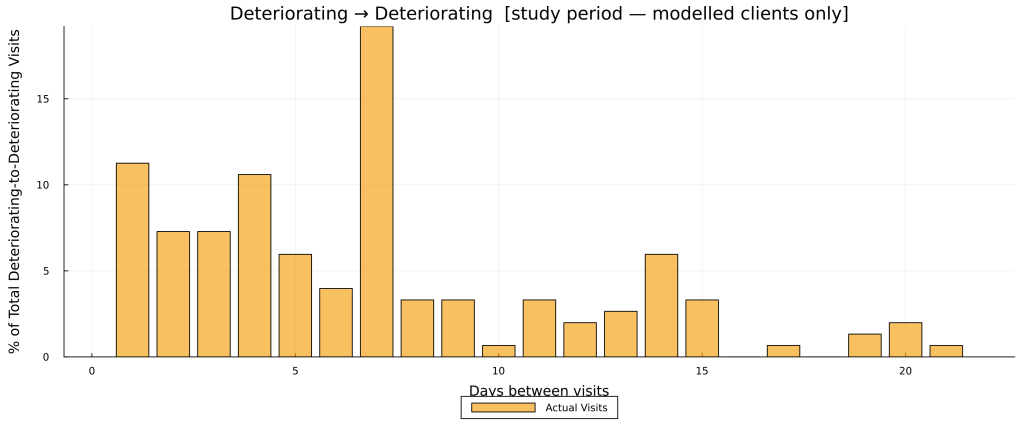

Figure 3: Deteriorating-to-Deteriorating Visit Histogram

Taken together, these figures demonstrate that the existing schedule does not closely follow the ideal clinical guidelines even for the most straightforward transition types. Therefore, we expect that a model scheduling to the ideal guidelines would lengthen the time between stable visits and shorten it for some deteriorating clients, decreasing and increasing the number of visits correspondingly. With this context established, we now turn to the development of the model itself.

6 Modelling Choices

Broadly put, with the aim of minimising the total number of late visits, the model assigns each client's visits to specific days across the nominated study period, subject to staffing, fairness, and visit frequency constraints (derived from the scheduling guidelines). A number of choices were made in translating this problem into a mathematical model, most of which are described and justified below.

6.1 Modellable Clients

Of the 487 clients in the dataset, not all were suitable for optimisation. Clients who were assessed at least once as Unstable or Terminal during the study period, who were discharged or died within the study window, or who had seven or more visits in the study period were excluded. These clients often require care that is more nuanced than what a frequency-based scheduling model can sensibly capture, and in some cases introduced technical complications related to the minimum spacing constraints discussed in Section 6.5. All clients excluded from optimisation had their visits fixed to the dates on which they actually occurred, so while not subject to optimisation, they still counted towards the daily capacity constraints and consumed nurse capacity on the relevant days.

After exclusions, 327 clients remain for optimisation. It should be noted that although the excluded group only contains a third of all clients from the study period, they account for approximately half of all visits in the actual schedule over the study period.

6.2 Phase Trajectories and Forward-Filling

A key characteristic of the model is that client phases are associated with visits rather than calendar dates, the limitations of such an approach are discussed in Section 10.While the model is free to determine when visits occur for modelled clients, the phase associated with each visit is fixed to what was recorded in the dataset. We refer to this sequence of phases across a client's visits in chronological order as their phase trajectory.

For pre-existing clients, the trajectory begins with their most recent prestudy visit, with remaining entries corresponding to their visits within the study period. For new clients, the trajectory begins with their first visit in the study period. In both cases, all modelled clients are treated as having exactly seven visits in their trajectory.

In practice many clients had fewer than seven visits in their trajectory. To avoid encoding the underservicing we observed in the real schedule directly into the model, the trajectory of every client was extended to seven visits by forward filling the last recorded phase. Seven was chosen to keep model runtime manageable, but still allow for clients deteriorating over the entire study period to have roughly weekly visits.

