Inventory-responsive donor-management policy: A tandem queueing network model

Taozeng Zhu,[a] Nicholas Teck Boon Yeo,[b] Sarah Yini Gao,[c] Gar Goei Loke[d,*]

a. Institute of Supply Chain Analytics, Dongbei University of Finance and Economics, Dalian 116025, China;
b. International Business School Suzhou, Xi'an Jiaotong-Liverpool University, Suzhou 215000, China;
c. Lee Kong Chian School of Business, Singapore Management University, Singapore 178899;
d. Rotterdam School of Management, Erasmus University Rotterdam, 3062PA Rotterdam, Netherlands

Published in Manufacturing & Service Operations Management. 25, (4), 1585-1602. DOI: 0.1287/msom.2023.1228

Abstract: Problem definition: In the blood-donor-management problem, the blood bank incentivizes donors to donate, given blood inventory levels. We propose a model to optimize such incentivization schemes under the context of random demand, blood perishability, observation period between donations, and variability in donor arrivals and dropouts. Methodology/results: We propose an optimization model that simultaneously accounts for the dynamics in the blood inventory and the donor's donation process, as a coupled queueing network. We adopt the Pipeline Queue paradigm, which leads us to a tractable convex reformulation. The coupled setting requires new methodologies to be developed upon the existing Pipeline Queue framework. Numerical results demonstrate the advantages of the optimal policy by comparing it with the commonly adopted and studied threshold policy. Our optimal policy can effectively reduce both shortages and wastage. Managerial implications: Our model is the first to operationalize a dynamic donor-incentivization scheme, by determining the optimal number of donors of different donation responsiveness to receive each type of incentive. It can serve as a decision-support tool that incorporates practical features of blood supply-chain management not addressed thus far, to the best of our knowledge. Simulations on existing policies indicate the dangers of myopic approaches and justify the need for smoother and forward-looking donor-incentivization schedules that can hedge against future demand variation. Our model also has potential wider applications in supply chains with perishable inventory.

Keywords: Pipeline Queues, blood collection, donor management, humanitarian operations, perishable inventory management

1. Introduction

Blood transfusion is an essential aspect of many medical treatments. In the United States alone, approximately 21 million blood components are transfused every year, saving over 4 million lives annually in the process (Satyavarapu and Wagle 2020). The global demand for blood is also on an upward trajectory, fueled primarily by increased complexities in medical procedures and aging populations. Consequently, blood donation and its management remain an integral part of the medical infrastructure. In addition to traditional difficulties in the management of blood, like its short shelf life and variability in its demand and supply, recent trends also exacerbate the challenges, such as shrinking eligible donor populations observed in many countries (Greinacher et al. 2010, Müller-Steinhardt et al. 2017). These pressures on both the demand and supply side call for more effective blood-supply management to meet growing demand, while minimizing wastage.

In managing the supply of blood, the decision maker possesses two main control levers: recruiting new donors and ensuring the regularity of donation amongst existing donors (i.e., donor retention) (Mugion et al. 2021). Our study is motivated by Singapore's blood-donation landscape, where there are significant challenges to the recruitment of new donors, rendering the first lever less effective (Health Sciences Authority 2020). This is also observed in other developed states (World Health Organization 2017). Consequently, the central blood bank will need to focus on the management of the existing donor pool. This is often conducted through donor incentivization, with the goal of increasing the regularity of donations and preventing dropouts.

Funding: This study was funded by the Singapore Management University through a research [Grant 20-C207-SMU-015] from the Ministry of Education Academic Research Fund Tier 1.

In general, incentivization can be implemented with or without monetary rewards. Despite ample discussion in both literature and practice on the adoption of economic rewards to motivate blood donations (e.g., Lacetera et al. 2013, 2014, and Sun et al. 2019), many nations still adopt the World Health Organization's (WHO) stance against monetary incentives (World Health Organization 2020). This is based on the fact that economic rewards may spur donations for the wrong reasons and, in some cases, the withholding of important health information that can compromise blood-inventory safety. Singapore adopts a similar stance. Regardless of whether the economic reward is disbursed, the management of blood supply continues to remain less controllable than other goods, and supply-demand matching depends on an effective incentivization scheme.

Designing an effective incentivization strategy is not a unique challenge in blood supply-chain management. It has been well studied in other domains, such as customer acquisition and retention (e.g., Afeche et al. 2017), queue management (e.g., Lingenbrink and Iyer 2019), service management (e.g., Choudhary et al. 2022), etc. Various forms of incentivization schemes are also adopted in practice. The goal of such incentivization schemes is to alter the behavior of the targeted individuals to take desired actions. In the context of donor management, formulating an effective incentivization scheme can be difficult. In practice, many blood banks today adopt threshold policies to manage their blood inventory (Lee et al. 2008, Shi et al. 2014, Sun et al. 2016); in other words, they start reaching out to their donors whenever the existing inventory falls below a certain point. At that point, the blood bank reaches out to as many donors as they can within limitations of manpower of volunteers and time. For example, the Singapore Red Cross will call up donors and encourage them to donate when there are shortages. Other indirect measures are also implemented, such as publishing real-time blood-inventory levels of different blood groups on their website, so as to incentivize donors of specific blood groups to donate. This also has the added effect of discouraging donations when supply is ample. Other forms of blood-donor-incentivization schemes have also been studied and discussed in the literature. For example, Aravindakshan et al. (2015) studied marketing tools to encourage blood donations. More recently, in Heger et al. (2020), the authors proposed a mechanism called a "Registry," where donors can volunteer to join and are called up when there is a shortage. The empirical study found that donors who chose to be placed on the Registry have a significantly different probability of donation, pre-incentivization and postincentivization.

We note that the overall effect of all the above-mentioned incentivization mechanisms, implemented in practice or studied in the literature, is to alter the donors' donation behavior via changing the *probability of donation*. Despite all the findings that illustrate how those proposed mechanisms achieve blood supply-demand matching outcomes, the question remains whether they are optimal (i.e., minimize blood shortage and wastage in the long run). To this end, we consider a general incentivization mechanism that changes the *probability of donation* and aim to provide a decision-support tool to optimize such a mechanism, so as to minimize blood shortage and wastage.

To design a decision-support tool that is effective in the practical context, the framework needs to address unique challenges and capture realistic features of the blood supply chain. Firstly, the demand and supply of blood can vary over time. For example, in Singapore, donation rates are observed to fall during the Lunar New Year period due to cultural beliefs. This leads to periodic shortages of supply during festivities and cultural celebrations. Planning for such variations is crucial, as seen in the scale-up of incentivization weeks before such periods (Shi et al. 2014). Secondly, there are critical hard time-based constraints that must be observed. Specifically, donors must observe a fixed duration of time between consecutive donations for health and safety reasons. This is termed the "observation window" and is commonly three months in most countries (Health Sciences Authority 2019, National Health Service 2020). However, blood units have an expiration window of 42 days, which is significantly shorter than the three-month observation period. The longer timescale in the supply than the demand indicates intrinsic bottlenecks in mustering enough supply to meet demand at short notice. As such, forward planning becomes critical. The combination of this and the previous challenge of time-variability dictates the need for a multiperiod model that is both forward-looking and reacts dynamically to the existing and expected inventory levels. Consequently, in our paper, we shall specifically study the donor-management problem in the transient, multiperiod setting that incorporates blood expiry windows and fixed observation windows between donations. Lastly, although the ultimate goal of incentivization optimization is to change the supply and inventory to the desired levels, it can be challenging to explicitly characterize the impact of incentivization on the eventual blood supply and inventory. This is because incentivization alters donors' donation patterns and their *probability of donation*. This uncertainty in how donors react to the incentivization, in turn, influences future incentivization decisions. The complex *endogenous* interaction between the decision to incentivize and the uncertainty in the donations needs to be carefully modeled to arrive at an accurate description of the blood inventory.

1.1. Key Approaches in the Literature

Blood supply-chain management has been studied widely and extensively (Karaesmen et al. 2011, Beliën

and Forcé 2012, Gunpinar and Centeno 2015). Here, we review the most related literature on both donor management and blood-inventory management, which will be relevant to our work, as we attempt to model both aspects. Broadly, we shall discuss the literature based on the adopted methodologies in three streams—namely, dynamic programming approach, queueing approach, and mechanism design—where we highlight the most related framework, Bayesian persuasion in queues (Lingenbrink and Iyer 2019).

1.1.1. Dynamic Programming Approach. Dynamic programming is frequently seen in the literature, as it can be applied to model the perishable inventory process. Nahmias and Pierskalla (1973) studied the setting of a perishable product with a two-period product life and characterized optimal ordering policy facing random demand. This was extended to m periods (Nahmias 1976). Prastacos (1981) developed an optimal policy to allocate perishable inventory to demand from n locations, minimizing both expected shortages and outdates under random supply and demand. Chen et al. (2019) considered a blood center facing two different demand streams with different freshness requirements and characterized the structure of the optimal blood-collection quantities and inventory policies. Ayer et al. (2019) studied blood-collection operations and obtained near-optimal blood-collection schedules through a two-interval Markov decision process (MDP) formulation.

In our setting of donor incentivization, it would be difficult to implement a dynamic programming approach. The presence of finite time windows necessitates the tracking of the age of blood and donor observation times in the state space, drastically increasing its complexity. Together with the complex dynamics of donor eligibility, it would be difficult to avoid tractability challenges.

1.1.2. Queueing Approach. The other common approach in the literature is to model the dynamics as a queueing system. Graves (1982) studied a single queue and server inventory system with constant replenishment and exponential demand stream. The authors postulate that the lifetime on the oldest unit in stock is equivalent to a virtual waiting time of an M/M/1 queue under a first-in-first-out (FIFO) policy. Goh et al. (1993) considered a model with two classes of demands, where each of the classes represents separate demand streams. Sarhangian et al. (2018) extended the literature by incorporating inventory freshness within the system and evaluating a threshold-based blood-allocation policy that considered the age of the blood units.

These studies mainly focus on characterizing and evaluating the queueing process. Hence, they are not readily amenable to the question of policy design. The assumption of the steady state, although useful in the analysis of the queueing system, does not gel with the fundamentally transient and state-dependent nature of the problem we intend to consider.

1.1.3. Mechanism-Design and Bayesian Persuasion Approach. The mechanism-design approach is one of the most classical methodologies widely used to model incentivization problems. Many other frameworks have been developed based on mechanism-design models to incorporate unique features of the problem of interest. One of the most relevant frameworks is Bayesian persuasion, proposed by Kamenica and Gentzkow (2011), which models a single informed principal choosing information to convey to an uninformed agent to motivate (persuade) him to take a certain action that favors the principal. This framework has been applied to studying different problems, such as price discrimination (Bergemann et al. 2015), medical testing or treatment (Schweizer and Szech 2018), and auctions (Bergemann et al. 2017). Although Bayesian persuasion might be promising to model donor incentivization in our study, to capture the inventory-dependent donor-incentivization feature, the Bayesian persuasion model must be embedded in the stochastic processes of blood donors and inventory. Along this line, the most related work, Lingenbrink and Iyer (2019), considered revealing information to participants within an unobservable queueing system to encourage or discourage participation. This proposed model integrates Bayesian persuasion with stochastic processes and studies behavioral responses to incentives and signals in queues and inventory networks. Relying on steady-state assumptions, the authors designed a state-dependent policy while considering behavioral patterns. However, this approach will be difficult for accommodating the blood-donation features, particularly the safety observation window and the perishable nature of blood inventory. Similar to the queueing literature, the reliance on steady-state assumption leaves little room for optimization in the transient setting. Furthermore, as our research aims to provide a practical decision-support tool for blood management, we do not intend to adopt a mechanism-design approach, given that it generally imposes strong assumptions on utility functions and behavioral rationality.

