Service Science

Publication details, including instructions for authors and subscription information:
http://pubsonline.informs.org

Dynamic Scheduling of Recurring Multisession Appointments with Heterogeneous Clients

Reut Noham, Karen Smilowitz

To cite this article:
Reut Noham, Karen Smilowitz (2026) Dynamic Scheduling of Recurring Multisession Appointments with Heterogeneous Clients. Service Science

Published online in Articles in Advance 23 Jun 2026

. https://doi.org/10.1287/serv.2026.0025

informs
https://pubsonline.informs.org/journal/serv

SERVICE SCIENCE
Articles in Advance, pp. 1–17
ISSN 2164-3962 (print), ISSN 2164-3970 (online)

Dynamic Scheduling of Recurring Multisession Appointments with Heterogeneous Clients

Reut Noham,[a,*] Karen Smilowitz[b]

[a] School of Industrial and Intelligent Systems Engineering, Tel-Aviv University, Tel Aviv-Yafo 6997801, Israel; [b] Department of Industrial Engineering and Management Sciences, Northwestern University, Evanston, Illinois 60208

*Corresponding author

Contact: reutno@tauex.tau.ac.il, https://orcid.org/0000-0002-0462-5097 (RN); ksmilowitz@northwestern.edu, https://orcid.org/0000-0003-2905-4566 (KS)

Received: February 26, 2026
Revised: May 11, 2026
Accepted: May 18, 2026
Published Online in Articles in Advance: June 23, 2026

https://doi.org/10.1287/serv.2026.0025

Abstract. Clients seeking paramedical and rehabilitation services require recurring treatment sessions over an extended period. Unlike single-visit problems, these services must assign each accepted client to a fixed, recurring day–time slot that remains occupied throughout the treatment program. This structure creates long-term capacity commitments that limit future scheduling flexibility and complicate acceptance decisions, particularly when clients differ in availability and required program durations. Motivated by an early intervention program for infants and toddlers with developmental delays, we study scheduling policies designed to address this combination of heterogeneity, uncertainty, and recurrence constraints. We model the multisession appointment scheduling problem as a Markov decision process in which requests arrive sequentially and decisions must consider both immediate feasibility and the long-term implications of blocking a slot across many periods. Our analysis identifies key structural elements of the scheduling decision, including a slot-selection guideline that assigns accepted clients to the least popular feasible slot and a duration-based threshold that characterizes acceptance behavior. These insights highlight the value of preserving flexibility and anticipating demand when scheduling recurring appointments under uncertainty. Building on these results, we develop a heuristic that groups schedule states into occupancy categories and applies simplified acceptance thresholds. Computational experiments show that this anticipatory approach outperforms first come, first served benchmarks, particularly when slot popularity is uneven or program durations vary widely. Whereas we do not model health outcomes directly, prior research links improved access, timely initiation, and continuity of care with better therapeutic results, underscoring the broader potential impact of more efficient scheduling.

Supplemental Material: The online appendix is available at https://doi.org/10.1287/serv.2026.0025.

Keywords: OR in health services • multisession appointment scheduling • recurring appointments • dynamic scheduling • Markov decision process

1. Introduction

In health services such as paramedical therapies, rehabilitation, or mental health services, heterogeneous clients with diverse needs attend frequent appointments over an extended period of time. These clients can be heterogeneous in various aspects, such as their availability, preferred time slots, and the duration of those services. Whereas single (one-time) visits for health services may be less sensitive to the timing of appointments, for health services requiring multiple recurring appointments, both clients and providers are much more sensitive to the timing of appointment slots. Because of their high frequency, such services become part of the clients' routine, and therefore, it is critical to account for their availability and timing preferences. From the providers' perspective, multiappointments also require long-term scheduling commitment, which limits their ability to meet the preferences and availability of future clients.

This paper considers multisession services in which the timely start of treatment, continuity of care, and persistence in treatment over time can contribute to better quality of care (Haggerty et al. 2003) and, thus, significantly affect clients' health outcomes. Early intervention (EI) services are a prime example of services characterized by these factors and are, therefore, the focus of this paper. Designed to support infants and toddlers with developmental delays or disabilities, EI services include interventions such as physical therapy, occupational therapy, speech therapy, and other services aimed at strengthening the acquisition of age-appropriate skills to enhance development, addressing the needs of both child and family. In the United States, as per Public Law 105–17 (1997), the services are provided for free or at a reduced cost and are publicly supervised. The early childhood years present a critical window within which brain plasticity more intensively facilitates the ability to learn new skills such that studies show that engaging EI therapies early and with regular frequency yields greater improvements than interventions at a later age for conditions known to adversely affect developmental progress such as cerebral palsy, Down syndrome, or autism (Majnemer 1998, Sheinkopf and Siegel 1998, Corsello 2005, Bonnier 2008, Spittle and Morgan 2025). Therefore, addressing developmental issues as early as possible is critical for significantly impacting the future quality of life for children receiving these services.

The scheduling process for EI services typically begins with an evaluation that culminates in a suggested treatment plan, that is, a set of required services (therapies), their frequency, and duration. Under ideal conditions, a given healthcare system would have sufficient capacity to meet the needs of each client and provide both evaluation and treatment in a timely manner. In this ideal scenario, sessions would also be scheduled at convenient times for the clients to interfere as little as possible with their other commitments and increase schedule adherence. In practice, however, system capacity is limited, and matching supply (available treatment slots) and demand (clients' requests for treatment) is difficult, especially when certain time blocks are more popular than others (Gupta and Wang 2008). The challenge of balancing clients' needs, operational costs, and health outcomes motivates our research, which is geared toward addressing and improving each of these aspects.

In this paper, we study the multisession appointment scheduling (MSAS) problem for a single provider, in which heterogeneous clients with different time preferences and treatment durations request service over time. Upon the arrival of each request, the provider must decide whether to accept the client and, if so, how to assign the client to a time slot. We investigate dynamic scheduling policies that adopt an anticipatory approach, using probabilistic information about future requests to guide current decisions. This contrasts with the largely myopic scheduling practices commonly used in paramedical and rehabilitation services, in which decisions are typically based only on current availability.

A defining characteristic of the MSAS setting is the requirement that each accepted client be assigned to a fixed, recurring day–time slot for the entire duration of the treatment program. This structure is common in long-term therapeutic and rehabilitation services, particularly in pediatric and early intervention settings, in which maintaining a consistent appointment time supports routine, adherence, and continuity of care. From an operational perspective, however, this hard recurrence requirement fundamentally limits scheduling flexibility: once a client is accepted, the same slot is blocked for multiple future periods and cannot be reallocated or reshuffled across periods. This persistent, slot-specific capacity commitment creates strong intertemporal coupling between scheduling decisions and fundamentally alters the nature of the acceptance problem. Unlike single-visit or flexible multiappointment settings, in which appointments are independent, sequential, or reschedulable across periods, accepting a client in the MSAS setting entails a long-term opportunity cost that depends jointly on the client's program duration and the future demand for the assigned slot. As a result, scheduling decisions must balance immediate feasibility with their long-term impact on access and capacity, leading to structural trade-offs that do not arise in more flexible appointment systems.

We focus on the single-provider MSAS problem as a core setting that captures these fundamental intertemporal trade-offs under uncertainty. Whereas many healthcare systems involve multiple providers, isolating the single-provider problem allows us to characterize the structural drivers of acceptance and assignment decisions before introducing additional layers of coordination. Our objective is to characterize and design scheduling policies that balance the key considerations inherent in multisession appointment systems, in particular, the trade-off between maximizing access for heterogeneous clients and efficiently utilizing long-term capacity. We further aim to understand how different prioritization objectives, such as client-oriented versus provider-oriented rewards, shape optimal scheduling behavior.

The contributions of this paper are as follows: First, we formulate the MSAS problem as a Markov decision process (MDP) that captures the long-term capacity commitments induced by recurring appointments together with heterogeneity in client availability and program duration. Second, we analyze two distinct reward structures, client dominance and provider dominance, and, under each, identify key structural components of the optimal policy. In particular, we characterize a duration-based threshold structure for acceptance decisions and establish a principled slot-selection guideline that prioritizes assigning clients to less frequently requested slots. Third, building on these structural insights, we develop a simple and implementable heuristic, the traffic light (TL) policy, that approximates the optimal decision logic, thereby significantly reducing computational complexity. Finally, through a comprehensive numerical study, we demonstrate that the proposed heuristic consistently outperforms commonly used first come, first served

(FCFS) policies, especially in settings with uneven slot popularity or high variability in program durations, and yields actionable managerial insights for anticipatory scheduling in healthcare services.

By explicitly modeling recurring appointments, our work highlights the operational consequences of long-term capacity commitments in healthcare services. In settings such as early intervention and pediatric rehabilitation, these commitments are closely tied to continuity of care, routine, and service accessibility. Beyond this specific context, the MSAS framework applies to a broader class of services characterized by long-lasting treatments and recurring appointments, for which providers face similar trade-offs between access, flexibility, and capacity utilization.

2. Related Literature

The literature on appointment scheduling in healthcare spans several related domains. In this section, we focus on three key streams that inform our work. The first concerns dynamic scheduling models, in which appointment decisions are made sequentially under uncertainty. Within this stream, we distinguish between models that assume homogeneous demand and those that incorporate patient-specific time preferences into the decision process. The third stream addresses multi-appointment scheduling, including repeated or interdependent appointments, which are central to many therapeutic and chronic care contexts. These bodies of work provide the methodological and conceptual foundations for our model.

