DYNAMIC INVENTORY ALLOCATION IN HUMANITARIAN OPERATIONS WITH MULTI-CLASS DEMAND AND FUNDING UNCERTAINTY

CEM YARKIN YILDIZ

ÖZYEĞİN UNIVERSITY

AUGUST, 2025

DYNAMIC INVENTORY ALLOCATION IN HUMANITARIAN OPERATIONS WITH MULTI-CLASS DEMAND AND FUNDING UNCERTAINTY

by

Cem Yarkın Yıldız

Thesis
Submitted in Partial Fulfillment of the
Requirements for the Degree of

Master of Science

in
Industrial Engineering

Advisor: Prof. Ömer Erhun Kundakcıoğlu
Coadvisor: Asst. Prof. Mehmet Önal

Graduate School of Science and Engineering
Özyeğin University
İstanbul

August, 2025

DYNAMIC INVENTORY ALLOCATION IN HUMANITARIAN OPERATIONS WITH MULTI-CLASS DEMAND AND FUNDING UNCERTAINTY

Approved by:

Prof. Ömer Erhun Kundakcıoğlu,
Advisor
Department of Industrial Engineering
Özyeğin University

Assoc. Prof. Dilek Günneç Danış
Department of Industrial Engineering
Özyeğin University

Asst. Prof. Mehmet Önal,
Co-advisor
Department of Industrial Engineering
Özyeğin University

Prof. Ali Ekici
Department of Industrial Engineering
Özyeğin University

Prof. Raha Akhavan-Tabatabaei
Department of Operations Management
and Business Analytics
Sabancı University

Approval Date: August 4, 2025

To my family...

DECLARATION OF ORIGINALITY

I hereby declare that I am the sole author of this thesis and that this is the true copy of my thesis, including the final revisions, approved by my thesis committee. All data and information have been obtained, produced, and presented in accordance with the rules of research ethics and principles of academic honesty. As required by these rules, to the best of my knowledge I have acknowledged ideas, thoughts, and any copyrighted material in accordance with the standard referencing rules. I certify that any part of this thesis has not been submitted for a degree or diploma in another educational institution.

Cem Yarkın Yıldız

ABSTRACT

This study develops a dynamic inventory control framework for humanitarian aid allocation under multi-class demand and uncertain funding. We consider a central planner responsible for distributing a critical aid item to camp based and surrounding urban beneficiaries. Demand arises stochastically from two groups: internal (camp) demand, which incurs time sensitive deprivation costs when unmet, and external (urban) demand, which incurs fixed referral penalties when unsatisfied. Inventory replenishment is constrained by a known total funding amount that arrives in uncertain installments over a finite horizon. To approximate real time fulfillment, we employ fine grained intra period discretization, enabling dynamic rationing decisions at the micro period level. The problem is formulated as a finite horizon stochastic dynamic programming that integrates stochastic funding arrivals, continuous time demand realization, and threshold based rationing policies. We examine both linear and exponential time sensitive deprivation cost structures. Numerical experiments and structural analysis demonstrate the cost effectiveness and adaptability of dynamic policies under varying funding regimes. The proposed framework provides rigorous analytical foundations and actionable insights for resource constrained humanitarian supply chains.

Keywords: stochastic inventory management, dynamic inventory rationing, multiple demand classes, funding uncertainty, humanitarian operations

ÖZET

Bu çalışma, çoklu talep sınıfları ve belirsiz finansman koşulları altında insani yardım tahsisine yönelik dinamik bir envanter kontrol çerçevesi geliştirmektedir. Çalışmada, kamp içindeki ve çevredeki kentsel bölgelerde yaşayan hak sahiplerine kritik bir yardım kaleminin dağıtımından sorumlu merkezi bir planlayıcı ele alınmaktadır. Talep, iki gruptan rassal olarak oluşmaktadır: Karşılanmadığında zaman duyarlı mahrumiyet maliyetine yol açan iç (kamp) talep ve karşılanmadığında sabit yönlendirme cezaları doğuran dış (kentsel) talep. Envanter yenilemesi, toplam tutarı bilinen ancak taksitlerinin zamanlaması belirsiz olan bir finansman akışı ile sınırlandırılmıştır. Gerçek zamanlı talep karşılamayı modelleyebilmek için dönemler, ince zaman dilimlerine ayrıştırılmış ve mikro dönem düzeyinde dinamik paylaştırma kararları yapılabilir hâle getirilmiştir. Problem, rassal finansman gelişlerini, sürekli zamanlı talep akışlarını ve eşik tabanlı paylaştırma politikalarını entegre eden sonlu ufuklu bir rassal dinamik programlama modeli olarak formüle edilmiştir. Hem doğrusal hem de üstel biçimli zaman duyarlı mahrumiyet maliyet yapıları incelenmiştir. Sayısal deneyler ve yapısal analizler, farklı finansman rejimlerinde dinamik politikaların maliyet etkinliğini ve uyarlanabilirliğini ortaya koymaktadır. Önerilen çerçeve, kaynak kısıtlı insani tedarik zincirleri için güçlü analitik temeller ve uygulanabilir içgörüler sunmaktadır.

Anahtar Kelimeler: rassal envanter yönetimi, dinamik envanter tahsisi, çoklu talep sınıfları, finansman belirsizliği, insani yardım operasyonları

ACKNOWLEDGEMENTS

I would like to express my deepest gratitude to my advisor, Prof. Ö. Erhun Kundakcıoğlu, for his invaluable guidance and support throughout my journey. His mentorship and patience have been pivotal in shaping both my academic and personal development.

I am sincerely thankful to my co-advisor, Dr. Mehmet Önal, for his insightful suggestions, engaging discussions, and continuous support, all of which have significantly enriched this study and enabled me to complete this thesis.

I extend my sincere appreciation to my family for their unconditional love, patience, and encouragement. Their sacrifices and persistent belief in me have been the foundation of my academic and personal growth.

Above all, I am profoundly grateful to my wife, İzel, for her endless understanding, encouragement, and emotional support during the most challenging moments of this journey. Her belief in me has been my greatest source of motivation.

Finally, I gratefully acknowledge the support of the Scientific and Technological Research Council of Turkey (TÜBİTAK) under grant 122M839, whose generous funding and trust made this research possible.

TABLE OF CONTENTS

LIST OF TABLES

LIST OF FIGURES

LIST OF ACRONYMS AND ABBREVIATIONS

ADP Approximate Dynamic Programming

BL Back-Loaded Funding

CPU Central Processing Unit

DRL Deep Reinforcement Learning

ES Equal Sized Funding

FL Front-Loaded Funding

IASC Inter-Agency Standing Committee

IDP Internally Displaced Person

INFORM Index for Risk Management

MDP Markov Decision Process

NGO Non-Governmental Organization

OCHA United Nations Office for the Coordination of Humanitarian Affairs

OR Operations Research

PPO Proximal Policy Optimization

UNHCR United Nations High Commissioner for Refugees

1. INTRODUCTION

Natural and human-made disasters continue to cause widespread disruption, often leaving affected populations in urgent need of humanitarian assistance. In 2025, the United Nations Office for the Coordination of Humanitarian Affairs (OCHA) projected that over 305 million people worldwide would require such support, underscoring a sharp rise in global vulnerability [1]. According to the OCHA Global Humanitarian Overview Report, the primary driver of the increasing humanitarian demand is forced displacement, which has reached historically unprecedented levels due to ongoing conflicts and the intensifying impacts of climate change [1]. As of mid-2024, the United Nations High Commissioner for Refugees (UNHCR) reported 122.6 million forcibly displaced individuals globally—the highest number on record and the twelfth consecutive year of increase [2].

Displaced populations can remain within their countries as *internally displaced persons* (IDPs) or cross international borders and be classified as *refugees* [3]. Governments and humanitarian organizations establish formal and informal temporary settlements in response to the humanitarian crises, commonly referred to as *humanitarian camps*. Displaced individuals reside and receive humanitarian assistance in those camps. These camps are typically managed by a centralized authority (e.g., host country or lead humanitarian agency) to coordinate service delivery [3]. They provide essential assistance, including shelter, food, water, sanitation, and medical aid, and play a crucial role in meeting the immediate needs of displaced communities. Individuals residing in humanitarian camps are referred to as *camp* or *internal* beneficiaries. They are considered solely dependent on the aid provided by the camp management.

While many displaced persons reside in camps, a substantial share of displaced individuals are integrating into host communities or settling in urban areas. Displaced individuals living outside formal camps are referred to as the *urban* or *external* beneficiaries. In regions where humanitarian aid is provided, the number of urban beneficiaries can reach significant proportions. For instance, according to the mid-year report of UNHCR, urban refugees comprised approximately 78% of all refugees as of 2024. This figure illustrates the increasingly decentralized and heterogeneous nature of displacement patterns [2].

Urban beneficiaries also rely heavily on humanitarian aid, complicating resource distribution due to their geographic dispersion and lack of centralized coordination [4]. Given their dependency, all the needs of camp beneficiaries should be met. At the same time, they should also strive to meet the needs of the urban population. According to the Sphere Handbook, humanitarian assistance should be allocated only based on need, regardless of location [5]. This principle emphasizes the importance of serving both camp-based and urban beneficiaries, with aid prioritized according to vulnerability and dependence [6]. Since resources are limited and beneficiary needs differ, humanitarian organizations face difficult prioritization decisions. In the face of these uncertainties and resource scarcity, central authorities should adopt robust and efficient policies to ensure that the most vulnerable populations are reached.

Financial uncertainty and scarcity remain among the most significant challenges facing humanitarian operations. International donors, including governments, multilateral organizations, and non-governmental organizations (NGOs) provide funding for humanitarian activities. Humanitarian funding is often delayed, fragmented, or restricted by donor conditions. These limitations reduce operational flexibility and hinder effective procurement and planning [7, 1]. In 2024, only 40% of global humanitarian funding requirements were met, resulting in considerable service gaps and unmet needs [8]. The UNHCR's supply chain strategy for 2024–2030 aims to achieve a 50% reduction in excess inventory and a 20% decrease in inventory aging. This vision highlights the importance of efficient inventory management in uncertain and resource-constrained environments. These constraints and uncertainties emphasize the need for robust policies that effectively allocate humanitarian aid to both camp and urban beneficiaries.

This research is motivated by the above-mentioned uncertainties and challenges. While several studies explore inventory management in humanitarian settings, most adopt simplifying assumptions such as single-class demand, deterministic funding regimes, or end-of-period fulfillment structures. These simplifications do not fully reflect the dynamics observed in practice. Our model addresses these limitations by introducing a dynamic allocation framework with real-time rationing and stochastic funding. We consider an inventory planning problem in which a central planner distributes a single aid item to a population affected by a humanitarian crisis. Demand arises from two beneficiary groups: (i) the camp (internal) and (ii) the urban (external) population. These groups differ in both urgency and the consequences of unmet demand. For camp residents, unsatisfied demand is assumed to be backordered and penalized through a time-sensitive *deprivation cost* that increases with delay. In contrast, unmet urban demand is treated as lost sales, incurring a fixed *referral cost* that reflects the cost of redirecting or compensating for unmet needs outside the primary distribution channel.

To allocate aid across beneficiary groups, the camp manager adopts the following policy: demand from internal residents is immediately satisfied if the aid item is available in stock. Otherwise, it is backlogged. The cost of backlogging a unit demand may be a convex function of the time during which it is backlogged. On the other hand, external demand may be deliberately unmet, even when inventory is available.

In multi-class demand systems, the deliberate deferral of fulfilling certain demand classes, in order to preserve inventory for anticipated future demands from higher-priority classes, is known as a *rationing policy*. The camp manager implements such a rationing policy to ensure that there is sufficient inventory for potential future demands from internal residents. In such a case, the urban resident is referred outside the camp, and a fixed penalty is incurred to account for unmet demand. We focus on an *immediate fulfillment* setting, where demand must be satisfied or rejected as it arises, to capture the operational realities of humanitarian logistics. This reflects the urgency-driven nature of aid distribution and the impracticality of periodic satisfaction.