1.2. Main Approach and Contributions

The review of existing methodologies underscores their limitations in terms of tractability and reliance on various assumptions, such as steady-state analysis and specifications of the utility function. In addition, the literature is particularly scant on how donor dynamics affect the supply of blood. Blood supply is often either modeled as stochastic, with a known distribution or as a decision variable. The omission of donor dynamics fundamentally renders it challenging to examine the impact of donor incentivization, and, hence, extending the existing models in the literature is untenable.

To this end, we instead adopt the paradigm recently introduced in Pipeline Queues (P-Queues, for short; Bandi and Loke 2018). In this paper, the authors introduced a new framework specifically targeting problems that involve the optimization of flows within a queueing network. The framework is introduced in the transient setting, leads to state-dependent policies, and is polynomial-time solvable. We believe this fits the needs of our intended problem. The technique is also adopted in two works within the literature—namely, Tang et al. (2020), which studied the vehicle-repositioning problem as a transient queueing network; and Zhou et al. (2022), which examined the patient scheduling problem under patient re-entry. In both cases, they involve feedback loops in the network, which is especially relevant to our work—where, similarly, donors return for multiple donations. Despite the complexity, both works arrived at tractable models that also perform strongly against benchmarks.

In our paper, we present a model for optimizing donor incentivization using a P-Queue model that aims to reduce the risk of blood shortages and wastage. We make the following contributions:

a. *Practical and tractable decision-support tool to blood-donor management*: To the best of our knowledge, our paper is the first to consider optimizing donor incentivization that responds to the inventory level dynamically. Unlike related works in the literature that present theoretical analysis with elegant structural results, but focus only either on (i) the inventory problem, where the supply is assumed to be exogenous, or (ii) the donor-management problem, where demand is ignored, we examine specifically how to make donor-management decisions as a function of the present inventory levels. Our approach allows us to build a decision-support tool, requiring fewer restrictive assumptions, incorporating challenging features usually omitted and simplified in other works, and analyzing the problem in a dynamic and transient setting. We adopt the novel P-Queue framework, which is tractable (polynomial-time-solvable; Theorem 1) and generates the optimal state-dependent donor-incentivization policy.

b. *Model flexibility and applicability*: Our base model can be easily adapted to address multiple incentivization schemes, multiclass donors, and multiple blood types (discussion in Section 2.7), or different time-dependent structures (Online Appendix B), without implications on the overall tractability. We also believe that our work has wider relevance to supply incentivization and perishability, such as in two-sided markets (e.g., home-sharing) and supply chains with social impacts (e.g., charity organizations and food donations). In particular, our coupled network model is able to handle both demand and supply dynamics separately, and the P-Queue technique can easily handle hard temporal cut-offs, such as perishability.

c. *Theoretical extensions*: Our paper advances the technique of P-Queue in three aspects. First, we propose a modification that allows relaxing some independence assumptions in the original framework (Propositions 6 and 7). Second, we introduce the novel concept of reducing the one-period delay that exists between the queues and servers (discussion after Equation 3). Finally, our model involves a coupled queueing system, for which is not immediately clear how it could be executed from the original framework.

d. *Superior performance*: We compare the optimal P-Queue policy with threshold-type policy benchmarks in the numerical experiments (Section 3). Our optimal policy can effectively reduce both shortages and wastage. Our optimal policy also exhibits features such as differentiated incentivization strategies for different classes of donors, forward-looking and risk-averse planning, and smoothing of interventions to maintain stability in the inventory levels, leading to its superior performance over the benchmarks. Specifically, the insights derived from the numerical study are summarized as follows:

a. Our model outperforms threshold policies because it is able to account for nonstationary future demand and uses more real-time information about the state; and

b. We highlight several features of a good donor-/inventory-management policy that can serve as high-level guidance for practitioners. The policies obtained from our model possess these three properties compared with the benchmark.

i. Maintaining a high level of inventory helps to fulfill demand, but entails better management of old blood to reduce wastage.

ii. Incentivizing low-responsive donors during times of surpluses helps maintain a high level of active donors and reduces dropouts amongst low-responsive donors.

iii. Maintaining a minimum pool of eligible high-responsive donors grants the greater capacity to react to prolonged shortages.

1.3. Organization of the Paper

The remainder of this paper is organized as follows. In Section 2, we introduce the main model setup and present tractable reformulation. We also briefly discuss potential extensions that capture more features in practical blood-donation settings. In Section 3, we design a numerical simulation study and showcase the novelty and advantages of our framework against benchmark policies. We conclude the paper in Section 4 and discuss possible future research. To keep the discourse succinct, we relegate all the proofs to Online Appendix A.

1.4. Convention

We adopt the convention $\min\emptyset = \infty$, where $\emptyset$ is the empty set and $\log 0 = \max\emptyset = -\infty$. For brevity, for a given indexed variable $p^{t,s}$, we abuse the notation

$(1-p)^{t,s}$ to mean $1-p^{t,s}$. Let $\mathcal{P}(X)$ be defined as the space of probability distributions on X.

2. Basic Blood-Donation Model

In this section, we present an integrated model of blood supply and demand with a particular focus on donation behavior and dynamics. Our model is broadly motivated by the technique of Pipeline Queues (Bandi and Loke 2018). To see the technique in action in other problem settings, readers are referred to Zhou et al. (2022) and Tang et al. (2020).

2.1. Sequence of Events

Our proposed model is a series of three queue-server dyads, representing the blood inventory, donors before incentivization, and incentivized donors, as illustrated in Figure 1. They organize into two queuing networks, one modeling the eligibility of donors for donation and the incentivization process, and the other, the blood inventory. The two networks are related via the process of blood donation—for each donor who successfully completes a donation, that corresponding blood packet triggers an inflow for the inventory network. We term this setting a tandem network.

Let T be the last modeling time and S be the largest "present delay" for index s, which we later introduce. Define $[T]$ as $\{1,\ldots,T\}$ and $[T]_0 := \{0\}\cup[T]$; $[S]$ and $[S]_0$ are defined analogously. Consider a finite time horizon $t \in [T]_0$, where $t=T$ is the last modeling time period, and $t=0$ represents the initial state. In the P-Queue framework, time-in-node is also tracked. We use the index $s \in [S]_0$ to track the present delay of each job in a server or a queue. It is assumed that $S > T$.

We first describe the network for donors. In the "eligible donors" server, each job represents a donor who has completed their safety observation window (i.e., is at present eligible for donation). Specifically, let $y^{t,s}$, $t \in [T]_0$, $s \in [S]_0$ be the random variable representing the number of eligible donors in the server, who have yet to donate by time t after s periods of completing their observation window. Donors can donate without incentivization. This is modeled by the direct feedback from the eligible donors server to the observation queue. For those donors who receive incentivization, they will be routed to the loop of the incentivized donors. More details on incentivization will be discussed later. Donation is modeled as service completion in the servers. After donation, the donor returns to the observation queue, where they must observe a $S_o(<S)$ safety observation period before they are next eligible for donation. We denote the state of the observation queue with the random variable $z^{t,s}$, $t \in [T]_0$, $s \in [S]_0$ representing the number of donors who have observed exactly s periods of the S_o observation window at time t. Once the observation window is completed, the donor may rejoin the pool of eligible donors in the server. The movement of donors who have observed s periods by time t into the server is modeled by the auxiliary decision variables $q^{t,s}$, $t \in [T]$, $s \in [S]_0$. Here, by definition, $q^{t,s} = 0, \ \forall s < S_o$. Note that the dynamics are not defined for $s > S$. Donors who stay past $s = S$ are assumed to have left the donor pool as *dropouts*.

We also model new donors to the system via an inflow into the observation queue. In practice, the blood bank only knows that the donor is a new donor when they make their first donation. As such, our assumption here is that every inflow has just made their first donation and, hence, enters the observation queue to complete their observation window before being allowed to donate again. We denote inflow as the random variable λ^t at time t.

Figure 1. Blood Donation Network with Incentivization

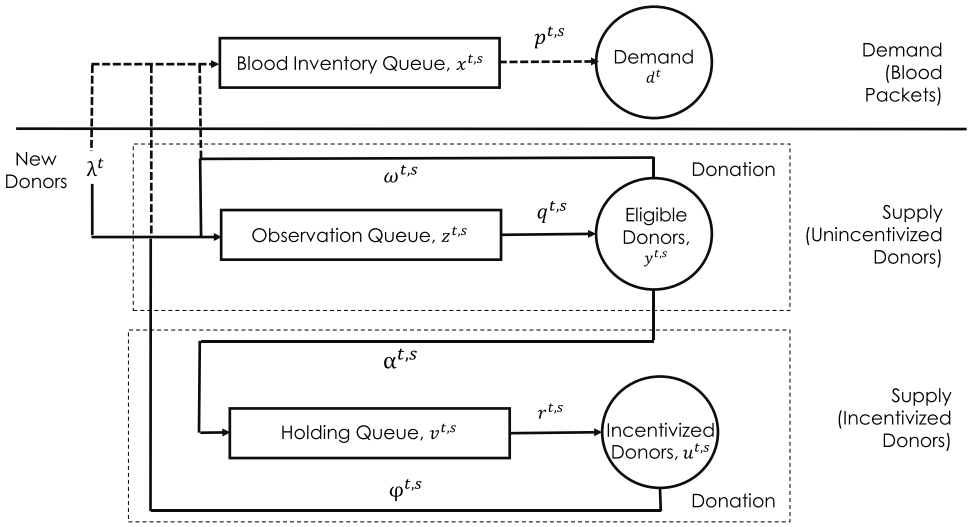

In our model, we describe incentivization as an outflow decision variable from the eligible donor server $y^{t,s}$. Keeping with the primitives of a P-Queue, we define this outflow as an adaptive decision: Let the *decision* variables $\alpha^{t,s} \in [0,1]$ represent the proportion of donors who have been eligible since s periods ago at time t, who are to receive the incentivization. After receiving this incentivization, the donors, represented by this outflow, would go into a new server, which we term the "incentivized donors" server, with state variables $u^{t,s}$, representing the number of donors who at time t have yet to donate, after s periods since incentivization. Adhering to the structure of P-Queues, we model an intermediate and auxiliary queue, termed the "holding queue," with state variables $v^{t,s}$, which receives the incentivized donors and dispatches them to the invited donors server via auxiliary decision variable, $r^{t,s}$, similar to the observation window. Like the other donors, upon successful donation, incentivized donors will return to the safety observation queue.

In the other network, we track the donated blood packets until they have expired or are used to fulfill the time-nonhomogeneous stochastic demand d^t. Here, each job in the network represents one blood packet. Denote the random variable $x^{t,s}$, $t \in [T]_0$, $s \in [S]_0$ as the number of blood packets at time t that have spent s periods in the blood inventory queue. Amongst these blood packets, the auxiliary decision variables $p^{t,s}$, $t \in [T], s \in [S]_0$ of them will be used to fulfill demand d^t. Let $S_e (< S)$ represent the shelf life of blood. Hence, $x^{t,s}$ for any $s \geq S_e$ would contribute to wastage as a result of expired blood. The parameters and notations used in the model are detailed in Table 1.

We remark that the push variables $q^{t,s}$ and $r^{t,s}$ are auxiliary variables to model the emptying of the queue into the servers. They are needed solely for model construction and do not correspond to any decisions in practice. As such, although they are solved as decision variables in the model, in reality, they will not be implemented.

2.2. Effect of Incentivization and Model Assumptions

Before defining the dynamics of our model, let us state clearly the assumptions that we make on the random variables—namely, the inflow, demand, and the service-time distributions of both incentivized and nonincentivized donors.