2.1. Dynamic Appointment Scheduling

Under a dynamic approach, scheduling decisions adapt to the current state of the system and to probabilistic information about future demand. Models for dynamic decision making can significantly enhance the agility and responsiveness of healthcare services. Studies in this area implement dynamic approaches under various assumptions, service characteristics, and objectives. Gupta and Denton (2008) provide an early survey of outpatient appointment scheduling, emphasizing the importance of modeling uncertainty and system responsiveness. A more recent and comprehensive review by Ahmadi-Javid et al. (2017) classifies appointment systems by decision levels (strategic, tactical, operational) and highlights MDPs as a key framework for modeling online decisions in the presence of stochastic arrivals, service durations, and capacity constraints. Several studies implement MDP-based models for dynamic slot assignment in healthcare. Patrick et al. (2008) develop a model for scheduling diagnostic patients with varying urgency levels, aiming to meet target waiting times. Liu et al. (2010) consider dynamic scheduling in the presence of no-shows and cancellations with the objective of maximizing the clinic's long-run net reward (revenue). Lu et al. (2018) propose a dynamic admission policy that balances provider revenue with patient waiting experience. Similar to their work, we also model the trade-off between provider utilization and client outcomes, which is captured in our model by the rejection rate.

Additional studies implement MDP-based formulations to support sequential appointment decisions under uncertainty. For example, Kolisch and Sickinger (2008) model the allocation of radiology slots among multiple patient classes with stochastic arrivals; Lin et al. (2011) focus on sequential scheduling under no-show uncertainty; and Parizi and Ghate (2016) propose a dynamic, multiresource scheduling model that incorporates cancellations and overbooking. These models address important operational complexities but are limited to single-appointment settings without explicit modeling of patient-specific availability. Our work builds on this literature by extending dynamic scheduling models to a setting that combines recurring appointment requirements and client heterogeneity in availability by considering their preferences and treatment duration.

2.2. Incorporating Patient Preferences

A subset of dynamic scheduling models explicitly incorporates patient preferences into the decision process. Gupta and Denton (2008) note in their survey that most models at the time ignored patient scheduling preferences despite the fact that such preferences affect adherence and satisfaction. Several works address this gap by developing dynamic models that explicitly consider patient choice. Gupta and Wang (2008) study the problem of allocating slots to appointment requests in primary care clinics, describing it as an MDP. Given patient choices (requests that specify physician, day, and slot), their model determines which requests the clinic should accept to maximize long-term revenue. Whereas their model assumes that patients observe the system state at the time of booking, Wang and Gupta (2011) relax this assumption and present a framework for designing an adaptive appointment system. In their model, patients indicate a set of acceptable appointment options, and the clinic decides whether to accept a request and which slot to assign. The model includes an acceptance probability estimation component and a booking decision tool, formulated as a stochastic dynamic programming problem with the objective of maximizing clinic revenue and improving patient–provider matching.

Feldman et al. (2014) also incorporate patient preferences in an interday scheduling problem. They model a booking system in which, upon request, a set of appointment dates is offered to the patient. When the set of open dates is revealed, the patient may choose one of the dates or leave the system without scheduling an appointment. The model is used to find the optimal set of offered days based on the state of the schedule and patient preferences, accounting for cancellations and no-shows. Recent papers that also include patient preferences in dynamic appointment scheduling are Wang and Fung (2015), Li et al. (2018), and Wang et al. (2018).

Among these, our model is most closely related to the work of Gupta and Wang (2008) in that both address dynamic slot acceptance with patient-specific preferences. Whereas their setting is limited to one-time appointments, their structure for state-dependent decision making under patient choice forms the methodological basis for our model. We extend this approach to accommodate multiple recurring appointments, incorporating client-specific time preferences, and treatment duration constraints.

2.3. Multisession and Chronic Care Scheduling

The literature on multiple appointments usually focuses either on a series of visits held on the same day across different resources as in same-day, multistation scheduling problems (Han et al. 2024) or on interdependent appointments that must follow a specific sequence with lead times between them, such as diagnostic tests followed by consultations (Moradi et al. 2025). In these settings, appointments are linked through precedence constraints, waiting-list dynamics, and coordination across multiple resources. Another related stream considers visits over the course of several days, during which patients may revisit the same resource of treatment (Diamant et al. 2018, Marynissen and Demeulemeester 2019). Whereas all of these studies address multiappointment scheduling, most focus on short-term sequences or tightly linked care episodes. In contrast, a related stream of work, emerging primarily in the context of chronic disease management, considers recurring visits. These models are closer in structure to our setting. In this context, Deo et al. (2013) develop a finite-horizon stochastic dynamic model that allocates a limited number of periodic visit slots to patients and explicitly considers health outcomes in the objective. Because of capacity constraints, not all patients can be monitored in each period, and disease control depends on the interval between visits. The decision problem focuses on selecting which patients to schedule over time given stochastic disease progression and limited monitoring capacity. However, the model does not incorporate slot-level assignment decisions or recurring time commitments; capacity is allocated at an aggregate level without binding patients to specific day–time slots across periods. Savelsbergh and Smilowitz (2016) extend this line of work in a static, single-batch scheduling setting by incorporating time-of-day preferences and patient-specific no-show probabilities. They construct population appointment schedules over a finite planning horizon, explicitly modeling disease progression and attendance behavior. However, appointments are assigned independently in each period: patients are not bound to a fixed recurring time slot across periods, and slot identity does not persist over time. As a result, capacity is allocated flexibly period by period rather than through long-term slot-specific commitments.

In a different healthcare context, Saure et al. (2012) study dynamic multiappointment scheduling for radiation therapy. Their model captures treatments that span multiple consecutive sessions and formulates the problem as an infinite-horizon MDP. Whereas treatments require capacity over multiple days, capacity is aggregated across identical treatment units, and appointments are not tied to a fixed recurring slot over time. As a result, resource consumption is time-aggregated rather than slot-specific. Building on these prior models, we focus on a setting characterized by fixed slot-level recurrence under heterogeneous availability. In our model, once a client is accepted, the same day–time slot is reserved over multiple future periods. This hard recurrence structure creates persistent, slot-specific capacity commitments and induces intertemporal coupling that alters feasible state transitions and the opportunity cost of acceptance decisions. Such a structure is central in early intervention and rehabilitation services, in which maintaining a consistent appointment time is operationally and clinically important.

The challenges of multisession scheduling also appear in home healthcare (HHC), in which patients often require repeated visits over a treatment period. However, because of the decentralized nature of service delivery and because appointments in this setting occur in different locations, the substantial majority of the HHC literature addresses scheduling and routing problems that deal with requirements such as preferences and continuity of care (Cissé et al. 2017, Fikar and Hirsch 2017). Bard et al. (2014) use a mixed-integer linear program to assign patients to therapists and schedule weekly visits under a set of practical constraints. This continues the earlier work by Shao et al. (2012), and both of which assume demand is known over the planning horizon. Bennett and Erera (2011) introduce consistency constraints to ensure that patients receive care once a week over several weeks. They use a rolling-horizon, myopic scheduling approach to maximize the number of patients served, incorporating time preferences. These models consider continuity of care but do not address dynamic decision making under uncertain future demand, nor do they focus on assigning recurring slots at fixed times.

Across the broader appointment scheduling literature, the dominant modeling paradigm remains single-visit scheduling under uncertainty. Although a growing body of work addresses multiappointment or chronic-care settings, these models typically consider short-term sequences of visits or allow flexibility in allocating appointments across periods without binding patients to a specific time slot over time.

In contrast, we study a dynamic scheduling problem in which each accepted client occupies the same day–time slot over multiple future periods subject to heterogeneous availability constraints. This hard recurrence structure induces slot-specific, intertemporal capacity commitments that fundamentally alter the opportunity cost of acceptance and assignment decisions. The resulting MSAS problem defines a structurally distinct setting that is not explicitly modeled in prior work.

3. The Multisession Appointment Scheduling Problem

In this section, we introduce the MSAS problem. The problem schedules recurring appointments over an extended period, and clients are assigned to a service provider whom they meet on a regular basis for a certain predefined term, for example, once a week over one year. An important component of the problem is the heterogeneity of the clients. We assume heterogeneity among clients in terms of their availability, meaning each client is available at different times and can only attend sessions during those specific times as well as in their required program length, meaning each client requires a different number of treatment sessions. Figure 1 describes the general scheduling process. A service coordinator is responsible for facilitating, coordinating, and monitoring the services. Once the coordinators receive a new client's request, they refer it to a service provider. The referral is based on professional factors, such as the provider's expertise, and operational parameters, such as geographical location and perceived availability. Then, the clients specify their availability to the provider. The provider reviews the current schedule and either offers the client a slot or rejects the request. The provider may reject the request if no suitable slots are available or to reserve slots for future (yet still unknown) clients. If a request is denied, the clients are referred back to the service coordinator, who redirects them to a different service provider, potentially causing a delay in the start of treatment. Thus, rejection in our model represents a local scheduling decision rather than a system-wide denial of service, still capturing the operational cost of delayed access.

This paper presents our analysis of the single provider problem, which means focusing on stages 3–5 of the process, (marked in Figure 1 in a dashed rectangle). In Section 7, we discuss some extensions to the multiprovider setting and future work that includes steps 1–3.

3.1. Problem Settings, Assumptions, and Objectives

We now focus on the dynamics between a single provider and the clients assigned to that provider, and this defines the scope of the problem we model. A central characteristic of the MSAS problem studied in this paper is that accepted clients are assigned to a fixed day–time slot that recurs in every period for the duration of the treatment program. This structure reflects long-term therapeutic settings in which care is delivered on a regular schedule and continuity over time is operationally important. As a result, accepting a client creates a persistent capacity commitment that occupies the same slot over multiple future periods. Time is discretized into periods. Within this setting, the sequence of events in each period is as follows:

i. A client request arrives, specifying a program duration and a set of acceptable appointment slots.

ii. The provider decides whether to accept the request and, if so, assigns the client to one of the slots in the client's acceptable set that are currently available.

iii. Treatment sessions are conducted for all active clients.

