Hospitalization versus Home Care: Balancing Mortality and Infection Risks for Hematology Patients

Mor Armony
NYU—Stern, marmony@stern.nyu.edu

Galit B. Yom-Tov
Technion—Israel Institute of Technology, gality@technion.ac.il

Problem definition: Previous research has shown that early discharge of patients may hurt their medical outcomes. However, in many cases the "optimal" length of stay (LOS) and the best location for treatment of the patient are not obvious. A case in point is hematology patients, for whom these are critical decisions. Patients with hematological malignancies are susceptible to life-threatening infections after chemotherapy. Sending these patients home early minimizes infection risk, while keeping them longer for hospital observation minimizes mortality risks if an infection occurs. We develop LOS optimization models for hematology patients that balance the risks of patient infection and mortality.

Methodology/results: We develop a Markov decision process formulation to explore the impact of the infection and mortality risks on the optimal LOS from a single-patient perspective. We further consider the social optimization problem in which capacity constraints limit the ability of hospitals to keep patients for the entirety of their optimal LOS. Using fluid models, we find that the optimal solution takes the form of a two-threshold policy. This policy may block some patients and immediately route them to home care, or speed up some patients LOS and send them to be home-cared early after an observation period at the hospital.

Managerial implications: Physicians can use our model to determine a personalised optimal LOS for patients according to their infection and mortality risk characteristics. Furthermore, they can adjust that decision according to the current hospital load. In a case study, we show that around 75% of the patient population need some observation period. If the hospital is overloaded, using a speedup only policy is optimal for 90% of the patient types; applying it to all patient types increases overall mortality risk by 0.5%.

1. Introduction

Cancer is one of the leading causes of death in the US. In addition to the enormous toll of cancer treatment on patients' health, the costs associated with such treatment are significant: according to AHRQ (2014), $87.8 billion was spent in 2014 in the US on cancer-related healthcare, with 27% spent on hospital inpatient stays. Indeed, cancer inpatients incur higher total hospital costs and longer length of stay (LOS) compared to non-cancer inpatients (Suda et al. 2006). Cancer patients are treated at the hospital for the disease itself as well as its potential complications.

For example, cancer inpatients are more susceptible to healthcare-associated infections (HAI) than non-cancer inpatients (Cornejo-Juárez et al. 2016). These infections increase patients' LOS and mortality. Mortality risks following an infection are so high for some types of cancer (e.g., acute leukemia) that patients may stay at the hospital just for the sake of being monitored to allow for a quick response in case an infection occurs (Carmen et al. 2019). Thus, a patient "optimal" LOS is a decision that depends on infection and mortality risks, among other factors. It is therefore important to understand the considerations associated with the LOS decision (or, put differently, the decision of when to send a patient home), especially for cancer patients.

Most of the current operational literature assumes that, everything else being equal, it is better to keep a patient at the hospital for as long as is medically justified. In particular, in the absence of cost considerations or capacity constraints, there should be no rush to send a patient home. One critical consideration that is often absent from these discussions is the risk of a patient developing an infection while at the hospital. According to Magill et al. (2018), in 2015, 3.2% of hospitalized patients developed HAI. This suggests that there is an important added benefit to sending a patient home early that is beyond cost or capacity savings.

Focusing on the question of when to send a patient home is closely related to recent trends in patient care of moving treatments out of the hospital and into outpatient clinics or patients' homes (van Tiel et al. 2005), given that both settings allow for physician and nurse visitations as well as various tests. These solutions have been shown to be safe for some patient populations (van Tiel et al. 2005); however, for appropriate implementation, one needs to consider the risks and benefits of the hospital versus out-of-hospital alternatives with respect to each individual patient. In this work, we develop a systematic approach that takes these factors and others into account to determine optimal patient length of stay, shedding light on how in-hospital and home care observations should be combined.

Our work builds on a recent empirical study (Carmen et al. 2019) conducted in a hematology ward (HW) of an Israeli hospital that indicates that the risk of developing an infection at the hospital is greater than the risk of developing an infection at home, and that both risks change over time. This is illustrated in Figure 1(a). On the flip side, provided that a patient has developed an infection, the survival rates are higher if the infection has occurred at the hospital due to the proximity of adequate care (see Figure 1(b)). In light of this tradeoff we are interested in gaining insight into the decision of when to send a patient home based on their location-dependent infection hazard rate and mortality risk. What makes this decision even more challenging is that the shape of the time-to-infection hazard rate is dependent upon patient-specific factors (such as their specific disease), as is illustrated in Figure 1(c). As argued above, due to the tradeoff between infection

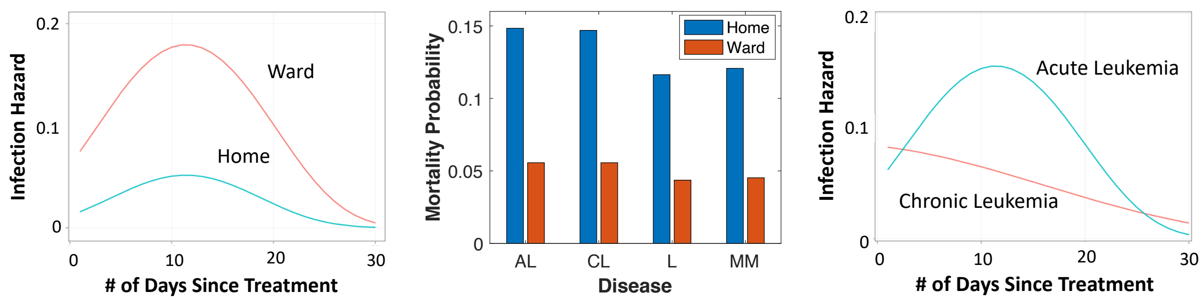

(a) Infection hazard (b) Mortality probability given infection for four disease types (c) Infection hazard as a function of disease type for in-ward patients

Figure 1 Infection and Mortality Hazard Rate Functions (Carmen et al. 2019) (Acute Leukemia (AL), Chronic Leukemia (CL), Lymphoma (L), and Multiple Myeloma (MM))

and mortality risks, this decision is relevant even in the absence of cost considerations or capacity constraints. Thus, we examine both the uncapacitated and the capacitated cases.

In the uncapacitated case (§3) our focus is on the individual patient assuming that the interaction between patients is minimal. We formulate the problem of choosing when to send a patient home as a Markov decision process (MDP) with the objective of maximizing total expected survival reward minus hospitalization cost. We find that the optimal policy is not necessarily of a threshold type. Nevertheless, under the realistic assumption that once a patient is sent home she will not be sent back to the hospital unless an infection ensues, any deterministic policy is effectively of a threshold type. Indeed, we make this assumption. Interestingly, even if the hospitalization cost is zero, this effective threshold may imply sending the patient home earlier than the maximal time allowed in order to reduce infection risk. Surprisingly, we find that the effective threshold is *not* monotone in the at-home infection hazard rate.

Turning our attention to the capacitated case (§4), the first-order question to address is how much capacity is needed to handle the system demand. This is equivalent to asking what is the system's offered load, i.e., what is the expected number of patients in the ward if capacity was unlimited. Unfortunately, in practice, due to high costs stemming from the fact that these patients need isolated rooms during treatment and highly proficient medical staff, and due to the large increase in demand in recent years stemming from improvements in treatment and patient survival (LLSC 2016), the system's capacity is typically not sufficient to handle all of this offered load. For example, data from the Technion SEELab show that the hematology unit at the Rambam Medical Center (a large tertiary hospital in Israel) has an average of 97% occupancy. It is therefore critical to consider systems with limited capacity and to come up with rules that determine which patients to send home and what is the best time to do so.

To address the question of offered load, note that traditionally studies on offered load in queueing systems take service time as exogenous (Whitt 2013). For example, for a stationary multiserver G/G/N queue the offered load is equal to the arrival rate times the expected service time. What makes our setting unique is that the service time here is a *result of an optimization*, determined as the optimal hospital length of stay and is an output of the MDP we formulated for the uncapacitated case. Moreover, the offered load can be adjusted as needed by changing the time that patients are sent home.

The optimization of the timing of sending patients home given an overloaded system addresses some important questions. Clearly, the optimal thing to do, if this was possible, would be to send every patient home at the optimal time prescribed by the analysis of the uncapacitated case. But given that the system does not have sufficient capacity to handle all of the demand, would it be best to use an equitable policy that sends everyone home earlier than the optimal time? Or would it be better to send some patients home at the optimal time while sending the rest of the patients home right away? Or would a much more complex policy be needed, with multiple thresholds applied to multiple groups of patients? We find that while the universe of possible policies is large, the capacitated problem reduces to the much simpler problem of dividing the patient population into up to *two* groups, with each group having its own threshold of when to send its patients home. We further explore the specifics of this two-class threshold policy as it applies to certain empirically driven risk functions.

Using a case study (§5) that is based on real patient data, we demonstrate how optimal LOS varies with patient characteristics and suggest that 75% of the patient population needs *some* observation period. We then examine how patient-discharge policies change if capacity constraints apply. Finally, we show that a variety of two-threshold policies should be used for some patient groups, but that a single-threshold speedup is by far the most common one.

2. Literature Review

Our paper is related to several research streams: optimization of medical decisions, optimization of patient flow and resource utilization, and asymptotic approximations of queueing systems.

Our research is motivated by cancer treatment. According to the World Health Organization, "Cancer is the second leading cause of death globally, and is responsible for an estimated 9.6 million deaths in 2018. Globally, about 1 in 6 deaths is due to cancer" (`https://www.who.int/news-room/fact-sheets/detail/cancer`). Within the various types of cancer, patients with hematological malignancies are known to be highly susceptible to infections, since the disease and/or therapy significantly weaken their immune system, leading to considerable infection-related mortality (Cornely et al. 2015, Halfdanarson et al. 2017, Taccone et al. 2009). While the past decade has witnessed significant advances in treatment strategies for hematological cancers, prevention of infections and adequate treatment for infected patients still pose a major challenge. In a large retrospective study of more than 41,000 cancer patients admitted to the hospital due to suspected infection, mortality rates among those who were treated for leukemia, lymphoma, and myeloma were as high as 14.3%, 8.9%, and 8.2%, respectively (Kuderer et al. 2006).

Our paper is inspired by empirical evidence accumulated in recent years on how the physical location of treatment or observation impacts health outcomes. Specifically, Carmen et al. (2019) showed that the location choice for post-treatment observation (dedicated ward, general ward (GW), or home) impacts infection and mortality among hematology patients. Chan et al. (2019) showed that the level of treatment of critical patients, as indicated by the type of ICU, impacts mortality. And, in a broader view, Song et al. (2020) showed that patient off-placement impacts LOS, mortality, and readmission. Overall, these studies indicate that choosing the right location is a key factor in treatment outcomes.

Patients' length of stay has been shown to be connected with health outcomes as well. Reducing LOS (speedup) in response to high load was shown to increase mortality of cardiothoracic surgery patients (Kc and Terwiesch 2009), and increase ICU patients' readmission (Kc and Terwiesch 2012). In contrast, increasing LOS was shown to reduce mortality of patients with acute myocardial infarction (Bartel et al. 2020), indicating that the influence of LOS on health outcomes is intricate. This raises the need to explicitly include this influence in models designed to optimize LOS. For example, Chan et al. (2012) formulated such a model to support decisions to discharge ICU patients to a step-down unit. More recently, a hospital-wide optimization method was proposed by Shi et al. (2020) to support patient-discharge decisions.

Most of the literature up to now has assumed that the best location for patients is at the hospital, and that early discharges are driven by limited capacity or high hospitalization cost. Our situation is different. In the hematology case, there is a delicate tradeoff between location, LOS, and health, due to the reduced risk of developing an infection at home. Therefore, discharge decisions are not driven only by capacity and cost considerations.

Patient early discharge (speedup) is one way for hospitals to deal with over-congestion. An alternative policy is to block patients from entering the hospital. For example, for ED services, this can be done by applying ambulance diversion (Allon et al. 2013) or, more subtly, by providing wait time information (Dong et al. 2019b). In our context, one can view the decision of sending a patient home immediately after treatment as a form of service denial. Finally, some literature explicitly deals with the tradeoff of admission control versus speedup. This tradeoff has been studied in the Markovian setting of a single-server queueing system (Adusumilli and Hasenbein 2010, Ata and Shneorson 2006), a multi-server queueing system (Lee and Kulkarni 2014, Yom-Tov and Chan

2021), and a multi-class queueing system (Ulukus et al. 2011, Turhan et al. 2012). Here, in our capacitated model, we generalize the distributional assumptions, by considering general service times and general risk functions in a multi-server setting. Indeed, allocating limited capacity when patients health condition may change dynamically over time was considered in the context of community-based treatments of chronic diseases (Deo et al. 2013), mass-casualty events (Mills et al. 2013) and hospital treatment (Nambiar et al. 2020, Ouyang et al. 2020).

Mathematically, our analysis of the capacitated model builds on the literature of heavy-traffic approximations for general queueing systems. Most closely related to our setting is the framework developed by Whitt (2006) to use a fluid model to in the analysis of the G/G/n+GI queueing system in overload. This framework was further utilized by Bassamboo and Randhawa (2016) who studied prioritization policies in a queueing model with general abandonment and general service times (G/G/n+GI). They focused on who to serve next given that some customers will abandon if they have to wait too long before their service starts; in particular, in their framework, once service starts it cannot be interrupted. By contrast, we focus on the decision of who to send home next, which in queueing theory terms translates to a decision of whose service time should be truncated and by how much.

3. The Uncapacitated Case: A Single-Patient Perspective

We start by studying the uncapacitated case, in which we assume a single patient and no capacity constraints, and determine the optimal length of stay that would maximize the patient's expected survival reward minus hospitalization cost. Our approach relies on an MDP formulation of the decision of when to send the patient home.

3.1. MDP Formulation for a Single Patient

We propose a discrete-time MDP formulation in which the physician makes a decision every period (day) on whether to keep the patient in the hospital or to discharge her. Cancer treatments are given in cycles, and our focus here is on a single arbitrary cycle. Each cycle starts with chemotherapy treatment followed by a recovery stage to allow the immunization system to recuperate. The treatment stage has a fixed duration according to the treatment protocol while the recovery stage can be done using hospital protective isolation or home care. Thus, the recovery stage is flexible and is the focus of our decision model. Our formulation assumes a finite decision horizon for the recovery stage of a specific cycle. This is consistent with the working assumption for hematological wards that if a patient has not gotten infected within the first 30 days after treatment, she will not get infected due to that treatment within that treatment cycle any longer (Carmen et al. 2019) (i.e., any infection afterwards is assumed to be unrelated to the specific chemotherapeutic treatment).

In our MDP formulation we focus on the decision of when to send a patient home after the treatment and *prior* to developing infection. If and when an infection occurs it is clear that the patient should be hospitalized. Thus, we refer to a state wherein a patient has developed an infection as an absorbing state. We also make the realistic assumption that once a patient has been sent home, she will not return to the hospital unless she has developed an infection. Thus, sending a patient home will also result in a transition to an absorbing state. The states and transitions are illustrated in Figure 2. According to our formulation, the MDP components are defined as follows:

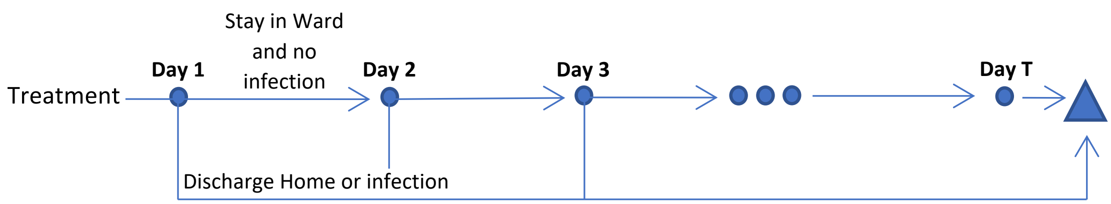

Figure 2 **MDP state transition**

- Define $\mathcal{S}$, the set of system states. We interpret the state s, for $s \in \mathcal{S}$, $s \in \{1,...,T\}$, as being at the beginning of day s after completing $s-1$ hospitalization days. In addition, we have an absorbing state, Δ (indicating that the patient was either discharged or infected). Hence, $\mathcal{S} = \{1,...,T,\Delta\}$. The initial state is $s=1$ and the final state is $s=\Delta$.
- Define $\mathcal{A}$ as the set of admissible actions. In general, $\mathcal{A} = \{w,h\}$, where w stands for *ward* (we use "ward" and "hospital" interchangeably) and h for *home*. We denote by $a(s)$ the action taken in state s. Hence, $a(s) = w$ means that the patient stays in the hospital ward in state s, and $a(s) = h$ means that the patient is discharged at state s. It is assumed that at the beginning of the horizon (time 0) the patient is at the hospital.
- Let $t = 0,...,T$ denote the time period. We use a subscript t to denote the state or action at time t.
- Define P as the probability transition matrix. The entry $P(s,s',a)$ in the matrix P describes the probability of moving from state s to s' given choice of action a. Hence, $P(s,s',a) = Pr(s_{t+1} = s'|s_t = s, a_t = a)$. If the patient is discharged (i.e., $a(s) = h$), she moves from state $s \in \{1,...,T\}$ to Δ. Hence, $P(s,\Delta,h) = 1$ for all $s \in \{1,...,T\}$.