Assumption 1. *Inflow $\lambda^t \sim \Lambda^t$ and demand $d^t \sim D^t$ distributions can be time-nonhomogeneous, but are independent across t. We assume that their moment-generating functions exist.*

These assumptions are relatively loose. Time nonhomogeneity can capture seasonal patterns, and if there is a need to observe nonindependent distributions across time, the model may still be solved on a rolling horizon manner, where the information of the new distribution is updated when new data are realized. We also consider the setting where inflow and demand are bounded, which guarantees that their moment-generating function exists.

Assumption 2. *The decision of each donor to donate is independent of any other donor. Furthermore, we assume that at any time $t \geq 0$, the probability of whether the donor will donate in the next time period (without incentivization), conditional on the fact that they had not donated $s \geq 0$ periods since they became eligible, is the same for all donors; and this is denoted by $\omega^{t+1,s+1}$.*

Independence is a reasonable assumption in our setting of blood donation. At first glance, requiring the conditional probabilities of donation to be the same appears restrictive. However, this representation is, in fact, general—any nonstationary discrete-time service-time distribution can be represented in the form of $\omega^{t,s}$ (Dai and Shi 2017). Moreover, as we shall see later, due to the structure of our model, its tractability is not influenced by adding an index i to the conditional probabilities to model heterogeneous donors; in other words, this assumption can be relaxed when we define different conditional probabilities $\omega_i^{t,s}$ on subpopulations i of donors.

Assumption 3.

a. *Incentivization can occur at any time point t or period since eligibility s;*

b. *For any donor, upon incentivization, the probability of donation is independent of the probability of donation prior to incentivization;*

c. *The probability of donation in the next time period, given incentivization, is independent across donors. This probability, at any time $t(\geq 0)$, is the same for all donors conditional on the time since they were incentivized at $s(\geq 0)$; and we denote this probability by $\phi^{t+1,s+1}$; and*

d. *The act of incentivization occurs prior to the act of donation at each time period.*

We wish to consider incentivization mechanisms that influence donors' probability of donation, as opposed to modeling incentivization as an inclusion or exclusion policy. This allows us to cover a large range of incentivization mechanisms, such as (i) incentivizing donors to return to donate or to hasten their donations, (ii) incentivizing donors to delay their donations, and (iii) dispensing different incentives to different subpopulations of donors. We discuss this more in Section 2.7 and also in Section 3. This is also in line with literature on mechanism design and Bayesian persuasion, where the effect of incentivization is modeled by the change in the agent's probability distribution of taking actions. For example, the Bayesian persuasion framework models agents as expected utility maximizers. The incentivization scheme works as a mechanism that alters agents' beliefs on unknown system states via a Bayesian updating process, subsequently changing the postincentivization expected utility and the probability of taking certain actions. Given

Table 1. List of Parameters and Variables

Parameter or variable	Definition
Dimensions	
T	Last modeling time
S_e	Shelf life of blood packets
S_o	Safety observation period of donors
S	Index upper bound of the present delay in each server or queue; also the point at which donors are considered dropouts
Parameters	
$\omega^{t,s}$	Probability of donation for donors who waited for s periods after completing safety observation period at time t
$\phi^{t,s}$	Probability of donation for donors who have yet to donate s periods after incentivization at time t
W	Target on total wastage by the last time period, T, that the decision maker hopes to keep within
Primary random variables	
λ^t	Random arrivals of first-time donors at time t
d^t	Demand for blood at time t
Random state variables	
$y^{t,s}$	Random variable of the number of eligible donors at time t, completed observation period, and yet to donate for s periods
$z^{t,s}$	Random variable of the number of donors observing safety observation for s periods at time t
$x^{t,s}$	Random variable of the number of blood units in inventory queue for s periods at time t
$u^{t,s}$	Random variable of the number of invited donors who have yet to donate for s periods at time t
$v^{t,s}$	Random variable of the number of registered and invited donors in holding queue for s periods at time t
Decision variables	
$q^{t,s}$	Push variable of donors dispatched to donor pool after spending s periods in the observation queue at time t
$p^{t,s}$	Push variable of blood units dispatched to fulfill demand after spending s periods in blood inventory queue at time t
$r^{t,s}$	Recourse variable of invited donors dispatched to the server after waiting for s periods at time t
$\alpha^{t,s}$	Decision variable of proportion of donors to approach

that our paper aims to provide a practical decision-support tool, we do not intend to incorporate utility models or belief-updating processes to detail the interaction between donors and an incentivization scheme to avoid stylized assumptions on utility functions and donors' rationality. Instead, we take the model-free, but *data-driven*, approach that treats the overall effect of the incentivization mechanism, captured by the difference between ω (i.e., probability of donation prior to incentivization) and ϕ (i.e., probability of donation after incentivization), as inputs, which can be estimated from the data. The benefit is that our model can work with any specific incentivization scheme, as long as the data support estimating the donation probabilities before and after the incentivization.

Assumption 3(a) allows the mechanism to be general in time. Assumption 3(b) sets up a setting similar to those in the mechanism design and Bayesian persuasion domain, where donors act differently postincentivization and independently of their prior decision to donate. This assumption is also analogous to, but more general than, the Markov property in MDPs, which require the process to be memoryless, hence requiring current states to capture all relevant information that affects the future. There is also empirical evidence from studies in Heger et al. (2020) supporting the assumption that the effect of events before incentivization is not as significant as the incentivization itself. Assumption 3(c) is the incentivized analog to Assumption 2. Finally, Assumption 3(d) just describes the staging of the events and does not have any real implications on the model. The setup of Assumption 3 motivates our approach to model incentivization as an outflow decision variable from the eligible donor server $y^{t,s}$.

One may question why the time since the last donation is not a factor that determines the probability of donation postincentivization. In Heger et al. (2020), the authors do not find empirical evidence that the time since the last donation is a significant factor. Nonetheless, to illustrate the flexibility of our framework and in other cases where the time since the last donation indeed matters, we have put in an extension where we consider time since the last donation as a factor in Online Appendix B.

2.3. Dynamics

We begin by describing the dynamics of the donors and start with the server consisting of eligible donors. Specifically, at time t, inflow into the server, which is equivalently all donors who have spent zero periods in the

server so far, is represented by $y^{t,0}$. This is fully made up of donors who have completed their safety observation period: $y^{t,0} = \sum_{s=0}^{S} q^{t,s}, \ \forall t \in [T]$.

For $s > 0$, the dynamics of $y^{t,s}$ only involve outflows. There are three key groups within the cohort of $y^{t,s}$. By Assumption 3(d), incentivization occurs before donation; in other words, we first identify $\alpha^{t,s}y^{t-1,s-1}$ of them as targets for incentivization and route them to the holding queue, whose dynamics we will describe in a bit. Of the remaining $(1-\alpha^{t,s})y^{t-1,s-1}$ of them, the identical and independence assumption in Assumption 2 justifies the definition for all $t \in [T]$, $s \in [S]$,

$$y^{t,s} = \text{Bin}((1-\alpha^{t,s})y^{t-1,s-1}, 1-\omega^{t,s}),$$

who are the donors that did not donate at time t and continue to remain in the server. Consequently, $f^{t,s} = \text{Bin}((1-\alpha^{t,s})y^{t-1,s-1}, \omega^{t,s})$ number of donors would complete their donation and return to the observation queue. For convenience later, we denote $\Omega^{t,s} := \prod_{\tau=0}^{\min\{s,t\}-1} (1-\omega^{t-\tau,s-\tau})$, the cumulative probability that the donor's donation time is at least s.

Notice here that, in keeping with the P-Queue framework, we have decided to model the incentivization decision as a proportion of the state, $\boldsymbol{\alpha}$, rather than as a deterministic number. This ensures that the incentivization is always an *adaptive* decision, reacting to the state $\boldsymbol{y}$.

Proposition 1 (Independence of Pipelines). *For a fixed time $t \in [T]$, for any $s, s' \in [S]$ such that $s \neq s'$, $y^{t,s}$ and $y^{t,s'}$ are independent.*

Proposition 1 is the culmination of the fact that the dynamics are organized by "pipelines" $\{y^{t-\tau,s-\tau}\}_{\tau=0}^{\min\{t,s\}}$, which trace back to either $y^{0,s-t}$, some initial condition, or inflow $y^{t-s,0}$, which is the sum of decision variables. In other words, the dynamics essentially play out within cohorts.

This concludes the dynamics for the eligible donors server. The dynamics in the incentivized donors server follow along similar lines. First, inflow is given by $u^{t,0} = \sum_{s=0}^{S} r^{t,s}, \ \forall t \in [T]$. Next, Assumption 3(c) being identical to Assumption 2 allows us also to define a Binomial distribution on the state $u^{t,s} = \text{Bin}(u^{t-1,s-1}, 1-\phi^{t,s})$. The absence of the decision variables here enables us to state a stronger result:

Proposition 2.

a. *Let $\Phi^{t,s} := \prod_{\tau=0}^{\min\{s,t\}-1}(1-\phi^{t-\tau,s-\tau})$; then, $\forall t \in [T]$,*

$$u^{t,s} = \begin{cases} \text{Bin}(u^{t-s,0}, \Phi^{t,s}) & \text{for } 0 < s < t, \\ \text{Bin}(u^{0,s-t}, \Phi^{t,s}) & \text{for } t \le s \le S. \end{cases}$$

b. ***Independence of Pipelines***: *For a fixed time $t \in [T]$, for any $s, s' \in [S]$ such that $s \neq s'$, $u^{t,s}$ and $u^{t,s'}$ are independent.*

We move on to the dynamics of the queues, starting with the holding queue. As previously mentioned, the $\alpha^{t,s}y^{t-1,s-1}$ donors earmarked for incentivization leave the eligible donors server and enter this queue. Thus, the inflow into the holding queue can be defined by

$$v^{t,0} = \sum_{s=1}^{S} \alpha^{t,s} y^{t-1,s-1} - r^{t,0}. \tag{1}$$

Here, the term $r^{t,0}$ represents the donors who may immediately be pushed into the incentivized donors server. This term is critical, as it is previously absent in the original P-Queue framework. Without it, the newly incentivized donors must minimally wait for at least one time period before donating, which defeats the purpose of incentivization. There are more serious consequences of this additional term for the other queues, which we explain when we arrive at that point.

The rest of the dynamics of the holding queue comprises the outflows, and the only outflows from the holding queue are the dispatches into the incentivized donors server. Hence, for $t \in [T]$,

$$v^{t,s} = v^{t-1,s-1} - r^{t,s} = \begin{cases} v^{t-s,0} - \sum_{\tau=0}^{s-1} r^{t-\tau,s-\tau} & \text{for } 0 < s < t, \\ v^{0,s-t} - \sum_{\tau=0}^{t-1} r^{t-\tau,s-\tau} & \text{for } t \le s \le S. \end{cases}$$

We also describe the dynamics of the observation queue, which is more complicated. Notice that the inflow into the observation queue is all donors who have just donated blood, and there are three sources of them—namely, the new donors, the donors who donated with incentivization, and those who did so without. Thus, one can respectively state the dynamics as

$$z^{t,0} = \lambda^t + \sum_{s=1}^{S} \text{Bin}(u^{t-1,s-1}, \phi^{t,s}) + \sum_{s=1}^{S} \text{Bin}((1-\alpha^{t,s})y^{t-1,s-1}, \omega^{t,s}). \tag{2}$$

Notice that here, we could have added a term $-q^{t,0}$ like in (1), but $q^{t,0} = 0$, by definition of the observation period. Similar to the case for v, the outflow is written as, $\forall t \in [T]$,

$$z^{t,s} = \begin{cases} z^{t-s,0} - \sum_{\tau=0}^{s-1} q^{t-\tau,s-\tau} & \text{for } 0 < s < t, \\ z^{0,s-t} - \sum_{\tau=0}^{t-1} q^{t-\tau,s-\tau} & \text{for } t \le s \le S. \end{cases}$$

Finally, we describe the dynamics of the blood packets in the blood inventory. Note that every inflow into the observation queue is marked by a donation. Hence, the inflow of blood packets into the blood-inventory queue is exactly equal to the inflow of donors into the observation queue. As such, the inflow dynamics into the blood inventory x would be the same as that for the blood donors z. Indeed, for all $t \in [T]$,

$$x^{t,0} = \lambda^t + \sum_{s=1}^{S} \mathrm{Bin}(u^{t-1,s-1}, \phi^{t,s}) + \sum_{s=1}^{S} \mathrm{Bin}((1-\alpha^{t,s})y^{t-1,s-1}, \omega^{t,s}) - p^{t,0}. \quad (3)$$

Here, the additional term $p^{t,0}$ is included to ensure that newly donated blood packets can be immediately dispatched to fulfill demand. This avoids the minimal wait time of one period, which is undesirable in a high-demand setting.