Importantly, the model is not required to use all visits. A visit that cannot be feasibly scheduled within the horizon without violating the minimum spacing requirements (discussed below) is left unassigned, meaning the forward fill gives the model freedom to schedule as many visits as the constraints permit rather than artificially inflating the final visit count.

6.3 Anchoring Pre-existing Clients

For pre-existing clients, the most recent prestudy visit is used as the reference point from which their study period visits are scheduled with respect to the frequency guidelines. Without this anchor, the model would have no awareness of prior visit history and would treat every pre-existing client as though they were new at the start of the study period producing schedules that do not reflect realistic client loads.

6.4 New Clients First Visits

For new clients, the date of their first visit is fixed to when it actually occurred during the study period. As with the prestudy anchor for pre-existing clients, this keeps the schedule grounded in reality, and avoids the model introducing new clients at arbitrary times without the referral data that would be needed to do so meaningfully. All their remaining visits are then freely optimised. Incorporating referral data to inform first visit timing is left as a potential extension, discussed further in Section 11.

6.5 Phase Transition Rule

The clinical scheduling guidelines specify a maximum desired interval to the next visit based on the client's current palliative care phase. A minimum interval is also applied, to prevent the model from scheduling visits unrealistically close together. These intervals are not applied per phase alone, but per phase transition. That is, the desired window (formed by the minimum and maximum intervals) for when the next visit should be depends on both the phase of the current visit and the next visit. These are summarised in Table [X] 1.

Table 1: Transition-Window Table

Transition Type	**Min days ($f^{\min}$)**	**Max days ($f^{\max}$)**
Stable → Stable	21 days	28 days
Stable → Deteriorating	14 days	21 days
Deteriorating → Deteriorating	3 days	7 days
Deteriorating → Stable	3 days	7 days

The motivation for using transition-specific windows is to better reflect the clinical reality. For example, a client who was last seen as Stable, but who has since deteriorated warrants a considerably earlier follow-up than the Stable guidelines alone would suggest. In practice, this follow-up is prompted by a crisis call, a mechanism not captured by the model or identifiable in the data.

Among the transition types, only Stable-to-Deteriorating has a meaningfully different window than what a per phase approach to the guidelines would produce. In the observed data, both Stable-to-Stable and Stable-to-Deteriorating transitions occur considerably earlier than what the guidelines suggest; Stable-to-Deteriorating visits are earlier still. Coupling this empirical support with the clinical justification of the crisis call mechanism, we find it reasonable to impose the tighter visit window on that transition. In contrast, Deteriorating-to-Stable was left unchanged, as there is no comparable clinical justification for visiting later given that phase

improvements only update during visits.

6.6 Buffer Zone

For some pre-existing clients, a small tolerance is added to the maximum desired interval for the first visit in the study period. This applies to clients whose gap from their last prestudy visit to the start of the study period is: at least 21 days for those with their last prestudy visit being Stable, 3 days for Deteriorating, and 1 day for all other phases. For these clients, rather than using the standard desired maximum interval, the model sets the effective desired maximum to the aforementioned gap plus a small additional tolerance of 3 days for clients whose prestudy visit was in the Deteriorating phase, and 7 days for those in the Stable Phase. No tolerance is applied for clients whose prestudy visit was in the other phases.

Without this adjustment, the model would often needlessly accumulate lateness penalties from the start of the schedule and be incentivised to support an unrealistic concentration of visits at the start of the schedule. A more appropriate implementation of this adjustment is discussed in Section 11.

6.7 Key Assumptions

Several simplifying assumptions underlie the model, some of which have been noted in earlier sections and are collected here for clarity. Each nurse is assumed to be capable of four in-person visits per shift, and nurses are treated as identical and interchangeable. New client's first visits are assumed to consume two visit slots rather than one, reflecting the additional time an initial assessment requires. Clients are assumed to be available on any day the model assigns them a visit. Finally, phases are assumed to be associated with visits rather than dates and only update at the point of a visit. With these choices and assumptions established, we now present the model formulation.