If the patient stays at the hospital for observation for another day (i.e., $a(s) = w$) then she may move to state $s+1$ if she develops no infection during that period, or move to Δ if she does. Hence, $P(s,s+1,w)$ and $P(s,\Delta,w)$ for all $s \in \{1,...,T\}$ are determined by the hazard rate of developing an infection in state s at the hospital. We define by $r_w(s)$ (by $r_h(s)$) the risk function[1] of developing an infection at time s given that the patient is in the ward (at home) at the beginning of that period and has not developed an infection. Then, $P(s,\Delta,w) = r_w(s)$ and $P(s,s+1,w) = 1 - r_w(s)$ for all $s \in \{1,...,T-1\}$. Formally, we assume that $r_w(T) = r_h(T) = 0$.

Once we get to state Δ, we stay there indefinitely. Hence, $P(\Delta,\Delta,a) = 1$ for all $a \in \{h,w\}$. At the end of the horizon, T, any uninfected patient who is still at the hospital is sent home. Thus, formally, we have that $P(T,\Delta,w) = P(T,\Delta,h) = 1$.

Note that the functions r_w and r_h are of general form. Figure 1(c) demonstrates the empirical shape of this function for two patient types: chronic leukemia and acute leukemia. As can be seen in the figure, the former exhibits a monotone decreasing infection hazard rate function and the latter patient type exhibits an increasing and then decreasing infection hazard rate function.

It is important to note that we assume that the hazard rate functions depend on the time that has elapsed since the completion of the treatment and on the location of the patient at that time. Importantly, the infection hazard rate function does not depend on the time that the patient was placed in that location. Thus, for example, $r_h(t)$ is the risk of developing an infection exactly t time units after treatment (provided that no infection has developed prior to that time) given that the patient is at home at that time and independently of when the patient was sent home from the hospital.

- Define a reward matrix R. Its elements $R(s,s',a)$ denote the reward gained from making the transition from state s to state s', given action a. To spell out the elements of the reward function we need to first define the specific gains and costs realized by each action in every state.

We normalize the rewards such that the patient receives a reward of 1 if she survives a cycle and 0 otherwise. A positive reward can be accumulated in one of two cases: either 1) the patient survived until state T without developing an infection (recall the assumption that one cannot develop an infection in state T or thereafter, i.e., $r_w(T) = r_h(T) = 0$.); or 2) the patient developed an infection and survived. A patient who has developed an infection survives with probability p_w (or p_h) if the infection developed while the patient was at the hospital (or home).

We also incur costs in each state. Denote by c ($c \geq 0$) the hospitalization cost (at the hospital) per day for days $1,...,T-1$. (Note that the patient does not incur a hospitalization cost on day T since in the last day discharge occurs automatically if the patient has not developed an infection). In the case where the patient incurs home care costs as well, c reflects the difference in cost per day.

Denote by c_I the cost of treating infections (I for infection). This infection treatment cost includes all hospitalization costs incurred during the period that the patient is hospitalized from

[1] We use risk function and hazard rate function interchangeably.

the beginning of the infection until recovery/death. The infection treatment cost is assumed to be independent of the time or the location at which the infection started. In the context of our formulation, the cost c_I is incurred once we move to state Δ, unless an infection did not occur at all up to time T. For convenience, we formulate this cost as a reward that is received if no infection occurred up to time T. We assume that $c_I >> c$. This is a reasonable assumption because, in expectation, treating an infection requires more than one day of hospitalization.

We next compute $R(s,\Delta,h)$ for all $s<T$, which is the total expected reward to go when the patient is *discharged* in state s. For ease of notation we denote this function by $R_h(s) := R(s,\Delta,h)$. This reward function takes into account the probabilities of getting an infection and of recovering at or after time s until the end of the horizon. This reward function can be computed by backward induction as follows:

$$\begin{aligned} R_h(T) \quad &= 1 + c_I \\ R_h(T-i) &= r_h(T-i)p_h + (1-r_h(T-i))R_h(T-i+1) \quad \forall i \in \{1,...,T-1\}. \end{aligned}$$

This can be rewritten as

$$R_h(T-i) = p_h\left(1 - \prod_{j=1}^{i}(1-r_h(T-j))\right) + (1+c_I)\prod_{j=1}^{i}(1-r_h(T-j)), \quad \forall i \in \{1,...,T-1\}. \tag{1}$$

It is easy to see that $R_h(s)$ is an increasing function of s, because it is a convex combination of p_h and $1+c_I$, where the latter term is clearly the greater of the two and its weight is increasing in s. It reflects the fact that if a patient is sent home later, the *cumulative* chance of getting an infection given that the patient has not developed an infection so far is decreasing.

To compute the reward gained from keeping a patient in the ward for an extra day, we need to consider two options: a) the patient gets infected on that day, or b) the patient remains uninfected for another day. In the former case, the expected reward is $p_w - c$: the system incurs a one-day hospitalization cost and receives a reward for the expected patient survival. In particular, $R(s,\Delta,w) = p_w - c$ for all $s \in \{1,...,T-1\}$. If the patient remains uninfected, the system incurs only the daily hospitalization cost, c. Thus, $R(s,s+1,w) = -c$, $\forall s<T$. If the patient reaches state T without getting an infection the reward is $1+c_I$. Hence, $R(T,\Delta,w) = 1+c_I$.

In sum, the immediate reward function R for all $s \in \{1,..,T,\Delta\}$ is

$$R(s,s',a) = \begin{cases} -c, & \text{if } \ s<T, \ s'=s+1, \ a(s)=w; \\ p_w - c, & \text{if } \ s<T, \ s'=\Delta, \ a(s)=w; \\ R_h(s), & \text{if } \ s<T, \ s'=\Delta, \ a(s)=h; \\ 1+c_I, & \text{if } \ s=T, \ s'=\Delta, \ \text{for all } a; \\ 0, & \text{if } \ s=\Delta, \ s'=\Delta, \ \text{for all } a. \end{cases} \tag{2}$$

We assume the following properties for the hazard rate function parameters:

Assumption 1. 1. *In-hospital infection risk is higher than home-care infection risk, i.e.,*

$$r_h(s) \leq r_w(s), \quad \forall s.$$

2. *Survival probability at the hospital is higher than survival probability at home, i.e.,*

$$p_h \leq p_w < 1.$$

As noted in the introduction, this assumption is consistent with the empirical findings of Carmen et al. (2019).

3.1.1. The MDP formulation. Define a policy π such that $\pi_t(s)$ is the action the physician takes at time t if the state is s. Denote by $V_T(\pi)$ the expected reward over the finite horizon T if policy π is used. In particular,

$$V_T(\pi) := E^{\pi}\left[\sum_{t=1}^{T-1} R(s_t, s_{t+1}, a_t) + R(s_T, \Delta, a_T)\right].$$

Let $v_t(s)$ be the optimal reward-to-go function from time t onward, given that the state at time t is s. (Note that, by definition, at time t, the state s can be either t or Δ.) By Equation (2), since no reward is gained once state Δ is reached we have that

$$v_t(\Delta) = 0, \quad \forall t = 1, ..., T.$$

For all $s \neq \Delta$, we have that

$$v_T(s) = 1 + c_I, \quad s = T,$$

and for $t = 1, ..., T-1$ and $s = t$, we have that

$$v_t(s) = \max \begin{cases} R_h(s), & a_t = h; \\ p_w \cdot r_w(t) + v_{t+1}(s+1) \cdot (1 - r_w(t)) - c, & a_t = w. \end{cases} \tag{3}$$

Let $R_w(s)$ be the reward-to-go if the patient is kept at the hospital at time $t = s$, $s \neq \Delta$, and the optimal action is taken from time $t+1$ onward. Then,

$$R_w(s) = p_w \cdot r_w(s) + v_{s+1}(s+1) \cdot (1 - r_w(s)) - c. \tag{4}$$

Now we can rewrite Equation (3) more compactly as

$$v_s(s) = \max\{R_h(s), R_w(s)\}, \quad s \neq \Delta. \tag{5}$$

Note that the problem at hand is an unconstrained MDP with finite state and action spaces. Thus, we can conclude that there exists an optimal policy that is non-randomized, using standard MDP theory. We next establish some structural properties of the MDP solution, and discuss their implications.

3.1.2. Structural properties of the optimal policy. Solving the single-patient MDP is *numerically* tractable given any set of system and patient parameters, due to the relatively small state space. The MDP formulation also allows us to uncover some important structural properties of the optimal solution such as threshold form and monotonicity with respect to various parameters. These structural properties are the focus of this section. Surprisingly, we find that the optimal solution is not necessarily of a threshold type. We also discover that, counter-intuitively, the optimal policy is not monotone in the risk of developing an infection at the hospital.

As is often the case for MDPs with a one-dimensional state space, one might expect the optimal solution for this MDP to be of a threshold type. It turns out that this is not necessarily true here. For the optimal policy to be of a threshold type it is necessary that if it is optimal to send a patient home at time t_0 then it is also optimal to send her home at times t, for all $t > t_0$, if she happens to still be at the hospital at that time. The counterexample in Table 1 shows that this is not always the case. Specifically, in this example, it is optimal to send a patient home at times 3 and 4 but not at time 5. Note that in this counterexample we assumed a ward-acquired infection risk function that is increasing and then decreasing, as is typical for some hematological maladies; see Figure 1(a).

Table 1 An Example of an Optimal MDP Solution that Is Not of a Threshold Type

State - s	$r_w(s)$	$R_h(s)$	$R_w(s)$	$v(s)$	Decision
1	0.29	2.07	2.30	2.30	ward
2	0.31	2.87	2.97	2.97	ward
3	0.33	4.00	3.97	4.00	home
4	0.35	5.59	5.37	5.59	home
5	0.3	7.84	7.90	7.90	ward
6		11.00	11.00	11.00	home by default

$c = 0.1, c_I = 10, p_h = 0.1, p_w = 0.7, r_h = 0.29, T = 6$

Although the optimal policy lacks a threshold structure, the system nevertheless has a well-defined *effective* threshold, which is the first time in which it is optimal to send a patient home. This is an effective threshold because it is assumed that once a patient is sent home she will not be readmitted to the ward during the same treatment cycle, unless she develops an infection. Formally, we denote the effective threshold as t_{opt}, where

$$t_{opt} := \min\{1 \leq t < T \mid R_h(t) \geq R_w(t)\},$$

and where, if $R_h(t) < R_w(t)$ for all $t < T$, then $t_{opt} := T$.

Next we examine the monotonicity properties of the effective threshold with respect to various system parameters. Intuitively, one expects the effective threshold to be monotonously increasing in p_w and r_h and monotonously decreasing in c, p_h, r_w, and c_I. It turns out that this intuition holds true for all of those parameters except for r_w, the risk of developing an infection at the hospital (for a counterexample see Table 2 below). The monotonicity statements given below are proved straightforwardly using the MDP formulation and backward induction.

Proposition 1. *Under Assumption 1, the effective threshold t_{opt} is:*

(a) monotone decreasing in c. That is, as the hospitalization cost, c, increases, the patient will be discharged home earlier.

(b) monotone decreasing in c_I. As the cost of hospitalization after infection, c_I, increases, the patient will be sent home earlier.

(c) monotone decreasing in p_h. As the survival probability in the case of infection at home, p_h, increases, the patient will be sent home earlier.

(d) monotone increasing in p_w. As the survival probability in the case of infection at the hospital, p_w, increases, the patient will be sent home later.

(e) monotone increasing in r_h. As the risk of infection at home, r_h, increases, the patient will be sent home later.

All the proofs for this section appear in Appendix A.

It is intuitive to expect the effective threshold to be monotone decreasing in r_w. That is, one would expect that the greater the risk of developing an infection in the ward, the sooner the patient will be sent home. It turns out that this intuition is incorrect, as is illustrated by the counterexample shown in Table 2. This table compares two scenarios, 1 and 2: in Scenario 1 the patient has a lower risk of infection at the hospital, r_w, compared to Scenario 2. At the same time it is optimal for the patient to be sent home sooner under Scenario 1. The example demonstrates that if the probability of recovering from an infection at the hospital, p_w, is large enough compared to the utility of sending a patient home, then as the risk of infection at the hospital, r_w, increases, the overall opportunity for recovery at the hospital increases and, therefore, keeping the patient longer at the hospital could be advantageous. We rigorously establish this intuition in Lemma 1 and Corollary 1, and then outline two simple sufficient conditions on the system primitives for monotonicity to hold in Corollaries 2 and 3.

We now establish sufficient conditions that support this latter explanation, by showing that under these conditions the effective threshold is indeed decreasing in the in-hospital infection probability.

Lemma 1 **(Monotonicity in r_w)**. *If $p_w \le v_s(s)$ for all s, then the effective threshold is monotone decreasing in r_w.*

Corollary 1. *If $p_w \le R_h(s)$ for all s, then the effective threshold is monotone decreasing in r_w.*

Table 2 An Example of Non-Monotonicity in the Hospital Infection Probability r_w

		Scenario 1			Scenario 2		
State (s)	$R_h(s)$	$r_w(s)$	$R_w(s)$	Decision	$r_w(s)$	$R_w(s)$	Decision
1	0.29	0.38	0.29	home	0.48	0.33	ward
2	0.37	0.37	0.36	home	0.47	0.38	ward
3	0.47	0.36	0.45	home	0.46	0.46	home
4	0.62	0.35	0.58	home	0.45	0.57	home
5	0.82	0.34	0.76	home	0.44	0.72	home
6	1.10		1.10	home		1.10	home

$c=0.2, c_I=0.1, p_h=0.1, p_w=0.7, r_h=0.28, T=6$

The following corollary allows us to establish a sufficient condition for monotonicity of the effective threshold in r_w as a simple expression of the system parameters. Note that both c and r_w play no role in this condition.

Corollary 2. *If the system parameters satisfy the condition*

$$\frac{p_w - p_h}{1 + c_I - p_h} \leq \prod_{j=1}^{T-1}(1 - r_h(T-j)), \tag{6}$$

then the effective threshold is monotone decreasing in r_w.

In particular, whenever c_I or p_h are large enough or p_w is small enough, monotonicity holds.

Finally, we establish, as another corollary to Lemma 1, that the effective threshold is monotone decreasing in the risk of developing infection in the hospital whenever the cost of hospitalization is 0.

Corollary 3. *If $c=0$, then the effective threshold is monotone decreasing in r_w.*

3.2. Characterizing the Effective Threshold

In this section we focus on the family of effective threshold policies. As per our discussion so far, the optimal solution to the MDP is, by assumption, an effective threshold, and thus, in looking for the optimal policy, it is sufficient to focus on this family. Here we study the value function associated with all effective threshold policies and specify some sufficient conditions under which this value function has has a simple, easy-to-characterize, maximum. As argued above, the maximizing effective threshold is in turn an optimal solution to the MDP.

To be concrete, we define an effective threshold (τ) policy to be such that a patient stays in the ward until time τ and is sent home at that time if an infection has not been developed by then. Denote by J_τ the total expected reward of the τ policy from time 1 to T. For example, J_1 is the total expected reward of a patient who is sent home immediately at time 1 (we refer to this strategy as blocking) and $J_{t_{opt}}$ is the total expected reward if the optimal effective threshold is used. By definition of t_{opt}, J_τ obtains is maximal value at $\tau = t_{opt}$.

For an arbitrary effective threshold τ, define $R_\tau(t)$ to be the reward-to-go under the τ policy from time t to T. From time τ to T, the reward gained from discharging the patient at time τ is equal to $R_\tau(t) = R_h(t)$, for $t \in \{\tau, ..., T\}$. From time 1 to $\tau - 1$, the patient is assumed to be in the ward. Hence,

$$R_\tau(t) = r_w(t)p_w + (1 - r_w(t))R_\tau(t+1) - c, \quad \forall t \in \{1, ..., \tau - 1\}. \tag{7}$$

By definition, the total expected reward satisfies

$$J_\tau = R_\tau(1). \tag{8}$$

We note that the optimal effective threshold is the one that maximizes J_τ as a function of τ. We next analyze specific cases in which J_τ may have a simple structure and where the resulting optimal policy can be simply inferred. The first two results refer to cases where the fundamental tradeoff between risk of developing infection and speed of access to treatment does not exist (i.e., when Assumption 1 does not hold).

Proposition 2. *If $r_w(\cdot) \le r_h(\cdot)$, $p_w \ge p_h$ and $c = 0$ then J_τ is monotone increasing. Therefore, the optimal policy is to stay hospitalized indefinitely (until time T).*

Proposition 3. *If $r_w(\cdot) > r_h(\cdot)$ and $p_w = p_h$ then J_τ is monotone decreasing. Therefore, the optimal policy is no observation regardless of the hospitalization costs.*

Note that the assumption that $p_w = p_h$ in Proposition 3 is strong, and may be realistic only if the patient lives in close proximity to the hospital, and the hospital has dedicated capacity in the emergency department for these patients. As we discussed in the introduction, it is more realistic to assume that $p_w > p_h$, as in Assumption 1. Consistent with that assumption, the case where $r_w(\cdot) > r_h(\cdot)$ and $p_w > p_h$ is more realistic, and is also more complicated. In that case, we are able to characterize the structure of the solution only when the risk functions are constant over time, and show that in that case J_τ is quasi-concave. To that end, assume that the risk functions $r_w(\cdot)$ and $r_h(\cdot)$ are constant over time and define $f(\tau)$ to be the difference in the reward-to-go at time $\tau - 1$ if the threshold is changed from τ to $\tau - 1$, i.e., $f(\tau) = R_\tau(\tau - 1) - R_{\tau-1}(\tau - 1)$. In particular,

$$f(\tau) = r_w(p_w - p_h) - (r_w - r_h)(1 - r_h)^{T-\tau}(1 + c_I - p_h) - c, \quad 1 < \tau \le T, \tag{9}$$

which can be seen by noting that $R_\tau(\tau) = R_{\tau-1}(\tau) = R_h(\tau)$ and by referring back to (1) and (7).