Finally, the outflow from the blood inventory corresponds to the blood packets utilized to fulfill the demand: For all $t \in [T]$,

$$x^{t,s} = \begin{cases} x^{t-s,0} - \sum_{\tau=0}^{s-1} p^{t-\tau,s-\tau} & \text{for } 0 < s < t, \\ x^{0,s-t} - \sum_{\tau=0}^{t-1} p^{t-\tau,s-\tau} & \text{for } t \le s \le S. \end{cases}$$

2.4. Constraints

Before listing the constraints, let us first caveat that we will be describing the constraints as if they were deterministic. However, the constraints essentially are functions of random variables, and, thus, we would need some measure to evaluate these constraints later. In the next subsection, we will provide a chance constraint interpretation and our rationale for proposing our model.

The main objective of a blood bank is to fulfill the demand for blood, while minimizing wastage. To handle this multiobjective problem, although it may be possible to write a combined objective that quantifies the trade-off for blood shortages against the cost of wastage, we avoid this, as it is generally difficult to prescribe such a trade-off. Instead, in the P-Queue framework, both objectives are modeled as constraints with the aim of finding a feasible policy that runs a high probability of meeting the demand and keeping wastage below a certain level.

At time t, demand is fulfilled by blood allocated from the inventory $p^{t,s}$ of different age s:

$$\sum_{s=0}^{S_e-1} p^{t,s} \ge d^t \quad \forall t \in [T].$$

Here, the limits run till $S_e - 1$, as expired blood cannot be used to fulfill the demand. This leads to $p^{t,s} = 0, \ \forall s \ge S_e$.

Demand fulfillment is modeled at every time period separately, as opposed to summed over all time periods t. This is because the latter can potentially result in large shortages at a particular time point in exchange for low or no shortages at other times. The blood bank also has to limit the total wastage of blood over the planning horizon under some level W:

$$\sum_{s=S_e}^{S} x^{T,s} + \sum_{t=1}^{T} \left(\sum_{s=0}^{S_e-1} p^{t,s} - d^t \right) \le W.$$

The first term represents the total expired blood packets, while the second term captures the total number of excess blood packets that were overcommitted to fulfilling the demand, summed over all time periods $t \in [T]$.

On servers, we can impose capacity constraints, whereas for queues, constraints can be imposed to ensure that the queue does not build up when there is spare capacity in the server. In both cases, such constraints can be written in Linear Forms (4). This is done by $a_n^s = 1, n \in \{x, y, z, u, v\}$, where the left-hand side (LHS) of (4) represents the total number of donors or blood packets in the server or queue at any given time t. The corresponding linear forms give rise to capacity constraints and queue-clearing constraints. Note that those linear forms in (4) can also represent other types of constraints required to control the system, such as waiting time in the queue and a blood-demand fulfillment policy. Specifically, when $a_n^s = \max\{0, s - S_o\}$, the LHS gives rise to the total waiting time after completing the observation window and before being dispatched to the eligible donors server, which are situations that we seek to reduce. Commonly adopted blood-demand fulfillment policies include the FIFO policy, where demand is fulfilled by the least fresh blood packet to ensure little wastage, or the last-in-first-out (LIFO) policy, which allocates the freshest blood packet to the demand (Sarhangian et al. 2018). The linear forms can specify both types of policies, and there may be many ways to model them. In our paper, we adopt the FIFO demand-fulfillment policy as a demonstration that is enforced by constraint $a_n^s = s^2$. LIFO can be similarly modeled (e.g., $a_n^s = 1/(s+1)^2$).

$$\sum_{s=0}^{S} a_n^s n^{t,s} \le b_n^t, \qquad n \in \{x, y, z, u, v\}. \quad (4)$$

We introduce a budget that restricts the maximum number of incentivized donors. Here, we can consider the constraint $\sum_{s=0}^{S} c^{t,s} y^{t,s} \le B^t$, where $c^{t,s}$ is the per-donor cost of dispensing the incentive to donors who have yet to donate for s periods by time t, and B^t is the budget at time t. For example, to impose a constraint on the maximum number of donors that can be called up in any

period, one can adopt $c^{t,s} \equiv 1$. Notice that this is also of the form (4).

We also have the feasibility constraints on push variables to ensure that there will not be more donors made eligible than there are donors already in the observation queue: $z^{t,s}, x^{t,s} \geq 0,\ \forall t \in [T], s \in [S]_0$, which are equivalent to $q^{t,s} \leq z^{t-1,s-1}, p^{t,s} \leq x^{t-1,s-1}$. This is a special case of (4), where $a^s = -1$ for some s, and zero otherwise. As we see in the next section, this does not hinder our attempts to arrive at a tractable formulation.

Lastly, after donation, donors must observe the safety period and cannot be pushed into the eligible donors server: $q^{t,s} = 0,\ \forall t \in [T], s < S_o$. For demand fulfillment, expired blood packets cannot be used: $p^{t,s} = 0,\ \forall t \in [T], s \geq S_e$.

Most critically, note that the constraints introduced above are all linear in the state variables $x^{t,s}, z^{t,s}, y^{t,s}$, decision variables $p^{t,s}, q^{t,s}$, and exogenous uncertainties λ^t, d^t. This will help us achieve tractable reformulations.

2.5. Model and Reformulation

In the P-Queue framework, stochastic constraints $\tilde{\zeta} \leq 0$ are modeled via their corresponding surrogates defined under the Aumann and Serrano (2008) riskiness index, $C_{k,\theta}[\tilde{\zeta}](:= k \log \mathbb{E}[\exp(\tilde{\zeta}/k\theta)]) \leq 0$. These surrogate constraints control the probability of violation of the original constraint, as Proposition 3 details.

Proposition 3. *Let $k, \theta > 0$. For a random variable $\tilde{\zeta}$, define $C_{k,\theta}[\tilde{\zeta}] = k \log \mathbb{E}[\exp(\tilde{\zeta}/k\theta)]$.*[1] *If $C_{k,\theta}[\tilde{\zeta}] \leq 0$, then,*

$$\mathbb{P}[\tilde{\zeta} \geq \Delta] \leq \exp(-\Delta/k\theta) \quad \forall \Delta > 0.$$

Proposition 3 is a standard result from the literature on *Satisficing* (Brown and Sim 2009, Brown et al. 2012) and is a direct application of Markov's inequality. From Proposition 3, we can see that both k and θ control the probability of constraint violation—the smaller k or θ, the sharper the guarantees. In the literature, k is treated as the global risk level, which we attempt to minimize, whereas θ is viewed as an idiosyncratic parameter that controls the tightness of each constraint. Here, we label each θ according to the constraint type "n" and time t, denoted as $\theta_{n,t}$. Indexing by t allows the decision maker to flexibly vary the tightness of the constraint in time t. For example, if demand shortages at early times are more critical than later times, the decision maker may make $\theta_{d,t}$ smaller for small t. In this paper, we consider two levels of tightness, θ_{hard}, for hard constraints that should never be violated (e.g., capacity constraints) and θ_{soft} for soft constraints that can be violated, but should only be done so infrequently (e.g., waiting time). The value of θ_{hard} is much smaller than that of θ_{soft} (e.g., $\theta_{\text{hard}} = 0.01$ versus $\theta_{\text{soft}} = 1$). In practice, the decision maker would calibrate θ based on the probability $\exp(-\Delta/k\theta)$ of incurring violation Δ.

Proposition 3 motivates the following optimization problem:

$$\begin{aligned}
\min_{k>0,\,p,q,r,\alpha} \quad & k \\
\text{s.t.} \quad & C_{k,\theta_{d,t}}\left[d^t - \sum_{s=0}^{S_e-1} p^{t,s}\right] \leq 0, \\
& \qquad \forall t \in [T] \qquad \textbf{[Demand]} \\
& C_{k,\theta_{W,t}}\left[\sum_{s=S_e}^{S} x^{T,s} + \sum_{t=1}^{T}\left(\sum_{s=0}^{S_e-1} p^{t,s} - d^t\right) - W\right] \leq 0, \\
& \qquad \textbf{[Wastage]} \\
& C_{k,\theta_{n,t}}\left[\sum_{s=0}^{S} a_n^s n^{t,s} - b_n^t\right] \leq 0,\ n \in \{x,y,z,u,v\}, \\
& \qquad \forall t \in [T] \qquad \begin{bmatrix} \textbf{Capacity, Budget} \\ \textbf{Queue-clearing} \\ \textbf{Waiting-time, FIFO} \end{bmatrix} \\
& C_{k,\theta_{n,t}}[-n^{t,s}] \leq 0, \qquad n \in \{x,z,v\}, \\
& \qquad \forall t \in [T],\ \ \forall s \in [S]_0 \qquad \textbf{[Push Constraints]} \\
& q^{t,s} = 0,\ \ \forall t \in [T],\ \forall s < S_o \quad \text{and} \quad p^{t,s} = 0 \\
& \qquad \forall t \in [T],\ \forall s \geq S_e \qquad \textbf{[Logic Constraints]} \\
& n^{t,s} \geq 0, n \in \{p,q,r\}, \quad \text{and} \quad \alpha^{t,s} \in [0,1], \\
& \qquad \forall t \in [T],\ \forall s \in [S] \qquad \textbf{[Decisions]}
\end{aligned} \tag{5}$$

In the rest of this section, we will show that the above surrogate constraints $C_{k,\theta_{n,t}}[\sum_{s=0}^{S} a_n^s n^{t,s} - b_n^t] \leq 0$ can be reformulated for each of $n \in \{x,y,z,u,v\}$ into computationally tractable forms that are jointly convex in the decision variables. The reformulations for each of them differ due to the different dynamics at each of the nodes. Because of the additive nature of the operator $C_{k,\theta}[\cdot]$, even if the right-hand-side terms b_n^t are decision variables, such as in the wastage constraint that is linear in x with push variables $p^{t,s}$ in b_n^t, they would lead to exactly the same expressions had we treated them as constants. Additionally, additivity also allows the demand terms d^t, which are conditionally independent of the rest of the dynamics, to be separated—for example, $C_{k,\theta_{d,t}}[d^t - \sum_{s=0}^{S_e-1} p^{t,s}] = C_{k,\theta_{d,t}}[d^t] + C_{k,\theta_{d,t}}[-\sum_{s=0}^{S_e-1} p^{t,s}]$—leaving a linear expression in the state variables. Thus, it suffices to provide reformulations for $C_{k,\theta_{n,t}}[\sum_{s=0}^{S} a_n^s n^{t,s} - b_n^t] \leq 0$ alone. As such, we omit their reformulations for brevity. Also note that in the following propositions, if index sets are empty, the corresponding terms or constraints

are understood to be omitted. Define the function $\rho_\pi^{t,s}(\zeta) := \log(1-\pi^{t,s}+\pi^{t,s}\exp(\zeta))$.