7 Model Formulation

7.1 Notation

Sets

C	set of modelled clients
V	set of visit indices $\{1, \ldots, \bar{v}\}$, where $\bar{v} = 7$ for all clients
D	set of schedulable days $\{1, \ldots, T\}$, spanning from the earliest prestudy anchor date across all pre-existing clients to the end of the study period
D_{pre}	subset $\{1, \ldots, 90\}$ of D corresponding to days before the study period begins

Parameters

$\bar{v}$	maximum number of visits per client ($\bar{v} = 7$)
T	total number of days in D ($T = 138$)

$f^{\min}_{c,v}$	minimum gap (days) between visit v and visit $v{+}1$ of client c
$f^{\max}_{c,v}$	maximum desired gap (days) between visit v and visit $v{+}1$ of client c
N_d	number of nurses available on day d (set arbitrarily large for $d \in D_{\text{pre}}$)
e_d	visit-slot capacity consumed by excluded clients on day d (new visits still count twice)
$k_{c,v}$	capacity weight of visit v for client c ($k_{c,1} = 2$ for new clients, 1 otherwise)
$w_{c,v}$	phase-based lateness penalty weight for visit v of client c, discussed in Section 8
M	sufficiently large scalar penalty for early termination

Variables

The primary decision variables are:

$y_{c,v,d} \in \{0,1\}$	1 if visit v of client c is scheduled on day d
$f^{+}_{c,v} \geq 0$	lateness penalty variable for the gap between visit v visit $v+1$ for client c
$p_{c,v} \in \mathbb{Z}_{\geq 0}$	penalty variable for early termination after visit v of client c

The following auxiliary variables are defined through constraints:

$x_{c,v} \in \{0,1\}$	1 if visit v of client c is assigned within the horizon
$t_{c,v} \in \mathbb{Z}_{\geq 0}$	day on which visit v of client c is scheduled
$l_{c,v} \in \{0,1\}$	1 if a feasible follow-up visit can occur after visit v before T

7.2 Formulation

$$\min \quad \sum_{c\in C}\sum_{v\in V} w_{c,v}\, f^{+}_{c,v} + \sum_{c\in C}\sum_{v\in V} M\,(\bar{v}-v+1)\,p_{c,v} \tag{1}$$

$$\text{s.t.} \quad x_{c,v} = \sum_{d\in D} y_{c,v,d} \qquad \forall\, c\in C,\ v\in V \tag{2}$$

$$t_{c,v} = \sum_{d\in D} d\cdot y_{c,v,d} \qquad \forall\, c\in C,\ v\in V \tag{3}$$

$$l_{c,v} = \sum_{d=1}^{T-f^{\min}_{c,v}} y_{c,v,d} \qquad \forall\, c\in C,\ v\in V \tag{4}$$

$$\sum_{d\in D} y_{c,v,d} \leq 1 \qquad \forall\, c\in C,\ v\in V \tag{5}$$

$$x_{c,v+1} \leq x_{c,v} \qquad \forall\, c\in C,\ v\in\{1,\dots,\bar{v}-1\} \tag{6}$$

$$y_{c,1,d^{*}_{c}} = 1 \qquad \forall\, c\in C \tag{7}$$

$$y_{c,v,d} = 0 \qquad \forall\, c\in C,\ v\in\{2,\dots,\bar{v}\},\ d\in D_{\text{pre}} \tag{8}$$

$$x_{c,v+1}=1 \ \Rightarrow\ t_{c,v+1}-t_{c,v} \geq f^{\min}_{c,v} \qquad \forall\, c\in C,\ v\in\{1,\dots,\bar{v}-1\} \tag{9}$$

$$\sum_{c\in C}\sum_{v\in V} k_{c,v}\cdot y_{c,v,d} \ \leq\ 4N_d - e_d \qquad \forall\, d\in D \tag{10}$$

$$f^{+}_{c,v} \geq \left(t_{c,v+1}-t_{c,v}\right) - f^{\max}_{c,v} \qquad \forall\, c\in C,\ v\in\{1,\dots,\bar{v}-1\} \tag{11}$$