Proposition 4. *Assume that $0 < r_h < r_w < 1$ and both are constant over time and that Assumption 1 holds. Let $\hat{\tau} := \max\{2 \le \tau \le T \mid f(\tau) > 0\}$, and $\hat{\tau} := 1$ if $f(\tau) \le 0$ for all $\tau \ge 2$. Then, J_τ is (strictly) increasing up to $\hat{\tau}$ and is decreasing afterwards. Specifically, if $\hat{\tau} = T$ (which is true if and only if $f(T) > 0$), then J_τ is monotone increasing. If $\hat{\tau} = 1$ (which is true if and only if $f(\tau) \le 0$ for all $\tau \ge 2$), then J_τ is monotone decreasing. Thus, it is optimal to send the patient home at time $\hat{\tau}$; i.e., $\hat{\tau}$ is an optimal effective threshold.*

We will see in Section 5.1, that although we cannot prove such a result for general risk functions, in most cases our date indicate that J_τ is indeed increasing-decreasing as in Proposition 4.

4. The Capacitated Model: A Hospital Ward Perspective

So far we have examined the problem of when to send a patient home assuming that the hospital has ample capacity. In practice, a hospital may be capacity-constrained, where patients may be forced to be sent home sooner than is optimal, due to lack of space, beds, equipment, or medical personnel. In such situations one might be forced to deviate from the optimal discharge policy in a meaningful way. This section focuses on the question of how to determine when to send a patient home when there is not enough hospital capacity to implement the uncapacitated solution without modifications.

The first question that comes to mind in this context is how to determine when the hospital is indeed capacity constrained. To address this question it is useful to model the hospital ward as a queueing system. Consider a hospital ward with n beds. We can think of these beds as multiple parallel servers in a queueing system. Let λ be the arrival rate of patients into the ward and let S be a random variable that represents the patient length of stay. Let $U = \lambda E[S]$ be the offered load inflicted on the system by its patients. Then, to determine whether or not the system is under- or over-loaded one looks at U versus n. Specifically, if $U << n$, then the capacity constraints are not very restrictive, whereas if $U \geq n$, then the capacity constraints are substantial. A specific challenge that arises in this context is that the service time, S, and more specifically its expected value $E[S]$ are not exogenous. In other words, the patient length of stay is an *outcome* of optimization and may in fact be impacted by the system capacity.

To disentangle the patient length of stay from the system capacity we start by recalling another definition of offered load, which is the expected number of busy servers in an infinite server queue. Recall that t_{opt} is the optimal effective threshold for a patient in the uncapacitated case. In particular, in a system with an infinite number of servers a patient will stay at the hospital until time t_{opt} and then will be sent home, unless the patient has developed an infection prior to that time. We will refer to the stay of such a patient as a *full stay*. Let $S_{t_{opt}}$ be the patient length of stay given such a threshold policy. Then, we define the system offered load as $U_{t_{opt}} = \lambda E[S_{t_{opt}}]$, and consider the system to be overloaded if $U_{t_{opt}} >> n$.

In the context of our study it is natural to consider the overloaded regime for two reasons:

1. It is exactly in the overloaded regime where the tradeoff between utilizing capacity and optimizing hospital length of stay is critical.

2. The hospital ward that this study is motivated by is indeed overloaded, with an average occupancy of 97%. This high load is typical due to both the high cost of hematology hospitalization

(stemming from the patient's need to be isolated during and after treatment) and increased demand (stemming from advances in cancer treatment in recent decades).

We assume that patients who cannot be hospitalized for observation in the ward due to capacity constraints are sent home immediately. Therefore, our model will have blocking dynamics. In practice, a patient who does not have an available bed in the ward may be hospitalized in another hospital ward rather than be sent home. However, in general, this is not a desirable practice for hematology patients because not only will the patient be further exposed to hospital-acquired infections (Carmen et al. 2019), but this off-placement may result in inferior patient care (Song et al. 2020, Dong et al. 2019a). In fact Carmen et al. (2019) showed that Internal Wards may be strictly inferior to home-care for hematology patients. For the more general case, we briefly discuss in Section 6 how one might utilize our results to consider the three-way decision of Hematology ward vs. General ward vs. home. Thus, assuming for the time being that no patients are sent to a ward other than the hematology ward and that the same threshold t_{opt} is used to determine when to send a patient home for all patients who are not blocked, we can model this system as a $G/G/n/n$ loss system in the overload regime, with service time having the same distribution as $S_{t_{opt}}$.

The underlying policy where a patient will either be blocked or stay in the ward according to an effective threshold of t_{opt} is one possible policy to use in the capacitated case. However, one may argue that this policy is unfair because it favors the patients who are lucky enough to find an available bed in the ward over patients who are blocked. An alternative policy that treats all patients the same (in distribution) is, upon arrival of a patient to a full system, to send the patient with the longest LOS home and admit the new patient instead. In this case the length of stay distribution would be kept shorter overall, to allow *all* patients to spend some time in the ward. Clearly, there are many other policies to consider that combine elements from both of these two extreme-case policies.

To determine the best policy to use in the capacitated case, we need to find a way to enumerate all possible policies and to evaluate the reward as a function of the policy so as to find the policy that maximizes the reward of the entire patient population. The complexity of the underlying process makes the use of exact analysis prohibitively complex. Instead, we revert to an approximation using fluid models in which the discrete flow of customers is modeled as a continuous flow of fluid and where the discrete-time is replaced with continuous-time. This fluid model approach has been used by Whitt (2006), Kang and Ramanan (2010), and Zhang (2013) for queueing models with general service times and abandonment distributions where it was also rigorously justified in these settings. We use it in a way that mimics the approach taken by Bassamboo and Randhawa (2016) to optimize scheduling in an overloaded queueing system with impatient customers, albeit in a very different context. Here we do not rigorously prove that indeed the fluid model is the limit for the underlying stochastic system, rather, we simply postulate it.

4.1. The $G/G/n/n$ Fluid Model

Consider a $G/G/n/n$ system with generally distributed service time (LOS) $S_{t_{opt}}$ and offered load $U_{t_{opt}} = \lambda E\left[S_{t_{opt}}\right]$. Assume that the ward has n beds, and let $\rho = \frac{U_{t_{opt}}}{n} = \frac{\lambda E\left[S_{t_{opt}}\right]}{n}$. Further assume that the system is overloaded, and thus $\rho > 1$.

Because the system is overloaded, the occurrence of an arriving patient finding a full ward is frequent. Therefore, it is meaningful to consider what should be done when a patient arrives to a full ward. Two obvious options to consider if there are no available beds upon a new patient's arrival are:

1. Blocking: Send the new patient home, or
2. Speedup: Send home early the patient who has been in the ward the longest (before their optimal time of t_{opt}) to make room for the new patient.

More generally, a policy can mix between blocking and speedup and can also mix the choice rule as to which patient to send home in the case of a speedup and when. The fluid model can be utilized to calculate the approximated reward function for any given policy. We proceed with formulating the fluid model.

Define $F(x)$ as the cumulative distribution function (CDF) of the time until infection if the patient remains at the hospital for the full horizon T, and assume that F is continuous. In addition, denote by $F_\tau(x)$ as a different CDF which is a truncated version of $F(x)$ (that is, $F_\tau(x) = F(x)$ for $x < \tau$ and $F_\tau(\tau) = 1$). If we use the policy that keeps all patients in the ward until time τ and then sends them home (unless an infection has occurred prior to that time) then F_τ describes the patient service-time (LOS) distribution at the hospital.

The continuous time distribution necessitates an adaptation of the definition of the value function J_τ—the total expected reward associated with a single patient given a threshold policy τ—to continuous-time. Consider the two continuous-time infection hazard rate functions, $r_w(\cdot)$ and $r_h(\cdot)$, of the time until infection in the ward and at home, respectively. Then, for a threshold τ, we have that the infection hazard rate of the time until infection is $r_\tau(t) := r_w(t)1_{\{t<\tau\}} + r_h(t)1_{\{t\geq\tau\}}$, and, correspondingly, the CDF, $G_\tau(\cdot)$, of the time until infection under a threshold policy τ satisfies

$$G^c_\tau(t) := 1 - G_\tau(t) = \exp\left(-\int_0^t r_\tau(u)du\right) = \exp\left(-\int_0^{\tau\wedge t} r_w(u)du - \int_\tau^{\tau\vee t} r_h(u)du\right).$$

Thus, the value function J_τ may be expressed as

$$J_\tau = p_w G_\tau(\tau) + p_h G^c_\tau(\tau)\left(1 - \exp\left(-\int_\tau^T r_h(u)du\right)\right) - c\int_0^\tau G^c_\tau(u)du + (1+c_I)G^c_\tau(T), \tag{10}$$

where the first term corresponds to the product of the probability of contracting an infection in the ward and the probability of recovery. The second term stands for the product of the probability of contracting an infection at home and the probability of recovery. The third term is the hospitalization cost and the fourth term is the reward obtained if no infection has occurred.

Similarly to the discrete time distribution, we define t_{opt} to be the effective optimal threshold with respect to the reward function J_τ. That is, t_{opt} is the smallest threshold τ that maximizes J_τ. Note that here we simply postulate the existence of t_{opt} without proof. If an effective optimal threshold does not exist, one can simply replace it by $\tau = T$. Note that by the optimality of the effective threshold t_{opt}, it is never optimal to keep a patient in the ward beyond t_{opt}, but one may need to discharge the patient earlier due to capacity constraints.

Consider an arbitrary discharge policy π with discharge times that are no later than t_{opt}, and let F_π be the CDF of the service-time distribution under the policy π. Let S_π be the service-time of patients at the hospital under the policy π, and let μ_π be the corresponding service rate. In particular, we have that

$$\frac{1}{\mu_\pi} = E[S_\pi] = \int_0^{t_{opt}} F_\pi^c(x)dx.$$

For a threshold policy we will slightly abuse the notation τ to describe not only the threshold itself but also the policy associated with it. For example, the speedup policy described above is a threshold policy with a single threshold at some time $\tau = \tau_{spd}$, and we will refer to this policy simply as τ_{spd}.

To describe the fluid model in steady state we adopt the following characterization from Whitt (2006), adapted to a blocking (loss) system.

The G/GI/n/n Fluid Model in steady state. The G/GI/n/n fluid model with service-time distribution F_π, under overloaded conditions $\rho_\pi := \frac{\lambda E[S_\pi]}{n} > 1$, has a unique steady state q, where

$$q(0^-) = \bar{\lambda},$$

$$q(0) = \mu_\pi,$$

and

$$q(s) = \mu_\pi F_\pi^c(s), \quad s \geq 0,$$

for $\bar{\lambda} := \lambda/n$. The fluid blocking rate b_π is $\bar{\lambda} - \mu_\pi$.

Loosely speaking, $q(s)$ describes the fluid content in steady state of all of the fluid in the system that arrived exactly s time units ago. Figure 3 depicts the fluid model for the two special policies described above, namely, blocking (combined with full stay for patients who do receive a bed) and speedup, for the specific CDF functions F_π. Note that while the speedup policy in the stochastic system is not necessarily a threshold policy, it becomes a threshold policy under the (deterministic) fluid model because, in steady state, all the fluid that has been in the system the longest, arrived there at *exactly* the same time. Denote by τ_{spd} the speedup policy threshold, i.e., the maximal LOS of patients under a single-threshold that can be used when capacity constraint prevents everyone from staying until t_{opt}. To compute the speedup threshold, τ_{spd}, one needs to solve the equation

$$1/\bar{\lambda} = \int_0^{\tau_{spd}} F^c(x)dx,$$

whose solution exists due to the intermediate-value theorem and the overloading assumption.

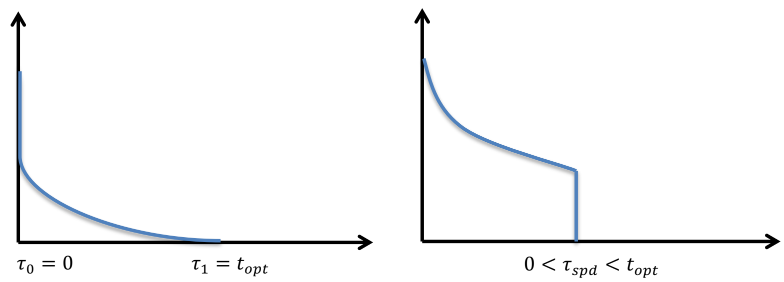

(a) Blocking or Full Stay (b) Speedup Only

Figure 3 Depiction of the Fluid Model in Steady State for the Blocking and Speedup Policies

4.2. General Policies

Up to now we have focused on a single-threshold policy where patients are sent home at the specified threshold unless they develop an infection beforehand. In general, we might consider a broader family of policies π in which, for all $x > 0$, a certain *fraction* $\psi_\pi(x)$ of the patients who have been in the ward for x time units are sent home at time x, where $0 \leq \psi_\pi(\cdot) \leq 1$. In addition, for $x = 0$, the policy blocks all patients who arrive to a full ward. A threshold policy τ ($\tau \geq 0$) is a special case of this family, with $\psi_\pi(\tau) = 1$ and $\psi_\pi(x) = 0$ for all $0 < x < \tau$. For simplicity we focus our attention on policies where the set of time points x such that $\psi_\pi(x) > 0$ is *finite*.[2] We refer to this finite set of $K + 2$ thresholds as $\vec{\tau} = \{\tau_i,\ i \in \{0, 1, ..., K, K+1\}\}, K < \infty$ where $0 = \tau_0 < \tau_1 < ... < \tau_K < \tau_{K+1} = t_{opt}$.

In the fluid model this policy may equivalently be described as a pair $(\vec{\tau}, \vec{\delta})$ of thresholds with a finite set of non-negative numbers $\vec{\delta} = \{\delta_i \geq 0,\ i \in \{0, 1, ..., K, K+1\}\}$, such that at time τ_i the fluid content is reduced by a mass of δ_i. Specifically, the fluid content in steady state for such a

[2] Given the continuity of the value function, J_τ, we may approximate any general policy to the required level of accuracy using a finite set of thresholds, hence, this assumption is not too restrictive (see Remark 1 in Bassamboo and Randhawa 2016).

general policy may be described as follows: the process starts at $q(0^-) = \bar\lambda$, then decreases instantly to $q(0) = \bar\lambda - \delta_0$, and, more generally, at time t we have that

$$q(t) = \left(\bar\lambda - \sum_{\{j:\tau_j \le t\}} \frac{\delta_j}{F^c(\tau_j)} \right) F^c(t), \quad \text{for} \quad t \ge 0. \tag{11}$$

From Equation (11) we see that for the policy to be admissible we must have that $\sum_{i=0}^{K+1} \frac{\delta_i}{F^c(\tau_i)} = \bar\lambda$.

To show that the two representations of the general policy—$(\vec\tau, \vec\delta)$ and ψ—are equivalent, note that we have that for all $x > 0$,

$$\psi(x) = \begin{cases} \frac{\delta_i}{q(\tau_i)+\delta_i} & \text{if} \quad x = \tau_i \text{ for some } i, \\ 0 & \text{otherwise.} \end{cases} \tag{12}$$

Given the above policy description it is natural to want to optimize over all the admissible pairs $(\vec\tau, \vec\delta)$ to find one that maximizes the overall value function in steady state. To do this, we next provide an alternative and equivalent description of such a policy. This alternative description lends itself well to evaluating the corresponding value function.

Following Bassamboo and Randhawa (2016) we observe that the admissible policy $\pi = (\vec\tau, \vec\delta)$ may alternatively be described as a partition of the patient population into $K+2$ classes with arrival rates $\bar\lambda_i = \frac{\delta_i}{F^c(\tau_i)}$, for $i = 0, 1, ..., K+1$, where for class i we apply the single threshold policy $\pi_i = \tau_i$. Figure 4 illustrates the equivalence between the two policy descriptions. A general admissible policy may be described, equivalently, as $\pi = (\vec\tau, \vec\lambda)$, since, given $\vec\tau$, there is a one-to-one correspondence between $\vec\delta$ and $\vec\lambda$.

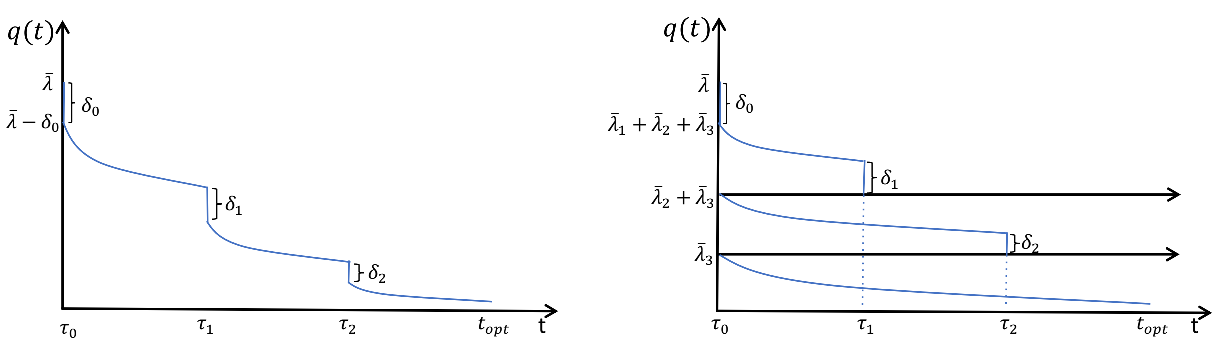

(a) Fluid dynamics (b) Fluid dynamics in a multi-class system

Figure 4 Fluid Dynamics with Multiple-Threshold Policy

Given this policy characterization, we are now in a position to state our fluid-level optimization problem.