Proposition 4. *For a given $t \in [T]$,*

$$C_{k,\theta_{u,t}}\left[\sum_{s=0}^{S} a_u^s u^{t,s} - b_u^t\right] = k\sum_{s=0}^{t-1}\sum_{s'=0}^{S} r^{t-s,s'}\rho_\Phi^{t,s}(a_u^s/k\theta_{u,t}) + k\sum_{s=t}^{S} u^{0,s-t}\rho_\Phi^{t,s}(a_u^s/k\theta_{u,t}) - b_u^t/\theta_{u,t}. \tag{6}$$

What we can see in this proposition is the effect of the risk correction at risk level k. Notice that if $\zeta > 0$, then $k\rho_\pi^{t,s}(\zeta/k) \to \pi^{t,s}\zeta$ as $k \to \infty$; thus, one recovers $\sum_{s=0}^{t-1}\sum_{s'=0}^{S} r^{t-s,s'}\Phi^{t,s}a_u^s/\theta_{u,t}$ and $\sum_{s=t}^{S} u^{0,s-t}\Phi^{t,s}a_u^s/\theta_{u,t}$ for the first and second terms in (6), which are just the expected value of the LHS of the constraint, split between contributions from the inflow (decision variables r) and the initial conditions, respectively. However, because $k\rho_\pi^{t,s}(\zeta/k)$ is convex and decreasing in k, a lower risk level k thus increases the right-hand side of (6), consequently tightening the constraint. Zhou et al. (2022) dive more deeply into this, and the interested reader is directed there.

Proposition 5. *For any $t \in [T]$, $C_{k,\theta_{y,t}}[\sum_{s=0}^{S} y^{t,s}a_y^s - b_y^t] \le 0$ is equivalent to the collection of the following constraints*:

$$y^{t,0}a_y^0/\theta_{y,t} + k\sum_{\tau=0}^{t-2}\xi^{t-\tau,1} + k\sum_{\tau=t}^{S}\xi^{1,\tau-t+1} \le b_y^t/\theta_{y,t}$$

$$\xi^{t,s} \ge \beta^{t,s}\rho_{1-\omega}^{t,s}(a_y^s/k\theta_{y,t}) \qquad \forall s \in [S]$$

$$\xi^{t-\tau,s-\tau} \ge \beta^{t-\tau,s-\tau}\rho_{1-\omega}^{t-\tau,s-\tau}(\xi^{t-\tau+1,s-\tau+1}/\beta^{t-\tau,s-\tau}) \qquad \forall \tau \in [t-1], s \in [S]\setminus[\tau].$$

Unlike Proposition 4, the dynamics of y involve the decision variables α. Hence, a closed form in the sense of (6) is no longer possible. In its place, we introduce the auxiliary variables ξ, which are dependent on the optimal auxiliary decision variables β, which we defined as a substitute for representing the original incentivization decision $\alpha^{t,s}$ via $1-\alpha^{t,s} = \beta^{t,s}/\beta^{t-1,s-1}$. These perspective functions $\beta\rho_{1-\omega}^{t,s}(\xi/\beta)$ are *jointly* convex in β and ξ, which is the point of the substitution. Because this representation uses more decision variables (β's) than there originally are (α's), it grants us the degree of freedom to determine the boundary values as $\beta^{t,0} = y^{t,0} := \sum_{s=0}^{S} q^{t,s}$ and initial conditions $\beta^{0,s} := y^{0,s}$.

Proposition 6. *For any $t \in [T]$, $C_{k,\theta_{v,t}}[\sum_{s=0}^{S} a_v^s v^{t,s} - b_v^t] \le 0$ is equivalent to the collection of the following constraints*:

$$\sum_{j=1}^{t-1}\frac{\beta^{j,0}-\beta^{j+1,1}}{\theta_{v,t}/a_v^{t-j-1}} + \sum_{j=t}^{S-1}\frac{\beta^{0,j-t}-\beta^{1,j-t+1}}{\theta_{v,t}/a_v^{t-1}} + \sum_{j=S-t}^{S-1}\frac{\beta^{0,j}-\beta^{1,j+1}}{\theta_{v,t}/a_v^{t-1}} + k\left(\sum_{j=2}^{t-1}\eta^{t-j+1,1} + \sum_{j=t}^{S-1}\eta^{1,j-t+1} + \sum_{j=S}^{t+S-2}\eta^{1,j-t+1}\right) - \frac{1}{\theta_{v,t}}\sum_{s=0}^{t-1}\sum_{\tau=0}^{s} a_v^s r^{t-\tau,s-\tau} + \frac{1}{\theta_{v,t}}\sum_{s=t}^{S} a_v^s\left(v^{0,s-t} - \sum_{\tau=0}^{t-1} r^{t-\tau,s-\tau}\right) \le b_v^t/\theta_{v,t}$$

$$\eta^{t-1,j-1} \ge \beta^{t-1,j-1}\rho_{1-\omega}^{t-1,j-1}\left(\frac{a_v^0(\beta^{t-1,j-1}-\beta^{t,j})}{k\theta_{v,t}\beta^{t-1,j-1}}\right) \qquad \forall j \in [S-1]\setminus\{1\}$$

$$\eta^{t-\tau,j-\tau} \ge \beta^{t-\tau,j-\tau}\rho_{1-\omega}^{t-\tau,j-\tau}\left(\frac{\beta^{t-\tau,j-\tau}-\beta^{t-\tau+1,j-\tau+1}}{k\theta_{v,t}\beta^{t-\tau,j-\tau}/a_v^{\tau-1}} + \frac{\eta^{t-\tau+1,j-\tau+1}}{\beta^{t-\tau,j-\tau}}\right) \qquad \forall \tau \in [t-2]\setminus\{1\}, j \in [t-1]\setminus[\tau]$$

$$\eta^{t-\tau,j-\tau} \ge \beta^{t-\tau,j-\tau}\rho_{1-\omega}^{t-\tau,j-\tau}\left(\frac{\beta^{t-\tau,j-\tau}-\beta^{t-\tau+1,j-\tau+1}}{k\theta_{v,t}\beta^{t-\tau,j-\tau}/a_v^{\tau-1}} + \frac{\eta^{t-\tau+1,j-\tau+1}}{\beta^{t-\tau,j-\tau}}\right) \qquad \forall \tau \in [t-1]\setminus\{1\}, j \in [S-1]\setminus[t-1]$$

$$\eta^{t-j,S-1} \ge \beta^{t-j,S-1}\rho_{1-\omega}^{t-j,S-1}\left(\frac{\beta^{t-j,S-1}-\beta^{t-j+1,S}}{k\theta_{v,t}\beta^{t-j,S-1}/a_v^{j-1}}\right) \qquad \forall j \in [t-1]$$

$$\eta^{t-j-\tau,S-\tau} \ge \beta^{t-j-\tau,S-\tau}\rho_{1-\omega}^{t-j-\tau,S-\tau}\left(\frac{\beta^{t-j-\tau,S-\tau}-\beta^{t-j-\tau+1,S-\tau+1}}{k\theta_{v,t}\beta^{t-j-\tau,S-\tau}/a_v^{j+\tau-1}} + \frac{\eta^{t-j-\tau+1,S-\tau+1}}{\beta^{t-j-\tau,S-\tau}}\right) \qquad \forall \tau \in [t-1]\setminus\{1\}, j \in [t-\tau-1]_0. \tag{7}$$

The reformulations for the queues are visibly much more complicated than the servers. This is because, unlike the servers, the inflows into the queues are themselves random variables, as opposed to decision variables. In the case of the dynamics of v, the inflow, as given in (1), involves the incentivized donors $\sum_{s=1}^{S}\alpha^{t,s}y^{t-1,s-1}$. Moreover, despite Proposition 1, because we are additionally summing over t over this expression, the summand does not simply decompose into mutually independent terms. To give an example, consider two donors who are in the holding queue at time $t=3$, where both arrived at the eligible donors server at time $t=1$, but one stayed for one time period before being incentivized, and the other for two time periods. Either way, donors from both of these routes belong to the same cohort $y^{1,0}$, and, hence, their contributions to $\sum_{s=0}^{S} a_v^s v^{t,s}$ cannot be independent. In the original P-Queue framework, this is ignored, and the framework instead assumes that cohorts arriving at $v^{t,0}$ are independent, which, as explained, is clearly violated. In Proposition 6 and subsequently also in

Proposition 7, we introduce a novelty in the technique to avoid this assumption and to evaluate the surrogate constraints. This is a significant improvement because the expressions obtained are very different, and it helps further refine the P-Queue framework in preserving its construct of having the greatest possible generality in terms of the class of queueing networks it can model.

Proposition 7. *For any* $t \in [T]$, $C_{k,\theta_{z,t}}[\sum_{s=0}^{S} a_z^s z^{t,s} - b_z^t] \le 0$ *is equivalent to*

$$C_{k,\theta_{z,t}}\left(\sum_{s=0}^{t-1} a_z^s \lambda^{t-s}\right) - \frac{1}{\theta_{z,t}} \sum_{s=0}^{t-1}\sum_{\tau=0}^{s-1} a_z^s q^{t-\tau,s-\tau} + \frac{1}{\theta_{z,t}} \sum_{s=t}^{S} a_z^s \left(z^{0,s-t} - \sum_{\tau=0}^{t-1} q^{t-\tau,s-\tau}\right) + Z_t(k) + k\left(\sum_{j=1}^{t-1} \psi^{t-j+1,1} + \sum_{j=t}^{t+S-1} \psi^{1,j-t+1}\right) \le b_z^t/\theta_{z,t}$$

$$\psi^{t,j} \ge \beta^{t,j} \rho_\omega^{t,j}(a_z^0/k\theta_{z,t}) \qquad \forall j \in [S-1]$$

$$\psi^{t-j,S} \ge \beta^{t-j,S} \rho_\omega^{t-j,S}(a_z^j/k\theta_{z,t}) \qquad \forall j \in [t-1]_0$$

$$\psi^{t-\tau,j-\tau} \ge \beta^{t-\tau,j-\tau} \sigma_\omega^{t-\tau,j-\tau}(a_z^\tau/k\theta_{z,t}, \psi^{t-\tau+1,j-\tau+1}/\beta^{t-\tau,j-\tau}) \qquad \forall \tau \in [t-2], j \in [t-1] \setminus [\tau]$$

$$\psi^{t-\tau,j-\tau} \ge \beta^{t-\tau,j-\tau} \sigma_\omega^{t-\tau,j-\tau}(a_z^\tau/k\theta_{z,t}, \psi^{t-\tau+1,j-\tau+1}/\beta^{t-\tau,j-\tau}) \qquad \forall \tau \in [t-1], j \in [S-1] \setminus [t-1]$$

$$\psi^{t-j-\tau,S-\tau} \ge \beta^{t-j-\tau,S-\tau} \sigma_\omega^{t-j-\tau,S-\tau}\left(\frac{a_z^{j+\tau}}{k\theta_{z,t}}, \frac{\psi^{t-j-\tau+1,S-\tau+1}}{\beta^{t-j-\tau,S-\tau}}\right) \qquad \forall \tau \in [t-1], j \in [t-\tau-1]_0,$$

where $Z_t(k) := k\sum_{j=1}^{t-1}\sum_{s'=0}^{S} r^{t-j,s'} \Upsilon_{1,j}^{t,j} + k\sum_{j=t}^{S-1} u^{0,j-t} \Upsilon_{j-t+1,j}^{t,j} + k\sum_{j=S}^{t+S-1} u^{0,j-t} \Upsilon_{j-t+1,S}^{t,j}$, $\Upsilon_{l,h}^{t,j} := \log(1 + \sum_{s'=l}^{h}(\exp(a_z^{j-s'}/k\theta_{z,t}) - 1)\overline{\Phi}^{t-j+s',s'})$, $\overline{\Phi}^{t,s} := \phi^{t,s}\prod_{\tau=1}^{\min\{s,t\}-1}(1-\phi^{t-\tau,s-\tau})$, *and* $\sigma_p^{t,s}(\zeta,\zeta') := \log(p^{t,s}\exp(\zeta) + (1-p^{t,s})\exp(\zeta'))$.