Inventory replenishment is constrained by limited donor funding, which arrives over a finite planning horizon. Humanitarian financing is often disbursed in uncertain installments, both in terms of timing and magnitude, unlike deterministic or fully observable funding regimes. This stochasticity complicates planning, as future budget availability cannot be precisely anticipated at the time of decision-making.

Contributions. This study contributes to the humanitarian inventory control literature by addressing key structural and operational uncertainties that remain underexplored in prior work. Our key contributions are:

- We formulate a finite-horizon inventory allocation model for humanitarian operations with two demand classes (internal backorders, external lost sales), incorporating stochastic funding installments and real-time fulfillment requirements.

- We introduce a discrete-time dynamic programming framework with intra-period discretization to approximate real-time responsiveness and immediate fulfillment.

- We show structural properties that the optimal replenishment policy is *modified base stock* policy, and the rationing policy follows *threshold* type structure for all model variants.

- We benchmark the proposed policy under deterministic and stochastic funding, and provide managerial insights on aid prioritization and different funding installment regimes.

The remainder of the paper is organized as follows. Literature Review reviews the relevant literature, emphasizing the unique aspects of the problem, related inventory control studies, and their applications in humanitarian contexts. Section 3 offers a detailed problem statement, followed by the mathematical developments and solution approach. In Section 3, we provide useful structural properties of the problem for replenishment and rationing decisions. Section 4 presents computational experiments and discusses managerial insights. Finally, Section 5 summarizes the study and outlines directions for future research.

2. LITERATURE REVIEW

The complexities that arise in the aftermath of disasters are classified as *post-disaster* scenarios in the humanitarian supply chain literature [9]. We consider a *complex emergency* setting in which the immediate disaster phase has passed, but prolonged humanitarian needs persist. Such crises often last over extended periods, requiring sustained humanitarian assistance. In the context of humanitarian aid allocation, we examine a multi-period inventory control problem characterized by stochastic demand. The system has two demand classes, incorporating a mix of backorders and lost sales, subject to uncertain funding disbursement schedules. While prior studies address some of these components individually, none integrates them within a unified modeling framework. To situate our contribution within the existing literature, we review prior work across three critical dimensions relevant to our setting: financial uncertainty, inventory rationing under mixed demand classes, and operational constraints in humanitarian logistics.

Financial constraints in inventory control are commonly modeled as either revenue-dependent capital or stockpoint capacity limits. Only a few studies capture financial uncertainty in the form of stochastic funding amounts and uncertain disbursement schedules [10]. One stream of humanitarian research focuses on developing policies for organizations to allocate their budgets effectively. For instance, [11] examine the distribution of nonprofit funding among program responses, fundraising, and administration. [12] explore budgeting decisions across different humanitarian programs. However, this stream primarily takes a high-level perspective, overlooking program-level decision-making. Our work complements this line of research by focusing on program-level decisions and their implications. Additionally, our findings offer insights that can inform higher-level budget allocation strategies, ultimately improving overall financial planning in humanitarian operations. For a broader discussion of funding structures in humanitarian settings, see [13].

[14] study a self-financing retailer who periodically replenishes stock from a supplier and sells it to the market. Replenishment decisions are constrained by cash flow, which is updated each period based on purchases and sales, with unmet demand treated as lost sales. The authors characterize the optimal policy as a capital-dependent base-stock policy, where the budget evolves endogenously through sales in a profit-driven setting. In contrast, [15] examine a single-item, finite-horizon inventory problem with a fixed, externally provided budget and lost sales. They formulate a two-stage decision problem: first, how to allocate the budget between stockpiling and shipment costs; and second, given the allocated shipping budget, how to dispatch relief items efficiently from the stockpile to the field. Using dynamic programming, they derive a replenishment policy that follows a (Q, R) structure.

[10] examine a stochastic funding structure in a humanitarian setting where there is no revenue generation. Authors develop a finite-horizon, periodic-review inventory model for a single product to analyze the effects of funding uncertainty, assuming all unmet demand is backlogged. [16] extend the model to a multi-treatment inventory allocation problem, introducing two demand states—less severe and more severe—where probabilistic transitions depend on allocation levels. While [10] consider only a single demand class, [16] incorporate two demand classes with state transitions between backlogged classes. Both models operate in a periodic-review setting and incorporate a general funding framework in which both the amount and timing of funding are stochastic. The authors show that the optimal replenishment policy follows a state-independent, modified base-stock structure.

In our setting, each humanitarian camp faces two demand classes: high-priority internal and lower-priority external demand. Since internal populations rely solely on camp management, their demand is backordered, whereas external demand is considered lost sales. Given this structure, developing an inventory-sharing mechanism that prioritizes internal demand while accounting for stochastic funding and future demand fluctuations is essential [17]. A trade-off arises between meeting external demand and maintaining sufficient inventory for future periods. Fulfilling external demand instead of holding inventory may reduce immediate costs. A naive policy prioritizes backorders, then internal, and finally external demands. It minimizes short-term costs and is optimal in scenarios without supply constraints, as discussed by [18]. However, under stochastic funding, where both timing and amount are uncertain, retaining inventory rather than allocating it immediately leads to better long-term outcomes. Given limited funding and unpredictable schedules, designing an efficient inventory-sharing mechanism becomes crucial.

Inventory rationing policies are employed in multi-class demand systems to allocate limited inventory efficiently. These policies are typically categorized based on the treatment of unmet demand (lost sales, backorders, or mixed scenarios), the review type (periodic or continuous), the number of demand classes (two or general N), and whether the rationing policy is static or dynamic between two replenishments. Our study focuses on the periodic inventory review problem under a mixed backorder and lost sales. Accordingly, we highlight a subset of studies closely related to our setting. For broader overviews of inventory optimization in multi-class demand systems, we refer readers to the comprehensive surveys by [18, 19, 20].

One widely studied approach in the literature is the critical-level policy [21]. In this policy structure, when on-hand inventory falls below a predefined threshold, high-priority demands are satisfied, while lower-priority demands are not. This approach effectively balances costs in multi-class demand systems, particularly when demand uncertainty and rationing decisions critically affect performance. [21] present an infinite-horizon periodic inventory control case, where there are two backordered demand classes and a deterministic lead time. They analyze the steady-state expected cost under base-stock and static critical-level policies. [22] propose a threshold-based base stock inventory policy for two backordered demand classes. The authors model the problem as a one-stage (S, T) model with regenerative cycles. In that model, demand is satisfied upon arrival, which plays a crucial role in cost derivation.

Two studies directly examine the impact of a mix of backorders and lost sales under a periodic inventory control policy. [23] examine an inventory system serving two demand classes, where the high-priority class is subject to lost sales and the low-priority class is subject to backorders. The core assumptions of the study are unlimited replenishment amounts and no on-hand inventory when backorders are available. Thus, this setting does not provide explicit threshold decisions. The authors characterize the structure of the optimal replenishment policy, proving that it follows a base-stock policy for nonnegative inventory levels and a non-decreasing order quantity for negative inventory levels. [24] study an inventory system with a mix of backorder and lost-sales demands, where priorities change dynamically. In simple terms, they discretize the period into N small intervals and solve the Markov Decision Process (MDP) for the period. They propose a dynamic rationing policy for the system, provide structural insights for the problem, and a heuristic policy for multi-period cases to mitigate the computational burden.

[25] propose a (Q, R) inventory model for a multi-supplier scenario, incorporating emergency and regular orders. They provide analytical insights and apply the Silver-Meal heuristic, demonstrating its applicability through a real-world case study in [26]. [27] investigate the allocation of aid to refugee camps while accounting for two distinct demand classes. Their study considers uncertainty arising from demand fluctuations and variability in replenishment cycle lengths. They propose a sharing threshold mechanism for external demand to optimize resource allocation. While conceptually related, their model differs from ours in several important respects. [27] model funding uncertainty through

uncertain replenishment cycles, implicitly assuming that funding amounts are sufficiently large to accommodate any allocation decision. In contrast, we consider a fixed funding amount with a stochastic funding schedule. In each period, the central decision-maker must decide whether to allocate the available budget based on the system state. Although [27] provide valuable managerial insights, their analysis does not explicitly model the operational effects of funding uncertainty. Furthermore, we employ a periodic inventory control framework, which aligns with the budgeting practices of humanitarian organizations, as financial commitments are typically made within a fiscal year and disbursed in tranches.

The primary distinction between humanitarian and commercial supply chains lies in their objective functions [28]. While commercial supply chains are profit-driven, the primary goal of humanitarian supply chains is to maximize social welfare [29]. Typical objectives in humanitarian operations include increasing demand coverage, reducing response time, minimizing shortages, preventing surplus inventory accumulation, and optimizing logistics costs [9]. [30] provide a comprehensive discussion on performance measurement in humanitarian relief chains. Beyond responsiveness and operational effectiveness, the timing of aid delivery is crucial. In this context, a more representative metric, deprivation, models the penalty for unmet demand non-linearly, explicitly accounting for the timeliness of demand satisfaction [6]. Unlike traditional approaches that assess unmet demand solely based on quantity over time, deprivation-based modeling incorporates waiting time, offering a more realistic representation of human suffering.

Despite extensive research on inventory control under financial constraints, existing studies primarily focus on commercial settings. They assume deterministic or capacity-driven budgets, or consider single-demand-class models. Our study advances the literature by introducing a multi-demand-class inventory control model within a humanitarian setting, where funding is stochastic. Unlike prior work that assumes deterministic replenishment schedules or fixed budgets, we explicitly model uncertainty in both funding amounts and timing. Furthermore, while previous studies have treated unmet demand as either entirely backordered or lost, this study incorporates a mixed demand structure, integrating rationing policies that dynamically prioritize demand classes. Although some similar discussions regarding demand satisfaction at the end of the period exist in the literature, no study assumes immediate satisfaction within a periodic inventory control model. Introducing a deprivation-based objective function further distinguishes this study, as it captures the human impact of delayed aid distribution—an essential yet often overlooked aspect of humanitarian operations. By bridging gaps in financial modeling, inventory control, and humanitarian logistics, this research lays the foundation for more effective decision-making in resource-constrained humanitarian operations. To the best of our knowledge, no existing study integrates immediate fulfillment with mixed backorder and lost sales demand, stochastic funding arrivals, and a deprivation-based objective in a periodic inventory control setting.

3. PROBLEM STATEMENT AND MATHEMATICAL MODEL

3.1 Problem Statement

We study a dynamic inventory allocation problem faced by a central planner responsible for distributing limited medical aid to a humanitarian camp and its surrounding urban population. Humanitarian supply chains in camp-based operations typically follow a centralized, top-down structure in which a central authority oversees the procurement, stockpiling, and distribution of essential aid [31]. This model aligns with widely adopted practices in humanitarian logistics and camp management [32]. Aid delivered to camps typically includes items such as blankets, clothing, shelter, and medical supplies. We focus on medical aid, as other essential goods are often locally sourced or donated in kind.

The objective is to dynamically allocate limited resources between camp and urban beneficiaries to minimize total expected cost. The cost structure includes inventory holding costs, time-sensitive deprivation penalties for fully backlogged internal demand, and fixed referral penalties for unmet external demand. The decision-maker dynamically adjusts replenishment and rationing policies in response to stochastic demand, uncertain funding, and class-specific consequences of unmet demand.

Population sizes for both groups are assumed to be known in each period. Demand is stochastic, driven by individual-level probabilities influenced by aid dependency [25, 33]. We adopt a probabilistic modeling approach due to the scarcity of historical data and the unpredictability of crises [34]. When individual probabilities are small and populations are large, demand can be approximated by a Poisson distribution, with the rate equal to the product of population size and per-capita demand probability. Each beneficiary requests at most one unit per period, and demands are independent across periods.

Camp residents rely solely on humanitarian aid, while urban beneficiaries may have access to local health systems. Ideally, all demands should be fulfilled, but limited inventory necessitates prioritization. Unmet internal demand is fully backordered and penalized using time-sensitive deprivation costs [6], while unmet external demand is treated as lost sales and penalized via fixed referral costs. Excess inventory presents an additional opportunity cost and results in holding costs. The UNHCR Supply Strategy emphasizes reducing excess inventory and minimizing inventory aging [35], highlighting the need for both service-level and financial optimization.