4.3. The Fluid-level Capacitated Optimization Problem

To find the optimal policy for hospital length of stay in the capacitated case at the fluid level, we seek to divide the fluid-level patient population $\bar{\lambda}$ into $K+2$ classes of sizes $\bar{\lambda}_i$ $(i=0,1,...,K+1)$, and determine a set of $K+2$ discharge thresholds that will maximize the total value gained by the entire patient population. The resulting optimization problem is given by

$$\begin{aligned} \sup_{K\in\mathbb{N};\ (\vec{\lambda},\vec{\tau})\in\mathbb{R}_+^{K+2}\times\mathbb{R}_+^{K+2}} \quad & \sum_{i=0}^{K+1} \bar{\lambda}_i J_{\tau_i} \\ \text{s.t.} \quad & \sum_{i=1}^{K+1} \frac{\bar{\lambda}_i}{\mu_{\tau_i}} \leq 1 \\ & \sum_{i=0}^{K+1} \bar{\lambda}_i = \bar{\lambda}, \\ & 0 = \tau_0 \leq \tau_1 \leq \tau_2 ... \leq \tau_K \leq \tau_{K+1} = t_{opt}. \end{aligned} \tag{13}$$

where J_{τ_i} is the value function associated with a threshold policy τ_i as defined in (10) and $1/\mu_{\tau_i}$ is the expected time in the ward for a patient of class i:

$$\frac{1}{\mu_{\tau_i}} = \int_0^{\tau_i} x f(x)dx + \int_{\tau_i}^{\infty} \tau_i f(x)dx = \int_0^{\tau_i} x f(x)dx + \tau_i F^c(\tau_i),$$

recalling that F is the CDF of the time until infection for a patient who remains in the hospital ward, and f is the corresponding PDF.

We next argue that in optimality the system is critically loaded; that is, there exists an optimal solution to (13) where the first constraint is obtained as an equality.

LEMMA 2. *If the system is overloaded when all patients are sent home according to the threshold policy t_{opt}, then if an optimal solution to (13) exists, then there exists a solution to (13) where the constraint $\sum_{i=1}^{K+1} \frac{\bar{\lambda}_i}{\mu_{\tau_i}} \leq 1$ is obtained as an equality.*

All the proofs for this section appear in Appendix B.

The proof for this lemma relies on the overload assumption and the optimality of the discharge threshold t_{opt}. Specifically, for any solution where the system is not critically loaded we can find another solution where more capacity is utilized and the objective function is at least as high.

By Lemma 2, problem (13) may be equivalently rewritten as

$$\begin{aligned} \sup_{K\in\mathbb{N};\ (\vec{\lambda},\vec{\tau})\in\mathbb{R}_+^{K+2}\times\mathbb{R}_+^{K+2}} \quad & \sum_{i=0}^{K+1} \bar{\lambda}_i J_{\tau_i} \\ \text{s.t.} \quad & \sum_{i=1}^{K+1} \frac{\bar{\lambda}_i}{\mu_{\tau_i}} = 1 \\ & \sum_{i=0}^{K+1} \bar{\lambda}_i = \bar{\lambda}, \\ & 0 = \tau_0 \leq \tau_1 \leq \tau_2 ... \leq \tau_K \leq \tau_{K+1} = t_{opt}. \end{aligned} \tag{14}$$

For problem (14), observe that for a fixed value of K and a fixed set of thresholds $\vec{\tau}$, the problem becomes a linear program in $\vec{\lambda}$, with two constraints. In particular, a basic solution will have at most two non-zero values of λ_i. Therefore, much like in Bassamboo and Randhawa (2016), we can conclude that there exists a solution to (14) in which at most two of the classes are non-empty. This is formalized in the following proposition.

Proposition 5. *There exists an optimal solution to* (14) *with at most two non-empty classes.*

The proof is essentially identical to the proof of Proposition 1 in Bassamboo and Randhawa (2016) and is hence omitted. By Proposition (5), optimization problem (14) may be reduced to the following:

$$\begin{aligned} \sup_{0\le\bar{\lambda}_l,\bar{\lambda}_h;0\le\tau_l\le\tau_h\le t_{opt}} \quad & \bar{\lambda}_l J_{\tau_l} + \bar{\lambda}_h J_{\tau_h} \\ \text{s.t.} \quad & \frac{\bar{\lambda}_l}{\mu_{\tau_l}} + \frac{\bar{\lambda}_h}{\mu_{\tau_h}} = 1, \\ & \bar{\lambda}_l + \bar{\lambda}_h = \bar{\lambda}. \end{aligned} \tag{15}$$

The problem can be further simplified, by setting $\bar{\lambda}_h = \bar{\lambda} - \bar{\lambda}_l$, to

$$\begin{aligned} \sup_{0\le\bar{\lambda}_l\le\bar{\lambda};0\le\tau_l,\tau_h\le\infty} \quad & \bar{\lambda}_l J_{\tau_l} + (\bar{\lambda} - \bar{\lambda}_l) J_{\tau_h} \\ \text{s.t.} \quad & \frac{\bar{\lambda}_l}{\mu_{\tau_l}} + \frac{\bar{\lambda} - \bar{\lambda}_l}{\mu_{\tau_h}} = 1. \end{aligned} \tag{16}$$

Notice that, given the values of τ_l and τ_h, $\bar{\lambda}_l$ is determined uniquely by solving the equation given by the constraint of (16), as $\bar{\lambda}_l = \frac{\bar{\lambda} - \mu_{\tau_h}}{1 - \mu_{\tau_h}/\mu_{\tau_l}}$ if $\tau_l \ne \tau_h$, and $\bar{\lambda}_l = 0$ otherwise. Furthermore, in order for $\bar{\lambda}_l$ to satisfy the conditions in the optimization of (16) we must have that

$$\mu_{t_{opt}} \le \mu_{\tau_h} \le \bar{\lambda} \le \mu_{\tau_l} \le \infty. \tag{17}$$

Let τ^* be such that $\mu_{\tau^*} = \bar{\lambda}$, where $\frac{1}{\mu_{\tau^*}} = \int_0^{\tau^*} x f(x) dx + \tau^* F^c(\tau^*)$. (Note that $\tau^* \equiv \tau_{spd}$.) Then, (17) implies that any feasible solution must have $\tau_l \le \tau^* \le \tau_h$. Therefore, since $\mu_{\tau_h} > 0$ it follows that $\bar{\lambda}_l \le \bar{\lambda}$ and (16) can be simplified to

$$\sup_{0\le\tau_l\le\tau^*\le\tau_h\le\infty} \bar{\lambda}_l(\tau_l,\tau_h) J_{\tau_l} + (\bar{\lambda} - \bar{\lambda}_l(\tau_l,\tau_h)) J_{\tau_h}, \tag{18}$$

where $\bar{\lambda}_l(\tau_l,\tau_h) = \frac{\bar{\lambda} - \mu_{\tau_h}}{1 - \mu_{\tau_h}/\mu_{\tau_l}}$, if $\mu_{\tau_l} \ne \mu_{\tau_h}$, and $\bar{\lambda}_l(\tau_l,\tau_h) = 0$, otherwise.

The culmination of the discussion thus far is in the conclusion that the solution of the capacitated problem reduces into *five* simple and mutually exclusive cases, as stated in the following corollary.

Corollary 4. *The optimal solution to the capacitated problem* (13) *has an up-to-two threshold structure as in* (18)*, with the following exhaustive and mutually exclusive cases:*

1. *Block or Speedup (Bl-Sp):* $\tau_l = 0,\ 0 < \tau_h < t_{opt}$,
2. *Two-level Speedup (2×Sp):* $0 < \tau_l < \tau_h < t_{opt}$,
3. *Speedup only (1×Sp):* $0 < \tau_l = \tau_h \le t_{opt}$,
4. *Speedup or Full Stay (Sp-FS):* $0 < \tau_l < \tau_h = t_{opt}$,
5. *Block or Full Stay (Bl-FS):* $\tau_l = 0,\ \tau_h = t_{opt}$.

The five possible cases are depicted in Figure 5.

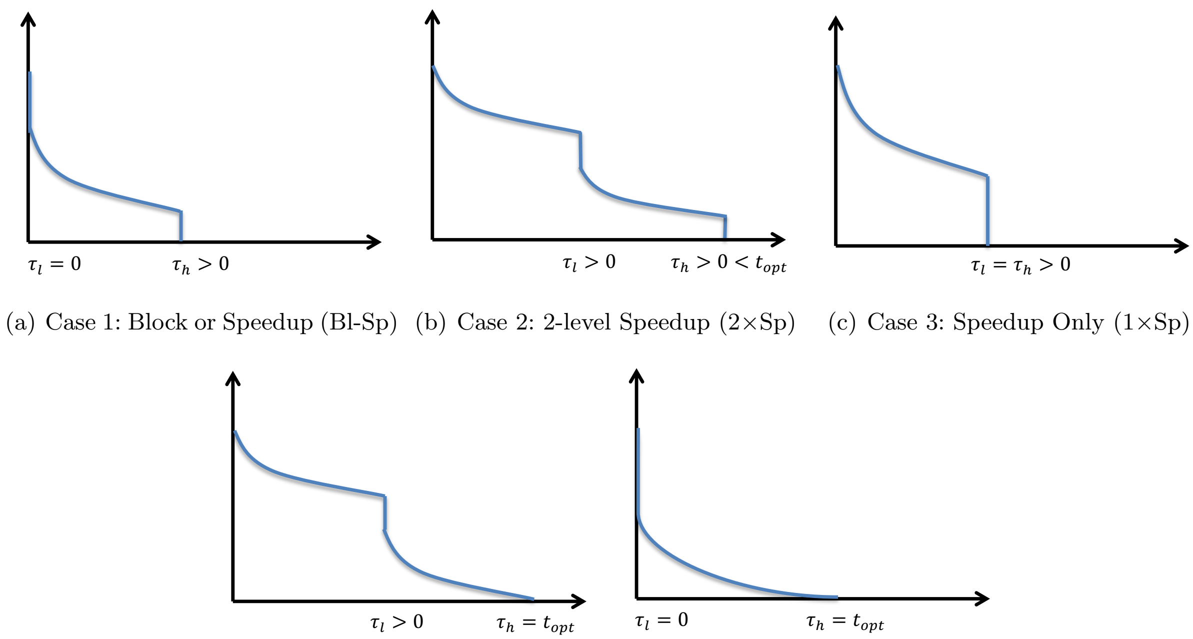

(a) Case 1: Block or Speedup (Bl-Sp) (b) Case 2: 2-level Speedup (2×Sp) (c) Case 3: Speedup Only (1×Sp)

(d) Case 4: Speedup or Full Stay (Sp-FS) (e) Case 5: Block or Full Stay (Bl-FS)

Figure 5 The five possible solutions to the capacitated optimization problem

Now that we have identified the five possible cases for the solution of the capacitated problem, we explore the question of whether one can identify the particular solution given some qualities of the problem primitives. We start with specifying some *necessary* conditions for the optimal solution of (18).

LEMMA 3 **(Necessary optimality conditions)**. *If the functions μ_τ and J_τ are both differentiable as a function of τ, with derivatives μ'_τ and J'_τ, respectively, then*

(a) An optimal solution to (18) of the form (τ_l, τ_h) with $0 < \tau_l < \tau^* < \tau_h < t_{opt}$ (the 2×Sp policy) must satisfy

$$\frac{J_{\tau_h} - J_{\tau_l}}{\mu_{\tau_h} - \mu_{\tau_l}} = \frac{\mu_{\tau_l}}{\mu_{\tau_h}} \frac{J'_{\tau_l}}{\mu'_{\tau_l}} = \frac{\mu_{\tau_h}}{\mu_{\tau_l}} \frac{J'_{\tau_h}}{\mu'_{\tau_h}}. \tag{19}$$

(b) An optimal solution to (18) of the form $(0,\tau_h)$ with $\tau^* < \tau_h < t_{opt}$ (Bl-Sp policy) must satisfy

$$\frac{J_0 - J_{\tau_h}}{\mu_{\tau_h}} = \frac{J'_{\tau_h}}{\mu'_{\tau_h}}. \tag{20}$$

(c) An optimal solution to (18) of the form (τ_l, t_{opt}) with $0 < \tau_l < \tau^* < t_{opt}$ (Sp-FS policy) must satisfy

$$\frac{(J_{t_{opt}} - J_{\tau_l})\mu_{t_{opt}}}{(\mu_{t_{opt}} - \mu_{\tau_l})\mu_{\tau_l}} = \frac{J'_{\tau_l}}{\mu'_{\tau_l}}. \tag{21}$$

Lemma 3 establishes that the ratio J'_τ/μ'_τ plays an important role in the characterization of the solution. This ratio expresses the marginal increase in value from increasing the threshold over the corresponding marginal decrease in service rate. While the former can potentially benefit the hospital, the latter hurts it because it increases the usage of capacity. This will be further formalized in Lemma 4 where we denote the negative value of this ratio as $\xi(\tau)$.

While Lemma 3 outlines necessary conditions for the optimality of the 2×Sp, Bl-Sp, and Sp-FS policies, it does not cover the two "boundary" policies of 1×Sp and Bl-FS. We are especially interested in characterizing conditions under which the *simplest* and most *equitable* policy that uses the same speedup threshold for all patients (1×Sp) is optimal. In the next lemma we will provide sufficient conditions for the optimality of these two boundary policies, as well as that of Bl-Sp, under the assumption that the function J is increasing up to the optimal threshold t_{opt}. In our empirically based numerical analysis in Section 5, we observe that virtually all the cases we encounter satisfy this assumption.

LEMMA 4 **(Sufficient conditions for optimality)**. *Assume that the functions μ_τ and J_τ are both differentiable as a function of τ*[3] *and that $J'_\tau > 0$ for all $0 \le \tau \le t_{opt}$. Define $\xi(\tau) := -\frac{J'_\tau}{\mu'_\tau}$. Then $\xi(\tau) > 0$. In addition, assume that $\xi(\tau)$ is monotone decreasing. Then,*

(a) The policies 2×Sp and Sp-FS are not optimal.

(b) If $\frac{J_\tau - J_0}{\mu_\tau} < \xi(\tau)$ for all $\tau^* \le \tau \le t_{opt}$ then the Bl-FS policy is optimal.

(c) If $\frac{J_{\tau^*} - J_0}{\mu_{\tau^*}} > \xi(\tau^*)$ then the policy 1×Sp is optimal.

(d) Finally, if there exists $\tilde{\tau}$ such that $\tau^* < \tilde{\tau} < t_{opt}$ with $\frac{J_\tau - J_0}{\mu_\tau} < \xi(\tau)$ for all $\tau^* \le \tau < \tilde{\tau}$ and $\frac{J_\tau - J_0}{\mu_\tau} > \xi(\tau)$ for all $\tilde{\tau} < \tau \le t_{opt}$, then the Bl-Sp policy is optimal with a speedup threshold $\tilde{\tau}$.

In examining the capacitated case one question that arises is what is the impact of the hospitalization cost c on the optimal policy. Specifically, since in the overloaded regime all beds are always occupied with probability 1, one may think that this cost has no impact on the optimal policy.

[3] We define the derivatives at 0 as the derivatives from the right and the derivatives at t_{opt} as the derivatives from the left.

However, we recall that the cost c is incorporated into the calculation of the function J_τ and hence does play a role in determining the optimal threshold t_{opt} (recall, for example, Proposition 1(a)). In particular, c plays a role in the optimal solution of the capacitated case, but only through its impact on J.

4.4. The Multi-Patient Type Case

In this section we extend the homogeneous-patient capacitated ward problem to heterogeneous patients. Consider a hospital ward with multiple types of patients, where each type has its own characteristics. Let type-j patients ($j = 1, ..., C$) have $r_h^j(\cdot)$ and $r_w^j(\cdot)$ infection hazard rate functions at home and in the ward, respectively, and let p_h^j and p_w^j be the corresponding probabilities of surviving that infection. These patient-type characteristics result in a type-dependent reward function J_τ^j, and a type-dependent optimal threshold t_{opt}^j. Additionally, consider the type-j arrival rate to be λ^j. Then, a hospital ward with limited capacity needs to determine when to send patients home while taking their types into account.

To address this problem we once again consider a system that operates in the overloaded regime. To define overload in this context we similarly let $S_{t_{opt}^j}^j$ be the length of stay of a type-j patient given a full stay with threshold t_{opt}^j. Let $U = \sum_{j=1}^{C} \lambda^j S_{t_{opt}^j}^j$ be the system's offered load. Assuming that the ward has n beds and, again, letting $\rho = \frac{U}{n}$, the system is considered overloaded if $\rho > 1$.

Similarly to the single-patient type case, we use a fluid-model approximation to evaluate the system's workload in steady state. We again consider a generic policy to be such that it divides each patient type j into $K^j + 2$ classes, where class i ($i = 0, ..., K^j + 1$) of patient type j has an arrival rate of $\bar{\lambda}_i^j$ and a single-threshold policy is applied to this class with a threshold of τ_i^j, with $0 = \tau_0^j \le \tau_1^j \le ... \le \tau_{K^j} \le \tau_{K^j+1} = t_{opt}^j$. The corresponding multi-type optimization problem can be written as

$$
\begin{aligned}
\sup_{\vec{K}\in\mathbb{N}^C;\ (\vec{\lambda}^j,\vec{\tau}^j)\in\mathbb{R}_+^{K^j+2}\times\mathbb{R}_+^{K^j+2}} \quad & \sum_{j=1}^{C}\sum_{i=0}^{K^j+1} \bar{\lambda}_i^j J_{\tau_i^j} \\
\text{s.t.} \quad & \sum_{j=1}^{C}\sum_{i=1}^{K^j+1} \frac{\bar{\lambda}_i^j}{\mu_{\tau_i^j}} \le 1, \\
& \sum_{i=0}^{K^j+1} \bar{\lambda}_i^j = \bar{\lambda}^j, \quad j = 1, ..., C, \\
& 0 = \tau_0^j \le \tau_1^j \le ... \le \tau_{K^j}^j \le \tau_{K^j+1}^j = t_{opt}^j, \quad j = 1, ..., C,
\end{aligned}
\tag{22}
$$

where the expected length of stay of a patient of type j and class i is

$$\frac{1}{\mu_{\tau_i^j}} = \int_0^{\tau_i^j} x f^j(x)dx + \int_{\tau_i^j}^{\infty} \tau_i^j f^j(x)dx = \int_0^{\tau_i^j} x f^j(x)dx + \tau_i^j(1 - F^j(\tau_i^j)).$$

For fixed values of $K^1, ..., K^C$ and of $\vec{\tau}^1, ..., \vec{\tau}^C$, the above is a linear program in $\vec{\lambda}^1, ..., \vec{\lambda}^C$ with $\sum_{j=1}^{C}(K^j+2)$ variables and $C+1$ constraints. Thus, similarly to Proposition 5, we can show that there exists an optimal solution to (22) with at most $C+1$ non-empty customer classes. This result is spelled out next.