In this proposition, we similarly introduce further novelties for avoiding independence assumptions, as earlier explained for Proposition 6. In this case, the innovations occur for the term $Z_t(k)$ and the nested ψ terms. Notice that all of these novel techniques introduce a new index of j, which splits into three summations. These summations actually arise out of a change of summation, which splits the space into three regions, depending on the boundary conditions. The reader may refer to the proofs in Online Appendix A for more details.

As mentioned previously, the dynamics of $x^{t,s}$ is analogous to $z^{t,s}$; consequently, its reformulation is also similar and is omitted for brevity.

2.6. Tractability

Note that in all of the above reformulations, we obtain constraints that are jointly convex in all decision variables (auxiliary or otherwise). This preserves model tractability.

Theorem 1 (Reformulation). *Problem* (5) *has a reformulation into a convex optimization problem with* $O(ST^3)$ *constraints. Moreover, it can be solved via a sequence of convex subproblems.*

We summarize in Table 6 in Online Appendix C.4 how each of the reformulations contribute to the eventual total number of constraints.

2.7. Practical Settings in Applications

At this point, we would like to present two ways that the model can be easily extended and subsequently discuss how this might be applicable to model practical settings better.

First, by adding an additional class index i on every state variable and decision variable, we can model the general situation of different donor classes, such as in the case of different blood groups. Because of the additivity (under independence) of the surrogate constraints, adding an index i for each class will lead to analogous reformulations and does not fundamentally alter the tractability guaranteed in Theorem 1. For the case of blood groups, the dynamics on the part of donors remain the same, except now divided into different blood groups. The only change occurs in the inventory, where we allow the blood of a particular type i to be used to fulfill the demand for another type j. This still leads to a set of dynamics that is linear in $x_i^{t,s}$ and $p_{i,j}^{t,s}$; hence, the analysis remains. We may further impose more penalties on using the wrong blood groups, as long as these penalties can be represented linearly in $p_{i,j}^{t,s}$. Another important example of a multiclass context is one where there are two distinct groups of high- and low-responsive donors, and the decision maker might be interested in incentivizing them differently. We explore this setting later in Section 3.

Second, our model also extends to the situation where the planner has multiple different incentives that would change the probability of donation in different ways and needs to decide on which incentives to use (or in combination or none at all). For example, a two-tiered system might include one option to incentivize donors to hasten their donation and the other to delay the donation. Our model can be altered trivially to handle such a situation by adding more layers, one for each type of incentive, similar to how the incentivization component of the donors' system is an added queue-server dyad. As the dynamics of the model are defined additively, doing so has no implications on the tractability of the model.

3. Numerical Studies

In this section, we present our numerical study, where we compare our model against benchmark policies in the context of Heger et al. (2020), which we shall refer to as the *reference literature* henceforth, and use their empirical results as the parameters for our model.

3.1. Context and Setup

In the reference literature, the authors conducted an empirical study, where in the first round, donors may sign up for a Registry and are informed that they would be contacted during times of blood shortages. In the subsequent round, the effect of participation in the Registry is tested in terms of their subsequent time-to-donation, based on whether the participants are exposed to a specific appeal for blood or not. The setting of their work gels with the setting of our model via a two-donor-class problem, which, as we have discussed in Section 2.7, does not lead to additional formulation complexities in our model. Here, donors who are in the Registry belong to one class, henceforth termed "high-responsive donors," and donors who are not in the Registry belong to another class, termed "low-responsive donors." Specifically, as the reference literature studies the donation likelihoods pre-incentivization and postincentivization of both classes of donors (as a function of time prior to donation or incentivization), their results naturally translate as the time-to-donation probabilities in our model, from which the probabilities $\omega^{t,s}$ and $\phi^{t,s}$ can be computed (Figure 8 in Online Appendix C.4). Other information, such as the relative proportions of the two classes of donors, can also be inferred from their study. For further details, please refer to Online Appendix C.2.

In this study, we consider the setting of a blood bank making donor-incentivization decisions under planning horizon $T = 5$, where one unit of time is three weeks. This choice is commensurate with the time interval between incentivization, which does not occur too frequently in practice, though the polynomial complexity of our model does allow us to model at finer time scales and retain tractability. Consequently, the safety observation period translates to $S_o = 4$ time periods (12 weeks), and the lifetime of blood packets is $S_e = 2$ time periods (6 weeks). The maximum donation time considered is $S = 12$ time periods (36 weeks), which is the definition used in the reference literature. In line with the model, if a donor has yet to donate within the 12 time periods since incentivization or eligibility, they are assumed to have dropped out of the donor pool.

The blood bank is assumed to possess information about the new donors inflow distribution Λ^t and the demand distribution D^t. These distributions pertain to the dynamics in the blood inventory and, thus, are not considered in the reference literature. As such, we had separately chosen these parameters. New donor inflow (set at mean 53) is chosen in such a way that the mean size of the donor pool remains roughly consistent over time. The mean demand was chosen to be slightly larger than the combined mean donation rate of new and repeat donors *without incentivization*, which forces the model to make incentivization decisions to avoid shortages. The decision maker is also assumed to know the initial states $x^{0,s}$, $y^{0,s}$, and $z^{0,s}$. The reference literature contains information on the proportion of the *total* donors in each class, but not its distribution across the time-in-node, s. To resolve this, we warm-started our simulations by running them over 100 time periods without incentivization and using their final states as the initial states for our simulations. This allows the system to arrive near the no-incentivization steady state, facilitating comparisons with the benchmark policies, which might perform poorly in transient regimes. This allows us to examine the reasons why the Pipeline Queues (PQ) model may result in a superior policy.

The goal of the blood bank is to determine the number of donors to be incentivized as a *decision rule* $\alpha^{t,s}$, for each of the classes of donors. They aim to keep the wastage low, while meeting demand. It is expected that the blood bank would implement only the decision at $t = 1$, allowing the new uncertainty to materialize, and then resolve the model thereafter, in a rolling-horizon manner. In this numerical study, the decision maker obtains the optimal incentivization decision rule via the proposed model, Problem (5). We shall not impose any capacity constraints on the servers, except only keeping the total number of incentivized donors under some budget, arising out of financial or capacity restrictions, of 80 donors per period. If the budget is set higher, forward planning is unnecessary; if it is set lower, the model might not be feasible in expectation. For any policy, if the intended number of incentivized donors exceeds 80, donors are incentivized on a FIFO basis, as donors with a longer time-since-last-donation have higher likelihood of dropping out.

We conduct numerical simulations on the performance of our model against benchmark policies (more details are in Online Appendix C.3). We run each simulation to 50 time periods, with the model solved on a rolling horizon basis. They are repeated a total of 160 times with the same initial condition, but different sample paths to compute the metrics at end of each time period: (i) occurrence of shortage or wastage and (ii) the degree of these shortages or wastage. Hence, results presented below will consist of the above metrics, averaged over $160 \times 50 = 8,000$ values.

3.2. Discussion of Benchmarks

For our simulations, we have chosen to benchmark against the threshold (TH) policy, which incentivizes donors whenever blood inventory falls below a particular threshold, Q. First, as explained in Section 1, this is the existing policy that blood banks often use today (Lee et al. 2008, Shi et al. 2014, Sun et al. 2016). Second, in many other inventory-control problems, threshold-type policies are known to be optimal under certain conditions. Third, to the best of our knowledge, we do not know of any other method in the literature that can be used to solve our model. The threshold policy is the closest adaptive policy that would constitute a good candidate.

Searching for an optimal TH policy over a general decision space would likely be computationally intensive. We focus on calibrating two parameters—namely, the threshold, Q (in step sizes of 10), and α (in step sizes of 0.1), which represents the proportion of eligible high-responsive donors incentivized. It is applied uniformly over all times-since-donation, s. Here, TH policy only incentivizes high-responsive donors because incentivization is triggered when the current inventory is low, and incentivizing low-responsive donors might not garner a sufficient response to address inventory shortages. We assume that the threshold policy, (Q, α), is held constant over time.

We iterate the TH policy over multiple values of Q and α and examine their performance on an efficiency frontier. We also identify two contender TH policies—SW policy (for "similar wastage") is chosen such that it has the lowest occurrence of shortage, given that it runs a similar occurrence of wastage as our Pipeline Queues policy; and SS policy (for "similar shortage"), which has the lowest occurrence of wastage, given that it runs a similar occurrence of shortage as our PQ policy.

The PQ policy was solved via interval bisection on the risk level k, the objective of our optimization model. The details can be found in Online Appendix C.1.

3.3. Performance Comparisons

We examine the performance of the various models under three demand patterns: (i) a time-homogeneous demand pattern; (ii) a periodic high- and low-demand pattern, alternating every three periods between a mean of 20 higher and then lower than the long-time average; and (iii) the same pattern as the first, except that with episodic instances of a large demand burst for a single period.

In Figure 2, we plot the efficiency frontiers of the TH policy (Q, α) in terms of the occurrences and magnitude of shortage and wastage, vis-à-vis the PQ policy, for all three different demand patterns. In Table 2, we zoom into the two benchmark cases selected based on their performance in the time-homogeneous demand setting (in Panel A). Panels B and C reflect their performance when the underlying demand is not time-homogeneous, but follows the other two patterns. Online Appendix C.4 provides the equivalent tables (Tables 7 and 8) for the cases where the two benchmark policies are selected based on their performance under the other demand patterns.

It is evident that the PQ policy lands squarely *outside* the efficiency frontier of the TH policy in all three cases—that is, the PQ policy systematically achieves a lower rate of shortage and wastage. Head-to-head comparisons against the benchmark TH policies indicate that this difference is significant. More precisely, the comparison against benchmark policies SW and SS shows that the TH policy needs to pay the price of at least a 16% increase in blood wastage or at least 32% as much blood shortage as wastage in order to match the PQ policy for the similar amounts of shortage and wastage, respectively, in the time-homogeneous case. For the other demand patterns, the difference is even more pronounced. Panels B and C indicate that the TH policies on the efficiency frontier for Panel A rapidly decay in performance once the demand pattern changes, indicating that there is no single stable threshold policy that would work well under an evolving environment. In comparison, the PQ policy will always be able to adapt to future demand information supplied to the model.

We attribute the effectiveness of the PQ policy to two reasons. First, the PQ policy is more general than the TH policy, as it solves a far greater number of decision variables, via an optimization model. The TH policy is similar to a PQ policy, where $\alpha^{t,s} = \alpha$ for all t, s when inventory is lower than Q; otherwise, zero. As such, the gains made by the PQ policy can be interpreted as the difference when we optimize for every $\alpha^{t,s}$. In contrast,

Table 2. Comparison of Performance for Best Candidate TH Policies under Time-Homogeneous Demand Patterns

Policies	Shortage occurrence (%)	Average shortage	Wastage occurrence (%)	Average wastage
		Panel A: Time-homogeneous demand		
SW	**0.73** (0.24)	**3.18** (1.08)	2.25 (0.24)	33.9 (1.12)
SS	**0.1** (0.08)	**0.46** (0.42)	**2.59** (0.25)	**36.3** (1.56)
PQ	0	0	2.23	33.9
		Panel B: Alternating periods		
SW	**0.83** (0.23)	**5.35** (2.11)	**4.36** (0.29)	**95.8** (1.97)
SS	**0.63** (0.17)	**3.49** (1.21)	4.84 (0.3)	98.5 (2.01)
PQ	0.04	0.04	5.1	99.7
		Panel C: Bursty demand		
SW	**3.56** (0.56)	**63.4** (10.7)	**2.26** (0.53)	**33.9** (3.98)
SS	**2.14** (0.4)	**31.8** (6.39)	**2.69** (0.54)	**37.3** (4.32)
PQ	0.31	3.24	4.48	49.3

Notes. The SW and SS policies used parameters $(Q = 100, \alpha = 0.2)$ and $(Q = 120, \alpha = 0.2)$, respectively. Bold indicates that PQ is significantly different from TH at significance level of 0.05, and the numbers in parentheses are half-widths of a 95% confidence interval for all measures.