The planner faces two decisions: (i) replenishment at the beginning of each period, subject to available funding, and (ii) rationing of on-hand inventory throughout the period. Inventory is always shared with camp residents. Each external demand is either satisfied or rejected upon arrival. This immediate-fulfillment structure increases responsiveness and supports operational realism. To manage these trade-offs, we propose a dynamic threshold-based rationing policy to determine how much inventory to allocate to urban demand versus reserve for future camp demand. The main challenge is to balance these allocations in the face of both demand and funding uncertainty.

Humanitarian organizations rely on funding from diverse sources—governments, NGOs, and private donors. Funding is typically pledged on an annual basis according to global needs assessments (e.g., the INFORM Severity Index), internal priorities, and donor expectations [12]. While the total committed amount is known, the disbursement schedule is stochastic, with installments arriving irregularly and often with delay [36, 37]. In our setting, the timing of funding tranches is stochastic, complicating long-term planning and equitable resource allocation.

To address these challenges, we formulate a dynamic inventory management framework that integrates the following components:

- ***Demand modeling:*** Beneficiary demand follows Poisson processes, driven by per-capita aid dependency and known population sizes.

- **Inventory rationing:** A dynamic sharing threshold policy adjusts allocations between internal and external demand classes in near real time, ensuring fairness and responsiveness.

- **Stochastic funding:** While funding vector (i.e., installments) is predetermined, its disbursement over time is stochastic. Replenishment decisions must respect the realized budget at each period.

- **Cost structure:** The model integrates holding costs, referral penalties for unmet external demand, and time-sensitive deprivation costs—modeled in both linear and exponential forms—for backordered internal demand.

This problem setting motivates the development of a stochastic dynamic programming model that jointly optimizes replenishment and rationing decisions across a finite planning horizon to ensure equitable and timely aid delivery.

3.2 Mathematical Model

The model captures key features of humanitarian supply chains, including multi-class stochastic demand, time-sensitive deprivation costs, and uncertain funding arrivals. We formulate the problem as a finite-horizon, discrete-time stochastic dynamic program. At the beginning of each period, a replenishment decision is made based on the available funding. Demand arises continuously within each period and must be fulfilled or rejected immediately upon arrival. To approximate real-time fulfillment within a discrete-time framework, we divide each period into fine-grained micro intervals, enabling immediate rationing decisions. We begin by outlining the system environment and decision timeline. A formal definition of the state space, decisions, transition dynamics, and cost structure for the base case follows this. We conclude the section by introducing structural variants that differ in cost formulation and funding dynamics.

3.2.1 Model Overview

We consider a finite-horizon inventory allocation problem in which a central planner distributes a single aid item to camp-based and urban populations. The goal is to allocate limited inventory in response to stochastic demand while minimizing the total expected cost. The cost function includes three components: (i) inventory holding costs, (ii) time-sensitive deprivation penalties for unmet internal (camp) demand, and (iii) fixed referral penalties for unmet external (urban) demand.

The model integrates two types of decisions: (i) replenishment decisions made periodically at the beginning of each period, subject to available funding, and (ii) real-time rationing decisions made at the arrival of each demand. This framework captures the operational complexity introduced by donor funding that arrives intermittently in known amounts and at unpredictable times.

Table 3.1: *Model sets.*

Notation	Description
H	Number of macro periods (replenishment periods)
N	Number of micro periods within each macro period
T	Total number of periods, defined as $T = HN + 1$
$\mathcal{S} = \{1, N+1, \ldots, N(H-1)+1\}$	Set of macro decision periods
$\mathcal{T} = \{1, 2, 3, \ldots, T\}$	Set of all micro periods (global time index)

The planning horizon is discretized into H macro periods, each comprising N micro periods of equal length as proposed in [24]. Macro periods represent the main decision epochs during which funding arrives (if any) and replenishment decisions are made. Let $\mathcal{S} \subset \mathcal{T}$ denote the beginning of each macro period, where $\mathcal{T}$ is the overall discrete time index. In between each macro period, demand arrives in continuous time but is approximated by a fine-grained discretization into N micro periods. This intra-period structure enables near-instantaneous rationing of inventory at each time step.

Let λ^I and λ^E denote the macro period Poisson arrival rates of internal and external demand, respectively, and define the total rate as $\lambda = \lambda^I + \lambda^E$ per macro period. The arrival rate per period becomes λ/N by discretizing each macro period into N micro steps of length $1/N$. Under the Poisson process assumption, the probability of observing more than one demand arrival in a single micro period is given by:

$$\mathbb{P}(D \geq 2) = 1 - \left(1 + \frac{\lambda}{N}\right) e^{-\frac{\lambda}{N}} \tag{3.1}$$

This discretization approximation ensures that at most one demand unit arrives per micro period, allowing the system to be modeled as a discrete-time Markov Decision Process (MDP). Each micro period serves as a decision epoch in which the system decides whether to share inventory in response to a single demand realization. To simplify the notation, we define the function $\eta(t)$ to map each micro period $t \in \mathcal{T}$ to its corresponding macro period index. This function identifies period t belongs to which macro period:

$$\eta(t) = \left\lceil \frac{t}{N} \right\rceil \tag{3.2}$$

Table 3.2: *Funding and decision variables.*

Notation	Description
Funding Variables	
K	Total number of installments
$\mathcal{K}$	Set of installments $\{1, 2, \ldots, K\}$
z_k	Funding amount of k^{th} installment $k \in \mathcal{K}$
F	Total committed funding $F = \sum_{k \in \mathcal{K}} z_k$
Decision Variables	
y_t	Replenishment quantity ordered at macro period $t \in \mathcal{S}$
a_t	Binary decision to share inventory with external demand at period $t \in \mathcal{T}$

Events and Decisions. The nomenclature for state, funding, and decision variables is defined in Table 3.2. At each period $t \in \mathcal{T}$, the following sequence of events and decisions occurs:

1. If $t \in \mathcal{S}$, corresponding funding installment is received (if it arrives), the available budget r_t is updated, and the number of remaining installments is decremented (if necessary);
2. If $t \in \mathcal{S}$, a replenishment quantity y_t is selected, subject to the current budget r_t, and the remaining budget is updated;
3. An inventory-sharing decision a_t is made, determining whether to allocate inventory to external demand;
4. A demand unit is realized, and the system state transitions accordingly.

The key problem parameters are summarized in Table 3.3. Replenishment decisions are made at macro periods based on the current system state under the zero lead time assumption. A limited budget constraint is explicitly incorporated, requiring that sufficient funds be available to execute any replenishment action, unlike most existing models, which often ignore this operational restriction. Following the replenishment decision, an inventory rationing decision is made, which is also dependent on the current state of the system.

We adopt a slightly modified funding framework inspired by [10, 16]. Funding is received in multiple installments over the planning horizon. Let $\mathbf{z} = (z_1, z_2, \ldots, z_K)$ denote the funding vector, where z_k is the amount of the k^{th} installment. Installments are received at the beginning of a macro period. Equivalently, funding becomes available at the end of the last micro period of the preceding macro interval—i.e., at period $t \notin \mathcal{S}$ such that $t + 1 \in \mathcal{S}$. This convention ensures that funding is fully observable prior to the replenishment decision. The total committed funding $F = \sum_{k \in \mathcal{K}} z_k$, the installment amounts, and the number of installments K are assumed to be known in advance. However, the arrival periods may vary under stochastic funding. This formulation also permits modeling deterministic funding schedules as a special case where $K = H$. For simplicity, we assume that one unit of funding corresponds to the procurement of one unit of inventory.

Table 3.3: *Problem parameters.*

Notation	Description
Demand Parameters	
λ^I	Poisson arrival rate of internal (camp resident) demand per macro period
λ^E	Poisson arrival rate of external (urban population) demand per macro period
p	Probability of a demand arrival in a micro period (λ/N)
p^I	Probability of internal demand arrives
p^E	Probability of external demand arrives
Cost Parameters	
c^E	Fixed referral penalty per unit of unmet external demand
c^H	Inventory holding cost per unit per macro period

Demand and Cost. We consider three possible demand realizations: (i) a single internal (camp resident) demand, (ii) a single external (urban) demand, or (iii) no demand at every period $t \in \mathcal{T}$ (including macro and micro periods). The probabilities of observing exactly one unit of demand are defined as follows:

$$p = 1 - e^{-\frac{\lambda}{N}}. \tag{3.3}$$

This formulation does not account for the negligible probability of multiple arrivals since fine-grained discretization ensures that such events have negligible probability mass and are thus omitted without loss of fidelity. Conditional on a demand arrival, the probability of it originating from internal (p^I) and external (p^E) demand is proportional to its share of the total rate:

$$p^I = \frac{\lambda^I}{\lambda} p, \quad p^E = \frac{\lambda^E}{\lambda} p. \tag{3.4}$$

To reflect the temporal granularity of inventory transactions, the per micro period holding cost is scaled as follows:

$$\tilde{c}^H = \frac{c^H}{N} \tag{3.5}$$

where c^H denotes the holding cost per macro period, and $\tilde{c}^H$ represents the effective holding cost incurred at each period. The system applies holding costs at the micro period level to penalize inventory keeping more precisely over time. In contrast, the referral cost c^E is not scaled, as it represents a one-time fixed penalty for unmet external demand, independent of the time step at which it occurs. The deprivation cost represents the time-sensitive penalty incurred between the arrival time $t \in \mathcal{T}$ of an internal demand and its satisfaction macro period $t \in \mathcal{S}$. Since internal demand is immediately satisfied whenever inventory is available, deprivation costs are incurred only by backlogged internal demands. The penalization is time-dependent and may take either a linear or exponential form. We provide detailed expressions and necessary modeling assumptions for both cases in the corresponding sections.

System State, Transition Dynamics, and Decisions. Let the system state at period t be denoted by $s_t = (x_t, r_t)$. This state representation serves as the common foundation across all model variants. Each formulation builds upon this base by incorporating additional logic into the transition dynamics and cost evaluation, depending on the assumed

Table 3.4: *State variables.*

Notation	Description
s_t	Full system state at time $t \in \mathcal{T}$
x_t	Net inventory level available at the beginning of period $t \in \mathcal{T}$
r_t	Available funding at the beginning of period $t \in \mathcal{T}$
$V_t(s_t)$	Minimum expected cost (i.e., cost-to-go or value function) at period $t \in \mathcal{T}$

deprivation cost structure and funding regime. x_t shows the net inventory level. We know that $x_t > 0$, there exists an inventory, and when $x_t < 0$, it represents the number of backlogged internal demands (i.e., the deprived population), and r_t is the available budget at the beginning of the period. The value function $V_t(s_t)$ denotes the minimum expected cost from period t onward, given the current state s_t.

At the end of the last micro period $(t-1) \notin \mathcal{S}$ of each macro period $t \in \mathcal{S}$, the funding installment is realized. This amount becomes available at the start of macro period t, where the available budget r_t is updated accordingly. Based on the observed state at the beginning of the first micro period of macro period t, the decision-maker selects the replenishment amount y_t. Replenishment decisions are made only at macro period boundaries. The demand of the deprived population is satisfied only at macro period boundaries. After the replenishment decision y_t is made, the available inventory is first allocated to meet the unmet internal demand. Due to the unconditional prioritization of internal demand, inventory is always first allocated to satisfy any backlogged demand. This immediate fulfillment structure of the model does not allow for the coexistence of a deprived individual and an inventory. In contrast, inventory-sharing decisions are made at each micro period, based on the observed system state and realized demand. We normalize the funding such that one unit of funding enables the procurement of one unit of aid.

In each micro period, fully dynamic rationing policies with no restrictions can decide whether to allocate resources to internal and external demand. This flexibility enables the decision-maker to adjust the rationing strategy at each micro step to achieve cost-optimal outcomes. However, such highly responsive policies may be impractical or perceived as unfair in real-world humanitarian settings, where stability and need-based assessment are also important. To enhance fairness and operational interpretability, we adopt a modified policy structure: internal demand is always prioritized and fulfilled unconditionally whenever inventory is available, without any further rationing decisions. This policy structure eliminates the need for explicit allocation decisions for internal requests, ensuring that all available inventory is automatically directed to the deprived population when possible, which guarantees consistent and priority-based fulfillment for those in greatest need. This simplification may result in slightly higher expected costs compared to fully flexible policies. However, it reflects priorities in humanitarian operations and yields a more stable and need-based policy.