PROPOSITION 6. *There exists an optimal solution to* (22) *that creates at most* $C+1$ *patient classes.*

Notwithstanding, any feasible solution to (22) will have at least one threshold associates with each of the C patient types due to the second set of constraints. The following corollary immediately follows:

COROLLARY 5. *There exists an optimal solution to* (22) *where* $C-1$ *of the patient types have a single type-specific threshold applied to them and up to one patient type has two type-specific thresholds applied to it.*

We thus conclude that handling the multi-patient type capacitated case is not as onerous as one might think. It involves figuring out what single patient-type should have two thresholds applied to it as well as identifying the relevant $C+1$ thresholds. Note that the $C-1$ classes that have a single threshold apply a speedup only policy (Figure 5(c)), where the speedup threshold in this case could also be 0.

5. Numerical Analysis

In this section we explore the practical implications of our theoretical results using real-world scenarios that are based on the empirical observations of Carmen et al. (2019). We start by exploring the uncapacitated, single-patient case and then continue by investigating the impact of imposing capacity constraints.

The paper Carmen et al. (2019) provides an empirical model that computes patient infection and mortality risks based on patient characteristics such as age, disease, chemotherapy treatment length (days), health condition at the end of treatment (based on white blood cell (WBC) counts), patient location, and more. Here we utilize the output of their model and feed these risk functions back into our model in order to glean insights into how our model might be used in practice.

5.1. The uncapacitated Case

For the uncapacitated case, we investigate how various factors affect the optimal observational hospital length of stay for hematology patients in the presence of ample capacity. Our analysis includes all four types of hematological cancer diseases: acute leukemia (AL), chronic leukemia (CL), lymphoma (L), and multiple myeloma (MM). For each disease type we vary some of the patient characteristics such as age and treatment length. We then run the MDP model presented in Section 3.1 to find the optimal effective threshold given these patient characteristics. Consistent with practice, for all cases we consider the maximum length of stay to be $T = 30$. In addition, for ease of exposition, we assume zero hospitalization costs (i.e., $c = c_I = 0$).

Figure 6 shows how the optimal effective threshold changes as a function of patient characteristics such as type (disease, age), current medical state (treatment protocol, WBC (white blood cell count) at the end of the protocol treatment), and medical history (number of past infections). We observe that, in general, as age increases, it is optimal for the patient to stay longer in the hematology ward (Figure 6(a)). Another difference is observed when looking at the length of protocol (Figure 6(b)): patients with a medium-length protocol (6–8 days), that are the most aggressive treatment protocols, should stay longer in the ward than patients with short- or long-length protocol. The state of the patient at the end of treatment is an important factor too, as observed in Figure 6(c): high-risk patients whose WBC at the end of their treatment is low need a much longer observation period and will need to be discharged later. Finally, patient history also impacts risk and the optimal observation time, as we observe in Figure 6(d). The non-monotonicity around the low number of past infections likely follows from the fact that the first couple of treatment cycles have a higher infection risk than later cycles (see Table 2 in Carmen et al. 2019).

The infection risk functions observed in virtually all of the above examples are of two main types (as is seen in Figure 1(c)): one type exhibits a monotone decreasing infection risk function and the other exhibits an increasing and then decreasing infection risk function (i.e., increasing-decreasing). For each patient type, we observed that the two infection risk functions—for home and hospital—have the same shape and differ only in the level of risk, i.e. the *shape* of the risk functions depends only on patient characteristics. We find that, in general, the increasing-decreasing risk functions result in no observation period unless the patient has a very long history of infection (more than 8 occurrences) and is above 55 years of age, while the monotone decreasing risk functions result in observation periods. We demonstrate the former alternative next.

Let us consider a specific case study, presented in Figure 7. The figure shows the policy change as a function of age for patients with a long history of infections (9), but who finished this particular treatment in good condition (WBC > 1000). In this case the increasing-decreasing infection risk function reaches a peak in day 12 (see Figure 7(a)) and the risk increases with age. Figure 7(b) illustrates an interesting phenomenon where up to age 40 no observation is recommended, while if the patient is above 45 years old observation is recommended until around the peak of the risk function. This drastic change is not apparent when we just look at the incremental effect of age on the risk function, which is not large. Things become more clear when we observe the expected reward function J_τ in Figure 7(c), which changes from a decreasing function for patients below

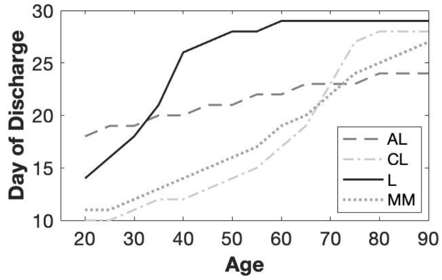

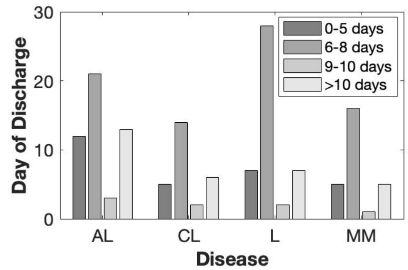

(a) Policy as a function of disease and age

(b) Policy as a function of disease and treatment protocol length

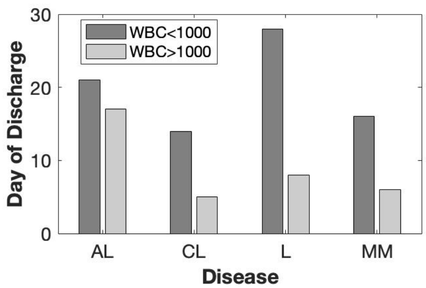

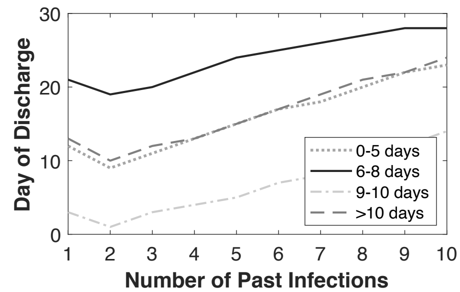

(c) Policy as a function of disease and WBC count at the end of treatment

(d) Policy as a function of treatment protocol length and history of past infections

Figure 6 **Optimal Effective Threshold (Based on Infection and Mortality Risks Given in Carmen et al. 2019)**

45 years old to an increasing-decreasing function for older patients. Thus, while the change in J_τ appears to be continuous in patients' age, the resulting optimal policy is not.

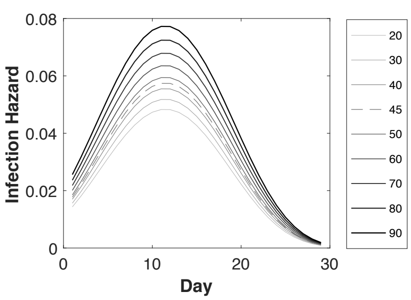

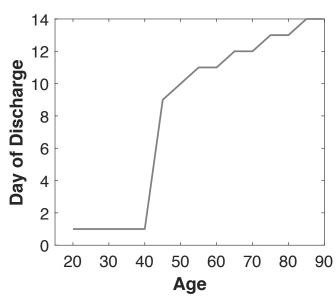

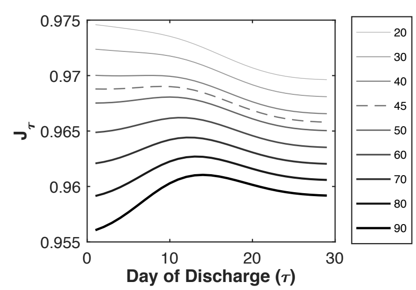

(a) Ward-acquired infection risk function of age

(b) Optimal day of discharge, t_{opt}, as a function of age

(c) J_τ as a function of age

Figure 7 **Policy as a Function of Age (AL, High WBC, 6–8-Days Protocol, 9 Past Infections)**

Figure 8 shows a histogram of the optimal observation time in days for a random sample of 1200 patients treated in the HW. The figure shows that, in the incapacitated case, around 25% of the patients should not stay for observation at all, while the rest would benefit from staying in the hospital for some period. In practice, it may be the case that the hospital does not have space for all of these patients and, therefore, will either cut their LOS short or block new patients from being observed. This capacitated case is the the focus of the next section.

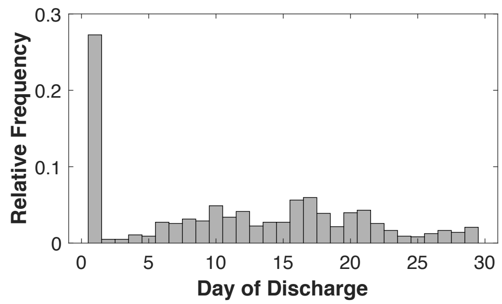

Figure 8 **Histogram of Optimal Observation Days (Random Sample of 1200 Patients Treated in HW)**

5.2. The Capacitated Case

In this section, we demonstrate how capacity constraints change the optimal policy and increase total patient mortality risk. Our main performance measure here is the expected reward, J, which in this particular case is equivalent to the total survival rate because of our assumption of zero hospitalization costs.

We start by examining the shape of different expected reward functions. As demonstrated in Figure 7(c), we observe that the reward function, J, is generally flat around the optimal threshold t_{opt}. J may change drastically as we get farther away from t_{opt}; therefore, strict capacity constraints will result in high patient risk (low total survival rate), while less stringent constraints will not change patients risk dramatically.

We next proceed to examine the policy structures that appear in our data as a function of system load, assuming a homogeneous patient population. According to Corollary 4 we expect the optimal policy to have up to two thresholds and belong to one of five possible types. Lemma 4 suggests that only three of the five policy types may be of relevance, under certain conditions. We wish to empirically test which policy types should be the most prevalent in practice. In this analysis we assume a range of system loads from moderate to high load where ρ equals 1.02, 1.05, 1.1, or 1.2 (i.e., the arrival rate $(\bar{\lambda})$ is $\mu_{t_{opt}} \times \rho$).

We solve the optimization problem for all 19,200 combinations of disease (4), WBC level (2), age(15), number of past infections (10), treatment protocol length (4), and system load (4). For these instances we identify parameter combinations that result in the optimality of all five aforementioned policies depicted in Figure 5. Table 3 presents the frequency of each such policy type. Because in all of those parameter combinations τ_{spd} is not necessarily an integer, a precise $1{\times}$Sp is rarely feasible, due to discretization. We therefore consider any policy in which $\tau_h - \tau_l = 1$ a

"1×Sp-type" policy, where the group that is sped up more is discharged one day ahead of the rest of the patients. The set of 1×Sp-type policies may be further divided into two subcases: one in which the upper threshold $\tau_h = t_{opt}$ and a second case in which the upper threshold $\tau_h < t_{opt}$; the former may also be considered as an Sp-FS policy and the latter as a 2×Sp policy. Those subcases are labeled in Table 3 as columns "1×Sp or Sp-FS" and "1×Sp or 2×Sp", respectively. Note that we have encountered no pure Sp-FS solutions, but we did encounter some that are pure 2×Sp solutions.

Our numerical analysis leads to the following observations: a) Policies of 1×Sp type cover around 90% of our examples. Only 7.2% are pure 2×Sp type, 2.5% are of Bl-Sp type, and 0.1% are Bl-FS. In particular, Bl-Sp and Bl-FS should be hardly used in practice. b) As the load in the system increases, policies of 1×Sp type become more prevalent than pure 2×Sp. However, Bl-Sp captures a higher percentage of extremely high loads (e.g., when $\rho = 1.2$). c) The difference in prevalence of policy types between the different types of cancer is small. The only exception is AL, which has a higher percentage of Bl-Sp and "1×Sp or 2×Sp". The dominance in the optimality of 1×Sp type over all other cases demonstrates that optimality and fairness largely go hand in hand so that one does not need to sacrifice one goal over the other.

Comparing the total reward (which in this case is equivalent to the total survival rate because of our zero hospitalization cost assumption) of all patients under the uncapacitated policy to the capacitated policy, we find that the decrease in the total survival rate due to limited capacity is less than 1.5% (when $\rho = 1.2$). (When $\rho = 1.1$ the maximal decrease is 0.7%, when $\rho = 1.05$ the maximal decrease is 0.25%, and when $\rho = 1.02$ the maximal decrease is 0.05%.)

The high frequency of occurrences where a policy of 1×Sp type is optimal raises the following question: how much would total survival rate decrease if we restricted our attention to the simple and fair 1×Sp policy. We thus compared the difference in expected reward (where the hospitalization cost is 0) when applying the 1×Sp solution instead of the optimal capacitated policy. For all the examples we examined above, the decrease in total survival rate (i.e., the reduction in expected reward) ended up being less than 0.5%. We can therefore suggest that, in practice, the fair policy of 1×Sp may be used exclusively, simplifying the problem significantly without compromising much in terms of optimality.

6. Conclusions

Our paper analyzed LOS optimization of hematology patients that balances between hospital-acquired infection risk (which drives hospitals to discharge high-risk patients as soon as possible) and home-acquired infection risk (which drives hospitals to keep high-risk patients as long as possible, to prevent delays in treating infection). Using MDP formulation we explore the connection

Table 3 **Frequency of solution types as a function of patient characteristics**

Case		Bl-Sp	2×Sp	1×Sp or 2×Sp	1×Sp or Sp-FS	Sp-FS	Bl-FS
All		2.5%	7.2%	45.1%	45.0%	0.0%	0.1%
Age	20	1.8%	4.6%	38.8%	54.9%	0.0%	0.0%
	25	2.2%	4.7%	39.5%	53.2%	0.0%	0.4%
	30	1.6%	4.7%	40.3%	53.0%	0.0%	0.4%
	35	2.0%	5.4%	40.9%	51.4%	0.0%	0.4%
	40	2.4%	5.5%	41.6%	50.1%	0.0%	0.4%
	45	2.5%	6.1%	42.5%	48.9%	0.0%	0.0%
	50	2.3%	6.6%	43.2%	47.9%	0.0%	0.0%
	55	2.3%	7.1%	44.8%	45.8%	0.0%	0.0%
	60	3.0%	7.6%	45.3%	43.5%	0.0%	0.7%
	65	2.5%	7.8%	46.8%	42.8%	0.0%	0.0%
	70	2.4%	8.4%	47.5%	41.7%	0.0%	0.0%
	75	2.7%	9.0%	48.6%	39.7%	0.0%	0.0%
	80	3.0%	9.1%	49.8%	38.2%	0.0%	0.0%
	85	2.7%	9.4%	50.7%	37.2%	0.0%	0.0%
	90	2.9%	10.1%	51.6%	35.3%	0.0%	0.0%
Disease	AL	7.9%	0.0%	62.9%	28.7%	0.0%	0.4%
	CL	0.0%	10.9%	32.5%	56.6%	0.0%	0.0%
	L	0.0%	17.4%	35.2%	47.4%	0.0%	0.0%
	MM	0.0%	2.5%	44.4%	53.1%	0.0%	0.0%
Load	1.02	0.0%	8.7%	16.9%	74.2%	0.0%	0.2%
	1.05	0.2%	11.2%	34.0%	54.4%	0.0%	0.2%
	1.1	0.9%	8.8%	56.6%	33.5%	0.0%	0.2%
	1.2	8.7%	0.2%	73.0%	18.1%	0.0%	0.0%

between the infection risk function and the optimal policy from the single patient perspective. We then consider the social optimization problem in which capacity constraints limit the ability of the hospital to work with the otherwise optimal solution. Our analysis covers a wide range of risk function dynamics that occur in practice and, therefore, it can guide hospitals in terms of observation policy decisions.

One important realistic aspect that our model captures is the patient heterogeneity. In particular, a hematology ward may have several types of patients being treated in parallel. Based on our numerical analysis of the single-patient type, we learned that limiting our attention to a speedup-only policy does not increase patient risk by much. Therefore, in the more complicated multi-patient type case, it may be sufficient to focus on the speedup only policies, instead of the whole range of possible policies. Even if that is not the case, our analysis of the capacitated case reveals that at most one patient type will need more than a single discharge threshold.

Hematology units are inherently small due to the need to keep patients isolated (which is expensive), and hence are typically overloaded. But one may also use our model for other types of medical conditions, where capacity may be less constrained. In such units the number of patients may vary over time and thus dependencies between load and infection risk may also play an important role in the decision between hospitalization and home care. Such dependencies are natural to consider because infection risk may increase with the number of people a patient encounters during her stay. Further work could explore these dependencies.