Figure 1. The Scheduling Process

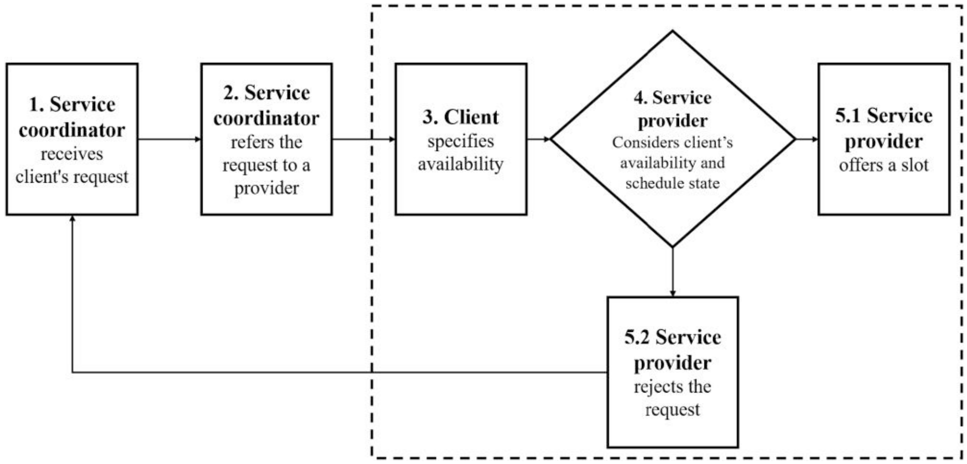

We adopt the following assumptions:

• Treatment frequency: Each accepted client attends one session per period. The period length is chosen to reflect the regular cadence of treatment (e.g., weekly sessions) so that program duration corresponds directly to the number of recurring commitments. More complex treatment frequencies and recurrence structures are discussed in Section 7.

• Client availability: Each request includes a subset of acceptable slots. Assignment decisions are restricted to this set.

• Capacity per slot: At most one client can be assigned to a given slot in each period. A slot represents an indivisible provider–time resource (e.g., a therapist-hour or a treatment room at a specific time); systems with multiple parallel resources can be represented by multiple slots.

• Arrivals: At most one client request arrives in each period. This reflects long-term therapeutic settings in which each provider maintains a relatively stable panel of ongoing clients and new requests arrive infrequently compared with the duration of treatment programs. This assumption enables a clean formulation of the decision process, preserving the core intertemporal trade-offs induced by recurring appointments. Allowing multiple arrivals per period requires sequential or batched decisions within a period but does not alter the fundamental structure of the problem; see Section 7 for further discussion.

• Cancellations and no-shows are negligible: In the service environments motivating this study, clients are assigned to recurring slots that align with their stated availability and become part of a stable weekly routine. Under such preference-based, fixed recurring assignments, attendance rates are typically high, and unexpected vacancies are relatively rare compared with open-access or flexible scheduling systems.

• Immediate start: Accepted clients begin treatment in the period immediately following their request. This reflects settings in which timely initiation of care is operationally important and avoids introducing additional state variables related to waiting times. Delayed starts add a waiting-state dimension to the model but do not alter the core capacity-allocation structure studied here.

The provider's decision problem is to determine, upon each request arrival, whether to accept or reject the client and, if accepted, which slot to assign, taking into account both immediate feasibility and the long-term impact of occupying capacity. Because of uncertainty in future client requests, accepting a client today may reduce the system's ability to accommodate future demand. In particular, assigning a client to a slot for an extended period creates an opportunity cost that depends on both the duration of the program and the future demand for that slot.

We consider two primary objectives. The first is to minimize rejected requests, reflecting the importance of timely access to care. The second is to maximize provider revenue, reflecting efficiency and financial sustainability. Whereas these objectives may appear aligned, the long-term capacity commitments induced by recurring appointments generate a natural trade-off between them. The following section formalizes this decision problem as an MDP.

3.2. Model Formulation

We model the MSAS problem as an MDP in which each period represents a new request arrival. The decision maker determines whether to accept or reject the request and, if accepted, to which slot it is assigned. The state captures the provider's schedule over multiple future periods, reflecting capacity commitments created by prior decisions. For simplicity, we assume one request arrival per period (see Remark 1).

Provider's slots: Let $\mathcal{J}$ denote the set of appointment slots within a single period (e.g., every hour from 9:00 to 17:00, five days a week). At most one client can be assigned to any slot in a given period.

Schedule state: The provider's schedule is represented by

$$Y = [y_j]_{j\in\mathcal{J}},$$

where $y_j \in \{0, 1, \ldots, N\}$ is the availability delay of slot j. Here, N denotes the maximum treatment horizon, that is, the longest possible duration (in decision epochs) of any recurring care program. If $y_j = 0$, slot j is currently available; if $y_j > 0$, the slot is occupied and next becomes available after y_j decision epochs. The set of all feasible schedules is

$$\mathcal{Y} = \{0, 1, \ldots, N\}^{|\mathcal{J}|}.$$

Client request (demand): A request is represented by

$$D = (L, J),$$

where $L \in \mathcal{L}$ is the program duration (number of recurring sessions) and $J \subseteq \mathcal{J}$ is the client's preferred set of time slots. If no request arrives at the current period, we set $D = (0, \varnothing)$.

Feasible slots: Given Y and $D = (L, J)$, define

$$\tilde{\mathcal{J}}(Y, D) := \{ j \mid j \in J, y_j \in \{0, 1\}\},$$

the set of slots that are available for immediate (or next period) assignment.

Actions: An action is denoted $a \in \mathcal{A}(Y, D)$, where a is the selected slot. The feasible action set is

$$\mathcal{A}(Y, D) = \begin{cases} \{0\}, & \text{if } \tilde{\mathcal{J}}(Y, D) = \varnothing, \\ \{0\} \cup \{ j \mid j \in \tilde{\mathcal{J}}(Y, D)\}, & \text{otherwise.} \end{cases}$$

Setting $a = 0$ denotes a rejection; otherwise, the request is accepted for its full duration L and assigned to a.

Arrival distribution. Let $p_{L,J} := \mathbb{P}(D = (L, J))$ denote a joint probability mass function (pmf) on $(\mathcal{L} \times 2^{\mathcal{J}}) \cup \{(0, \varnothing)\}$ with $p_{0,\varnothing} = p_\varnothing$ and $\sum_{L,J} p_{L,J} = 1$. For $L > 0$, we assume that arrivals of positive length have independent marginals so that

$$p_{L,J} = (1 - p_\varnothing)\, \alpha_L^+ \beta_J^+, \qquad \sum_{L>0} \alpha_L^+ = 1, \quad \sum_J \beta_J^+ = 1.$$

Here, $\{\alpha_L^+\}_{L \in \mathcal{L}}$ and $\{\beta_J^+\}_{J \subseteq \mathcal{J}}$ are pmfs on positive program lengths and preference sets, respectively.

Immediate reward: A positive reward is obtained only when a request is accepted. Specifically,

$$R(a) = \begin{cases} R(L), & \text{if } a \neq 0, \\ 0, & \text{otherwise.} \end{cases}$$

Note that $R(a)$ depends only on the program length L whenever the request is accepted and not on the assigned slot. Accepting a client for a program of length L yields a reward that is nondecreasing in L. In particular, $R(\cdot)$ may be constant (i.e., independent of L) or increasing. We assume $R(\cdot)$ is convex, which captures that longer programs can generate greater and increasingly efficient value because of stable revenue and reduced per-session overhead.

State transitions: Each accepted request blocks the selected slot for the next L decision epochs, whereas all other occupied slots progress one step toward availability. All vector operations below are understood component-wise. Let $\mathbf{1}$ denote the $|\mathcal{J}|$-dimensional vector of ones. For $Y = (y_j)_{j \in \mathcal{J}}$, we define

$$Y - \mathbf{1} := (y_j - 1)_{j \in \mathcal{J}}, \qquad \max(Y, 0) := (\max\{y_j, 0\})_{j \in \mathcal{J}}.$$

We define the deterministic transition function, $f : \mathcal{Y} \times \mathcal{A} \times \mathcal{L} \to \mathcal{Y}$, as follows:

$$f(Y, a, L) = \begin{cases} \max(Y - \mathbf{1}, 0), & \text{if } a = 0, \\ \max(Y - \mathbf{1}, 0) + L\boldsymbol{e}_{\mathbf{a}}, & \text{if } a \in \mathcal{J}, \end{cases} \tag{1}$$

where $\boldsymbol{e}_{\mathbf{a}}$ is a unit vector with a single one at entry a and zero elsewhere. The first case corresponds to rejecting the request, after which the system evolves only by the natural progression of time. The second case corresponds to accepting a client whose program of length L occupies slot a for L subsequent decision epochs.

Remark 1. For clarity of exposition, we assume that at most one client request arrives in each decision epoch. This simplifying assumption allows a clean analytical formulation of the MDP and facilitates the structural analysis that follows. In empirical implementations, the model can be extended to allow multiple requests to arrive within the same calendar period (e.g., within a week). Such an extension preserves the core structure of the model, in which ongoing programs advance at a fixed periodic cadence, but this requires handling several sequential or batch scheduling decisions within each period.

Value function and Bellman equation: Let $V(Y|D)$ denote the maximum expected value given the current state Y and the observed request $D = (L, J)$, and let $V(Y)$ denote the unconditional value function when only the distribution of future requests is known. We consider a discounted infinite-horizon setting with discount factor $\gamma \in (0, 1)$.