Figure 2. (Color online) Comparison with Different Threshold Policies with Call Budget 80

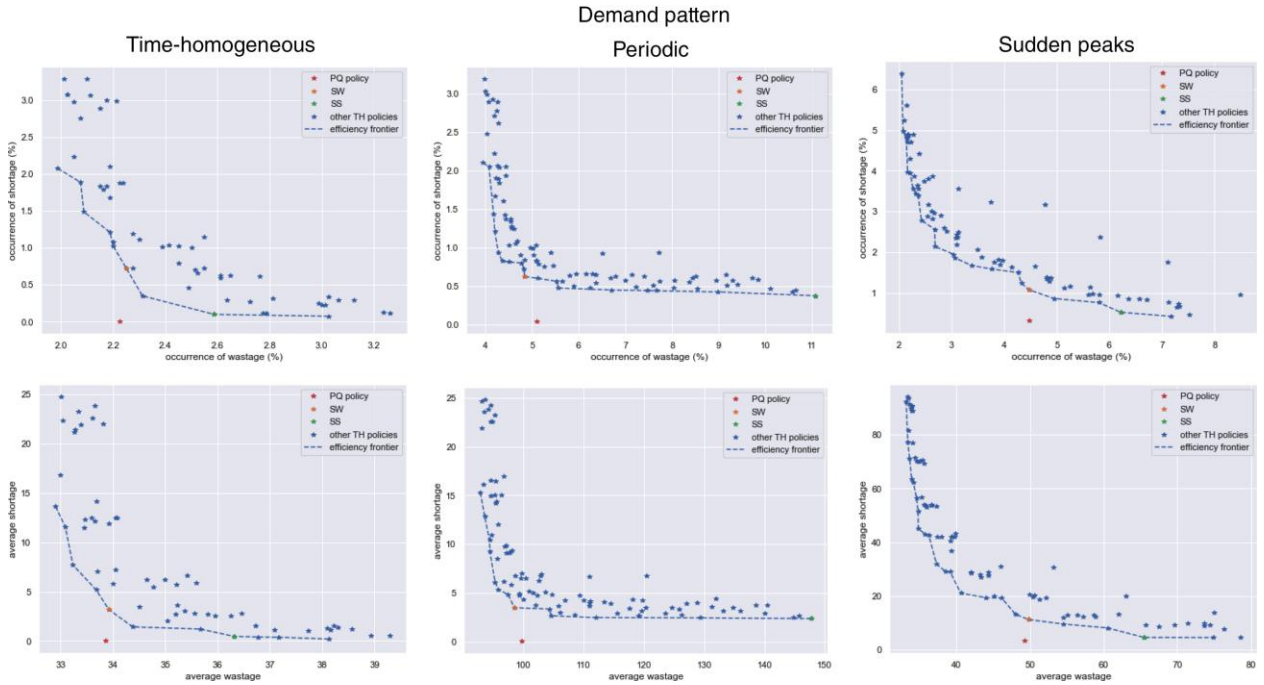

to the best of our knowledge, we do not know how to construct a formulation to solve for the optimal threshold-policy parameters using an optimization model that avoids the curse of dimensionality in the transient setting. Second, the TH policy instead only uses information about the current total inventory. In contrast, the PQ policy is privy to information at a more granular level of the distribution of time-to-expiry of the blood packets. Moreover, it also has the current time-since-donation information of its donors. Once again, it will be impractical to expect the TH policy to be extended to incorporate such information, as it rapidly increases the dimensions of the problem.

By incorporating future demand information, the PQ policy is forward-looking. A closer scrutiny across the different demand patterns in Figure 2 reveals that the scales of shortages and wastages increase drastically when the demand is no longer time-homogeneous—that is, the scale of improvement of the PQ policy grows more prominent. This difference arises because the TH policy is passive and reactionary—it incentivizes donors only when there is a shortage, which is likely a result of the onset of higher demand. The PQ policy instead is forward-looking and is able to make use of future information or anticipated changes in the inventory to decide on its incentivization strategy.

3.4. Policy Structure

In this subsection, we attempt to explain how the different policies behave in situ and relate these differences to policy performances. In general, however, it is difficult to make such analyses because the PQ policy is a function mapping from a large dimensional space to the decision space also of a large dimension. It is infeasible to analyze how each dimension contributes to the differences in policy performances. Instead, we decide to generate the eventual number of donors incentivized over each sample path and compare this single metric across different policies. For this particular analysis, we have chosen one particular sample path at random. We have checked and confirmed that the insights remain the same for other sample paths.

In Figure 3, we plot the number of incentivized donors over a particular sample path. Each chart in the figure refers to a different TH policy benchmark (SW and SS). Under the PQ policy, which incentivizes both high- and low-responsive donors, the number of high-responsive donors is plotted as "PQ high"; the total number of incentivized donors is plotted as "PQ policy"; the difference is the number of low-responsive donors incentivized. The sample path of the blood demand at every period is given in Figure 9 in Online Appendix C.4 to facilitate the discussion.

First, notice that the PQ policy starts incentivizing donors earlier and stops incentivizing donors later *vis-à-vis* the TH policy, but incentivizes donors in lower numbers at every one of these time points. We can see this as a sign that the PQ policy is attempting to smooth the demand fluctuations to reduce wastage. Moreover, the fact that the PQ policy starts incentivizing before the TH policy indicates that the former is forward-looking, which is the benefit when solving a multiperiod model, as opposed to the TH policy, which is passive and reactionary.

Figure 3. (Color online) Comparison of a Number of Incentivized Donors over Time between Different TH Policies SW (Top) and SS (Bottom) Against the PQ Policy in the Time-Homogeneous Case

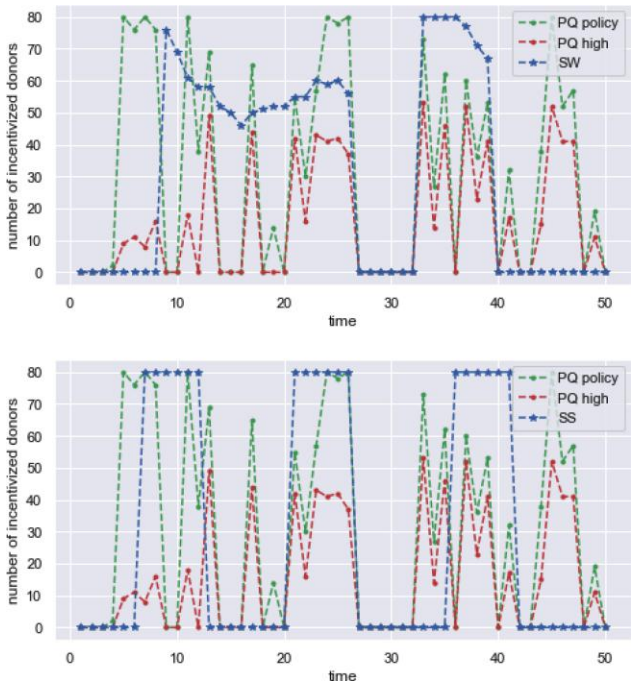

Note. The budget is 80; thus, incentivized donors are capped at 80 donors, even if the stipulated number is higher.

Secondly, the PQ policy incentivizes a mixture of low- and high-responsive donors. This proportion varies depending on the inventory state. At early times (first 15 time periods), the PQ model believes that the initial inventory is relatively full and conserves the high-responsive donor pool and starts by incentivizing the low-responsive donors. As the inventory levels begin to stabilize, the PQ policy starts to incentivize a higher proportion of high-responsive donors. In other words, the PQ policy is risk-pooling between incentivizing low-responsive donors in larger numbers and preserving high-responsive donors to tackle future urgent shortages. Also, high-responsive donors are more likely to donate blood of their own accord, and, thus, the relative gains from incentivizing them are lower in the longer term.

3.5. Analysis of the Average State

In this subsection, we compare the average inventory levels (Figure 4); the average number of donors in the observation queue—that is, noneligible donors (Figure 5); and the average number of eligible donors (Figure 6) under each policy for the time-homogeneous demand setting. Table 3 further shows the ratios between the number of noneligible donors and eligible donors. Here, the average states are computed by averaging the states

Figure 4. (Color online) Average State of Blood Inventory, x

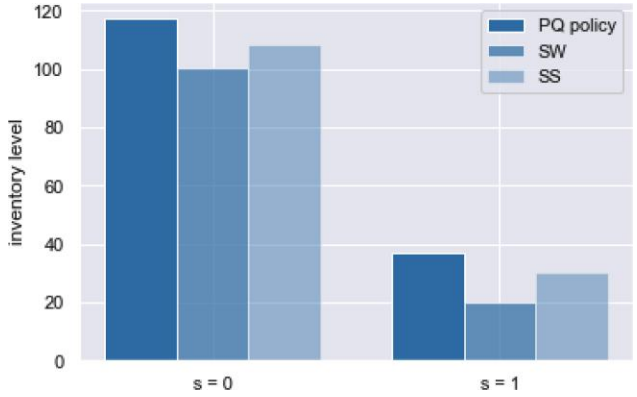

Note. $s = 1(s = 0)$ refers to "old (new) blood."

in the last 25 time periods of the simulations across all 80 iterations.