Another interesting direction to explore as further research is the decision of when to send a patient home when information on their medical condition is dynamically updated over time. We note that models for predicting infection risk of Hematology patients dynamically over time do not exist in the literature yet. Hence, following Carmen et al. (2019), we assume that the infection hazard rate is known as soon as the patient finishes treatment, and is not updated (expect for the elapsed time and whether or not the patient got an infection) while the patient is under observation. For all practical purposes, our current models are the best that can be implemented with current technology. But we believe that new information that is received during observation may change the infection hazard rate of a specific patient and thus the optimal observation time may need to be updated dynamically. Our analysis can be straightforwardly generalized to handle this case in the uncapacitated single-patient context, but more research is needed to determine how to handle this case when capacity constrains are in effect. This will be important to address as new models for predicting patient risks dynamically become available.

We assume in our paper that patients either stay in the hematology ward or are sent home, while in practice a patient who does not have an available bed in the hematology ward may be hospitalized in the general ward rather than be sent home. Yet, the dedicated hematology ward is always better than the general ward for hematology patients (Carmen et al. 2019), both in terms of infection risk and in terms of mortality risk. This suggests that our analysis may be utilized to incorporate the option of general-ward hospitalization versus home care in a sequential fashion.

To summarize, our paper explored the question of where a patient should be observed following cancer treatment, considering that both the hospital care and home care have their pros and cons. Our framework allows one to find a solution that strikes the right balance between the two locations and offers the "best of both worlds." In particular, we show that in many cases, one does not necessarily have to choose between hospital care and home care and, in fact, a combination of the two is optimal. Our framework may be used to study other service systems where there are inherent tradeoffs between professional service and self-service and a fine balance needs to be achieved.

References

Adusumilli KM, Hasenbein JJ (2010) Dynamics admission and service rate control of a queue. *Queueing Systems* 66:131–154.

AHRQ (2014) Medical expenditure panel survey, URL https://meps.ahrq.gov/data_stats/tables_compendia_hh_interactive.jsp, last accessed Feb, 18, 2019.

Allon G, Deo S, Lin W (2013) The impact of hospital size and occupancy of hospital on the extent of ambulance diversion: Theory and evidence. *Operations Research* 61(3):544–562.

Ata B, Shneorson S (2006) Dynamic control of an M/M/1 service system with adjustable arrival and service rates. *Management Science* 52(11):1778–1791.

Bartel AP, Chan CW, Kim SH (2020) Should hospitals keep their patients longer? The role of inpatient care in reducing post-discharge mortality. *Management Science* 66(6):2326–2346.

Bassamboo A, Randhawa RS (2016) Scheduling homogeneous impatient customers. *Management Science* 62(7):2129–2147.

Carmen R, Yom-Tov GB, Nieuwenhuyse IV, Foubert B, Ofran Y (2019) The role of specialized hospital units in infection and mortality risk reduction among patients with hematological cancers. *PLoS ONE* 14(3):e0211694.

Chan C, Green L, Lu L, Lekwijit S, Escobar G (2019) Assessing the impact of service level when customer needs are uncertain: An empirical investigation of hospital step-down units. *Management Science* 65(2):459–954.

Chan CW, Farias VF, Bambos N, Escobar G (2012) Optimizing ICU discharge decisions with patient readmissions. *Operations Research* 60(6):1323–1342.

Cornejo-Juárez P, Vilar-Compte D, García-Horton A, López-Velázquez M, Ñamendys-Silva S, Volkow-Fernández P (2016) Hospital-acquired infections at an oncological intensive care cancer unit: Differences between solid and hematological cancer patients. *BMC Infectious Diseases* 16(1):274.

Cornely O, Gachot B, Akan H, Bassetti M, Uzun O, Kibbler C, et al (2015) Epidemiology and outcome of fungemia in a cancer cohort of the infectious diseases group (IDG) of the European Organization for Research and Treatment of Cancer (EORTC 65031). *Clinical Infection Diseases* 61(3):324–331.

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. *Operations Research* 61(6):1277–1294.

Dong J, Shi P, Zheng F, Jin X (2019a) Off-service placement in inpatient ward network: Resource pooling versus service slowdown, working paper.

Dong J, Yom-Tov E, Yom-Tov GB (2019b) The impact of delay announcements on hospital network coordination and waiting times. *Management Science* 65(5):1969–1994.

Halfdanarson T, Hogan W, Madsen B (2017) Emergencies in hematology and oncology. *Mayo Clinic Proceedings* 92(4):609–641.

Kang W, Ramanan K (2010) Fluid limits of many-server queues with reneging. *The Annals of Applied Probability* 20(6):2204–2260.

Kc D, Terwiesch C (2009) Impact of workload on service time and patient safety: An econometric analysis of hospital operations. *Management Science* 55(9):1486–1498.

Kc D, Terwiesch C (2012) An econometric analysis of patient flows in the cardiac intensive care unit. *Manufacturing & Service Operations Management* 14(1):50–65.

Kuderer N, Dale D, Crawford J, Cosler L, Lyman G (2006) Mortality, morbidity, and cost associated with febrile neutropenia in adult cancer patients. *Cancer* 106(10):2258–2266.

Lee N, Kulkarni V (2014) Optimal arrival rate and service rate control of multi-server queues. *Queueing Systems* 76:37–50.

LLSC (2016) Facts and statistics: The Leukemia & Lymphoma Society of Canada, URL https://www.llscanada.org/disease-information/facts-and-statistics.

Magill SS, OLeary E, Janelle SJ, Thompson DL, Dumyati G, Nadle J, Wilson LE, Kainer MA, Lynfield R, Greissman S, et al. (2018) Changes in prevalence of health care–associated infections in US hospitals. *New England Journal of Medicine* 379(18):1732–1744.

Mills AF, Argon N, Ziya S (2013) Resource-based patient prioritization in mass-casualty incidents. *Manufacturing & Service Operations Management* 15(3):361–377.

Nambiar S, Mayorga M, Capan M (2020) Resource allocation strategies under dynamically changing health conditions, URL https://papers.ssrn.com/sol3/papers.cfm?abstract_id=3525866, working paper.

Ouyang H, Argon NT, Ziya S (2020) Allocation of intensive care unit beds in periods of high demand. *Operations Research* 68(2):591–608.

SEELab (2020) Service enterprise engineering laboratory, Technion, URL https://web.iem.technion.ac.il/en/service-enterprise-engineering-see-lab/general-information.html.

Shi P, Helm JE, Deglise-Hawkinson J, Pan J (2020) Timing it right: Balancing inpatient congestion versus readmission risk at discharge, Forthcoming in Operations Research.

Song H, Tucker AL, Graue R, Moravick S, Yang JJ (2020) Capacity pooling in hospitals: The hidden consequences of off-service placement. *Management Science* 66(9):3825–3842.

Suda KJ, Motl SE, Kuth JC (2006) Inpatient oncology length of stay and hospital costs: implications for rising inpatient expenditures. *Journal of Applied Research* 6(2):126–132.

Taccone F, Artigas A, Sprung C, Moreno R, Sakr Y, Vincent J (2009) Characteristics and outcomes of cancer patients in European ICUs. *Critical Care* 13(1):R15.

Turhan A, Alanyali M, Starobinski D (2012) Optimal admission control in two-class preemptive loss systems. *Operations Research Letters* 40(6):510–515.

Ulukus MY, Güllü R, Örmeci L (2011) Admission and termination control of a two-class loss system. *Stochastic Models* 27(1):2–25.

van Tiel FH, Harbers MM, Kessels AG, Schouten HC (2005) Home care versus hospital care of patients with hematological malignancies and chemotherapy-induced cytopenia. *Annals of Oncology* 16(2):195–205.

Whitt W (2006) Fluid models for multiserver queues with abandonments. *Operations Research* 54(1):37–54.

Whitt W (2013) OM forum: Offered load analysis for staffing. *Manufacturing & Service Operations Management* 15(2):166–169.

Yom-Tov GB, Chan CW (2021) Balancing admission control, speedup, and waiting in service systems. *Queueing systems* 97:163–219.

Zhang J (2013) Fluid models of many-server queues with abandonment. *Queueing Systems* 73:147–193.

Appendix A: Proofs for Section 3

We prove Proposition 1 by breaking it into smaller more specific lemmas each concentrating on the monotonicity with respect to a different parameter.

LEMMA 5 (**Monotonicity in** c). *Under Assumption 1, the effective threshold t_{opt} is monotone decreasing in c. That is, as the hospitalization cost, c, increases the patient will be discharged home earlier.*

PROOF: Let $R_w(s;c)$ be the hospital (ward) reward-to-go, as defined in (4), when the hospitalization cost is c. Since $R_h(s)$ is independent of c, it is sufficient to prove that $R_w(s;c)$ is monotone decreasing in c. To see why this is sufficient, suppose that $R_w(s;c)$ is decreasing in c and define $t_{opt}(c)$ as the effective threshold when the hospitalization cost is equal to c. Suppose that $c_1 < c_2$. By definition, $R_h(t_{opt}(c_1)) \geq R_w(t_{opt}(c_1);c_1)$. By the monotonicity of $R_w(s;c)$ in c, this implies that $R_h(t_{opt}(c_1)) \geq R_w(t_{opt}(c_1);c_2)$. But, by definition, $t_{opt}(c_2) = \min\{1 \leq t \leq T \mid R_h(t) \geq R_w(t;c_2)\}$. Thus, $t_{opt}(c_2) \leq t_{opt}(c_1)$.

We next use backward induction to show that $R_w(s;c)$ is indeed decreasing in c. In fact, we show that $R_w(s;c)$ and $v_s(s;c)$ are *both* decreasing in c, where $v_s(s;c)$ is the optimal reward-to-go function, as defined in (5), when the hospitalization cost is c.

We first establish the monotonicity of $R_w(s;c)$ and $v_s(s;c)$ for the end of the horizon at $s = T-1$. Indeed, $R_w(T-1;c) = p_w r_w(T-1) + v_T(T;c)(1-r_w(T-1)) - c = p_w r_w(T-1) + (1+c_I)(1-r_w(T-1)) - c$. Hence, $R_w(T-1;c)$ is decreasing in c, and $v_{T-1}(T-1;c) = \max\{R_h(T-1), R_w(T-1;c)\}$ is also decreasing in c.

For the backward induction step, suppose that the monotonicity of $R_w(s;c)$ and $v_s(s;c)$ holds for state $s = T-i+1$; then, we need to establish that it also holds for state $s = T-i$. Indeed, for state $s = T-i$, we have that $R_w(T-i;c) = p_w r_w(T-i) + v_{T-i+1}(T-i+1;c)(1-r_w(T-i)) - c$ is a sum of three functions, two that are decreasing in c and one is independent of c, and hence the total is decreasing in c. In addition, $v_{T-i}(T-i;c)$ is decreasing in c, as it is equal to the maximum between two functions, one that is independent of c and the other one that is decreasing in c. $\square$

LEMMA 6 (**Monotonicity in** c_I). *Under Assumption 1, the effective threshold t_{opt} is monotone in c_I. As the cost of hospitalization after infection, c_I, increases, the patient will be sent home earlier so as to reduce the probability of developing an infection.*

PROOF: We first argue that it is sufficient to prove that

$$R_h(s;c_I) - R_w(s;c_I) \text{ is increasing in } c_I \text{ for all } s, \tag{23}$$

where $R_h(s;c_I)$ is defined in (1) and $R_w(s;c_I)$ is defined in (4), both assuming that the cost of contracting an infection is c_I. Indeed, if (23) holds, assume that $c_I^1 < c_I^2$ and let $t_{opt}(c_I)$ be the effective threshold associated with the cost c_I. By the definition of effective threshold, we have that $R_h(t_{opt}(c_I^1);c_I^1) - R_w(t_{opt}(c_I^1);c_I^1) \geq 0$. Together with (23), it follows that $R_h(t_{opt}(c_I^1);c_I^2) - R_w(t_{opt}(c_I^1);c_I^2) \geq 0$. This implies that the first point t such that $R_h(t;c_I^2) - R_w(t;c_I^2) \geq 0$ has to satisfy $t \leq t_{opt}(c_I^1)$.

We now turn to the proof of (23). Specifically, we will show by backward induction on s that for $c_I^1 < c_I^2$,

$$R_h(s;c_I^2) - R_h(s;c_I^1) \geq R_w(s;c_I^2) - R_w(s;c_I^1) \geq 0. \tag{24}$$

In order for the induction step to go through, we need to also establish for each value of s that

$$R_h(s;c_I^2) - R_h(s;c_I^1) \geq v_s(s;c_I^2) - v_s(s;c_I^1) \geq 0, \tag{25}$$

where (25) is an auxiliary statement needed in the proof of (24).

Initiation: For state T both (24) and (25) hold, as $R_h(T;c_I) = R_w(T;c_I) = v_T(T;c_I) = 1 + c_I$, by definition.

Induction: For general s, suppose that (24) and (25) hold for $s+1$. By definition, $R_w(s;c_I) = p_w r_w(s) + v_{s+1}(s+1;c_I)(1-r_w(s)) - c$, and $R_h(s;c_I) = p_h r_h(s) + R_h(s+1;c_I)(1-r_h(s))$. Therefore, $R_w(s;c_I^2) - R_w(s;c_I^1) = (v_{s+1}(s+1;c_I^2) - v_{s+1}(s+1;c_I^1))(1-r_w(s)) \geq 0$, and $R_h(s;c_I^2) - R_h(s;c_I^1) = (R_h(s+1;c_I^2) - R_h(s+1;c_I^1))(1-r_h(s)) \geq 0$. To complete the proof of (24) for state s, using Assumption 1, it is sufficient to establish that $R_h(s+1;c_I^2) - R_h(s+1;c_I^1) \geq v_{s+1}(s+1;c_I^2) - v_{s+1}(s+1;c_I^1)$, which follows by the induction hypothesis on (25).

To show (25) at state s, we need to consider four cases depending on the utility-maximizing action in (3):

1. $v_s(s;c_I) = R_h(s;c_I)$, for all $c_I \in \{c_I^1, c_I^2\}$,
2. $v_s(s;c_I) = R_w(s;c_I)$, for all $c_I \in \{c_I^1, c_I^2\}$,
3. $v_s(s;c_I^1) = R_h(s;c_I^1)$ and $v_s(s;c_I^2) = R_w(s;c_I^2)$, and
4. $v_s(s;c_I^1) = R_w(s;c_I^1)$ and $v_s(s;c_I^2) = R_h(s;c_I^2)$.

Note that in case 1, (25) at state s trivially holds. In case 2, (25) at state s follows directly from (24) at state s. Case 3 cannot be realized since it contradicts (24) (unless there is a tie between $R_h(s;c_I^j)$ and $R_w(s;c_I^j)$ for $j = 1,2$, but then case 3 reduces to cases 1 and 2). Finally, in case 4, the first inequality in (25) reduces to $R_w(s;c_I^1) \geq R_h(s;c_I^1)$, which is true in case 4, since assuming that $v_s(s;c_I^1) = R_w(s;c_I^1)$ means that keeping the patient in the ward is the utility-maximizing action in the case where $c_I = c_I^1$. It remains to show that, in case 4, $R_h(s,c_I^2) - R_w(s,c_I^1) \geq 0$. But note that $R_h(s,c_I^2) - R_w(s,c_I^1) \geq R_w(s,c_I^2) - R_w(s,c_I^1)$ by the assumption on the optimality of sending a patient home when $c_I = c_I^2$, and that $R_w(s,c_I^2) - R_w(s,c_I^1) \geq 0$ by (24). Thus, we have established (25) for a general value of s. □

Lemma 7 (**Monotonicity in p_h**). *Under Assumption 1, the effective threshold is monotone decreasing in p_h. As the survival probability in case of infection at home, p_h, increases, the threshold will be lower; hence, the patient will be sent home earlier.*

Proof: We first argue that it is sufficient to prove that

$$R_h(s;p_h) - R_w(s;p_h) \text{ is increasing in } p_h \text{ for all } s. \tag{26}$$

Indeed, if (26) holds, assume that $p_h^1 < p_h^2$ and let $t_{opt}(p_h)$ be the effective threshold associated with the recovery probability p_h. By the definition of the effective threshold, we have that $R_h(t_{opt}(p_h^1);p_h^1) - R_w(t_{opt}(p_h^1);p_h^1) \geq 0$. By (26), we have that $R_h(t_{opt}(p_h^1);p_h^2) - R_w(t_{opt}(p_h^1);p_h^2) \geq 0$. This implies that the first point t such that $R_h(t;p_h^2) - R_w(t;p_h^2) \geq 0$ has to satisfy $t \leq t_{opt}(p_h^1)$, hence, $t_{opt}(p_h^2) \leq t_{opt}(p_h^1)$.

To prove (26) we will use backward induction. Specifically, we will show by backward induction on s that $R_h(s;p_h^2) - R_w(s;p_h^2) \geq R_h(s;p_h^1) - R_w(s;p_h^1)$. In order for the induction step to go through, we need to establish for each value of s that

$$R_h(s;p_h^2) - R_h(s;p_h^1) \geq R_w(s;p_h^2) - R_w(s;p_h^1) \geq 0, \tag{27}$$

and that

$$R_h(s;p_h^2) - R_h(s;p_h^1) \geq v_s(s;p_h^2) - v_s(s;p_h^1) \geq 0, \tag{28}$$

where (28) is an auxiliary statement needed in the proof of (27).

Initiation: For state T the statement is trivial, as $R_h(T;p_h) = R_w(T;p_h) = v_T(T;p_h) = 1 + c_I$, by definition.