The conditional Bellman equation is

$$V(Y|D) = \max_{a \in \mathcal{A}(Y,D)} \{R(a) + \gamma V(f(Y, a, L))\}, \tag{2}$$

where $R(a) = R(L)$ if the request is accepted ($a \neq 0$)) and zero otherwise. Because each feasible action a either rejects the request or accepts and assigns it to one of the feasible slots $j \in \tilde{\mathcal{J}}(Y, D)$, the Bellman equation can be written explicitly as

$$V(Y|D) = \begin{cases} \gamma V(f(Y, 0, L)), & \text{if } \tilde{\mathcal{J}}(Y, D) = \varnothing, \\ \max\{\gamma V(f(Y, 0, L)), \\ \quad \max_{j \in \tilde{\mathcal{J}}(Y,D)} \{R(L) + \gamma V(f(Y, j, L))\}\}, & \text{otherwise.} \end{cases} \tag{3}$$

Infinite-horizon objective: The unconditional value function integrates over all possible realizations of requests:

$$V(Y) = \sum_{L \in \mathcal{L}_0} \sum_{J \subseteq \mathcal{J}} p_{L,J}\, V(Y|D = (L, J)), \quad \mathcal{L}_0 := \{0\} \cup \mathcal{L}. \tag{4}$$

The optimal policy π^* maps each observed state and request to a feasible action that maximizes the expected value:

$$\pi^*(Y, D) \in \arg\max_{a \in \mathcal{A}(Y,D)} \{R(a) + \gamma V(f(Y, a, L))\}. \tag{5}$$

The policy, therefore, decides, for each arriving request D, whether to reject it ($a = 0$) or accept and assign it to one of the feasible slots $j \in \tilde{\mathcal{J}}(Y, D)$. The resulting mapping π^* yields the maximum discounted expected reward over an infinite horizon.

4. Structural Components of the Scheduling Policy

In this section, we analyze the structural elements that shape scheduling decisions in the MSAS model. Some of these elements serve as principled structural guidelines that simplify the decision space and enable a tractable characterization of the policy (Section 4.1), whereas others follow directly from the Bellman optimality equations given this structural simplification and yield formally provable acceptance behavior (Section 4.2). This combined perspective allows us to isolate the key drivers of scheduling decisions and understand how different operational objectives influence the resulting policy structure.

To examine these effects, we consider two reward regimes: (i) Client dominance: this regim reflects a system that prioritizes minimizing client rejections. We model this by setting the reward function to be constant, $R(L) = r$, for all accepted requests. Although rejections do not incur an explicit penalty, the loss of reward relative to acceptance effectively imposes a cost of r for each rejected request. (ii) In provider dominance, this regime reflects a system that prioritizes maximizing provider revenue. Here, $R(L)$ is nondecreasing in L, either linear or strictly convex, representing the total revenue from all sessions of the treatment program. A linear reward captures constant per-session value, whereas a strictly convex reward reflects increasing marginal value for longer engagements.

Our analysis proceeds in two stages. First, in Section 4.1, we examine the slot-selection component. When a request is accepted and multiple slots in the client's preference set are feasible, the provider must determine which of these slots to occupy throughout the program duration. Ideally, one could select the slot that maximizes the Bellman continuation value, that is,

$$\hat{j} \in \arg \max_{j \in \tilde{\mathcal{J}}(Y,D)} \{R(L) + \gamma V(f(Y, j, L))\}. \tag{6}$$

However, because recurring assignments induce complex, high-dimensional state transitions, directly characterizing the maximizer in (6) is analytically difficult. By recurring assignment, we mean that once a client is accepted into slot j, the same slot is reserved for L consecutive decision epochs. Because the system state is the vector $Y \in \{0, \ldots, N\}^{|\mathcal{J}|}$ describing the remaining occupancy of all slots, each acceptance modifies one coordinate of this vector and influences its evolution over multiple future periods. As a result, the transition dynamics exhibit intertemporal, slot-specific coupling across periods. Section 4.1, therefore, develops a simple and principled slot-ranking guideline, motivated by structural intuition and supported by numerical evidence, that selects a unique slot among all feasible options and thereby reduces the slot-selection stage to a single choice.

Second, in Section 4.2, conditional on this slot-selection structure, we analyze the acceptance decision: whether to schedule the client into $\hat{j}$ or reject the request. Unlike slot selection, acceptance behavior follows directly from the Bellman optimality equation and, therefore, differs across the two reward regimes. In both the client-dominance and provider-dominance cases, we show that optimal acceptance behavior follows a threshold structure with respect to program duration. Combining a structurally guided slot-selection rule with analytically derived acceptance thresholds yields a transparent and tractable representation of the scheduling policy in the MSAS model.

4.1. Slot Selection: A Structural Ranking Guideline

When a request arrives and multiple slots in the client's preference set are feasible, the first step in the scheduling decision is to determine which of these feasible slots should be used if the provider chooses to accept the request. Because accepting a client blocks the chosen slot for several future periods, identifying the most appropriate slot is an essential first step before evaluating whether acceptance is beneficial.

To guide this choice, we introduce the following notation. For each slot $j \in \mathcal{J}$, let ψ_j denote its slot inclusion probability, defined as the probability that j appears in the preference set of a randomly arriving request: $\psi_j = \Pr(j \in J)$. These probabilities, derived from the empirical distribution $\{\beta_J^+\}$ introduced in Section 3.2, quantify how frequently each slot is likely to be requested in the future.

This observation motivates a simple pairwise ranking principle: given two feasible slots $j, k \in \tilde{\mathcal{J}}(Y, D)$ with $\psi_j < \psi_k$, assigning the client to the less popular slot j preserves future flexibility by keeping the more frequently requested slot k available for later arrivals. Similar intuition appears in the single-session setting of Wang and Gupta (2011), in which selecting the least popular feasible slot improves the likelihood of accommodating future demand.

Guided by this structural insight, we adopt the following ranking rule whenever multiple feasible slots exist: among all feasible slots in the preference set, the system assigns the client to the slot with the smallest inclusion probability.

Formally, let

$$\hat{j} \in \arg \min_{j \in \tilde{\mathcal{J}}(Y,D)} \psi_j, \tag{7}$$

denote the least popular feasible slot. Under this structural guideline, if the request is accepted, it is always assigned to $\hat{j}$. The Bellman Equation, (3), therefore, simplifies to

$$V(Y|D) = \begin{cases} \gamma V(f(Y, 0, L)), & \tilde{\mathcal{J}}(Y, D) = \emptyset, \\ \max\{\gamma V(f(Y, 0, L)), R(L) \\ \qquad + \gamma V(f(Y, \hat{j}, L))\}, & \text{otherwise}, \end{cases} \tag{8}$$

reducing the scheduling decision in each period to a binary choice: accept the request into $\hat{j}$ or reject it. Whereas a full proof of slot-ranking optimality in the general multisession setting is analytically challenging because of the cross-period coupling induced by recurring commitments, the structural guideline above provides a consistent and computationally effective basis for policy design. Section 4.2 builds on this slot-selection rule to characterize optimal acceptance behavior under each reward regime.

4.2. Acceptance Policy

We now shift our focus from the question of slot selection to that of request acceptance: when is it optimal to schedule a client at all? This decision is driven by the reward function $R(L)$, which reflects the system's operational objective. As discussed earlier, a constant reward corresponds to a client-dominance setting, in which all accepted requests are valued equally and the goal is to maximize throughput. An increasing and convex reward reflects a provider-dominance setting, in which longer treatment programs generate higher cumulative value.

These objectives induce different acceptance behaviors. In the client-dominance case, treating all clients as having equal value incentivizes accepting as many requests as possible. However, longer treatment programs consume capacity over multiple periods and may block future arrivals. Consequently, the optimal policy may reject long-duration clients to preserve flexibility and maintain high acceptance rates. In contrast, under provider-dominance the priority shifts toward maximizing cumulative reward. Here, long-duration clients, despite occupying capacity for extended periods, may be preferred because their higher total reward outweighs the opportunity cost of rejecting several shorter clients.

In both settings, the optimal acceptance policy admits a threshold structure in the client's program duration L; however, the direction of the threshold reverses. Under client dominance, shorter programs are more attractive; under provider dominance, the policy favors longer programs. This structural inversion highlights how acceptance priorities evolve in response to the underlying reward objective.

4.2.1. The Client-Dominance Case. To prove the optimality of a threshold-type policy, we show that, if it is optimal to reject a request for slot j and program duration L when the provider's schedule is Y, then it is also optimal to reject in schedule Y a request for slot j and program duration $L+1$. In this setting, in which all accepted requests yield the same reward regardless of duration, we denote this constant reward by r, that is, $R(L) = r\ \forall L$.

Lemma 1. *If $V(f(Y,j,L))$ is nonincreasing in L for any fixed $Y \in \mathcal{Y}$ and $j \in \mathcal{J}$, then there exists an optimal policy of threshold type in L, in which requests are accepted if and only if their duration does not exceed a threshold.*

The proof is provided in Online Appendix A.1.

Lemma 2. *The maximum expected value, $V(f(Y,j,L))$, is nonincreasing in L. That is,*

$$V(f(Y,j,L)) \geq V(f(Y,j,L+1)) \quad \forall L \in \mathcal{L},\ \forall Y \in \mathcal{Y},\ \forall j \in \mathcal{J}.$$

Lemma 2 is intuitively clear: the optimal value should not increase when fewer scheduling options are available. A formal proof is provided in Online Appendix A.2.

Lemmas 1 and 2 together establish that the optimal acceptance policy, conditional on the slot-selection rule of Section 4.1, has a threshold structure in the client-dominance setting with respect to the client's program duration L. Specifically, when the provider's schedule is Y, a request for slot j is accepted if and only if the client's program duration is at most $L^C_{Y,j}$; otherwise, the request is rejected. Here, $L^C_{Y,j}$ denotes the client-dominance threshold for accepting a request of slot j when the current schedule is Y. Note that the threshold may vary across both slots and schedule states; that is, there may be a different threshold $L^C_{Y,j}$ for each combination of $Y \in \mathcal{Y}$ and $j \in \mathcal{J}$. In practical terms, this means that, when the schedule becomes more constrained, the provider should favor shorter treatment programs to maintain future scheduling flexibility.