In the model formulation, we consider three distinct scenarios based on the nature of demand. In the no-demand case, the system evolves forward, incurring immediate costs. When internal demand arises, it must be fulfilled immediately if inventory is available. We need to decrease the inventory position by one when internal demand arrives. For external demand, the decision-maker faces a choice: either fulfill the request using available inventory or reject it, incurring a referral cost. These decision points are only meaningful when $x_t \geq 1$; otherwise, the system has no discretion. Internal demand results in increased deprivation, and external demand leads to a referral cost. If inventory is available and the sharing condition defined in (3.7) holds, the decision-maker may allocate inventory to meet demand. To capture this dependency, we define the binary *sharing decision* variable $a_t \in \mathcal{A}(x_t)$, where:

$$\mathcal{A}(x_t) = \begin{cases} \{0,1\} & \text{if } x_t \geq 1 \\ \{0\} & \text{if } x_t \leq 0 \end{cases} \tag{3.6}$$

The inequality in (3.7) formalizes the logic behind the sharing decision. It states that the cost savings from retaining one additional unit of inventory are less than the combined cost of referring the demand externally and holding that unit. In other words, if it is optimal to share one unit of aid, then the inequality in (3.7) must hold:

$$V_{t+1}(x_t - 1, r_t) - V_{t+1}(x_t, r_t) \leq c^E + \tilde{c}^H \tag{3.7}$$

In the following sections, we introduce variant formulations based on the deprivation cost structure and funding regime. In the case of linear deprivation, the deprivation cost is time-sensitive and penalizes time linearly. Exponential cases highlight the importance of waiting time, as it is a function of time. We assume that the funding vector is known in advance for each variant. Meaning that, installment amounts are known; however, the exact timing of their arrivals is uncertain. In the deterministic case, we certainly know the exact arrival period of each funding installment.

3.2.2 Linear Deprivation

In the linear time-sensitive deprivation case, the cost of deprivation accumulates uniformly over time. That is, the penalty incurred per unit time is constant across all micro periods. This structure implies a *memoryless* cost function, meaning that the system does not need to track the exact arrival times of individual backlogged demands. Let $\tilde{c}^I$ denote the per-period deprivation cost rate. Then, for a backlogged internal demand arriving at micro period $t \in \mathcal{T}$ and satisfied at macro period $\tau \in \mathcal{S}$, the total deprivation cost incurred is:

$$\tilde{c}^I(\tau - t) \tag{3.8}$$

Since the cost accrues at a constant rate, it suffices to track only the total number of deprived individuals at any time. Their arrival times are irrelevant for computing the total cost. Under this structure, the system incurs a cost of $\tilde{c}^I$ for each deprived individual in every micro period until they are satisfied. Therefore, the total deprivation cost in any period is proportional to the size of the deprived population at that time.

In both deterministic and stochastic funding scenarios, the installment size in the final period is assumed to be known. Therefore, the terminal cost at $t = T$ is evaluated as follows:

$$V_T(x_T, r_T) = \begin{cases} c^H x_T, & \text{if } x_T \geq 0, \\ c^I(-x_T - r_T)^+, & \text{otherwise.} \end{cases} \tag{3.9}$$

Note that, operator $(\cdot)^+$ is $(x)^+ = \max(0, x)$.

3.2.2.1 Deterministic Funding

In the deterministic funding setting, committed funding arrives in predefined installments at known macro periods. This formulation allows modeling budget cycles with full certainty. The recursive structure distinguishes between micro period to micro period and macro period transitions. We first provide the immediate cost functions and state transformations under three distinct demand realizations. The immediate cost function is shaped as:

$$
\begin{aligned}
V_t(x_t, r_t) =&(1-p)\left(\tilde{c}^H(x_t)^+ + \tilde{c}^I(-x_t)^+ + V_{t+1}(x_t, r_t)\right) \\
&+ p^I\left(\tilde{c}^H(x_t-1)^+ + \tilde{c}^I(-x_t+1)^+ + V_{t+1}(x_t-1, r_t)\right) \\
&+ p^E \min_{a_t \in \mathcal{A}(x_t)} \left(\tilde{c}^H(x_t-a_t)^+ + \tilde{c}^I(-x_t)^+ + c^E(1-a_t) + V_{t+1}(x_t-a_t, r_t)\right)
\end{aligned}
\tag{3.10}
$$

Here, in the first term, we show the no-demand arrival. Our system iterates one more period without changing the state. In the second term, internal demand arises and consumes one inventory. In the third term, we see the trade-off between the

In the replenishment state, the problem follows the same transition functions after replenishment. Meaning that, we will apply the same transition logic with post-replenishment state variables. The post-replenishment state variables are given as:

$$
x'_t = x_t + y_t, \quad r'_t = r_t - y_t. \tag{3.11}
$$

Depending on the realization of the demand, the immediate cost (or stage cost) is shaped as provided. In between two micro periods, we only decide on sharing inventory with externals. Recognize that, if $x_t = 0$, the sharing decision also becomes degenerate (i.e., not sharing).

At period $t \notin \mathcal{S}$ such that $t+1 \in \mathcal{S}$, a funding installment becomes available. In the deterministic setting, we assume that the system receives an installment at every macro period. However, the installment amount may be zero (which is the same as no funding arrival). Accordingly, the funding installment at period $t+1$, denoted by $z_{\eta(t+1)}$, is incorporated by increasing the budget in the next-period state:

$$
r'_t = r_t + z_{\eta(t+1)}, \quad x'_t = x_t. \tag{3.12}
$$

In summary, the only structural difference between macro and micro transitions lies in the arrival of funding and the replenishment decision. The remaining transition mechanics are identical, determined by the post-replenishment inventory state. The system is in a non-decision (frozen) state when $x_t \leq 0$, which simplifies the recursion by reducing the policy space in these periods.

3.2.2.2 Stochastic Funding

In the stochastic funding setting, installments arrive randomly over the planning horizon. The total number of installments, denoted by $K \leq H$, and installment amounts (i.e., funding vector) are known in advance, but their timing is uncertain. To capture this, we introduce an additional state variable m_t, representing the number of remaining installments at period t. Now state space becomes $s_t = (x_t, r_t, m_t)$. The number of remaining installments is initialized with $m_1 = K-1$ since we assume the first installment is available at the beginning of the planning horizon. Natural result of this assumption is $r_1 = z_1$.

Funding arrivals occur at macro period boundaries and are modeled as a Bernoulli process. The probability of receiving an installment at period $t \in \mathcal{T}$ where $t+1 \in \mathcal{S}$ is determined by the ratio of remaining installments to remaining macro periods. For $m_t \geq 1$, the possible macro periods of next funding arrival are given by $\Gamma(s_t) = \{1, 2, \ldots, H - m_t\}$, i.e., before the end of the planning horizon, all the installments are received. Those $\Gamma(s_t)$ sets represent the values of macro period indices. The corresponding probabilities for τ macro periods ahead—valid for $\tau \in \Gamma(s_t)$—are provided as:

$$p_\tau(s_t) = \begin{cases} \left(\prod_{s=1}^{\tau-1} \left(1 - \frac{m_t}{H-\eta(t)-s}\right)\right) \frac{m_t}{H-\eta(t)-\tau} & \text{if } \tau \geq 2, \\ \frac{m_t}{H-\eta_t(t)} & \text{if } \tau = 1, \\ 1 - \frac{m_t}{H-\eta_t(t)} & \text{if } \tau = 0. \end{cases} \tag{3.13}$$

Here, $\tau = 0$ means the probability of funding not arriving in the upcoming period, which is the complement of the probability that funding will arrive in the first macro period beginning ($p_1(s_t)$, ($p_0(s_t) = 1 - p_1(s_t)$). We define the transition equations for periods $t \notin \mathcal{S}$ to $t+1 \in \mathcal{S}$—i.e., from micro to macro periods based on whether a funding installment arrives. The corresponding installment received at this transition is denoted by $z := z_{K-m_t+1}$. So the value functions are nothing but calculated as:

$$
\begin{aligned}
V_t(x_t, r_t, m_t) =& (1-p)\left(\tilde{c}^H(x_t)^+ + \tilde{c}^I(-x_t)^+ + \sum_{\tau\in\{0,1\}} p_\tau(s_t)V_{t+1}(x_t, r_t+\tau z, m_t-\tau)\right) \\
&+ p^I\left(\tilde{c}^H(x_t-1)^+ + \tilde{c}^I(-x_t+1)^+ + \sum_{\tau\in\{0,1\}} p_\tau(s_t)V_{t+1}(x_t-1, r_t+\tau z, m_t-\tau)\right) \\
&+ p^E \min_{a_t\in\mathcal{A}(x_t)}\left(\tilde{c}^H(x_t-a_t)^+ + \tilde{c}^I(-x_t)^+ + c^E(1-a_t) + \sum_{\tau\in\{0,1\}} p_\tau(s_t)V_{t+1}(x_t-a_t, r_t+\tau z, m_t-\tau)\right)
\end{aligned}
\tag{3.14}
$$

The recursive formulation for micro periods in the stochastic funding case closely mirrors the deterministic case, with the key difference being the addition of the state variable m_t, which tracks the number of remaining funding installments. The main distinction lies in the expectation operator outside the state transition, which captures the two possible realizations in the next macro period—whether a funding installment arrives or not. At macro periods $t \in \mathcal{S}$, a replenishment decision is made (if applicable), and the system transitions forward accordingly. Importantly, the post-replenishment dynamics follow the same recursion as in the deterministic case. This ensures that the same state transition functions apply uniformly across both micro and macro periods, unified through state updating. While the funding arrivals introduce stochasticity at macro boundaries, the overall decision logic and structural mechanics of the dynamic program remain consistent with the deterministic setting.

3.2.3 Exponential Deprivation

Traditional backlogged inventory models often impose a linear penalty for delayed fulfillment. However, in humanitarian contexts, the urgency of unmet internal demand typically intensifies exponentially over time [6]. To capture this behavior more accurately, we adopt an exponential deprivation cost structure, where the cost of delay grows exponentially with waiting time, reflecting the compounding human impact of deferred aid.

Under a linear deprivation structure, cost evaluation is relatively straightforward due to its memoryless property, which allows the system to track only the total number of deprived individuals. In contrast, exponential deprivation requires precise arrival time information for each unmet internal demand, as the penalty depends critically on the duration of deprivation. The cost incurred for fulfilling an internal demand that arrives at period $t \in \mathcal{T}$ and is satisfied at macro period number t' is given by:

$$c^I(t, t') = \alpha \left(e^{\beta(t' - \frac{t+N}{N})} - 1 \right), \tag{3.15}$$

where α is the deprivation cost coefficient, β is the exponential growth rate, and $(t' - \frac{t+N}{N})$ represents the fulfillment delay in macro period units. Remember, t' is the satisfaction macro period.

One potential solution for calculating deprivation cost is to augment the state space to retain individual arrival times for all backlogged internal demands. State augmentation could be achieved through extending the state space to record waiting times for up to HN periods. However, this approach becomes computationally intractable easily as the time horizon grows. Instead, we exploit the problem structure to enable exact cost computation without state augmentation. This is achieved through the following assumptions:

- All internal backorders are fulfilled in the order of arrival. This assumption is both ethically justified and mathematically consistent, as the exponential cost function is strictly increasing in waiting time. For any fixed fulfillment period t', we have $c^I(t_1, t') > c^I(t_2, t')$ when $t_1 < t_2$.

- At every macro period $t \in \mathcal{S}$, the decision-maker must replenish at least minimum of deprived population and budget units to satisfy as much of the deprived population as the budget allows. This policy ensures that estimated satisfaction times are exact due to the funding vector visibility.