Induction: For general s, suppose that (27) and (28) hold for $s+1$. By definition, $R_w(s;p_h) = p_w r_w(s) + v_{s+1}(s+1;p_h)(1 - r_w(s)) - c$, and $R_h(s;p_h) = p_h r_h(s) + R_h(s+1;p_h)(1 - r_h(s))$. Therefore, $R_w(s;p_h^2) - R_w(s;p_h^1) = (v_{s+1}(s+1;p_h^2) - v_{s+1}(s+1;p_h^1))(1 - r_w(s))$, and $R_h(s;p_h^2) - R_h(s;p_h^1) = (p_h^2 - p_h^1) r_h(s) + (R_h(s+1;p_h^2) - R_h(s+1;p_h^1))(1 - r_h(s))$. To complete the proof of (27) for state s, using Assumption 1, it is sufficient to establish that $R_h(s+1;p_h^2) - R_h(s+1;p_h^1) \geq v_{s+1}(s+1;p_h^2) - v_{s+1}(s+1;p_h^1) \geq 0$, which follows by the induction hypothesis on (28).

To establish (28) for state s, we need to consider four cases depending on the utility-maximizing action in (3). The four cases are represented by the value of $v_s(s;p_h) = \max\{R_h(s;p_h), R_w(s;p_h)\}$:

1. $v_s(s;p_h) = R_h(s;p_h)$, for all $p_h \in \{p_h^1, p_h^2\}$,
2. $v_s(s;p_h) = R_w(s;p_h)$, for all $p_h \in \{p_h^1, p_h^2\}$,
3. $v_s(s;p_h^1) = R_h(s;p_h^1)$ and $v_s(s;p_h^2) = R_w(s;p_h^2)$, and
4. $v_s(s;p_h^1) = R_w(s;p_h^1)$ and $v_s(s;p_h^2) = R_h(s;p_h^2)$.

Note that in case 1, (28) at state s trivially holds. In case 2, (28) at state s follows directly from (27) at state s. Case 3 is in contradiction to (27) (unless there is a tie between $R_h(s;p_h^j)$ and $R_w(s;p_h^j)$ for $j = 1,2$, but then case 3 reduces to cases 1 and 2). Finally, in case 4, (28) reduces to $R_w(s;p_h^1) \geq R_h(s;p_h^1)$ and $R_h(s;p_h^2) - R_w(s;p_h^1) \geq 0$. The former statement is true due to the assumption that keeping the patient in the ward is the utility-maximizing action in this case when $p_h = p_h^1$. For the latter we have that $R_h(s;p_h^2) - R_w(s;p_h^1) \geq R_w(s;p_h^2) - R_w(s;p_h^1) \geq 0$, where the first inequality is due to the optimality of sending a patient home at s when $p_h = p_h^2$ and the second inequality is due to (27). Thus, we have established (28) for all s. $\square$

LEMMA 8 **(Monotonicity in p_w)**. *Under Assumption 1, the effective threshold is monotone increasing in p_w. As the survival probability in case of infection at the hospital, p_w, increases, the threshold will be higher; hence, the patient will be sent home later.*

PROOF: We first argue that it is sufficient to prove that

$$R_w(s;p_w) \text{ is increasing in } p_w \text{ for all } s. \tag{29}$$

To see that (29) is indeed sufficient, notice first that $R_h(s)$ is independent of p_w. Now, assume that (29) holds and define $t_{opt}(p_w)$ as the effective threshold as a function of p_w. Also, let $R_w(t;p_w)$ be the hospital (ward) reward-to-go, as defined in (4), when the recovery probability given that the patient got an infection at the hospital is p_w. Suppose that $p_w^1 < p_w^2$. By definition, $R_h(t_{opt}(p_w^2)) \geq R_w(t_{opt}(p_w^2);p_w^2)$. By the monotonicity of $R_w(s)$ in p_w, this implies that $R_h(t_{opt}(p_w^2)) \geq R_w(t_{opt}(p_w^2);p_w^1)$. But, by definition, $t_{opt}(p_w^1) = \min\{0 < t \leq T \mid R_h(t) \geq R_w(t;p_w^1)\}$. Thus, $t_{opt}(p_w^1) \leq t_{opt}(p_w^2)$.

To prove (29) we will use backward induction. Specifically, we will show by backward induction on s that $R_w(s;p_w^2) \geq R_w(s;p_w^1)$. In order for the induction step to go through, we need to establish for each value of s that

$$R_w(s;p_w^2) \geq R_w(s;p_w^1), \tag{30}$$

and that

$$v_s(s;p_w^2) \geq v_s(s;p_w^1), \tag{31}$$

where (31) is an auxiliary statement needed in the proof of (30).

Initiation: For state T the statement is trivial, as $R_w(T;p_w) = v_T(T;p_w) = 1 + c_I$, by definition.

Induction: Assume that (30) and (31) hold at state $s+1$. We will establish that they are also true for state s. Note that $R_w(s;p_w) = p_w r_w(s) + v_{s+1}(s+1;p_w)\cdot(1-r_w(s+1)) - c$. It follows from the induction hypothesis that $R_w(s;p_w)$ is increasing in p_w. As for the value function, note that $v_s(s;p_w) = \max\{R_h(s), R_w(s;p_w)\}$, and hence is also increasing in p_w. □

Lemma 9 (**Monotonicity in r_h**). *Under Assumption 1, the effective threshold is monotone in r_h. As the risk of infection at home, r_h, increases, the threshold will be higher; hence, a patient will be sent home later.*

Proof: We first argue that it is sufficient to prove that

$$R_h(s;r_h) - R_w(s;r_h) \text{ is decreasing in } r_h \text{ for all } s. \tag{32}$$

To see that (32) is indeed sufficient, assume that (32) holds and define $t_{opt}(r_h)$ as the effective threshold as a function of r_h. Suppose that $r_h^1 > r_h^2$. By definition of the effective threshold, we have that $R_h(t_{opt}(r_h^1);r_h^1) - R_w(t_{opt}(r_h^1);r_h^1) \geq 0$. Due to the conjectured monotonicity, this implies that $R_h(t_{opt}(r_h^1);r_h^2) - R_w(t_{opt}(r_h^1);r_h^2) \geq 0$. This implies that the first point t such that $R_h(t;r_h^2) - R_w(t;r_h^2) \geq 0$ has to satisfy $t \leq t_{opt}(r_h^1)$. Hence, we have that $t_{opt}(r_h^2) \leq t_{opt}(r_h^1)$.

We prove (32) by backward induction. Specifically, we will show by backward induction on s that $R_h(s;r_h^2) - R_w(s;r_h^2) \geq R_h(s;r_h^1) - R_w(s;r_h^1)$. In order for the induction step to go through, we need to establish for each value of s that

$$R_h(s;r_h^2) - R_h(s;r_h^1) \geq R_w(s;r_h^2) - R_w(s;r_h^1) \geq 0, \tag{33}$$

and that

$$R_h(s;r_h^2) - R_h(s;r_h^1) \geq v_s(s;r_h^2) - v_s(s;r_h^1) \geq 0, \tag{34}$$

where (34) is an auxiliary statement needed in the proof of (33).

Initiation: For state T the statement is trivial, as $R_h(T;r_h) = R_w(T;r_h) = v_T(T) = 1 + c_I$ by definition.

Induction: For general s, suppose that (33) and (34) hold for $s+1$. By definition, $R_w(s;r_h) = p_w r_w(s) + v_{s+1}(s+1;r_h)(1-r_w(s)) - c$, and $R_h(s;r_h) = p_h r_h(s) + R_h(s+1;r_h)(1-r_h(s))$. Therefore, $R_w(s;r_h^2) - R_w(s;r_h^1) = (v_{s+1}(s+1;r_h^2) - v_{s+1}(s+1;r_h^1))(1-r_w(s))$, and $R_h(s;r_h^2) - R_h(s;r_h^1) = p_h r_h^2(s) + R_h(s+1;r_h^2)(1 - r_h^2(s)) - p_h r_h^1(s) - R_h(s+1;r_h^1)(1-r_h^1(s)) \geq p_h r_h^1(s) + R_h(s+1;r_h^2)(1-r_h^1(s)) - p_h r_h^1(s) - R_h(s+1;r_h^1)(1 - r_h^1(s)) = (R_h(s+1;r_h^2) - R_h(s+1;r_h^1))(1-r_h^1(s))$, where the inequality follows from the fact that by assumption $r_h^1(s) \geq r_h^2(s)$, and from the fact that, by (1), $p_h \leq R_h(s+1;r_h^2)$. That is, by replacing $r_h^2(s)$ with $r_h^1(s)$ in the expression above, we end up replacing a convex combination of p_h and $R_h(s+1;r_h^2)$ with another convex combination of these two terms that gives a higher weight to the lower term p_h. To complete the proof of (33) for state s, using Assumption 1, it is sufficient to establish that $R_h(s+1;r_h^2) - R_h(s+1;r_h^1) \geq v_{s+1}(s+1;r_h^2) - v_{s+1}(s+1;r_h^1) \geq 0$, which follows by the induction hypothesis on (34).

To show (34) for state s, we need to consider four cases depending on the utility-maximizing action in (3). The four cases depend on the maximizing action in $v_s(s;r_h) = \max\{R_h(s;r_h), R_w(s;r_h)\}$ for r_h^1 and r_h^2:

1. $v_s(s;r_h) = R_h(s;r_h)$, for all $r_h \in \{r_h^1, r_h^2\}$,
2. $v_s(s;r_h) = R_w(s;r_h)$, for all $r_h \in \{r_h^1, r_h^2\}$,
3. $v_s(s;r_h^1) = R_h(s;r_h^1)$ and $v_s(s;r_h^2) = R_w(s;r_h^2)$, and
4. $v_s(s;r_h^1) = R_w(s;r_h^1)$ and $v_s(s;r_h^2) = R_h(s;r_h^2)$.

Note that, in cases 1 (34) at state s trivially holds and in case 2 (34) follows directly from (33). Case 3 is in contradiction in (33) (unless there is a tie between $R_h(s;r_h^j)$ and $R_w(s;r_h^j)$ for $j = 1,2$, but then case 3 reduces to cases 1 and 2). Finally, in case 4, (34) reduces to $R_w(s;r_h^1) \geq R_h(s;r_h^1)$ and $R_h(s;r_h^2) - R_w(s;r_h^1) \geq 0$. The former statement is true due to the assumption about the utility-maximizing action in the case $r_h = r_h^1$. For the latter we have that $R_h(s;r_h^2) - R_w(s;r_h^1) \geq R_w(s;r_h^2) - R_w(s;r_h^1) \geq 0$, where the first inequality is due to the optimality of sending a patient home at s when $r_h = r_h^2$ and the second inequality is due to (33). Thus, we have established (34) for all s. □

Proof of Lemma 1 **(Monotonicity in r_w)**: We need to establish that if $p_w \leq v_s(s)$ for all s, then the effective threshold is monotone decreasing in r_w. Since $R_h(s)$ is independent of r_w, we argue that it is sufficient to prove that, under the Lemma's conditions,

$$R_w(s) \text{ is decreasing in } r_w \text{ for all } s. \tag{35}$$

To see that this is indeed sufficient, suppose that $R_w(s)$ is decreasing in r_w and define $t_{opt}(r_w)$ as the effective threshold as a function of r_w. Also, let $R_w(t;r_w)$ be the hospital (ward) reward-to-go, as defined in (4), when the infection hazard rate function at the hospital is r_w. Suppose that $r_w^1 < r_w^2$. By definition, $R_h(t_{opt}(r_w^1)) \geq R_w(t_{opt}(r_w^1);r_w^1)$. By the monotonicity of $R_w(s)$ in r_w, this implies that $R_h(t_{opt}(r_w^1)) \geq R_w(t_{opt}(r_w^1);r_w^2)$. But, by definition, $t_{opt}(r_w^2) = \min\{0 < t \leq T \mid R_h(t) \geq R_w(t;r_w^2)\}$. Thus, $t_{opt}(r_w^2) \leq t_{opt}(r_w^1)$.

To prove (35) we will use backward induction. Specifically, we will show that for each value of s that

$$R_w(s;r_w^2) \leq R_w(s;r_w^1), \tag{36}$$

and that

$$v_s(s;r_w^2) \leq v_s(s;r_w^1), \tag{37}$$

where (37) is an auxiliary statement needed in the proof of (36).

Initiation: For $s = T$: $R_w(T) = v_T(T) = 1 + c_I$, by definition, hence both (37) and (36) hold.

Induction: Assume that the monotonicity of $R_w(s+1;r_w)$ and $v_{s+1}(s+1;r_w)$ in r_w holds for state $s+1$. We show that in that case it also holds for state s. By definition, $R_w(s;r_w) = p_w r_w(s) + v_{s+1}(s+1;r_w) \cdot (1 - r_w(s)) - c$, which is decreasing in r_w, by assumption, since $R_w(s;r_w)$ is a convex combination of p_w and $v_{s+1}(s+1;r_w)$ (minus c), $p_w \leq v_{s+1}(s+1;r_w)$ by the Lemma's assumption, the coefficient of p_w is increasing, and from the induction hypothesis. In particular, $R_w(s;r_w^1) = p_w r_w^1(s) + v_{s+1}(s+1;r_w^1)(1 - r_w^1(s)) - c \geq p_w r_w^2(s) + v_{s+1}(s+1;r_w^1)(1 - r_w^2(s)) - c \geq p_w r_w^2(s) + v_{s+1}(s+1;r_w^2)(1 - r_w^2(s)) - c = R_w(s;r_w^2)$. Finally, the monotonicity of $v_s(s;r_w)$ in r_w, i.e. (37), follows straightforwardly by (5). □

Proof of Corollary 1: By definition $v_s(s) \geq R_h(s)$. Thus, if $p_w \leq R_h(s)$, the conditions of Lemma 1 are satisfied and therefore the effective threshold is monotone decreasing in r_w. □

PROOF OF COROLLARY 2: The first thing to note is that if (6) holds, then

$$\frac{p_w - p_h}{1 + c_I - p_h} \le \prod_{j=1}^{i} (1 - r_h(T-j)), \quad \forall i = 1, ..., T-1.$$

This is because removing terms from the product on the RHS of (6) will not decrease this product. The corollary then follows from Corollary 1 and Equation (1). □

PROOF OF COROLLARY 3: We will show by backward induction that if $c = 0$, then $p_w \le v_s(s)$ for all s. The proof will then follow from Lemma 1. For $s = T$, $v_T(T) = 1 + c_I \ge 1 > p_w$. Suppose that at time $s+1$ we have that $v_{s+1}(s+1) \ge p_w$. When $c = 0$, $R_w(s)$ is a convex combination of p_w and $v_{s+1}(s+1)$ — a term that is greater than or equal to p_w by the induction hypothesis. Thus, $R_w(s) \ge p_w$. But, by definition, $v_s(s) \ge R_w(s)$. □

PROOF OF PROPOSITION 2: Assume that $r_w(\cdot) \le r_h(\cdot)$, $p_w \ge p_h$, and $c = 0$. Let τ_1 and τ_2 be two thresholds such that $\tau_1 < \tau_2$. To prove Proposition 2 we need to show that $J_{\tau_1} \le J_{\tau_2}$. By (8), this is the same as showing that $R_{\tau_1}(1) \le R_{\tau_2}(1)$. In fact, we will prove a stronger claim in which $R_{\tau_1}(t) \le R_{\tau_2}(t)$ for all $t \in \{1, ..., T\}$ (where recall that $R_\tau(t)$ is the reward-to-go from t to T under a threshold τ, see (7)). The proof is based on backward induction.

For all $t \ge \tau_2$ we have that $R_{\tau_1}(t) = R_{\tau_2}(t)$ since the patient is at home from time τ_2 to T under both policies. We next consider times t such that $\tau_1 \le t < \tau_2$. Assume that the inequality holds for time $t+1$, i.e., $R_{\tau_2}(t+1) - R_{\tau_1}(t+1) \ge 0$, and show that it also holds for t.

$$\begin{aligned}
&R_{\tau_2}(t) - R_{\tau_1}(t) \\
&= r_w(t)p_w + (1 - r_w(t))\, R_{\tau_2}(t+1) - (r_h(t)p_h + (1 - r_h(t))R_{\tau_1}(t+1)) \\
&\ge r_w(t)p_w - r_h(t)p_h + R_{\tau_1}(t+1)(r_h(t) - r_w(t)) \\
&\ge r_w(t)p_h - r_h(t)p_h + R_{\tau_1}(t+1)(r_h(t) - r_w(t)) \\
&= (r_h(t) - r_w(t))(R_{\tau_1}(t+1) - p_h) \ge 0.
\end{aligned}$$

The first equality follows from plugging in Equations (1) and (7), and the assumption that $c = 0$. The third-line inequality follows from the inductive assumption. The fourth-line inequality is due to the proposition's assumptions regarding p. The last inequality is due to the proposition's assumptions regarding r and since $R_\tau(t) \ge \min\{p_w, p_h\} = p_h$, which can be shown by backward induction using Equations (1) and (7).

For $t < \tau_1$ the expressions are the same and so the inequality remains true.

Since J_τ is monotone increasing in τ, it gets its maximal value at $\tau = T$, which means it is optimal to hospitalize the patient until the end of the horizon. □

PROOF OF PROPOSITION 3: Assume that $r_w(\cdot) > r_h(\cdot)$ and $p_w = p_h$. Let τ_1 and τ_2 be two thresholds such that $\tau_1 < \tau_2$. To prove Proposition 3 we need to show that $J_{\tau_1} \ge J_{\tau_2}$. By definition, this is the same as showing that $R_{\tau_1}(1) \ge R_{\tau_2}(1)$. We will prove the stronger claim that $R_{\tau_1}(t) \ge R_{\tau_2}(t)$ for all $t \in \{1, ..., T\}$. The proof is based on backward induction.