$$x_{c,\bar{v}}=1 \ \Rightarrow\ f^{+}_{c,\bar{v}} \geq \left(T-t_{c,\bar{v}}\right) - f^{\max}_{c,\bar{v}} \qquad \forall\, c\in C \tag{12}$$

$$p_{c,v} \geq x_{c,v} - x_{c,v+1} - (1-l_{c,v}) \qquad \forall\, c\in C,\ v\in\{1,\dots,\bar{v}-1\} \tag{13}$$

$$\sum_{v\in V} f^{+}_{c,v} \leq 15 \qquad \forall\, c\in C \tag{14}$$

$$y_{c,v,d} \in \{0,1\}, \quad f^{+}_{c,v} \geq 0, \quad p_{c,v} \geq 0, \quad x_{c,v} \in \{0,1\}, \quad t_{c,v} \in \mathbb{Z}_{\geq 0} \quad \forall\, c, v, d \tag{15}$$

where d^{*}_{c} denotes the fixed first visit date for client c, the most recent prestudy visit date for pre-existing clients, and the date of the first study period visit for new clients.

7.3 Discussion of Formulation

Objective (1)**.** The objective has two terms. The first minimises the total lateness penalty across all clients and visits, weighted by $w_{c,v}$, a phase-based weight that reflects the clinical severity of a late visit. Specific values are given in Section 8. The second term penalises early termination, cases where a client is assigned fewer than $\bar{v}$ visits despite a follow-up being feasible within the horizon. The factor $(\bar{v} - v + 1)$ ensures that losing an earlier visit, which forfeits all subsequent visits, is penalised more heavily than losing a later one. M is set sufficiently large that the solver will always prefer any reduction in early termination over any reduction in lateness.

Auxiliary variable definitions (2)–(4)**.** Constraints (2) and (3) define $x_{c,v}$ and $t_{c,v}$ in terms of the primary decision variables $y_{c,v,d}$. Constraint (4) defines $l_{c,v}$ as an indicator of whether visit v is scheduled sufficiently early that a follow up visit satisfying the minimum interval requirement could still occur before the horizon ends. This is used in the early termination constraint (13) to distinguish between avoidable and unavoidable early terminations.

Visit assignment (5)–(6)**.** Each visit index is assigned to at most one day (5), and visit indices are assigned consecutively (6), a later visit cannot be assigned unless all earlier visits are also assigned.

First visit fixing (7)–(8)**.** The first visit for each client is fixed to d^{*}_{c} as described in Section 6. For pre-existing clients this is their most recent prestudy visit; for new clients it is the date of their first study period visit. Constraint (8) ensures that all subsequent visits are confined to the study period, preventing the model from scheduling visits 2 through $\bar{v}$ in D_{pre}.

Minimum interval (9)**.** When two consecutive visits are both assigned, the gap between them must meet the transition-specific minimum interval (discussed in Section 6.5. This is enforced as a hard constraint to ensure the model does not place visit unrealistically close together.

Capacity (10)**.** The total visit load on any day must not exceed available nursing capacity net of visits already consumed by excluded clients. The factor of 4 reflects the assumption that each nurse conducts at most four visits per day, and admissions take two visit slots. Staffing profiles across the different scenarios considered are described in Section 8.

Lateness (11)–(12)**.** Constraint (11) defines the lateness variable for consecutive assigned visits. Constraint (12) handles the study period tail: for the last assigned visit of each client, lateness is measured against the end of the horizon rather than a subsequent visit.

Early termination (13)**.** $p_{c,v}$ is activated when visit v is the last assigned visit for client c and $l_{c,v}$ confirms a follow-up would have been feasible. This prevents the model from artificially truncating schedules that would not be captured in lateness accumulation.

Lateness cap (14)**.** Total lateness penalty per client is bounded at 15 to ensure the model distributes lateness reasonably across the client base rather than concentrating it on a small number of clients.

8 Implementation

The model was implemented in Julia using the JuMP modelling framework and solved with the Gurobi optimiser.