Under these assumptions, we can infer the exact satisfaction period for any newly arriving internal demand. In contrast to the linear case, no additional penalty is needed for the remaining deprived population at the terminal period because the full time-weighted deprivation cost is already assigned at the time of each demand arrival. However, one should observe that the cost diminishes when the problem arrives at the end of the horizon and there is no motivation for satisfying the internal demand. This behavior is not intended; for this reason, we penalize the unmet demand at the end of the period with a deprivation cost of one full macro period. The terminal cost function at $t = T$ is therefore:

$$V_T(s_T) = \begin{cases} c^H x_T, & \text{if } x_T \geq 1, \\ (-x_T - r_T)^+ c^I(T, H+2) & \text{otherwise.} \end{cases} \tag{3.16}$$

We next formulate the transition dynamics under the exponential cost structure for both deterministic and stochastic funding scenarios.

3.2.3.1 Deterministic Funding

In the deterministic funding setting, the timing and amount of future funding installments are known in advance. When internal demand cannot be satisfied immediately due to inventory shortage, it is possible to determine the exact period in which it will be fulfilled. To determine the fulfillment time of a newly arriving internal demand, we define a fulfillment index function $d(x_t, r_t)$ that maps the current backlog and funding status to the index of the earliest installment sufficient to meet demand. The detailed construction is provided for both deterministic and stochastic funding regimes in Appendix A.1.

Illustrative Example. Suppose the funding vector over four macro periods is given by $\mathbf{z} = (3, 2, 2, 0)$, and each macro period consists of five micro periods. Then, the macro period boundaries are $(1, 6, 11, 16, 21)$ and the global time set is $1, 2, \ldots, 21$. Consider system state $s_3 = (x_t = -1, r_t = 1)$ at period $t = 3$. In the next macro period (starting at $t = 6$), two units of funding will arrive. Given that the current budget can be used, the projected inventory at $t = 6$ becomes $1 + 2 - 1 = 2$. Hence, the unsatisfied internal demand at $t = 3$ is guaranteed to be fulfilled at $t = 6$ which means $d(x_t - 1, r_t) = 2$.

The corresponding deprivation cost is:

$$c^I(3, d(x_t - 1, r_t)) = \alpha \left(e^{\beta(2 - \frac{3+N}{5})} - 1 \right) \tag{3.17}$$

This example demonstrates how, under deterministic funding and valid fulfillment assumptions, exact exponential deprivation costs can be computed without maintaining the full backlog timeline. We now investigate the transitions under two cases: $x_t \geq 1$ and $x_t \leq 0$. First lets provide the $x_t \geq 1$ case, which is shaped as:

$$\begin{aligned} V_t(x_t, r_t) =& (1-p) \left(\tilde{c}^H x_t + V_{t+1}(x_t, r_t) \right) \\ &+ p^I \left(\tilde{c}^H (x_t - 1) + V_{t+1}(x_t - 1, r_t) \right) \\ &+ p^E \min_{a_t \in \mathcal{A}(x_t)} \left(\tilde{c}^H (x_t - a_t) + c^E (1 - a_t) + V_{t+1}(x_t - a_t, r_t) \right) \end{aligned} \tag{3.18}$$

It is obvious that the equations are reduced to the micro period transition versions. Calculations are the same as previous calculations provided for the linear deprivation version of the problem, except that the deprivation cost is not accumulated and penalized once on arrival. For this reason, we drop the micro step deprivation cost and instead rest is the same. However, when $x_t \leq 0$, things are changed as follows:

$$
\begin{aligned}
V_t(x_t, r_t) =&(1-p)\left(V_{t+1}(x_t, r_t)\right) \\
&\quad + p^I\left(c^I(t, d(x_t - 1, r_t)) + V_{t+1}(x_t - 1, r_t)\right) \\
&\quad + p^E(c^E + V_{t+1}(x_t, r_t))
\end{aligned} \tag{3.19}
$$

In the no demand arrival case and the external demand case, we only proceed with the system with one step. We drop the holding costs and the sharing decision since there is no inventory. The significant change is that now we need to add the exact cost of deprivation: it is not progressive but a lump sum. We determine the exact arrival date, looking at the funding vector and determine the exact macro period when it will be satisfied. From micro period $t \notin \mathcal{S}$ to micro period $t+1 \notin \mathcal{S}$, equations are as given.

Transition at micro period $t \notin \mathcal{S}$ to macro period $t+1 \in \mathcal{S}$ is the same as with a linear variant of the problem. Exogenous information is revealed ($z := z_{\eta(t+1)}$, and the sharing decision is made based on the state. Meaning that right after the decision is made, funding is added to the value functions, and we calculate the value function cost accordingly.

One important nuance exists in this formulation. As can be observed, the arrival funding level is not considered in the immediate cost. The reason is that the arrived funding is not usable in the current period (i.e., observed after the decision). The cost function considers the future funding arrivals and the current budget jointly. Since the funding is not available during the period t but only arrives and can be used in the next period, the budget inside the function is not updated. Updated states are given as:

$$
r'_t = z_{\eta(t+1)} + r_t, \quad x'_t = x_t. \tag{3.20}
$$

In the macro to micro period transition, we again nearly follow the same calculation, however, with one different consideration. Without explicitly tracking the arrival dates of the rejected internal demand, the exact satisfaction period is calculated based on the earliest time backorder satisfaction assumption. Meaning that, cost is added beforehand, assuming the exact satisfaction date is known. For this reason, the model needs to ensure that the lower bound of the replenishment amount is $y_t \geq \min\{r_t, (-x_t)^+\}$. This constraint ensures the model tends to satisfy deprivations as much as possible. Depending on the replenishment decision, expected costs are calculated.

The primary difference between the linear and exponential cases under deterministic funding lies in the structure of the penalty. In the exponential case, the model does not apply the deprivation cost incrementally. Instead, when the inventory is depleted and internal demand is rejected, the model directly penalizes the deprivation based on the time of future satisfaction and the arrival of sufficient funding. The cost is applied as a lump sum function of the delay rather than being accumulated over time. Thus, we need to add a replenishment constraint to satisfy as much of the deprivation as possible to calculate the deprivation cost under the assumption of the satisfaction period.

3.2.3.2 Stochastic Funding

Under stochastic funding, we extend the state space by including m_t to track the number of remaining installments, as in the linear deprivation version of the setting. We introduce an additional variable κ_t, which denotes the exact macro period that the last unmet internal demand will be satisfied. Note that the remaining installments may be insufficient to meet the needs of the entire deprived population. In such cases where $\kappa_t = H$, the system enters a frozen state where no further action is taken until the end of the planning horizon.

Funding vector is fully observable, so it can be easily calculated in which installment system can satisfy all deprived populations as provided in Appendix A.1 as

$k(x_t, r_t, m_t)$. Note that, different from the deterministic version of this problem, the arrival time is not exactly known. However, we need to condition on how much installment ahead is enough to satisfy deprivation so that we can calculate the costs based on the expectation of funding arrival. The exact calculations are provided and discussed in detail.

Recall that internal demand is only rejected when $x_t = 0$. Once the inventory level reaches zero, the system enters a freeze state (i.e., no further action is required until the inventory becomes positive). When the system reaches the state $x_t = 0$, it must wait for the arrival of the first internal demand. Upon the arrival of this demand, we determine κ_t, which denotes the macro period in which the system will be able to satisfy the unmet internal demand (i.e., the next funding arrival time). This timing depends on the stochastic funding arrival process. Rather than relying on the expected cost (which may yield significant under or over estimation), we explicitly take into consideration the possible funding arrival scenarios.

Realization of no demand and external demand scenarios stays the same, and calculations are straightforward. However, internal demand requires special treatment for certain conditions. For any micro period $t \notin \mathcal{S}$ to $t+1 \notin \mathcal{S}$, five distinct cases may arise:

1. **Case (i)**: Inventory is available, i.e., $x_t \geq 1$.

2. **Case (ii)**: No inventory is available ($x_t \leq 0$), but the deprived population and new internal arrival can be satisfied using the available budget at the first macro period ahead (i.e., $-x_t + 1 \leq r_t$).

3. **Case (iii)**: There is enough available budget on hand to satisfy the deprived population. However, the new arrival cannot be satisfied with the current budget level ($-x_t = r_t$).

4. **Case (iv)**: Deprived population already cannot be satisfied with the current budget

($\kappa_t \geq 1$) but a new internal arrival does not alter the earliest installment macro period required to satisfy the deprived population, i.e., $k(x_t, r_t, m_t) = k(x-1, r_t, m_t)$.

5. **Case (v)**. A new internal arrival changes the earliest installment macro period $k(x_t, r_t, m_t) < k(x-1, r_t, m_t)$.

The calculations are provided case by case for regular micro period transitions, replenishment, and funding arrival considerations.

Case (i): $x_t \geq 1$**.** Condition $x_t \geq 1$ suggests two things: when inventory is positive, it is obvious that $\kappa_t = 0$, and we do not need to consider anything, and satisfy the internal demand. For the micro to micro period, transitions are shaped as:

$$\begin{aligned} V_t(x_t, r_t, m_t, 0) =& (1-p)\left(\tilde{c}^H x_t + V_{t+1}(x_t, r_t, m_t, 0)\right) \\ &+ p^I\left(\tilde{c}^H(x_t - 1) + V_{t+1}(x_t - 1, r_t, m_t, 0)\right) \\ &+ p^E \min_{a_t \in \mathcal{A}(x_t)}\left(\tilde{c}^H(x_t - a_t) + c^E(1-a_t) + V_{t+1}(x_t - a_t, r_t, m_t, 0)\right) \end{aligned} \tag{3.21}$$

Simply, calculations are the same as the deterministic funding counterpart of the problem. When $t \in \mathcal{S}$, replenishment decision updates the state with:

$$x'_t = x_t + y_t, \quad r'_t = r_t - y_t, \quad m'_t = m_t, \quad \kappa'_t = 0 \tag{3.22}$$

Under this post-replenishment state, we proceed to the state with the given calculations. Observe that there is no predetermined funding arrival. For this reason, when $t \notin \mathcal{S}$ and $t+1 \in \mathcal{S}$, we need to consider whether funding will arrive in the next period or not. Remember that $p_\tau(s_t)$ is the probability of a funding installment arriving τ macro period ahead. Thus, the calculation for micro periods stays the same, but we need an expectation on the arrived state, given that the new funding is z_{K-m_t+1}, which is:

$$(1 - p_1(s_t)) \times V_{t+1}(x_t, r_t, m_t, 0) + p_1(s_t) \times V_{t+1}(x_t, r_t + z_{K-m_t+1}, m_t - 1, 0) \quad (3.23)$$

Case (ii): $-x_t + 1 \leq r_t$**.** In this case, the available budget is sufficient to satisfy the new internal demand arrival. Thus, internal demand should be penalized according to the next macro period. This condition implies $\kappa_t = 0$ and one hidden condition, which is $x_t \leq 0$ since we already consider the scenario where $x_t \geq 1$. There is no inventory on hand; for this reason, the sharing decision and holding cost are not relevant in this case and all cases after this (remember that for $x_t \leq 0$ value of $\mathcal{A}(x_t) = \{0\}$). Equations for regular micro period transitions are given as:

$$\begin{aligned} V_t(x_t, r_t, m_t, 0) =& (1-p)\left(V_{t+1}(x_t, r_t, m_t, 0)\right) \\ &+ p^I \left(c^I(t, \eta(t)+1) + V_{t+1}(x_t - 1, r_t, m_t, 0)\right) \\ &+ p^E \left(c^E + V_{t+1}(x_t, r_t, m_t, 0)\right) \end{aligned} \quad (3.24)$$

At period $t \in \mathcal{S}$, replenishment decision updates and the funding arrival are exactly the same with Case (i) calculations.