4.2.2. The Provider-Dominance Case. In contrast to the client-dominance case, in the provider-dominance case, to prove the optimality of a threshold-type policy conditional on the slot-selection guideline, we show that, if it is optimal to accept a request for slot j and program duration L in schedule state Y, then it is optimal to accept in state Y a request for slot j and program duration $L+1$ as well.

Lemma 3. *If $V(f(Y,j,L))$ is nonincreasing and convex in L for any fixed $Y \in \mathcal{Y}$ and $j \in \mathcal{J}$, then there exists an optimal policy of threshold type in L, in which requests are accepted if and only if their duration exceeds a threshold.*

The argument builds on the convexity of the value function in L, which implies that the marginal loss from increasing program duration is nondecreasing. This structure ensures that the acceptance condition is preserved from L to $L+1$. For the proof, see Online Appendix A.3.

Lemma 4. *The maximum expected value, $V(f(Y,j,L))$, is nonincreasing and convex in L for any fixed $Y \in \mathcal{Y}$ and $j \in \mathcal{J}$.*

The proof for Lemma 4 can be found in Online Appendix A.4.

Lemmas 3 and 4 together establish that the optimal acceptance policy, given the slot-selection rule of Section 4.1, has a threshold structure in the provider-dominance setting with respect to the client's program duration L. Specifically, when the provider's schedule is Y, a request for slot j is accepted if and only if the client's program duration is at least $L^P_{Y,j}$, otherwise, the request is rejected. Here, $L^P_{Y,j}$ denotes the provider-dominance threshold for accepting a request for slot j when the current schedule is Y. Note that the threshold may vary across both slots and schedule states, that is, there may be a different threshold $L^P_{Y,j}$ for each combination of $Y \in \mathcal{Y}$ and $j \in \mathcal{J}$. Intuitively, this implies that, when longer programs generate higher marginal value, the provider should prioritize clients with extended treatment durations even if it means rejecting several shorter requests.

These results establish that, despite the differing objectives in the client-dominance and provider-dominance settings, the structure of the optimal policy in both cases conforms to a threshold form with respect to the client's program duration. However, the direction of the threshold and, therefore, the prioritization of requests reverses depending on which objective is being optimized. This contrast highlights the dependence of the optimal scheduling policy's structure on the specific objective being optimized. The threshold-based characterization derived here provides a concrete foundation for constructing heuristic policies tailored to each setting. Moreover, it offers a principled starting point for reasoning about more general scenarios in which access and revenue considerations are jointly weighted, leading to intermediate policies that balance competing priorities. Together, these structural insights form the analytical foundation for our heuristic development in Section 5, bridging the theoretical and practical aspects of the MSAS problem.

5. Traffic Light Heuristic Approach

In the previous section, we gain valuable structural insights into the scheduling decision, showing that, under our slot-selection guideline, the optimal acceptance behavior exhibits a threshold structure with respect to the client's program duration. Solving the MDP exactly, however, remains computationally intractable because of the size of the state space. The number of possible schedule states equals $(N+1)^{|\mathcal{J}|}$, which grows rapidly with the maximal program duration and the number of slots. In addition, the heterogeneity of client requests further increases the dimensionality of the problem. Whereas the structural results significantly simplify the characterization of the optimal policy, they still require computing state- and slot-dependent thresholds across a large number of states. This creates an additional computational challenge even after accounting for the analytical structure. Rather than attempting to compute these thresholds explicitly or approximate the value function over the full state space, we adopt a structure-driven approach and design a heuristic policy that builds directly on the key analytical insights derived above. We, therefore, introduce the traffic light (TL) heuristic, which leverages these insights to provide a simple and practical decision rule.

The TL heuristic groups states and slots with similar characteristics, effectively reducing the state space and applying the same threshold to all states within each group, trading off detailed state information for a simpler and more practical decision rule. The heuristic framework consists of three main steps. First, we categorize all schedule states into three possible states: green, orange, and red, reflecting increasing levels of congestion. The green state corresponds to low occupancy and is associated with minimal restrictions on request acceptance. The orange state reflects moderate congestion and calls for more selective assignment decisions. The red state represents high occupancy and necessitates stringent acceptance constraints to preserve future scheduling flexibility. The classification relies on an occupancy-rate measure, defined as the fraction of slot–period pairs occupied over a finite look-ahead horizon.

Second, we classify the slots into two groups: popular and less popular. In practice, this could be done by analyzing historical data and consulting with the service providers who can identify which slots are more in demand and which are less popular. Here too, a more careful assignment is advised for the popular slots because they appear in future requests with a higher probability. Consequently, the policy requires a total of $3 \times 2 = 6$ thresholds with one designated for each combination of state group and slot group.

Lastly, we determine these threshold values based on the categorized schedule states and slot groups. We denote the thresholds using the following notation: $L^{case}_{state,slot}$, where $case = \{C(client\ dominance), P(provider\ dominance)\}$, $state = \{G(green), O(orange), R(red)\}$ and $slot = \{p(popular), np(non\ popular)\}$. For example, $L^C_{O,p}$ represents the client-dominance case threshold for popular slots in the orange state, whereas $L^P_{R,np}$ represents the threshold for nonpopular slots in the red state under the provider dominance case. Building on the explanation provided earlier regarding the restrictive nature of each state, in the client-dominance case, we anticipate the following relationship: $L^C_{R,i} \le L^C_{O,i} \le L^C_{G,i}$, where i represents either p or np indices. This expectation aligns with our prioritization of clients with shorter program durations as the schedule becomes busier. Conversely, in the provider-dominance case, we expect $L^P_{G,i} \le L^P_{O,i} \le L^P_{R,i}$ as we prioritize clients with longer program duration. Similarly for the client-dominance case, we expect $L^C_{k,p} \le L^C_{k,np}$, whereas for the provider-dominance case we expect $L^P_{k,np} \le L^P_{k,p}$, where k represents G, O, R indices. For clarity of exposition, we summarize the implementation of the TL heuristic in Algorithm 1.

Algorithm 1 (Traffic Light Policy)

1. Observe the current schedule state Y and the incoming request (L, J).
2. Classify the schedule state as $state \in \{G, O, R\}$ according to the occupancy-rate measure.
3. Determine the set of feasible slots $\tilde{J}(Y, D)$. If $\tilde{J}(Y, D) = ø$, reject the request.
4. Among feasible slots, select the least-popular slot $\hat{j}$. Classify it as $slot \in \{p, np\}$ (popular or nonpopular).
5. Identify the corresponding threshold $L^{case}_{state,slot}$, where $case \in \{C, P\}$ denotes client- or provider-dominance.
6. Client-dominance ($case = C$): Accept if $L \le L^{C}_{state,slot}$; otherwise reject.
7. Provider-dominance ($case = P$): Accept if $L \ge L^{P}_{state,slot}$; otherwise reject.

In the following section, we present a comprehensive numerical study and describe the specific implementation used to calibrate state classification and threshold values. It is important to emphasize that the methodology adopted here for classifying schedule states and slot types represents only one possible instantiation of the TL framework. The heuristic itself is defined by the structural logic outlined above and can be implemented using alternative congestion measures, state partitions, or slot groupings that adhere to the same principles. Moreover, the framework naturally extends to a larger number of aggregated states or slot categories, and this requires determining a greater number of threshold parameters. Whereas increasing such differentiation may improve performance by capturing finer system dynamics, it also entails additional computational effort, highlighting a trade-off between the level of aggregation and tractability. As demonstrated in the next section, even a modest aggregation yields substantial improvements over first come, first served policies, maintaining practical implementability.

6. Computational Study

We now present a numerical study conducted to evaluate the performance of the TL heuristic and examine the sensitivity of its performance to various problem conditions, such as slot demand patterns and variability in client program durations. In addition to assessing the heuristic's performance, we also generate managerial insights into the structure and behavior of optimal scheduling policies under different operational scenarios.

6.1. Experimental Setting

In this experiment, we systematically generate problem instances comprising sequences of requests. The experimental design reflects core operational features of recurring treatment settings, including weekly appointment structures, heterogeneous program durations, and asymmetries in slot demand. Each instance simulates requests arriving throughout 1,000 time periods with at most one request arriving per period. Within these requests, each includes a preference list and the duration of the client program. We consider 20 available slots per period, that is, $|\mathcal{J}| = 20$, representing periods of one week with five days per week and four slots per day (early morning, late morning, afternoon and late afternoon). We assume a capacity of one client per slot.

The preference list for each request was randomly generated as follows: first, we generated a number, denoted as m, ranging from zero to five. Here, zero indicates no arrival of a request during that period, whereas values from one to five represent the length of the preference list if a request has arrived. The probability of having no request ($P(m = 0)$) was set to 5%, meaning there is a 5% chance of no request arriving during a given period. For the other values of m (one to five), we set the following probabilities: $P(m = 1, 5) = 0.1$, $P(m = 2, 4) = 0.23$, and $P(m = 3) = 0.29$. These probabilities were arbitrarily determined based on the assumption that clients are most likely to provide three optional slots with a smaller probability that they provide two or four options, and with the lowest probability, they provide only one option or are flexible with five options. Then, according to the value of m, we generated the slots comprising the preference list. For a given m, slots were sampled sequentially until m distinct slots were obtained. Here, we considered two types of preference behaviors: (i) equal slot preference and (ii) weighted slot preference. In the equal-preference setting, each slot was sampled with equal probability $1/|\mathcal{J}|$, so all slots have the same probability of inclusion in the preference list. In the weighted-preference setting, we divided the set of 20 slots into three groups: 5 popular slots, 5 less popular slots, and 10 least popular slots. Higher sampling weights were assigned to the popular group than to the less popular and least popular groups. Consequently, slots do not share the same inclusion probability, and popular slots are more likely to appear in the preference list. This grouping reflects common practice in which certain time blocks (e.g., early morning or late afternoon) are more convenient for clients and, therefore, more frequently requested.