All scenarios were solved to within a 1% MIP gap.

The model was run across five weekday staffing scenarios: Actual, Planned, 11 Nurse, 12 Nurse, and Shadow. The Shadow scenario does not correspond to a staffing level in the traditional sense. Rather, the daily capacity on each day is set to match the total number of visits that actually occurred in reality. This is achieved by setting $4N_d$ equal to the adjusted visit count (where new visits count twice), and then subtracting e_d as usual. This can loosely be interpreted as the actual nursing workforce plus any additional staff required to deliver the visits that occurred. By holding daily capacity fixed to what reality used, the Shadow scenario isolates the effect of optimisation alone. Any improvement in schedule quality relative to the actual schedule can be attributed to a more appropriate allocation of the same resources, though it is not without its limitations as a benchmark, as discussed in Section 10.

The Planned and Actual scenarios use the PCSE nursing staff roster, the former being the intended roster and the latter being who actually turned up on the day, with N_d set to the corresponding staffing level. The 11 Nurse and 12 Nurse staff scenarios are hypothetical, with a fixed number (11 or 12) of nurses available every weekday across the entire study period, used to explore the effect of increased nursing capacity on schedule quality. Unless otherwise stated, weekend capacity is fixed to the adjusted visit count observed on each weekend day.

The lateness penalty weight w_c, v is phase-based, taking the more severe of the phases at visit v and $v+1$. Taking values of 1 for Stable, 10 for Deteriorating, and 100 for Unstable/Terminal, ensuring the model prioritises clients with greater clinical need. For the final visit of each client, where no subsequent phase exists, the penalty is based on the current phase alone. Moreover, the first visit carries an additional multiplier of 50 to break the model's indifference to accumulating lateness on the first gap versus later gaps. The early termination penalty M was set to 10^9. Lastly, the frequency parameters $f_{c,v}^{\min}, f_{c,v}^{\max}$ are set per the transition rules, described in Section 6.5. With the model and implementation fully described, we now present the results across the five staffing scenarios.

9 Discussion of Results

Unless otherwise stated, all results in this section are reported for modelled clients only. Schedule quality is reported as average days late per client, and staffing levels are reported as average nurses per weekday. All results correspond to optimised model schedules across the five staffing scenarios, except for the real schedule from the data, which is denoted as Shadow on the graphs (as this staffing scenario had the same resources that reality used).

9.1 Effect of Optimisation at Same Resource Level

Figure 4 presents the Shadow staffing scenario, plotting average days late per client against average nurses per weekday for both the real schedule (shadow) and the optimised model schedule (optimised shadow).

The optimised Shadow schedule reduces average days late per client from 9.03 to 6.36, a substantial improvement achieved with the same resources reality used. This improvement is likely driven by two complementary

effects: Firstly, a redistribution of deteriorating client visits towards the weekly guideline, mostly bringing forward those seen too infrequently and slightly delaying those seen more often than required. Secondly, a shift in stable client visits toward or beyond the recommended three-to-four-week interval, which frees up capacity that can be directed towards clients with greater clinical need. The magnitude of the improvement suggests, as seen in Section 5, a notable portion of the lateness in the real schedule may be attributable to the allocation of existing resources rather than staffing levels alone.