Case (iii): $-x_t = r_t$**.** This case requires special treatment since we have sufficient funding for current deprived population, $\kappa_t = 0$; however, new internal demand triggers κ_t to set predetermined funding, which needs to be greater than current macro period ($\kappa_t \geq \eta(t) + 1$). The reason is that we know the next internal cannot be satisfied with the available budget, and we need to determine when it is going to be satisfied. It means we need to determine when the next funding will arrive. $\tau \in \Gamma(s_t)$ where $\Gamma(s_t) = \{1, 2, ..., H - \eta(t) - m_t + 1\}$. The term $(H - \eta(t) - m_t + 1)$ is the upper bound for how many macro periods ahead the first installment can arrive. The calculations are given as:

$$
\begin{aligned}
V_t(x_t, r_t, m_t, 0) =& (1-p)\left(V_{t+1}(x_t, r_t, m_t, 0)\right) \\
&+ p^I \left(\sum_{\tau \in \Gamma(s_t)} p_\tau(s_t) \times \left(c^I(t, \eta(t)+\tau) + V_{t+1}(x_t - 1, r_t, m_t, \eta(t)+\tau) \right) \right) \\
&+ p^E \left(c^E + V_{t+1}(x_t, r_t, m_t, 0) \right)
\end{aligned}
$$

In this condition, observe that replenishment is a degenerate decision. The reason is the lower bound of the replenishment decision. Please remember that one of the two assumptions for the calculation of the exact satisfaction period of the internal demand is $\min\{(-x_t)^+, r_t\} \leq y_t \leq r_t$. In this scenario, $-x_t = r_t$ then $r_t \leq y_t \leq r_t$ which means, we need to replenish as much of the available budget. The after replenishment state is given by:

$$
x'_t = x_t + r_t, \quad r'_t = 0, \quad m'_t = m_t, \quad \kappa'_t = 0 \tag{3.25}
$$

In micro period to macro period transitions for no demand and external demand scenarios, we need to consider whether funding will arrive in the next period or not. However, for internal arrival, if there is one, the system automatically decides on the next arrival. For this reason, we do not need to consider that scenario. Expectation is given by:

$$
(1 - p_1(s_t)) \times V_{t+1}(x_t, r_t, m_t, 0) + p_1(s_t) \times V_{t+1}(x_t, r_t + z_{K-m_t+1}, m_t - 1, 0) \tag{3.26}
$$

Case (iv): $k(x_t - 1, r_t, m_t) > k(x, r_t, m_t)$. We complete the cases where $\kappa_t = 0$. We consider a scenario where the system already enters the freeze state in this and the next case. The condition for the case shows that new internal demand cannot be satisfied with the budget on hand and the upcoming funding installment. Note that κ_t always shows the next funding installment arrival time. So, suppose the next arrival is not sufficient to satisfy the next arrival. In that case, it is safe to behave as funding (which is $z :=$

z_{K-m_t+1}) is already realized, and conditioning on the κ_t, we can determine when the next installment will arrive. Meaning that, given the funding arrival in the κ_t, we need to observe when the next funding will arrive (if applicable). In the light of this discussion, consider the state transitions :

$$
\begin{aligned}
V_t(x_t, r_t, m_t, \kappa_t) =& (1-p)\left(V_{t+1}(x_t, r_t, m_t, \kappa_t)\right) \\
&+ p^I \left(\sum_{\tau \in \Gamma(s_t)} p_\tau(s_t) \times \left(c^I(t, \kappa_t + \tau) + V_{t+1}(x_t - 1 + r_t + z, 0, m_t - 1, \kappa_t + \tau) \right) \right) \\
&+ p^E \left(c^E + V_{t+1}(x_t, r_t, m_t, \kappa_t) \right)
\end{aligned}
$$

We discuss that the replenishment decision is degenerate for those cases where $r_t \leq -x_t$. For this reason, post-decision states are the same as Case (iii), where $y_t = r_t$. Note that $\kappa_t > 0$ does not imply that $r_t = 0$ strictly. Situation for the funding arrival is quite similar to the replenishment case when $\kappa_t \neq 0$ conditions. There are two cases:

- $\kappa_t = 0$: No deprivation is present, or the existing deprivation does not require a funding installment for fulfillment (i.e., there is no predetermined funding arrival needed). This is only valid for Case (i) and Case (ii).

- If $\kappa_t = \eta(t+1)$, it means the predetermined arrival time is reached and funding arrives with certainty.

- If $\kappa_t > \eta(t+1)$, then we know that funding is not going to arrive with certainty.

For this reason, there is no expectation here. We update the state for external and no demand scenarios if $\kappa_t = \eta(t+1)$ as:

$$
V_{t+1}(x_t, r_t + z_{K-m_t+1}, m_t - 1, \kappa_t) \tag{3.27}
$$

and for $\kappa_t > \eta(t+1)$, no update is required.

Case (v): $k(x_t - 1, r_t, m_t) = k(x, r_t, m_t)$**.** In the final case, it is known that the current calculated κ_t is also sufficient for the new internal. Then, we need to penalize the internal demand depending on the κ_t. Equations are:

$$
\begin{aligned}
V_t(x_t, r_t, m_t, \kappa_t) =& (1-p)\left(V_{t+1}(x_t, r_t, m_t, \kappa_t)\right) \\
& + p^I \left(c^I(t, \kappa_t) + V_{t+1}(x_t - 1 + r_t + z, 0, m_t - 1, \kappa_t)\right) \\
& + p^E \left(c^E + V_{t+1}(x_t, r_t, m_t, \kappa_t)\right)
\end{aligned} \tag{3.28}
$$

Value function calculations for replenishment periods and funding arrival stay the same. The only difference with Case (iv) is that we do not update the (κ_t) value on the internal arrival. This concludes the modeling section. We present four distinct problem variants, each characterized by a different update mechanism and cost structure. In the subsequent section, we analyze their structural properties.

3.2.4 Structural Properties

This section establishes key structural properties of the model. We prove two main results: (i) the optimal replenishment decision at each macro period follows a *modified base-stock policy*, and (ii) the optimal rationing (sharing) decision for external demand follows a *threshold policy*. Throughout this section, we assume that holding costs are linear in inventory, deprivation costs are convex (either linear or exponential), referral costs are fixed per unit, lead times are zero, and internal demand always has unconditional priority. Funding affects replenishment only at macro periods. These structural results hold across all model variants considered.

Let the full state be $s_t = (x_t, r_t, m_t, \kappa_t)$. For the purpose of analyzing structural properties, we suppress the κ_t component because whenever $\kappa_t \geq 1$, no decision is made: replenishment becomes trivial (degenerate), and no sharing decision is applicable due to the absence of available inventory. The two primary state variables affecting convexity and threshold results are the inventory level x_t and the budget r_t. We write $V_t(x_t, r_t, m_t)$ to denote the value function when m_t is fixed. Convexity and monotonicity properties are proven pointwise with respect to these remaining state variables. To simplify the analysis, we define the net inventory position and note that the model never allows $x_t \geq 1$ and $\kappa_t \geq 1$ to occur simultaneously. Hence, whenever the inventory position is positive, we can assert that no freeze state is present.

Lemma 3.1 (Discrete Convexity in Inventory)**.** *Let $V_{t+1}(x_{t+1}, r_{t+1})$ be convex in x_{t+1} for all feasible funding levels r_{t+1}. Then, under the proposed model structure, the value function $V_t(x_t, r_t)$ remains convex in x_t at period t for each of the transition types.*

Proof. The proof appears in Appendix B.1. □

When the value function is convex in inventory, the marginal benefit (i.e., marginal reduction in expected cost) of holding an additional unit of inventory is non-increasing. Intuitively, as inventory increases, the risk of stockouts decreases, and additional inventory units provide less benefit while incurring holding costs. Therefore, the marginal value of inventory diminishes, which is consistent with the convex structure.

Theorem 3.2 (Existence of Threshold-Based Sharing Policy)**.** *For each period t, there exists a threshold $\Omega_t(r_t)$ such that the optimal sharing policy follows:*

$$a_t^*(x_t, r_t) = \begin{cases} 1, & \textit{if } x_t \geq \Omega_t(r_t), \\ 0, & \textit{otherwise}. \end{cases} \tag{3.29}$$

Proof. The proof appears in Appendix B.2. □

The threshold-based sharing policy provides two key benefits. First, instead of evaluating every possible inventory level, we only need to determine a single threshold value for the given state variables. This significantly reduces computational complexity, as we can identify the switch point where the sharing decision changes from not share to share, which is more efficient than performing a classical linear search. Second, for fixed state variables other than time, the optimal sharing threshold is non-increasing as the period t increases. Intuitively, as we approach the end of the planning horizon, the opportunity value of holding inventory for future periods diminishes, making it optimal to share more aggressively (i.e., the threshold does not increase over time).

Theorem 3.3 (Existence of Base-stock Replenishment Policy)**.** *For each macro period $t \in \mathcal{S}$, there exists an optimal base-stock level $\mathcal{X}_t(r_t)$ such that the optimal replenishment decision is budget-capped:*

$$y_t^* = \begin{cases} 0 & \textit{if } \mathcal{X}_t(r_t) - x_t \leq 0, \\ r_t & \textit{if } \mathcal{X}_t(r_t) - x_t \geq r_t, \\ \mathcal{X}_t(r_t) - x_t & \textit{otherwise}. \end{cases} \tag{3.30}$$

Proof. The proof appears in Appendix B.3. □

Convexity in inventory is sufficient to exploit the problem structure. When the value function is convex in inventory, there exists an optimal base-stock level that minimizes the cost function. Moreover, due to the relationship between inventory and budget (i.e., the combined purchasing capacity available for future periods), we can perform the optimization in terms of the total purchasing power $(x_t + r_t)$ while keeping other variables fixed. This facilitates an even faster exploration of optimal actions. The practical implications of structural properties are further discussed in Section 4.

4. COMPUTATIONAL STUDY

In this section, we present our computational analysis. We begin by addressing the computational challenges posed by the curse of dimensionality and describe how we overcome them using structural properties and dynamic programming enhancements. We present a real-world inspired humanitarian response case study with discussions of parameters. To evaluate the scalability and performance of the proposed model, we conduct a factorial experiment under varying demand settings. For this case study, we perform a detailed sensitivity analysis to highlight how funding uncertainty, and cost characteristics affect optimal policy decisions and overall system performance.

4.1 Computational Challenges

Dynamic programming problems are commonly affected by three types of the so-called curse of dimensionality: the size of the state space, the size of the outcome space, and the size of the action space [38]. In our problem, all three dimensions pose significant challenges—especially due to the long planning horizon and the multiple state variables introduced by funding uncertainty and deprivation tracking.

Consider the exponential deprivation variant under stochastic funding. The space complexity of storing the full value function can be approximated as:

$$\mathcal{O}\big(HN \times x^b \times r^b \times m \times \kappa\big),$$

where x^b and r^b represent the bound of the variables. The value function itself already exceeds tractable memory limits for typical computing environments. Replenishment decisions contribute additional memory requirements at a scale of

$$\mathcal{O}\big(H \times x^b \times r^b \times m \times \kappa\big),$$

while sharing decisions further increase the burden by another

$$\mathcal{O}\big(HN \times x^b \times r^b \times m \times \kappa\big).$$

Even after pruning unreachable or redundant states, the effective size of the decision space can easily become intractable. The resulting state space is so large that storing even a value function becomes infeasible on most commercial computers. For this reason, while calculating the value functions, we use slicing idea: instead of holding whole history, we keep only the last two i.e., current calculated period $V_t(\cdot)$, and one period ahead $V_{t+1}(\cdot)$, thus holding $2/(HN)$ portion of the value function becomes enough.

Another state space enhancement is on sharing policy. It is provided in the structural properties, rationing decision follows *threshold structure*, meaning that there exists a threshold level for doing rationing. Instead of tracking the rationing decision at every inventory level, we compute a threshold level for each combination of the remaining state variables. Rather than storing a full decision table, the problem is solved by determining and applying this threshold-based policy. Similarly, this happens in the inventory replenishment as well. We provide that, *modified-base stock policy* is optimal for our problem. This situation implies that, there exists a inventory replenishment target value for the given other state variables, so that we can calculate target value for fixed other variables. It benefits us on two issue: firstly, replenishment policy requires significantly less memory, and it can be calculated fast, instead of calculating for each state variables. Under monotonicity assumptions, we can use binary search which yields a $\mathcal{O}(\log N)$ complexity comparing it to searching linearly which results $\mathcal{O}(N)$ complexity. It is especially useful for large instances, which makes contribution more significant.