The duration of the clients' program was uniformly generated from a discrete list of values with increments of five, spanning from 10 to 30, 50, or 100 periods. This represents three distinct scenarios: (i) All clients have short-term programs, spanning from 10 to 30 periods. Consequently, there is a low variance between the durations of different clients. (ii) Client durations vary from short to medium term, ranging from 10 to 50 periods. This results in a medium variance between the durations of different clients. (iii) Client durations vary from short to relatively long term, spanning from 10 to

100 periods. As a result, there is a high variance between the durations of different clients. The third scenario reflects the conditions commonly observed in EI programs, in which clients may enter the program at any point from birth until close to age three, depending on the timing of diagnosis and referral. As a result, there is naturally high variation in how long each child remains in the program. The lower variance scenarios serve as part of a sensitivity analysis and may reflect environments with a narrower range of possible treatment durations.

For each combination of preference behavior (equal slot preference/weighted slot preference denoted respectively as E or W) and client program duration scenario (low/medium/high variance denoted respectively as L, M, or H) totaling six combinations (EL, EM, EH, WL, WM, WH), we randomly generated 20 instances according to the combination characteristics. As mentioned earlier, each instance simulated requests arriving throughout 1,000 time periods with at most one request arriving per period. Across all runs, the coefficient of variation of the objective values did not exceed 0.037, and the 95% confidence intervals around the means were consistently narrow, indicating stable and reliable simulation outcomes.

Before applying the TL heuristic, we set its decision thresholds, which determine whether to accept a request based on the current occupancy level. The setting process involved classifying states into green, orange, or red occupancy levels using an occupancy rate measure and identifying the optimal threshold values through exhaustive search. Full details of the occupancy rate calculation, the threshold definitions, and the parameter selection process as well as insights into the acceptance thresholds obtained under different scenarios, are provided in Online Appendix B.

6.2. Results

After fine-tuning the heuristic parameters, we implemented our solution method and calculated the objective value when following the TL policy. Each instance was solved three times: once for the client-dominance case with $R(L) = r = 1\ \forall L$ and twice for the provider-dominance case, first using a linear cost function $R(L) = L$ and then using a strictly convex cost function $R(L) = L^2/100$. These two forms allow us to examine the effects of different convex cost structures on the provider's revenue. Note that, under a linear cost structure, the specific values of r and the coefficient of L do not influence the optimal solution or the resulting policy. Instead, they scale the objective function proportionally without affecting the decision-making process. Therefore, for simplicity and interpretability, we set both parameters to one. In contrast, under a strictly convex cost structure such as $R(L) = L^2/100$, the specific form and scaling of the function can influence the optimal policy. In this case, the coefficient 1/100 was chosen to keep the magnitude of the cost comparable to the linear case, allowing for meaningful comparisons, preserving the qualitative effects of convexity.

To obtain an upper bound on the achievable performance, we solved a deterministic version of the problem for each simulated sample path, assuming full knowledge of all future arrivals, program durations, and preference sets. The resulting optimization problem was formulated as an integer linear program (Online Appendix C). Because this model assumes perfect information and optimizes globally over the entire horizon, its optimal value provides a valid upper bound on the performance of any online policy under uncertainty. In the computational study, all objective values are computed over the finite simulation horizon as the (undiscounted) total accumulated reward. The integer programs were solved to optimality using Gurobi 10.0.0.

The deterministic version of the model, whereas not intended to approximate the optimal solution of the stochastic problem, nevertheless provides useful intuition about how program duration influences acceptance decisions. To examine this effect, we evaluated the acceptance rate for each program duration across all deterministic scenarios. The acceptance rate represents the percentage of requests of a given duration that were accepted and scheduled out of all such requests that appeared in the generated instances. Figures 2 and 3 present the average acceptance rate by program duration across all instances for each of the six scenarios (each scenario combining preference behavior, equal or weighted, and program duration variance, low, medium, or high). Figure 2 corresponds to the client-dominance setting, whereas Figure 3 corresponds to the provider-dominance setting with $R(L) = L$.

The results reveal a consistent and interpretable trend. Under client dominance, acceptance rates decrease as program duration increases. Under provider dominance, acceptance rates increase with duration (a pattern that also appears when using the strictly convex reward function $R(L) = L^2/100$). These empirical patterns from the deterministic setting are consistent with (though not reflective of the optimality of) the analytical results obtained for the stochastic model. Specifically, whereas the deterministic solution is not optimal for the stochastic case, the observed monotonic trends mirror the threshold-type acceptance structure we characterize analytically: shorter programs are favored under client dominance, whereas longer programs are preferred under provider dominance.

Another observation concerns the difference in acceptance rates between the equal slots preference case and the weighted slot preference case. In both client-dominance and provider-dominance settings, we observe that acceptance rates are higher when the

Figure 2. Acceptance Rate Under Deterministic Demand: Client-Dominance Case

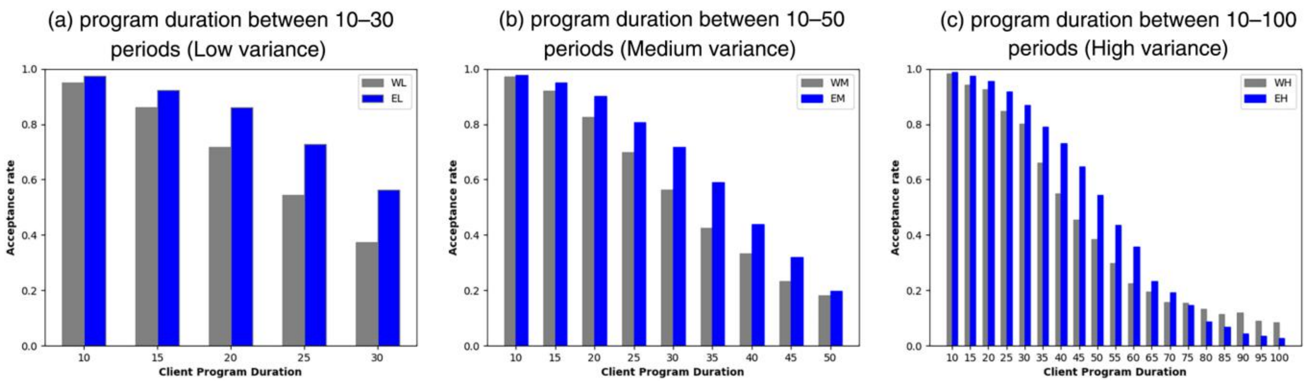

slots are evenly popular compared with when some slots are more popular than others. This is not surprising because evenly popular slots distribute demand more uniformly, reducing scheduling conflicts and increasing the likelihood of acceptance for each slot. These observations are further supported by additional small-scale experiments, in which the optimal policy can be computed exactly using dynamic programming. The results exhibit a threshold-type behavior consistent with our analytical findings (see Online Appendix D).

As a benchmark policy and to establish a lower bound on the value of the objective function, we implemented an FCFS policy on our generated instances. Under this policy, any request is accepted as long as there is at least one available slot in the client's preference list at the time of the request. The FCFS policy is commonly used in practice because of its simplicity and ease of implementation, making it an ideal baseline for comparison. When multiple slots are available, we examine two options for selecting the slot. The first option is to randomly select an available slot, denoted as $FCFS^R$. The second option, motivated by the slot-ranking guideline introduced in Section 4.1, is to choose the least-popular feasible slot. We denote the latter as $FCFS^{LP}$. Note that, when all slots are equally popular, both policies operate identically and yield the same results.

Table 1 presents the results obtained in the client-dominance setting. The table shows the relative gaps from the benchmark policy ($FCFS^R$) and from the optimal solution to the deterministic problem (OPT^D) for both the $FCFS^{LP}$ and traffic light (TL) policies. Additionally, the last column indicates the extent to which the relevant solution closes the gap between these lower and upper bounds, that is, taking the difference between the $FCFS^R$ solution and the solution provided by the examined policy and then dividing it by the difference between the OPT^D and $FCFS^R$ solutions. Details of the heuristic parameters and threshold values for each scenario are provided in Online Appendix B.

We observe that, in the weighted slots preference case, in which some slots are more popular than others, the TL policy improves the solution compared with the $FCFS^R$ policy by an average of 6.2% when the program duration variance is low (WL scenario). This improvement increases as the program duration variance rises, reaching an average improvement of 10.3% for the high variance (WH scenario). Because the objective function is linear in the number of rejections, the improvement value reflects the relative increase in the number of accepted requests. For instance, in the WM