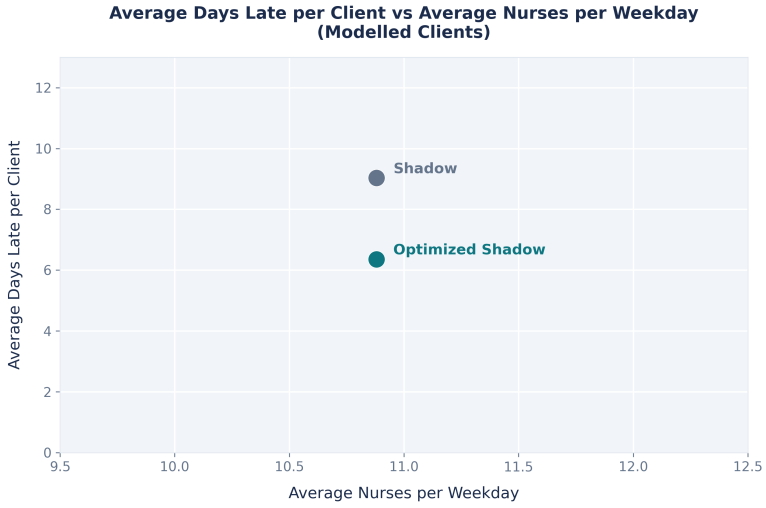

Figure 4: Real and Optimised Schedules (Shadow)

9.2 Effect of Staffing Level

Figure 5 extends the analysis to all staffing scenarios. Unsurprisingly, staffing level has a significant effect on average days late per client, with the relationship across scenarios approximately linear, suggesting each additional nurse per weekday yields a reduction of roughly 4 to 5 average days late.

When the model is run at the Actual (who showed up) and Planned staffing levels, average days late increases sharply relative to the Shadow scenario (same resources as reality). It should also be noted at these staffing levels, approximately 50 and 20 modelled clients respectively received only a single visit early in the schedule and were not seen again, suggesting the model ran out of capacity to service them. This aligns with discussion in Section 5, these staffing levels already required backfilling in reality, and the results here suggest that even when visits are better aligned with the ideal scheduling guidelines, they remain insufficient to support the active client base seen during the study period.

In contrast, increasing staffing by one to two nurses per weekday beyond planned levels, 11 and 12 nurse scenarios, produces significant reductions in average days late. Suggesting that a modest increase in staffing alongside scheduling better aligned with the clinical guidelines could go a considerable way toward relieving the capacity pressure identified in Section 3.

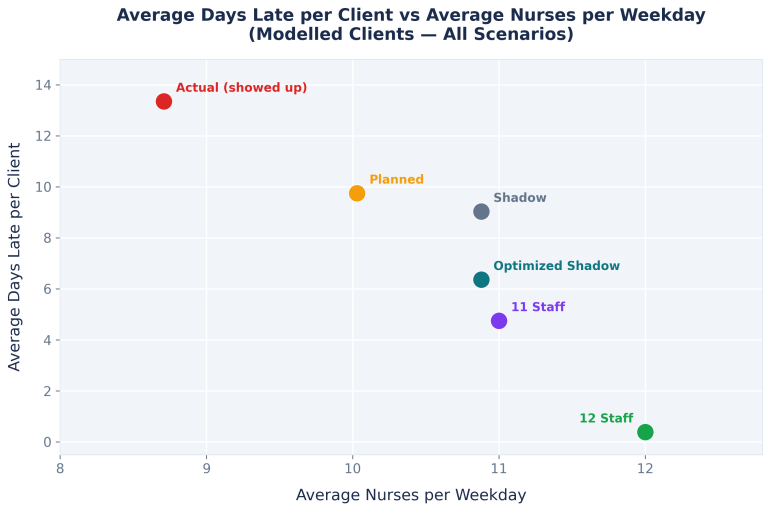

Figure 5: All Optimised Staffing Scenarios

9.3 Effect of New Client Capacity

Figure 6 illustrates the effect of reducing the capacity cost of new client's first visits from two slots to one, a change also applied to excluded clients via e_d. The reduction in average days late per client is substantial, highlighting new admissions as a meaningful capacity bottleneck. This is consistent with the intuition that initial assessment visits are placing a significant burden on daily nurse activities, and suggests that steps towards streamlining new admissions, alongside scheduling better aligned with guidelines, could have a serious impact on relieving capacity pressures.