4.2 Baseline Setting and Discussion on Parameters

Before discussing the other components of the problem, first discuss the H parameter value. Remember that the H controls the probability mass ignored by the system. It means, increasing H value helps the system to reflect the continuous approximation better. The Figure 4.1 illustrates the relationship between the probability of observing multiple demand arrivals in a micro period and the ratio of the number of micro periods between two macro periods (N) and total lambda $\frac{N}{\lambda}$. As this ratio increases (i.e., as the temporal resolution becomes finer), the probability of simultaneous arrivals diminishes rapidly. We select N sufficiently large such that the probability of multiple arrivals within a micro period is negligible (i.e., $\mathbb{P}(D \geq 2) \approx 0$). The precision of this approximation (i.e., the probability mass being ignored) is visualized in Figure 4.1.

Figure 4.1: *Relationship between N/λ and the probability of multiple arrivals in a micro period.*

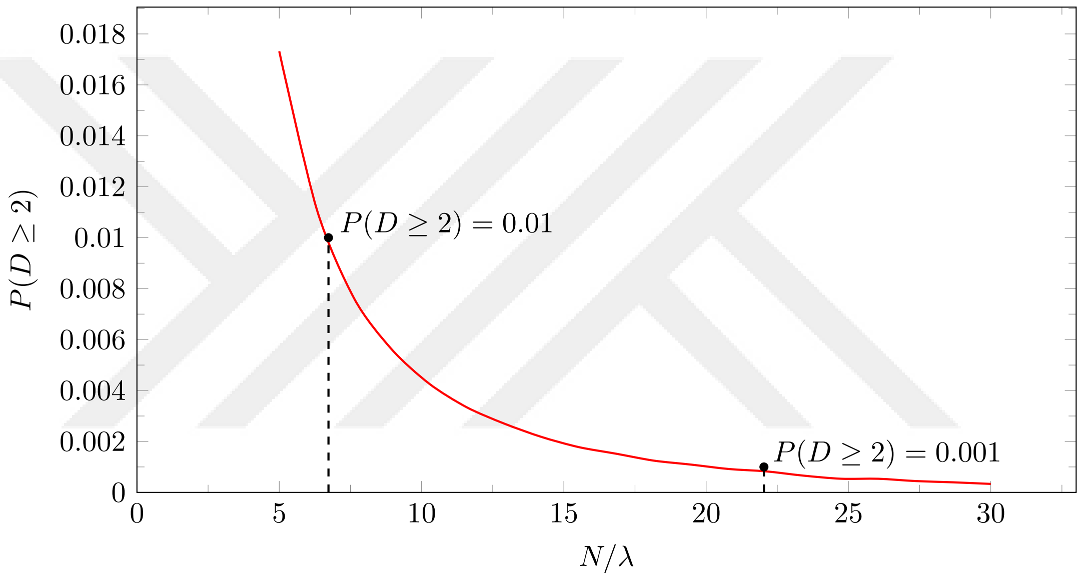

Ideally, we aim to set N to a large value. However, there is a trade-off between the value of N and computational time. Increasing N linearly increases the number of required calculations. In other words, the computational burden grows proportionally with N. Therefore, selecting a reasonable value for N ensures computational efficiency without compromising the resolution of the system. In our setting, we use $N/\lambda = 7$, which approximately results in ignoring only 1% of the demand. Note that, if the value of the λ and ratio of $N/\lambda = 7$ are not changing, then total micro period count is not changing.

Another important dimension is the planning horizon length. Since we operate under an uncertain funding environment, it is reasonable to assume that donors allocate budgets on a fiscal-year basis. Consequently, we set the planning horizon to one year for each scenario. Extending the horizon length beyond this period does not provide additional insights, as our focus is on understanding the trade-offs across different settings. However, varying H—that is, reviewing inventory at monthly, bi-weekly, or weekly intervals—can influence system performance. We therefore investigate the impact of horizon resolution.

The model considers three cost components, defined per macro period. The holding cost, c^H, can be viewed as the opportunity cost of not deploying funds for alternative responses. Thus, instead of merely providing numerical values, we emphasize the relative ratios among costs: c^E/c^H, c^I/c^H, and c^E/c^I, which express the criticality of serving beneficiaries compared to alternative responses. For simplicity, we assume $c^H = 1$ per macro period. Note that, when definition of the macro period changes, the cost parameters need to be scaled accordingly. For this reason, in all analysis, we assume that the costs are for per month. We define the criticality ratio for external demand (CR^E) and internal demand (CR^I) as the ratios of referral and deprivation costs to the holding cost. In general, we use $CR^E = 3$ and $CR^I = 5$, which implies referral cost is 3 (one time) and linear deprivation cost of 5 per month. For the exponential cost case, the waiting cost for one month is also set to 5, however scaled with α, leading to:

$$\alpha(e^{\beta} - 1) = CR^I, \quad \beta = ln\Big(\frac{CR^I + \alpha}{\alpha}\Big). \tag{4.1}$$

We need to define the costs per month to align with mathematical model and use it interchangeably. For this reason, from per month costs, we scale the cost accordingly for the each period length to fair comparison between the scenarios. Only difference is for exponential deprivation cost, where the time is given as macro period format. For this reason, we only scale the cost with normalizing β by given horizon information.

The total committed funding is defined as

$$F = \lambda^I + \gamma\,\lambda^E, \tag{4.2}$$

where $\gamma \in [0,1]$. This formulation ensures that the committed funding is always sufficient to cover the expected internal demand, while the parameter γ determines the proportion of the expected external demand that can be satisfied. For very low funding levels, there is no incentive to allocate resources to external demand. Therefore, to effectively analyze the impact of sharing, we restrict γ to the interval $[0,1]$.

We utilize realistic data from [27] to obtain meaningful results. The dataset provides demand statistics for refugee camps in Türkiye. In this study, we focus on the camp in Hatay, where the annual internal demand rate is $\lambda^I = 1071$, and the annual external demand rate is $\lambda^E = 2818$. Table 4.1 summarizes the main problem parameters.

Table 4.1: *Problem parameters for numerical analysis.*

Parameter	Symbol	Value
Number of micro periods per macro period	N	$N/\lambda = 7$
Planning horizon	H	1 year
Internal demand rate (annual)	λ^I	1071
External demand rate (annual)	λ^E	2818
Holding cost per month	c^H	1
Criticality ratio (external)	CR^E	3
Criticality ratio (internal)	CR^I	5
Exponential deprivation rate	β	$\ln\left(\frac{CR^I+\alpha}{\alpha}\right)$
Total committed funding	F	$\lambda^I + \gamma\lambda^E$
Funding ratio	γ	$[0,1]$

We use *fill rates* as a key performance statistic to compare the results. The fill rate reflects the ability of the camp manager to respond to demand in a timely manner. For internal (camp) demand, unmet requests are eventually fulfilled; thus, the fill rate can be interpreted as a measure of the timeliness of demand satisfaction. In contrast, external demand cannot be served in future periods once missed. Therefore, the external fill rate represents the proportion of arrivals that are satisfied.

4.3 Factorial Design for Computational Performance

To assess the computational scalability of the proposed model, we examine the total CPU time required to solve instances under increasing demand intensity. Let the number of macro periods be fixed at $H = 12$, and consider a funding parameter $\gamma = 0.3$. The internal-to-external demand ratio is set as $\frac{\lambda^I}{\lambda^E} = \frac{2}{5}$. We vary the total demand rate λ across four levels: λ, 2λ, 4λ, 8λ, where each value reflects the total expected monthly demand from both classes.

We assume that the stochastic funding process is realized over $K = 6$ installments, and compare four model variants: linear vs. exponential deprivation cost, and deterministic vs. stochastic funding. Table 4.3 reports the total CPU time (in seconds) for each configuration. The results demonstrate that computational effort increases significantly with the demand scale, especially under exponential deprivation and stochastic funding, due to the increased size of the state and action spaces and the complexity of the value function transitions.

Table 4.2: *Effect of per macro period λ on CPU Time ($\lambda = 70$ and $\gamma = 0.3$).*

		λ	2λ	4λ	8λ
CPU Times in sec	Linear Deterministic	2.62	3.97	14.95	191.48
	Linear Stochastic	4.98	24.09	213.71	922.55
	Exponential Deterministic	3.37	10.76	68.29	512.69
	Exponential Stochastic	169.78	431.32	3152.27	25106.43

The results in Table 4.2 reveal a clear relationship between the total demand intensity and computational effort. For all model variants, CPU time increases with demand scale, but the rate of increase differs substantially depending on the cost structure and funding regime. Under linear deprivation costs, both deterministic and stochastic funding scenarios remain computationally tractable across all demand levels. In particular, the deterministic case shows a modest increase in run time, suggesting that the model scales linearly with demand in this setting. The stochastic case exhibits greater sensitivity due to the additional uncertainty in funding timing, which enlarges the effective state space.

In contrast, the exponential deprivation cases demonstrate significantly higher computational burden, particularly when funding is stochastic. While the deterministic-exponential variant remains manageable up to 8λ, the stochastic-exponential case becomes computationally expensive even at moderate demand levels. This is expected, as exponential cost functions amplify the impact of delay, requiring finer value function resolution and deeper recursion, which in turn slows down convergence. These findings highlight the importance of model structure in computational performance. When both cost functions and funding mechanisms are nonlinear or uncertain, the computational complexity grows rapidly, necessitating careful tuning of approximation methods or pruning techniques in large-scale applications.

4.4 Numerical Results and Policy Insights

4.4.1 Analysis of Horizon Resolution

We now analyze the effect of time resolution by varying the number of replenishment opportunities H within a fixed total planning horizon. Specifically, we solve the problem for $H = \{2, 4, 12, 26, 52\}$, which correspond to coarser (e.g., biannual) to finer (e.g., weekly) discretizations of the same time span. For each value of H, we adjust the Poisson demand rates and cost parameters accordingly to maintain consistency with the total horizon length.

Table 4.3 summarizes the simulation results for the linear deprivation setting with $\gamma = 0.3$. As expected, increasing the temporal resolution (i.e., higher H) enables more frequent replenishment and rationing decisions, which in turn reduces deprivation costs and provide more stability (i.e., lowers variability) in total cost. Internal fill rates also

Table 4.3: *Costs and fill rates for different period resolution (linear deprivation, $\gamma = 0.3$, 100 replications).*

		H = 52	H = 26	H = 12	H = 4	H = 2
Fill Rate (External)	Max	0.38	0.39	0.43	0.49	0.48
	Mean	0.35	0.35	0.35	0.35	0.35
	Min	0.32	0.30	0.28	0.26	0.27
	StdDev	0.01	0.02	0.03	0.04	0.04
Fill Rate (Internal)	Max	0.86	0.87	0.90	0.93	0.96
	Mean	0.82	0.82	0.82	0.84	0.85
	Min	0.79	0.77	0.73	0.71	0.73
	StdDev	0.01	0.02	0.03	0.05	0.04
Holding Cost	Max	565.91	550.15	640.48	705.86	633.28
	Mean	531.78	517.56	546.79	555.95	563.39
	Min	505.74	477.97	471.90	460.50	475.84
	StdDev	12.51	16.35	28.82	39.38	33.45
Deprivation Cost	Max	166.10	226.47	324.56	572.01	945.75
	Mean	116.36	114.20	130.13	155.34	155.24
	Min	73.01	58.92	40.99	12.76	16.56
	StdDev	20.62	32.03	56.62	117.09	157.69
Referral Cost	Max	5469	5700	6015	6360	6300
	Mean	5135.58	5132.04	5129.70	5173.50	5118.30
	Min	4803	4446	4245	3540	3570
	StdDev	165.72	218.97	361.35	548.10	526.11
Total Cost	Max	6140.44	6341.37	6653.97	7012.33	7100.85
	Mean	5783.71	5763.80	5806.63	5884.79	5836.93
	Min	5459.36	5099.07	4907.99	4265.61	4197.48
	StdDev	166.26	221.06	361.69	549.34	554.20

improve moderately, reflecting more responsive inventory allocation to internal demand.

Interestingly, the external fill rate remains largely stable across different values of H, suggesting that the underlying rationing policy consistently prioritizes internal demand. The rationing policy is likely designed to protect internal demand first, regardless of horizon resolution This confirms that the external class receives a fixed share of resources regardless of decision frequency, as long as the internal demand can be met.