Figure 3. Acceptance Rate Under Deterministic Demand: Provider-Dominance Case

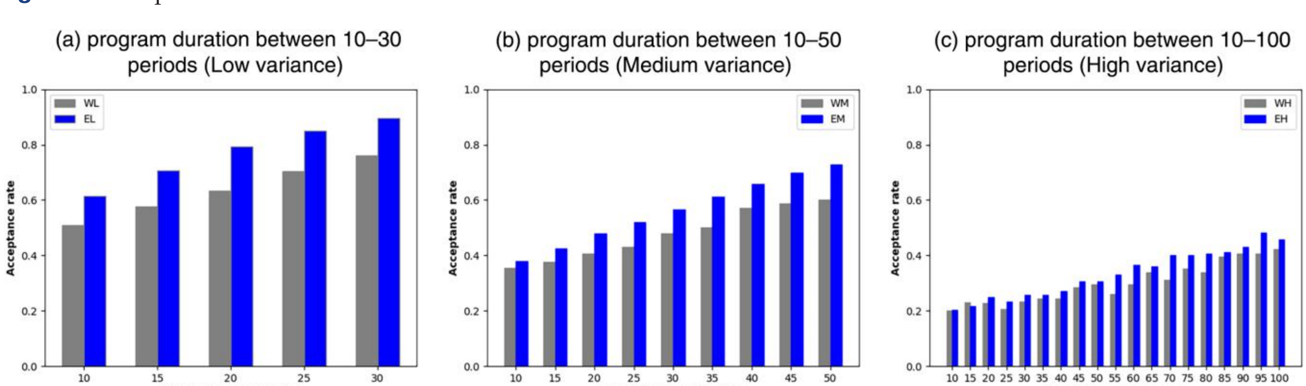

Table 1. Results Obtained in the Client-Dominance Setting

			Gap from $FCFS^R$, %		Gap from OPT^D, %		Percentage gap closed	
Scenario	Preference	Max duration	TL	$FCFS^{LP}$	TL	$FCFS^{LP}$	TL	$FCFS^{LP}$
WL	Weighted	30	6.2	5.7	22.8	23.5	26.3	24.1
WM	Weighted	50	8.1	2.3	14.3	21.5	41.5	11.9
WH	Weighted	100	10.3	0.9	9.9	21.3	56.1	5.1
EL	Equal	30	0.0	NA	51.6	51.6	0.0	NA
EM	Equal	50	2.4	NA	25.3	28.5	11.0	NA
EH	Equal	100	8.1	NA	14.2	24.4	41.6	NA

case, the TL policy results in an average increase of 8.1% in the total number of scheduled requests compared with the $FCFS^R$ policy. The improvement trend is also reflected in the gap from the OPT^D and the percentage of the gap closed. We note that a gap of 9.9% from the OPT^D indicates that the *TL* policy achieves approximately 90% of the maximum possible value. Given that the maximum possible value is unachievable, we consider these results to be satisfactory.

When all slots are equally preferred, we also observe improvements though they are more moderate. In the EL scenario, the policy converged with the FCFS policy, resulting in no relative improvement. For the EM scenario, we achieved an average improvement of 2.4%. The biggest improvement, 8.1%, was in the EH scenario, in which the program duration variability was high. This corresponds to achieving approximately 86% of the maximum possible value and reducing the gap between the upper and lower bounds by 41.6%.

Examining the performance of the $FCFS^{LP}$ policy, we observe an average improvement of 5.7% compared with the $FCFS^R$ in the WL scenario. However, as the variance in program duration increases, the performance difference between the $FCFS^{LP}$ and $FCFS^R$ policies decreases with only a 0.9% average improvement in the high variance (WH) scenario. Thus, we conclude that, as the variance in program duration increases, the advantage of the TL policy becomes more pronounced. This can be explained by the fact that, with higher variance, the threshold values in the TL policy differ more significantly from the maximum program duration. Because the FCFS policy can be seen as a special case of the TL policy in which the threshold values are set to the maximum program duration (in the client-dominance case), we observe larger differences between these policies and greater opportunities for improvement. The TL policy's ability to adjust its thresholds based on varying program durations allows it to manage the schedule more effectively, leading to substantial performance gains compared with the $FCFS^R$ and $FCFS^{LP}$ policies.

In the provider-dominance case, we first examine a linear revenue function $R(L) = L$, where the revenue gained from scheduling a session is proportional to the client's program duration. As shown in Table 2, an interesting outcome is that, for both the WL and WM scenarios, all thresholds identified in the TL policy match the lower bound of the program duration, effectively aligning with the $FCFS^{LP}$ policy. In the WH scenario, a slight improvement is observed with $L^P_{R,p} = 25$, whereas all other thresholds remain at the lower bound. When slot preferences are equal, selecting the least popular slot becomes irrelevant, and consequently, the $FCFS^R$ policy yields the best results in these cases. Overall, the gap from the deterministic optimal solution remains modest, not exceeding 12.1% and in most cases is below 10%. This suggests that, under a linear cost structure, the FCFS policy performs well.

Next, we examine a strictly convex revenue function $R(L) = L^2/100$, in which clients with longer program durations yield, on average, higher revenue. As shown in Table 3, the TL policy consistently outperforms the FCFS policies across all scenarios. In the weighted slots preference cases, the improvement of the TL policy over the $FCFS^R$ policy increases with the variance in program durations: from 4.5% in the WL scenario to

Table 2. Results Obtained in the Provider-Dominance Setting, $R(L) = L$

			Gap from $FCFS^R$, %		Gap from OPT^D, %		Percentage gap closed	
Scenario	Preference	Max duration	TL	$FCFS^{LP}$	TL	$FCFS^{LP}$	TL	$FCFS^{LP}$
WL	Weighted	30	3.8	3.8	7.6	7.6	30.8	30.8
WM	Weighted	50	2.7	2.7	6.9	6.9	26.0	26.0
WH	Weighted	100	1.9	1.7	6.8	6.9	20.0	18.3
EL	Equal	30	0.0	NA	12.1	12.1	0.0	NA
EM	Equal	50	0.0	NA	9.4	9.4	0.0	NA
EH	Equal	100	0.0	NA	7.6	7.6	0.0	NA

Table 3. Results Obtained in the Provider-Dominance Setting, $R(L) = L^2/100$

			Gap from $FCFS^R$, %		Gap from OPT^D, %		Percentage gap closed	
Scenario	Preference	Max duration	TL	$FCFS^{LP}$	TL	$FCFS^{LP}$	TL	$FCFS^{LP}$
WL	Weighted	30	4.5	3.8	13.6	14.1	21.3	18.2
WM	Weighted	50	6.8	2.8	14.0	17.3	28.2	11.5
WH	Weighted	100	10.5	1.6	14.1	21.0	36.7	5.6
EL	Equal	30	0.6	NA	16.1	16.5	2.8	NA
EM	Equal	50	3.6	NA	15.9	18.8	15.5	NA
EH	Equal	100	11.6	NA	14.0	22.9	38.9	NA

10.5% in the WH scenario. Conversely, the $FCFS^{LP}$ improvement over $FCFS^R$ decreases as the variance increases, dropping from 3.8% in WL to 1.6% in WH. This trend indicates that the advantage of the TL policy becomes more pronounced with higher variance in program durations. The effectiveness of the TL policy is also evident when looking at the gap from the deterministic optimal solution and the percentage of the gap closed. For instance, in the WH scenario, the TL policy closes 36.7% of the gap compared with only 5.6% by the $FCFS^{LP}$ policy. Notably, even in scenarios in which all slots are equally preferred, the TL policy demonstrates significant advantages when program duration variance is high. In the EH scenario, the TL policy achieves an 11.6% improvement over $FCFS^R$, closing 38.9% of the gap from the optimal solution. These results highlight the effectiveness of the TL policy in environments with high variability in program durations, especially under a strictly convex revenue structure.

These results can be attributed to the structure of the revenue function, in which clients with longer program durations contribute more to the overall revenue. As the variance in program durations increases, the adaptive threshold mechanism of the TL policy becomes more effective in prioritizing these high-revenue clients. Unlike the FCFS policies, which schedule without considering program duration, the TL policy adjusts its thresholds to favor longer programs. Similar to the client-dominance case, this adaptability allows the TL policy to better align scheduling decisions with the objective function in mind, leading to more significant performance improvements over the FCFS policies, particularly in environments with high variability in program durations.

7. Discussion

Appointment scheduling is a central operational challenge in healthcare, particularly when clients require a series of recurring treatment sessions. In such settings, scheduling decisions must balance client convenience, which supports adherence and continuity of care, with the provider's need to manage capacity efficiently over extended periods. A distinctive feature of our problem is that each client must be assigned to a fixed day–time slot that recurs throughout the entire treatment program. This realistic but restrictive requirement creates long-term capacity commitments and substantially increases the complexity of acceptance decisions. It also differentiates our setting from most multiappointment scheduling models, in which appointments can be placed flexibly across periods.

Because scheduling occurs sequentially under uncertainty with only probabilistic information about future arrivals, we model the multisession appointment scheduling problem as a Markov decision process. This framework captures the dynamic nature of the problem and the persistent consequences of each acceptance decision. Our aim was to design policies that accommodate heterogeneous client characteristics, differences in time preferences, availability, and program durations, respecting providers' operational constraints and the structure imposed by recurring appointments.

We examined two objectives: minimizing rejected requests (client dominance) and maximizing provider revenue (provider dominance). Analyzing each objective separately and conditional on the slot-selection structure adopted in Section 4.1, we establish that the optimal acceptance policy exhibits a threshold with respect to program duration. However, the direction of this threshold depends on the objective: under client dominance, shorter programs should be prioritized to preserve future flexibility, whereas under provider dominance, longer programs yield higher cumulative value. These structural insights illustrate how prioritization shifts when the operational objective changes. The two reward structures represent stylized extremes that isolate different prioritization logics. In practice, healthcare systems often balance access considerations with financial sustainability, leading to intermediate objective functions that mitigate exclusion at either end of the duration spectrum. The threshold structures identified here clarify how such trade-offs arise and provide a foundation for designing more balanced policies that account for both efficiency and access considerations.

To address the large state space of the MDP, we introduced the traffic light heuristic, an easily implementable policy informed by the structural insights developed in our analysis. The heuristic groups states with similar characteristics and applies a common threshold within each group, thereby reducing the dimensionality of the decision space. Our computational experiments show that $FCFS^{LP}$, which accepts a request if a preferred slot is available but assigns the client to the least popular slot, outperforms the standard $FCFS^{R}$ benchmark in both dominance regimes, highlighting the benefit of incorporating preference information. Moreover, the value of threshold-based policies, and of the traffic light heuristic in particular, increases with program duration variance. In such cases, the heuristic yields substantial performance gains relative to both *FCFS* benchmarks and performs competitively with the deterministic upper bound.