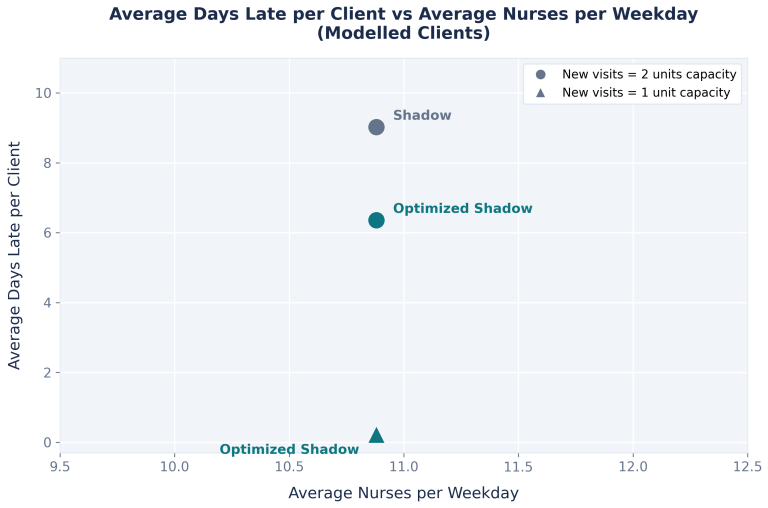

Figure 6: Impact of Streamlining New Admissions

9.4 Weekend Capacity

Figure 7 shows the daily visit count on weekends alongside available baseline nursing capacity (who showed up), analogous to the weekday figure in Section 3. In contrast to weekdays, weekend capacity is frequently underutilised, with spare slots available on most days. Figure 8 shows the effect of allowing the model to utilise this spare weekend capacity across all staffing scenarios.

Figure 7: Baseline Staffing Capacity On Weekends

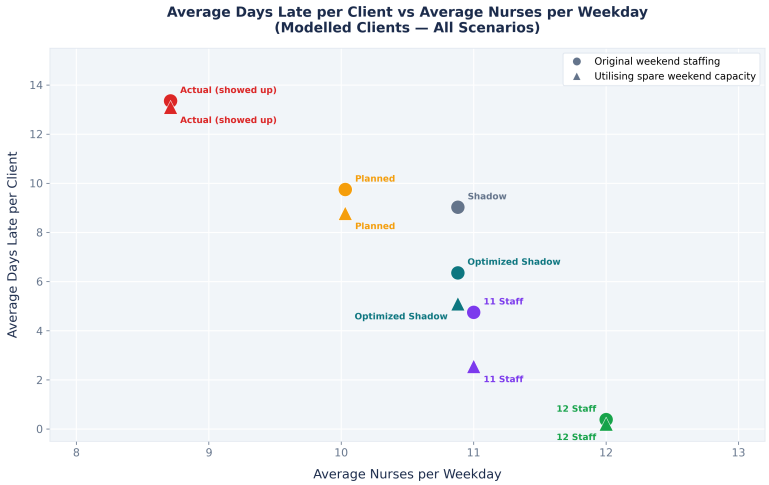

Figure 8: Effect of Utilising Spare Weekend Capacity

The benefit of additional weekend capacity is most pronounced for the middle level staffing scenarios, where weekday capacity is tight but not exhausted. In the Actual scenario, weekday capacity is so constrained that a small number of additional weekend visits does little to address this underlying shortfall. Conversely, the 12 Nurse scenario is already performing well under the original weekend capacity, leaving little lateness for more weekend visits to absorb. For the intermediate scenarios, weekday capacity to produce a workable schedule, but not sufficient to eliminate lateness entirely, making the additional weekend slots more effective. Unsurprisingly,

this suggests that making full use of existing weekend capacity is often an easy avenue for improving schedule quality. While the results suggest several avenues for operational improvement, they should nonetheless be interpreted in the context of the model's limitations.

10 Limitations

Several of the model's core assumptions warrant scrutiny. Both the assumptions that nurses are identical and interchangeable, and that each can conduct four in-person visits per shift, are reasonable and achievable based on consultation with PCSE. However, because the model takes no account for travel time, it is possible that some schedules are not operationally feasible. Clients are assumed to be available on any assigned day. In practice, this is not always true: clients may cancel, reschedule, or even be in hospital. However, without the relevant data, these factors could not be incorporated into the model. Lastly, phases are tied to visits rather than calendar dates and are assumed to update upon visits. By rescheduling visits to different days, the model implicitly assumes the phases recorded at each visit still reflects the client's condition on the rescheduled day, which obviously may not hold in practice.