For all $t \ge \tau_2$ we have that $R_{\tau_1}(t) = R_{\tau_2}(t)$ since the patient is at home from time τ_2 to T under both policies. We next consider times t such that $\tau_1 \le t < \tau_2$. We start from $t = \tau_2 - 1$ and move backward (toward

τ_1), by induction. Assume that the inequality holds for time $t+1$, i.e., $R_{\tau_2}(t+1) - R_{\tau_1}(t+1) \leq 0$, and show that it also holds for t.

$$
\begin{aligned}
&R_{\tau_2}(t) - R_{\tau_1}(t) \\
&= r_w(t)p_w + (1 - r_w(t))R_{\tau_2}(t+1) - c - (r_h(t)p_h + (1 - r_h(t))R_{\tau_1}(t+1)) \\
&\leq r_w(t)p_w - r_h(t)p_h + R_{\tau_1}(t+1)(r_h(t) - r_w(t)) \\
&= r_w(t)p_h - r_h(t)p_h + R_{\tau_1}(t+1)(r_h(t) - r_w(t)) \\
&= (r_h(t) - r_w(t))(R_{\tau_1}(t+1) - p_h) \leq 0.
\end{aligned}
$$

The first equality follows from plugging in Equations (1) and (7). The third line inequality is due to the induction hypothesis and the fact that $c \geq 0$. The fourth line equality is due to the proposition's assumption regarding p. The last inequality follows from the fact that $R_\tau(t) \geq \min\{p_w, p_h\} = p_h$, which can again be shown by backward induction on (1) and (7).

Once $t < \tau_1$ the expressions are the same and so the inequality remains true.

Since J_τ is monotone decreasing in τ it obtains its maximal value at $\tau = 1$, which means all patients are discharged immediately; i.e., the no-observation policy is optimal. □

Proof of Proposition 4: Assume that $0 < r_h < r_w < 1$, that both are constant over time, and that Assumption 1 holds. Recall the definition of $f(\tau)$ in (9). Specifically,

$$f(\tau) := R_\tau(\tau - 1) - R_{\tau-1}(\tau - 1) = r_w(p_w - p_h) - (r_w - r_h)(1 - r_h)^{T-\tau}(1 + c_I - p_h) - c, \quad 1 < \tau \leq T.$$

Let $\hat{\tau} := \max\{2 \leq \tau \leq T \mid f(\tau) > 0\}$, and $\hat{\tau} := 1$ if $f(\tau) \leq 0$ for all $\tau \geq 2$. Then, we need to prove that J_τ is (strictly) increasing in τ for all values up to $\hat{\tau}$ and is decreasing afterwards.

We first note that $f(\tau)$ is decreasing in τ. In particular, it crosses 0 at most once. Therefore, $f(\tau) > 0$ for all $\tau \leq \hat{\tau}$ and $f(\tau) \leq 0$ for all $\tau > \hat{\tau}$. Thus, to show that J_τ obtains its maximal value at $\hat{\tau}$ it is sufficient to show that (i) J_τ is locally strictly increasing in τ (namely, $J_\tau > J_{\tau-1}$), for any τ such that $f(\tau) > 0$, and that (ii) that $J(\tau)$ is locally (weakly) decreasing in τ (namely, $J_\tau \leq J_{\tau-1}$), for any τ such that $f(\tau) \leq 0$.

(i) Suppose that $f(\tau) > 0$. We will show that this implies that

$$R_\tau(t) > R_{\tau-1}(t) \quad \text{for all } 1 \leq t \leq \tau - 1. \tag{38}$$

This, in particular, will imply that $R_\tau(1) > R_{\tau-1}(1)$, which, by definition, is equivalent to showing that J_τ is locally strictly increasing in τ. We prove (38) by backward induction on t. For $t = \tau - 1$, (38) directly follows from the definition of $f(\tau)$ as $f(\tau) = R_\tau(\tau - 1) - R_{\tau-1}(\tau - 1)$ and the assumption on f being strictly positive at τ. Let $t < \tau - 1$ and suppose that $R_\tau(t) - R_{\tau-1}(t) > 0$. We want to show that this implies that $R_\tau(t-1) - R_{\tau-1}(t-1) > 0$. Indeed, by (7),

$$R_\tau(t-1) - R_{\tau-1}(t-1) = r_w p_w + (1 - r_w)R_\tau(t) - (r_w p_w + (1 - r_w)R_{\tau-1}(t)) > 0,$$

where the last inequality follows from the inductive assumption.

(ii) The proof of the sufficient condition for J_τ to be locally decreasing in τ is analogous to locally increasing proof. Details are omitted. □

Appendix B: Proofs for Section 4

Proof of Lemma 2: Suppose that the system is overloaded under the threshold t_{opt} and consider an optimal solution $(K, \vec{\lambda}, \vec{\tau})$ to (13). Suppose that $\sum_{i=1}^{K+1} \frac{\bar{\lambda}_i}{\mu_{\tau_i}} < 1$. Then, by the overloaded system assumption, there is at least one class k_0 such that $\tau_{k_0} < t_{opt}$. Recall that $\tau_{K+1} := t_{opt}$. Clearly, $\mu_{k_0} > \mu_{K+1}$, and, by the optimality of t_{opt}, we have that $J_{t_{opt}} \geq J_{\tau_{k_0}}$. We now construct a modified solution $(K, \tilde{\lambda}, \tilde{\tau})$ as follows:

- $\tilde{\tau}_k = \tau_k$, for all k.
- $\tilde{\lambda}_k = \bar{\lambda}_k$, for all $k \neq k_0$ and $k \neq K+1$.
- For $0 < \epsilon \leq \bar{\lambda}_{k_0}$, set $\tilde{\lambda}_{k_0} = \bar{\lambda}_{k_0} - \epsilon$ and $\tilde{\lambda}_{K+1} = \bar{\lambda}_{K+1} + \epsilon$.

Then as long as ϵ is small enough we have that for the new solution the two constraints of (13) are still satisfied and the objective function is not smaller than in the original solution. Repeat this process until $\sum_{i=1}^{K+1} \frac{\bar{\lambda}_i}{\mu_{\tau_i}} = 1$. □

Proof of Lemma 3: The proof follows a straightforward approach relying on first-order necessary conditions for optimality. First note that the Lagrangian corresponding to the optimal solution to problem (16) (which is equivalent to (18)) is

$$L(\tau_l, \tau_h, \bar{\lambda}_l, \alpha) = \bar{\lambda}_l J_{\tau_l} + (\bar{\lambda} - \bar{\lambda}_l) J_{\tau_h} + \alpha \left(\bar{\lambda}\mu_{\tau_l} + \bar{\lambda}_l \mu_{\tau_h} - \bar{\lambda}_l \mu_{\tau_l} - \mu_{\tau_h}\mu_{\tau_l} \right), \tag{39}$$

where $0 \leq \tau_l \leq \tau_h \leq t_{opt}$, $\bar{\lambda}_l \in [0, \bar{\lambda}]$ and $\alpha > 0$. Moreover, in this Lemma we focus on solutions of the form $\tau_l < \tau_h$. We obtain necessary conditions for the optimality of the internal solutions by taking the derivative of (39) with respect to each variable $(\tau_l, \tau_h, \bar{\lambda}_l, \alpha)$.

$$\frac{\partial L}{\partial \tau_l} = \bar{\lambda}_l J'_{\tau_l} + \alpha \left(\bar{\lambda}\mu'_{\tau_l} - \bar{\lambda}_l \mu'_{\tau_l} - \mu_{\tau_h}\mu'_{\tau_l} \right) = 0, \tag{40}$$

$$\frac{\partial L}{\partial \tau_h} = (\bar{\lambda} - \bar{\lambda}_l) J'_{\tau_h} + \alpha \left(\bar{\lambda}_l \mu'_{\tau_h} - \mu'_{\tau_h}\mu_{\tau_l} \right) = 0, \tag{41}$$

$$\frac{\partial L}{\partial \bar{\lambda}_l} = J_{\tau_l} - J_{\tau_h} + \alpha \left(\mu_{\tau_h} - \mu_{\tau_l} \right) = 0, \tag{42}$$

$$\frac{\partial L}{\partial \alpha} = \bar{\lambda}\mu_{\tau_l} + \bar{\lambda}_l \mu_{\tau_h} - \bar{\lambda}_l \mu_{\tau_l} - \mu_{\tau_h}\mu_{\tau_l} = 0. \tag{43}$$

By (40), we get

$$\alpha = \frac{\bar{\lambda}_l J'_{\tau_l}}{(\bar{\lambda}_l + \mu_{\tau_h} - \bar{\lambda})\mu'_{\tau_l}}.$$

Further, by (41), we get

$$\alpha = \frac{(\bar{\lambda}_l - \bar{\lambda}) J'_{\tau_h}}{(\bar{\lambda}_l - \mu_{\tau_l})\mu'_{\tau_h}}.$$

Finally, by (42), we get

$$\alpha = \frac{J_{\tau_h} - J_{\tau_l}}{\mu_{\tau_h} - \mu_{\tau_l}}.$$

Thus, item (a) in the lemma follows by replacing $\bar{\lambda}_l$ by $\frac{\bar{\lambda} - \mu_{\tau_h}}{1 - \mu_{\tau_h}/\mu_{\tau_l}}$ (by (43)), in two of the above terms and comparing them.

For item (b), notice that $\tau_l = 0$ and thus $\mu_0 = \infty$. Hence, $\bar{\lambda}_l = \bar{\lambda} - \mu_{\tau_h}$, and so the Lagrangian becomes

$$L(\tau_h, \bar{\lambda}_l, \alpha) = \bar{\lambda}_l J_0 + (\bar{\lambda} - \bar{\lambda}_l) J_{\tau_h} + \alpha \left(\bar{\lambda} - \bar{\lambda}_l - \mu_{\tau_h} \right). \tag{44}$$

We obtain the necessary conditions for optimality by taking the derivative of (44) with respect to each variable $(\tau_h, \bar{\lambda}_l, \alpha)$ as follows:

$$\begin{aligned}
\frac{\partial L}{\partial \tau_h} &= (\bar{\lambda}-\bar{\lambda}_l)J'_{\tau_h} - \alpha\mu'_{\tau_h} = 0,\\
\frac{\partial L}{\partial \bar{\lambda}_l} &= J_0 - J_{\tau_h} - \alpha = 0,\\
\frac{\partial L}{\partial \alpha} &= \bar{\lambda}-\bar{\lambda}_l - \mu_{\tau_h} = 0.
\end{aligned}$$

The statement of this item follows straightforwardly by extracting α from the first two derivatives and comparing the two, while assigning $\bar{\lambda}-\bar{\lambda}_l = \mu_{\tau_h}$ from the third equation.

Analogously, we can prove item (c). Note that $\tau_h = t_{opt}$ is an assumption here. Thus, the Lagrangian becomes

$$L(\tau_l, \bar{\lambda}_l, \alpha) = \bar{\lambda}_l J_{\tau_l} + (\bar{\lambda}-\bar{\lambda}_l)J_{t_{opt}} + \alpha\left(\bar{\lambda}\mu_{\tau_l} + \bar{\lambda}_l\mu_{t_{opt}} - \bar{\lambda}_l\mu_{\tau_l} - \mu_{t_{opt}}\mu_{\tau_l}\right). \tag{45}$$

We obtain the necessary conditions for optimality by taking the derivative of (45) with respect to each variable $(\tau_l, \bar{\lambda}_l, \alpha)$ as follows:

$$\begin{aligned}
\frac{\partial L}{\partial \tau_l} &= \bar{\lambda}_l J'_{\tau_l} + \alpha\left(\bar{\lambda}\mu'_{\tau_l} - \bar{\lambda}_l\mu'_{\tau_l} - \mu_{t_{opt}}\mu'_{\tau_l}\right) = 0,\\
\frac{\partial L}{\partial \bar{\lambda}_l} &= J_{\tau_l} - J_{t_{opt}} + \alpha\left(\mu_{t_{opt}} - \mu_{\tau_l}\right) = 0,\\
\frac{\partial L}{\partial \alpha} &= \bar{\lambda}\mu_{\tau_l} + \bar{\lambda}_l\mu_{t_{opt}} - \bar{\lambda}_l\mu_{\tau_l} - \mu_{t_{opt}}\mu_{\tau_l} = 0,
\end{aligned}$$

from which the statement of this item is easily extracted. □

Proof of Lemma 4:

(a) Suppose that $\xi(\tau)$ is monotone decreasing; then, we show that the necessary conditions for the optimality of 2×Sp and Sp-FS—i.e., Eqs. (19) and (21)—cannot hold. To see this, note first that because μ_τ is decreasing in τ,

$$-\frac{\mu_{\tau_l}}{\mu_{\tau_h}}\frac{J'_{\tau_l}}{\mu'_{\tau_l}} = \frac{\mu_{\tau_l}}{\mu_{\tau_h}}\xi(\tau_l) > \frac{\mu_{\tau_l}}{\mu_{\tau_h}}\xi(\tau_h) > \frac{\mu_{\tau_h}}{\mu_{\tau_l}}\xi(\tau_h) = -\frac{\mu_{\tau_h}}{\mu_{\tau_l}}\frac{J'_{\tau_h}}{\mu'_{\tau_h}},$$

which is in contradiction to (19). Now consider (21).

If $\xi(\tau)$ is monotonic decreasing then

$$\begin{aligned}
&\frac{(J_{t_{opt}} - J_{\tau_l})\mu_{t_{opt}}}{(\mu_{t_{opt}} - \mu_{\tau_l})\mu_{\tau_l}} > \frac{(J_{t_{opt}} - J_{\tau_l})\mu_{\tau_l}}{(\mu_{t_{opt}} - \mu_{\tau_l})\mu_{\tau_l}} = \frac{J_{t_{opt}} - J_{\tau_l}}{\mu_{t_{opt}} - \mu_{\tau_l}} = \frac{\int_{\tau_l}^{t_{opt}} J'_\tau d\tau}{\int_{\tau_l}^{t_{opt}} \mu'_\tau d\tau} = \frac{\int_{\tau_l}^{t_{opt}} J'_\tau \frac{\mu'_\tau}{\mu'_\tau} d\tau}{\int_{\tau_l}^{t_{opt}} \mu'_\tau d\tau}\\
&= \frac{\int_{\tau_l}^{t_{opt}} \xi(\tau)(-\mu'_\tau)d\tau}{\int_{\tau_l}^{t_{opt}} \mu'_\tau d\tau} > \frac{\int_{\tau_l}^{t_{opt}} \xi(\tau_l)(-\mu'_\tau)d\tau}{\int_{\tau_l}^{t_{opt}} \mu'_\tau d\tau} = -\xi(\tau_l) = \frac{J'_{\tau_l}}{\mu'_{\tau_l}}.
\end{aligned}$$

This is again because J_τ is increasing in τ and ξ_τ and μ_τ are both decreasing in τ. The resulting inequality is in contradiction to (21).

(b) We start by observing that by Corollary 4 an optimal solution to (18) exists and it belongs to one of the five cases outlined in the corollary. In the case of the current lemma, since the policies 2×Sp and Sp-FS have been eliminated in part (a), the only viable options are policies of type Bl-Sp, 1×Sp, and Bl-FS. What these cases have in common is that in all of them there is a speedup threshold τ, such that $\tau^* \le \tau \le t_{opt}$ and any patient who is not assigned a bed with the intention to be released at time τ is blocked. Thus, the key to identifying the optimal solution is to find the optimal speedup threshold τ.

Suppose that $\frac{J_\tau - J_0}{\mu_\tau} < \xi(\tau)$ for all $\tau^* \le \tau \le t_{opt}$ and consider τ such that $\tau^* < \tau < t_{opt}$. Then, we have that $\frac{J_\tau - J_0}{\mu_\tau} < \xi(\tau)$ if and only if

$$\mu'_\tau(J_\tau - J_0) + J'_\tau \mu_\tau > 0. \tag{46}$$

Now note that for policies of type Bl-Sp, 1×Sp, and Bl-FS, the objective function of (18) at a speedup threshold of τ is equal to

$$V_{cap}(\tau) := \lambda_l(\tau) J_0 + (\bar{\lambda} - \lambda_l(\tau)) J_\tau,$$

where $\lambda_l(\tau) = \bar{\lambda} - \mu_\tau$ and, in particular, $\lambda_l = 0$ in 1×Sp (because then $\bar{\lambda} = \mu_\tau$). Then, the derivative of $V_{cap}(\tau)$ is

$$V'_{cap}(\tau) = \lambda'_l(\tau) J_0 + \bar{\lambda} J'_\tau - \lambda'_l(\tau) J_\tau - \lambda_l(\tau) J'_\tau = (\bar{\lambda} - \lambda_l(\tau)) J'_\tau + \lambda'_l(\tau)(J_0 - J_\tau) =$$

$$\mu_\tau J'_\tau + \lambda'_l(\tau)(J_0 - J_\tau) = \mu_\tau J'_\tau + \mu'_\tau (J_\tau - J_0),$$

where the latter is positive, by (46). Therefore, the τ that maximize $V_{cap}(\tau)$ is the largest threshold possible in the range $[\tau^*, t_{opt}]$, i.e., t_{opt}, which implies that the Bl-FS policy is optimal here.

(c) In this case, first note that if $\frac{J_{\tau^*} - J_0}{\mu_{\tau^*}} > \xi(\tau^*)$ then $\frac{J_\tau - J_0}{\mu_\tau} > \xi(\tau)$ for all $\tau^* \le \tau \le t_{opt}$. Because

$$\frac{J_\tau - J_0}{\mu_\tau} > \frac{J_{\tau^*} - J_0}{\mu_{\tau^*}} > \xi(\tau^*) > \xi(\tau).$$

The rest of the proof is analogous to the proof of (b), expect for the reverse inequality of Eq. (46) which implies that $V_{cap}(\tau)$ is decreasing in τ.

(d) The proof of this item is straightforward given the proofs of the above two items and hence is omitted.

□