From Figure 4, we can see that the PQ policy maintains the highest amount of blood inventory (about 20% more than the next highest) compared with the TH policies. This is largely due to the PQ policy's forward-looking nature that it can maintain higher inventory levels without leading to higher wastage, even with a higher level of old blood. In contrast, the TH policy is reactionary and cannot actively increase the inventory until a shortage occurs. The PQ policy also maintains a larger number of active eligible donors (as in Figure 6(a)). This is coupled with the fact that it also has a lower proportion of high-responsive donors than the TH policy, as seen in Figure 6(b). This implies that the larger number of eligible donors arises under the PQ policy because of its ability to retain more low-responsive donors as active donors. It happens because incentivization serves as a reminder to the low-responsive donors to donate, which is part of the strategy of keeping them in the donor pool. Given that high-responsive donors naturally have a high chance of donating without incentivization, the PQ policy sees

Figure 5. (Color online) Average State of Observation Queue, z

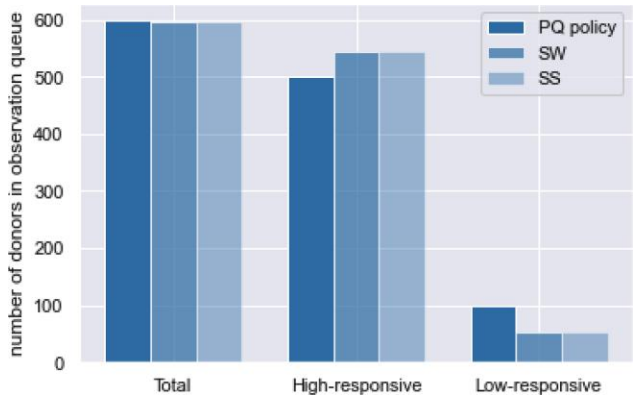

Figure 6. (Color online) Average State of Eligible Donors, y

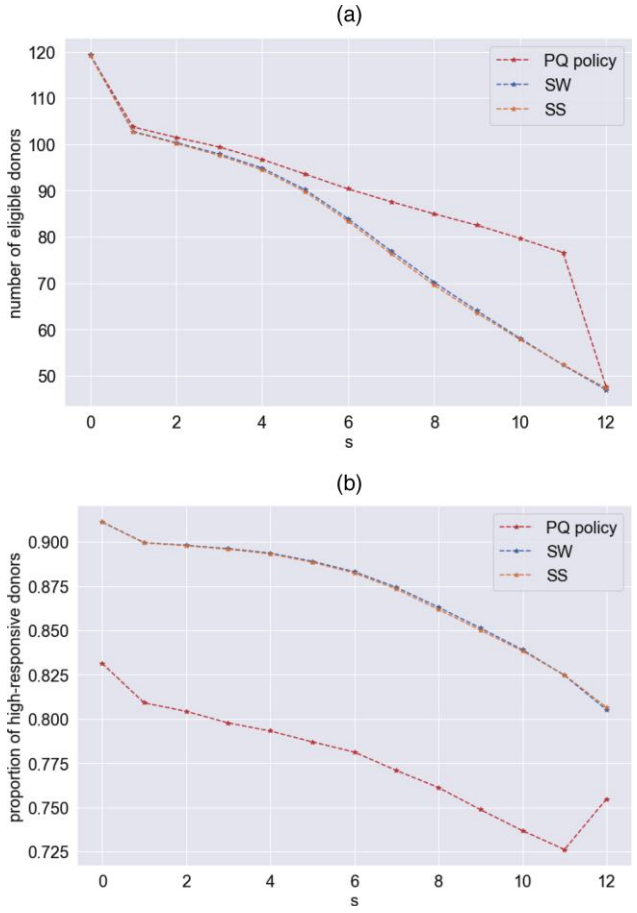

Notes. (a) Number of eligible donors. (b) Proportion of high-responsive donors.

fewer donor dropouts. This observation is also corroborated by Figure 5, which shows twice as many low-responsive donors in the observation queue than the TH policies. As entry to the observation queue directly corresponds to blood donation, the PQ policy receives twice as many donations from low-responsive donors as the TH policies. This cross-subsidizes the donations required from the high-responsive donors; consequently, more high-responsive donors are eligible under the PQ policy (see Table 3). Therefore, the PQ policy can better react to short-term shortages by having a larger group of eligible high-responsive donors.

These insights reveal reasons behind the superiority of the PQ policy and shed light on what constitutes a good inventory and incentivization policy. Namely, (i) maintaining a high level of inventory helps to fulfill demand, but entails better management of old blood to reduce wastage; (ii) incentivizing low-responsive donors during times of surpluses helps maintain a high level of active donors and reduces dropouts amongst low-responsive donors; and (iii) maintaining a minimum pool of eligible high-responsive donors grants greater capacity to react to prolonged shortages.

Table 3. Ratio of Noneligible to Eligible Donors—That Is, $\sum z / \sum y$

Policies	All	High-resp.	Low-resp.
PQ	0.514	0.550	0.386
SW	0.563	0.584	0.410
SS	0.566	0.587	0.410

Note. Resp., responsive.

4. Conclusion

In this paper, we proposed an optimization framework to solve the donor-incentivization policy in the blood-donor-management problem. Our framework is novel in that it simultaneously models the dynamics of both the blood inventory and the donor-flow process in a manner that can tractably solve the incentivization decisions. The numerical experiments demonstrate the advantages of our policy compared with structured policy benchmarks, such as the threshold policy, in reducing both shortages and wastage. The optimal policy gives rise to a smoother incentivization schedule that plans forward to avoid future starvation in case of demand surges. It also supports the implementation of a flexible and dynamic multiclass donor-incentivization policy. Moreover, our framework can be easily extended to practical situations, particularly in operationalizing an optimization decision-making support system for the proposed strategy of managing high- and low-responsive donors.

Work has begun in Singapore to classify blood donors into high- and low-responsive classes. It has opened the opportunity for blood-donor-management and is the basis for the application of our proposed framework. We hope to collaborate with the local authorities as part of our future work. Another area is to examine ways to establish the donation probabilities postincentivization. Future work can examine whether modern machine-learning methods can be applied for this purpose.

Acknowledgments

The authors thank Prof. Mor Armony, the associate editor, and the two anonymous reviewers for the insightful comments that have culminated in the current state of this paper.

Endnote

[1] We adopt the convention in the literature: $C_{k,\theta}[\tilde{\zeta}] = \text{ess sup}\{\tilde{\zeta}/\theta\}$ as $k \to 0$ and $\mathbb{E}[\tilde{\zeta}/\theta]$ as $k \to \infty$.

References

Afeche P, Araghi M, Baron O (2017) Customer acquisition, retention, and service access quality: Optimal advertising, capacity level, and capacity allocation. *Manufacturing Service Oper. Management* 19(4):674–691.

Aravindakshan A, Rubel O, Rutz O (2015) Managing blood donations with marketing. *Marketing Sci.* 34(2):269–280.
Aumann RJ, Serrano R (2008) An economic index of riskiness. *J. Polit. Econom.* 116(5):810–836.
Ayer T, Zhang C, Zeng C, White CC III, Joseph VR (2019) Analysis and improvement of blood collection operations. *Manufacturing Service Oper. Management* 21(1):29–46.
Bandi C, Loke GG (2018) Exploiting hidden convexity for optimal flow control in queueing networks. Preprint, submitted June 20, https://dx.doi.org/10.2139/ssrn.3190874.
Beliën J, Forcé H (2012) Supply chain management of blood products: A literature review. *Eur. J. Oper. Res.* 217(1):1–16.
Bergemann D, Brooks B, Morris S (2015) The limits of price discrimination. *Amer. Econom. Rev.* 105(3):921–957.
Bergemann D, Brooks B, Morris S (2017) First-price auctions with general information structures: Implications for bidding and revenue. *Econometrica* 85(1):107–143.
Brown DB, Sim M (2009) Satisficing measures for analysis of risky positions. *Management Sci.* 55(1):71–84.
Brown DB, De Giorgi E, Sim M (2012) Aspirational preferences and their representation by risk measures. *Management Sci.* 58(11):2095–2113.
Chen S, Li Y, Zhou W (2019) Joint decisions for blood collection and platelet inventory control. *Production Oper. Management* 28(7): 1674–1691.
Choudhary V, Shunko M, Netessine S, Koo S (2022) Nudging drivers to safety: Evidence from a field experiment. *Management Sci.* 68(6):4196–4214.
Dai JG, Shi P (2017) A two-time-scale approach to time-varying queues in hospital inpatient flow management. *Oper. Res.* 65(2):514–536.
Goh C-H, Greenberg BS, Matsuo H (1993) Two-stage perishable inventory models. *Management Sci.* 39(5):633–649.
Graves SC (1982) The application of queueing theory to continuous perishable inventory systems. *Management Sci.* 28(4):400–406.
Greinacher A, Fendrich K, Hoffmann W (2010) Demographic changes: The impact for safe blood supply. *Transfusion Med. Hemotherapy* 37(3):141–148.
Gunpinar S, Centeno G (2015) Stochastic integer programming models for reducing wastages and shortages of blood products at hospitals. *Comput. Oper. Res.* 54:129–141.
Health Sciences Authority (2019) Types of blood donations. Accessed January 6, 2021, https://www.hsa.gov.sg/blood-donation/types-of-blood-donations.
Health Sciences Authority (2020) The big blood picture 2020. Accessed January 6, 2021, https://www.hsa.gov.sg/docs/default-source/bsg/big-blood-picture-2020.pdf.
Heger SA, Slonim R, Garbarino E, Wang C, Waller D (2020) Redesigning the market for volunteers: A donor registry. *Management Sci.* 66(8):3528–3541.
Kamenica E, Gentzkow M (2011) Bayesian persuasion. *Amer. Econom. Rev.* 101(6):2590–2615.
Karaesmen IZ, Scheller-Wolf A, Deniz B (2011) Managing perishable and aging inventories: Review and future research directions. Kempf K, Keskinocak P, Uzsoy R, eds. *Planning Production and Inventories in the Extended Enterprise, International Series in Operations Research & Management Science*, vol. 151 (Springer, New York), 393–436.
Lacetera N, Macis M, Slonim R (2013) Economic rewards to motivate blood donations. *Science* 340(6135):927–928.
Lacetera N, Macis M, Slonim R (2014) Rewarding volunteers: A field experiment. *Management Sci.* 60(5):1107–1129.
Lee CK, Hong J, Hung ATF (2008) An update of blood donor recruitment and retention in Hong Kong. *Asian J. Transfusion Sci.* 2(2):47–50.
Lingenbrink D, Iyer K (2019) Optimal signaling mechanisms in unobservable queues. *Oper. Res.* 67(5):1397–1416.
Mugion RG, Pasca MG, Di Di Pietro L, Renzi MF (2021) Promoting the propensity for blood donation through the understanding of its determinants. *BMC Health Services Res.* 21(1):1–20.
Müller-Steinhardt M, Weidmann C, Klüter H (2017) Changes in the whole blood donor population in south-west Germany: 2010 vs. 2016. *Transfusion Med. Hemotherapy* 44(4):217–223.
Nahmias S (1976) Myopic approximations for the perishable inventory problem. *Management Sci.* 22(9):1002–1008.
Nahmias S, Pierskalla WP (1973) Optimal ordering policies for a product that perishes in two periods subject to stochastic demand. *Naval Res. Logist. Quart.* 20(2):207–229.
National Health Service (2020) Who can give blood. Accessed January 6, 2021, https://www.blood.co.uk/who-can-give-blood/.
Prastacos GP (1981) Allocation of a perishable product inventory. *Oper. Res.* 29(1):95–107.
Sarhangian V, Abouee-Mehrizi H, Baron O, Berman O (2018) Threshold-based allocation policies for inventory management of red blood cells. *Manufacturing Service Oper. Management* 20(2):347–362.
Satyavarapu A, Wagle D (2020) Improving the fragile US supply of blood. Accessed January 8, 2021, https://www.mckinsey.com/industries/public-and-social-sector/our-insights/improving-the-fragile-us-supply-of-blood.
Schweizer N, Szech N (2018) Optimal revelation of life-changing information. *Management Sci.* 64(11):5250–5262.
Shi L, Wang J, Liu Z, Stevens L, Sadler A, Ness P, Shan H (2014) Blood donor management in China. *Transfusion Med. Hemotherapy* 41(4):273–282.
Sun T, Gao G, Jin GZ (2019) Mobile messaging for offline group formation in prosocial activities: A large field experiment. *Management Sci.* 65(6):2717–2736.
Sun T, Lu SF, Jin GZ (2016) Solving shortage in a priceless market: Insights from blood donation. *J. Health Econom.* 48:149–165.
Tang Q, Zhang Y, Zhou M (2020) Robust vehicle repositioning with entropic risk measure. Preprint, submitted June 17, https://dx.doi.org/10.2139/ssrn.3612626.
World Health Organization (2017) The 2016 global status report on blood safety and availability. World Health Organization. Accessed January 8, 2021, https://apps.who.int/iris/handle/10665/254987.
World Health Organization (2020) Blood safety and availability. Accessed December 28, 2022, https://www.who.int/news-room/fact-sheets/detail/blood-safety-and-availability.
Zhou M, Loke GG, Bandi C, Zi QGL, Wang W (2022) Intraday scheduling with patient re-entries and variability in behaviours. *Manufacturing Service Oper. Management* 24(1):561–579.