Taken together, these computational findings also suggest several operational insights. First, explicitly accounting for recurring capacity commitments is most valuable in environments with high variability in program durations or uneven slot demand. In such settings, treating sessions as independent appointments may lead to systematic misallocation of highly requested time blocks or underestimation of long-term opportunity costs. Second, simple structural rules, such as prioritizing less popular slots and applying duration-based thresholds, provide an anticipatory mechanism that mitigates these effects without requiring full dynamic optimization. Importantly, these benefits do not depend on a particular calibration of the heuristic. Accordingly, the traffic light framework can be adapted by introducing more refined scheduling-state or slot-popularity categories though this comes at the cost of additional computation.

Our analysis focuses on a single-provider setting and on client heterogeneity expressed through availability and program duration. This scope allows us to isolate the dynamic scheduling implications of recurring appointments and long-term capacity commitments under uncertainty without conflating them with provider selection, referral, or coordination decisions. Even within this focused setting, the MSAS problem exhibits rich intertemporal structure, giving rise to threshold-based acceptance behavior and slot-dependent opportunity costs.

7.1. Modeling Scope, Limitations, and Extensions

The model developed in this paper focuses on a stylized setting designed to isolate the core intertemporal trade-offs induced by recurring capacity commitments. In particular, we assume a single session per period, at most one request arrival per period, and immediate treatment initiation. These assumptions enable a tractable MDP formulation and allow us to derive clear structural insights regarding acceptance and assignment decisions. In practice, treatment plans may involve more complex patterns, such as multiple sessions per period, less frequent visits, or multiple request arrivals within a given period. Incorporating such features requires a richer representation of the scheduling process.

In the current formulation, each accepted client occupies the same slot in consecutive future periods, so the system state can be represented through the remaining duration of each recurring commitment. More general treatment frequencies fundamentally alter this recurrence structure. For example, multiple sessions per period require simultaneous assignment of several recurring slots, whereas less frequent visits generate nonconsecutive future slot occupancies. As a result, the model needs to explicitly track future occupancy patterns across periods, substantially enriching the state representation, action space, and transition dynamics. Similarly, multiple arrivals per period require sequential or batch decision making within each period.

A related extension concerns the timing of treatment initiation. In the current model, accepted clients begin treatment immediately. Allowing delayed starts introduces additional decision variables governing admission timing and requires tracking scheduled future starts beyond immediate availability. Taken together, these extensions enrich the modeling framework but preserve the fundamental structure of the problem, in which each accepted client induces a persistent capacity commitment over time. From a practical perspective, some of these features can often be incorporated at the heuristic level, for example, by linking multiple slot assignments or applying sequential decision rules within a period. At the same time, these richer scheduling environments also motivate the development of more scalable solution methodologies. In particular, approximate dynamic programming and reinforcement learning approaches may provide useful alternatives for high-dimensional settings in which exact structural analysis becomes less tractable. Exploring such methods and comparing them with the structure-based heuristic proposed here, represents an important avenue for future research.

Beyond extensions to the scheduling structure itself, the proposed framework also naturally gives rise to richer objective formulations and operational settings. One promising direction involves mixed objective functions combining constant and nonconstant reward components. Preliminary deterministic results suggest that threshold structures persist but shift depending on the relative weight of each reward term: dominance of the constant component favors shorter programs, whereas the increasing component promotes longer programs. Such formulations may also be interpreted as incorporating an implicit penalty for rejection, simultaneously providing a natural way to represent different levels of treatment urgency or intensity through the reward structure. Fully characterizing these interactions in the stochastic setting remains an open challenge.

Additional extensions include programs that involve time-dependent slot changes, models in which clients choose among provider-generated offers, and systems with multiple providers. In a multiprovider environment, the recurring-slot structure continues to induce intertemporal capacity commitments at the provider-slot level, but the decision space expands to include provider selection and cross-provider coordination. Whereas capacity pooling across providers may mitigate some local bottlenecks, the core trade-off identified in this paper, between program duration and future slot-specific demand, remains central. Thus, we expect the structural insights regarding slot ranking and duration-based thresholds to extend, operating within a higher dimensional allocation framework. In multiprovider contexts, both centralized and decentralized organizational structures merit investigation as they may lead to different coordination mechanisms, flexibility levels, and system-wide performance outcomes.

Acknowledgments
The authors thank the Editor-in-Chief, the Associate Editor, and the anonymous referee for their insightful and constructive comments and suggestions, which have improved the quality of this paper.

References
Ahmadi-Javid A, Jalali Z, Klassen KJ (2017) Outpatient appointment systems in healthcare: A review of optimization studies. *Eur. J. Oper. Res.* 258(1):3–34.
Bard JF, Shao Y, Qi X, Jarrah AI (2014) The traveling therapist scheduling problem. *IIE Trans.* 46(7):683–706.
Bennett AR, Erera AL (2011) Dynamic periodic fixed appointment scheduling for home health. *IIE Trans. Healthcare Systems Engrg.* 1(1):6–19.
Bonnier C (2008) Evaluation of early stimulation programs for enhancing brain development. *Acta Paediatrica* 97(7):853–858.
Cissé M, Yalçındağ S, Kergosien Y, Şahin E, Lenté C, Matta A (2017) Or problems related to home health care: A review of relevant routing and scheduling problems. *Oper. Res. Health Care* 13:1–22.
Corsello CM (2005) Early intervention in autism. *Infants Young Children* 18(2):74–85.
Deo S, Iravani S, Jiang T, Smilowitz K, Samuelson S (2013) Improving health outcomes through better capacity allocation in a community-based chronic care model. *Oper. Res.* 61(6):1277–1294.
Diamant A, Milner J, Quereshy F (2018) Dynamic patient scheduling for multi-appointment health care programs. *Production Oper. Management* 27(1):58–79.
Feldman J, Liu N, Topaloglu H, Ziya S (2014) Appointment scheduling under patient preference and no-show behavior. *Oper. Res.* 62(4):794–811.
Fikar C, Hirsch P (2017) Home health care routing and scheduling: A review. *Comput. Oper. Res.* 77:86–95.
Gupta D, Denton B (2008) Appointment scheduling in health care: Challenges and opportunities. *IIE Trans.* 40(9):800–819.
Gupta D, Wang L (2008) Revenue management for a primary-care clinic in the presence of patient choice. *Oper. Res.* 56(3):576–592.
Haggerty JL, Reid RJ, Freeman GK, Starfield BH, Adair CE, McKendry R (2003) Continuity of care: A multidisciplinary review. *BMJ* 327(7425):1219–1221.
Han Y, Johnson ME, Shan X, Khasawneh M (2024) A multi-appointment patient scheduling system with machine learning and optimization. *Decision Anal. J.* 10:100392.
Kolisch R, Sickinger S (2008) Providing radiology health care services to stochastic demand of different customer classes. *OR Spectrum* 30(2):375–395.
Li X, Wang J, Fung RY (2018) Approximate dynamic programming approaches for appointment scheduling with patient preferences. *Artificial Intelligence Medicine* 85:16–25.
Lin J, Muthuraman K, Lawley M (2011) Optimal and approximate algorithms for sequential clinical scheduling with no-shows. *IIE Trans. Healthcare Systems Engrg.* 1(1):20–36.
Liu N, Ziya S, Kulkarni VG (2010) Dynamic scheduling of outpatient appointments under patient no-shows and cancellations. *Manufacturing Service Oper. Management* 12(2):347–364.
Lu Y, Xie X, Jiang Z (2018) Dynamic appointment scheduling with wait-dependent abandonment. *Eur. J. Oper. Res.* 265(3):975–984.
Majnemer A (1998) Benefits of early intervention for children with developmental disabilities. *Seminars in Pediatric Neurology* 5(1):62–69.
Marynissen J, Demeulemeester E (2019) Literature review on multi-appointment scheduling problems in hospitals. *Eur. J. Oper. Res.* 272(2):407–419.
Moradi S, Sauré A, Patrick J (2025) Express: Advance multi-priority, multi-appointment patient scheduling with dependent demand and lead times. *Production Oper. Management* 35(3):855–875.
Parizi MS, Ghate A (2016) Multi-class, multi-resource advance scheduling with no-shows, cancellations and overbooking. *Comput. Oper. Res.* 67:90–101.
Patrick J, Puterman ML, Queyranne M (2008) Dynamic multipriority patient scheduling for a diagnostic resource. *Oper. Res.* 56(6):1507–1525.
Saure A, Patrick J, Tyldesley S, Puterman ML (2012) Dynamic multi-appointment patient scheduling for radiation therapy. *Eur. J. Oper. Res.* 223(2):573–584.
Savelsbergh M, Smilowitz K (2016) Stratified patient appointment scheduling for mobile community-based chronic disease management programs. *IIE Trans. Healthcare Systems Engrg.* 6(2):65–78.
Shao Y, Bard JF, Jarrah AI (2012) The therapist routing and scheduling problem. *IIE Trans.* 44(10):868–893.
Sheinkopf SJ, Siegel B (1998) Home-based behavioral treatment of young children with autism. *J. Autism Developmental Disorders* 28(1):15–23.
Spittle AJ, Morgan C (2025) Early intervention for children with cerebral palsy. Panteliadis CP, ed. *Cerebral Palsy: From Childhood to Adulthood* (Springer, Berlin), 197–205.
Wang J, Fung RY (2015) Adaptive dynamic programming algorithms for sequential appointment scheduling with patient preferences. *Artificial Intelligence Medicine* 63(1):33–40.
Wang WY, Gupta D (2011) Adaptive appointment systems with patient preferences. *Manufacturing Service Oper. Management* 13(3):373–389.
Wang J, Chen YF, Xu M (2018) Optimization and approximation methods for dynamic appointment scheduling with patient choices. *Comput. Oper. Res.* 92:65–76.