The model operates under perfect information: phase trajectories and staffing levels are all known in advance and the model exploits this knowledge. Therefore, the results only represent an upper bound on achievable schedule quality. The Shadow scenario is a clear illustration of this, while useful for isolating the effect of optimisation, the model knows the total daily capacity used in advance, whereas in reality that capacity is only achieved reactively, giving it a significant planning advantage as the benchmark scenario.

The study covers a single period, which by PCSE's account, was unusually busy. Furthermore, palliative care also exhibits seasonal variation in both client phases and staffing levels. The results should therefore be interpreted with some caution in terms of their generalisability to other operating periods. A broader analysis across multiple study periods would strengthen findings considerably.

Finally, a substantial proportion of clients had their visits fixed to what actually occurred and were not optimisable. Consequently, the model only reflects the best achievable schedule for a subset of clients rather than the full population. Several of these limitations point naturally to directions for future work.

11 Next Steps

The most natural extension is to move from perfect-information model to a probabilistic one. A straightforward implementation within the existing framework would be to use a Markov approach to simulate client phase trajectories, estimating phase transition probabilities from the observed visit data. For pre-existing clients, a prestudy visit could serve as the anchor from which a full trajectory is simulated, while referral data could guide the expected trajectory of new clients based on their urgency classification. This approach would particularly improve forward filling. However, excluded clients would still need to be fixed, leaving room for further refinement.

Referral data could be incorporated as a zero-capacity visit with a tightly penalised frequency window for the next visit, calibrated to the urgency classification in the referral. This would allow the model to determine first visit timing for new clients endogenously rather than fixing it to observed dates.

Richer visit data, capturing cancellations, rescheduling and hospitalisations, would also considerably refine the model. Furthermore, identifying when visits were changed for clinical reasons, though potentially difficult to obtain, would allow for much more nuanced analysis on overservicing.

Sensitivity analysis on client load would also be valuable. During this project, one approach was attempted but abandoned due to methodological and time constraints: randomly sampling clients and either removing them or adding them (with slightly shifted visit dates) to simulate higher or lower client loads at fixed staffing levels. In practice this frequently caused infeasibilities, as adding enough clients pushed many excluded client visits onto weekends, exhausting weekend nursing capacity and rendering the model infeasible. A more principled simulation of client load variation remains an important direction for future work.

Finally, an immediate improvement would be a refinement of the buffer zone applied to pre-existing clients' first visits. The current implementation uses a phase-specific minimum threshold as the trigger condition rather than the maximum interval, which is inconsistent with the intent of the buffer. Moreover, the buffer thresholds currently use the per phase rule instead of the transition-window rules applied elsewhere in the model. Correcting both would make the treatment of pre-existing clients more internally consistent. While these directions primarily refine the existing architecture, alternative model formulations might produce more fruitful results.

12 Summary

This report set out to understand what an ideal visit schedule would look like for PCSE under the clinical guidelines, and what staffing levels would be required to deliver it. The analysis finds that the existing schedule departs from those guidelines, and that reallocating the same resources more appropriately noticeably reduces average days late per client, demonstrating that scheduling practice alone has a meaningful impact on schedule quality, independent of the staffing levels available. At current staffing levels, the model is unable to fully support the active client base under the guidelines, though increasing 1 to 2 nurses per weekday over planned levels produces schedules that adhere much more closely to them. Making better use of existing weekend capacity and streamlining new client admissions are also identified as avenues for low-cost improvement in schedule quality.

The results are best understood as an upper bound on what is achievable under the given guidelines and staffing levels. This reflects the model's reliance on perfect information and the fixing of a substantial proportion of clients to their actual visit dates. Accordingly, Section 11 offers a path toward a more operationally realistic model, with a probabilistic treatment of phase trajectories and richer referral and visit data as the most consequential directions for future work.