From a managerial perspective, increasing the decision frequency from low to moderate (e.g., from $H = 2$ to $H = 12$) yields noticeable improvements in service levels and cost. However, further increases beyond this point (e.g., $H = 26$ to $H = 52$) provide only marginal gains while substantially increasing computational complexity (i.e., diminishing returns). Based on this trade-off, we adopt a monthly control policy with $H = 12$ in subsequent analyses, which balances performance, tractability, and realism — particularly in humanitarian contexts where weekly inventory adjustments may be impractical.

4.4.2 Total Committed Funding Amount

One another dimension is funding amount. We need to understand, how γ value effects the system performance. In this case, we provide both linear and exponential deprivation case, and how the costs are changing for $H = 12$.

Figure 4.2: *Effect of γ on system performance under linear deprivation: (a) cost components, (b) fill rates.*

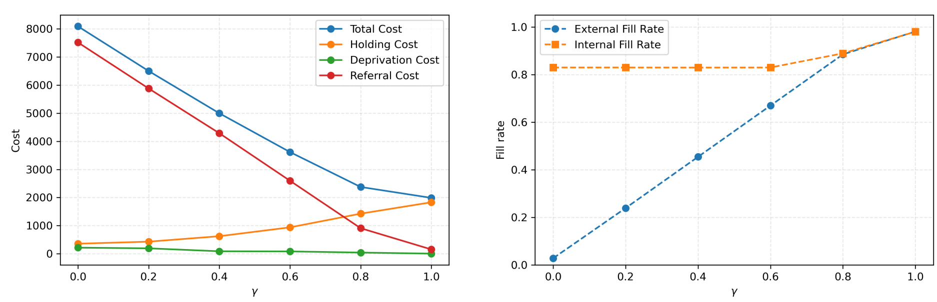

(a) (b)

Figure 4.3: *Effect of γ on system performance under exponential deprivation: (a) cost components, (b) fill rates.*

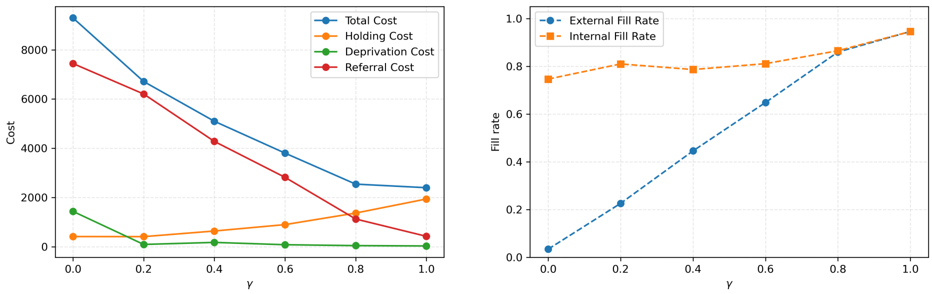

(a) (b)

The results presented in Figure 4.2 and Figure 4.3 are largely in line with expectations but reveal an important distinction. As γ increases, we would intuitively expect the system to fulfill internal demand at a higher rate. However, the internal fill rate remains nearly constant from $\gamma = 0.0$ to $\gamma = 0.6$. This behavior arises from a key difference compared to classical periodic inventory control problems: we penalize deprivation continuously over time, rather than only at the end of a period. Consequently, the benefit of satisfying internal demand early in a macro period diminishes as time progresses toward the end of period. As a result, the model strategically prioritizes external demand by setting the sharing thresholds close to zero, even at moderate values of γ. This prioritization stems from the high penalty associated with unserved external demand, which dominates the cost structure. This behavior persists across a wide range of γ, only shifting meaningfully when γ approaches higher values (e.g., 0.8 or 1.0), where internal and external fulfillment rates converge due to abundant inventory sharing.

4.4.3 Impact of Stochastic Funding Frequency

Stochastic funding scenarios introduce uncertainty in the timing of installment arrivals, while the total funding amount remains fixed. A key consideration is how the frequency of installment arrivals—reflected in their expected inter-arrival times—influences overall system performance. In this setting, the total promised funding is equally divided among the installments, so as the number of installments increases, the system receives smaller amounts more frequently.

We investigate the effect of stochastic funding frequency on system performance, holding the total committed funding amount fixed. The installment sizes are assumed to be equal, and funding arrivals occur stochastically over the planning horizon. To evaluate this effect, we solve the model for varying numbers of installments $\mathcal{K} = \{2, 4, 6, 8, 10\}$ under a fixed funding ratio $\gamma = 0.3$, considering both linear and exponential deprivation cost structures.

Figure 4.4: *Total cost under linear deprivation for different numbers of installments.*

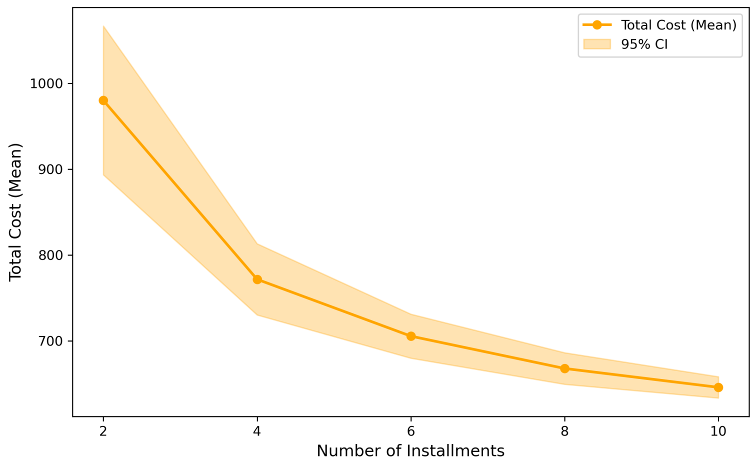

As illustrated in Figure 4.4, increasing the number of installments yields lower total costs, indicating improved system robustness against funding uncertainty. The most notable improvement occurs when the number of installments increases from 2 to 4. Beyond this point, the marginal gains become less pronounced. This diminishing return suggests that overly frequent funding provides limited additional benefit under linear deprivation costs.

Figure 4.5 presents the results for the exponential deprivation case. A similar

Figure 4.5: *Total cost under exponential deprivation for different numbers of installments.*

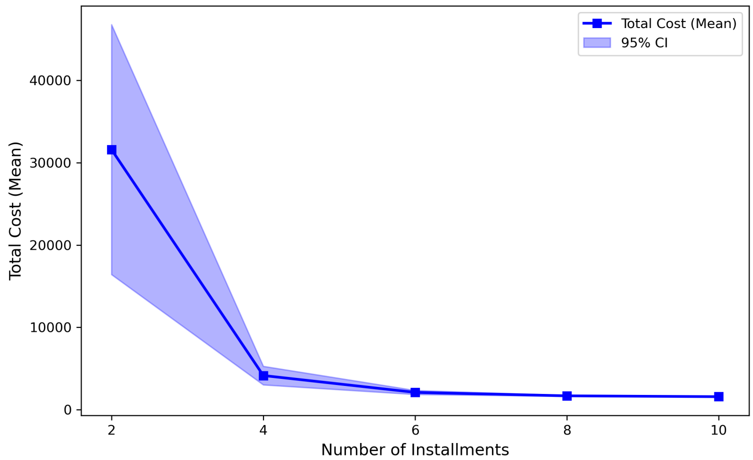

downward trend is observed with increasing installment frequency, but the magnitude of improvement between 2 and 4 installments is substantially larger. This difference arises from the convex nature of the exponential cost structure, where delays in fulfilling demand incur disproportionately higher penalties. As a result, higher installment frequencies significantly reduce expected deprivation by shortening the duration of funding shortages. Therefore, in the presence of exponential deprivation costs, the system exhibits a stronger preference for more frequent, smaller funding disbursements.

4.4.4 Impact of Funding Pattern

This analysis investigates the effect of the temporal structure of funding—referred to as the funding pattern—on system performance. Specifically, we examine how varying the timing of installment amounts impacts total cost across different model variants. In all cases, we assume that the number of installments K and their total sum are fixed. The only difference lies in how the total funding is distributed across the installments.

We consider the following three funding patterns:

- **Equal Sized Funding (ES)**: Each installment delivers an identical amount, serving as the baseline scenario.

- **Front-Loaded Funding (FL)**: The first half of the installments deliver $1.25\times$ the equal amount, while the second half deliver $0.75\times$ the equal amount.

- **Back-Loaded Funding (BL)**: The reverse of FL; the first half of installments deliver $0.75\times$ and the second half $1.25\times$ the equal amount.

All three regimes are calibrated such that the total committed funding over the horizon remains constant. This ensures a fair comparison that isolates the impact of timing alone.

Figure 4.6: *Relative total cost under different funding patterns for each model variant.*

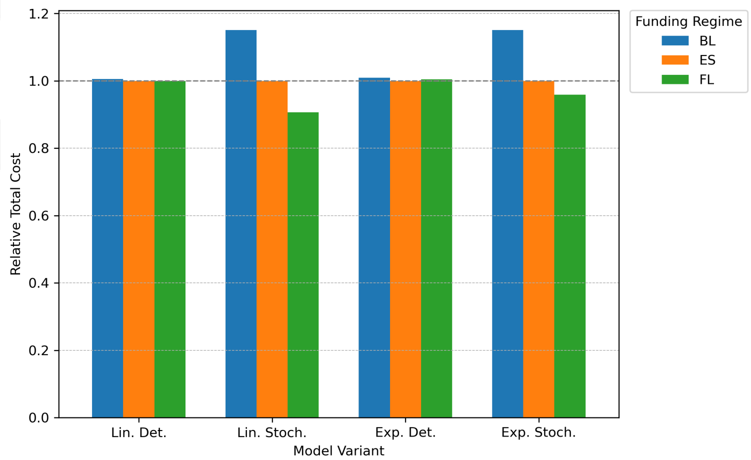

Figure 4.6 presents the relative total costs for each model variant (linear vs. exponential deprivation; deterministic vs. stochastic funding) under the three funding patterns. All values are normalized by the cost observed in the ES baseline.

The results highlight that stochastic funding scenarios are more sensitive to the timing of installments. In particular, both the linear and exponential stochastic cases incur significantly higher costs under BL, whereas FL leads to noticeable improvements in total cost, especially for the linear-stochastic case. This is intuitive: in the presence of funding uncertainty and deprivation cost accumulation over time, receiving larger amounts earlier reduces the expected duration of shortage, thereby lowering deprivation-related penalties.

In contrast, the deterministic scenarios exhibit relatively minor deviations from the baseline. Since the timing of funding is known with certainty, the system can plan replenishment accordingly, rendering the temporal pattern of funding less impactful. These findings emphasize the importance of early disbursement in stochastic humanitarian settings. When funding uncertainty is present, front-loading resources may serve as a practical and low-cost strategy to mitigate deprivation, even without increasing the total budget.

5. CONCLUSIONS

This study develops a dynamic inventory allocation framework tailored to humanitarian operations, where inventory must be allocated between internal demand (fully backordered, incurring time-sensitive deprivation costs) and external demand (lost sales, incurring fixed referral costs) under funding constraints. By explicitly modeling uncertainty in both demand arrivals and funding timing, the framework captures the core operational complexities observed in humanitarian contexts. Three key features characterize the model: (i) joint replenishment and sharing decisions under multi-class demand with finite inventory, (ii) a discrete-time approximation of real-time decision-making through micro periods within periodic review, and (iii) the use of both linear and exponential deprivation cost structures to reflect varying urgency profiles.

The problem is formulated as a finite-horizon stochastic dynamic program and solved via exact backward induction. Structural insights—such as base-stock-type replenishment and threshold-based sharing—enable tractable and interpretable policy structures. To manage the computational complexity, the model employs dynamic slicing of the value function and leverages monotonicity properties to accelerate recursion.

Numerical results reveal several key insights. First, the total amount and temporal distribution of funding significantly influence system performance. Early and frequent funding disbursements improve cost efficiency—particularly under exponential deprivation—by mitigating delays in serving internal demand. Second, increasing the planning resolution (via micro periods) enhances responsiveness and reduces deprivation, though gains diminish beyond a certain threshold. Third, the model consistently prioritizes internal demand when resources are scarce, reflecting the asymmetry in cost and urgency between internal and external demand classes.

The framework also suggests several avenues for future research. Non-stationary demand can be accommodated by allowing the number of micro-periods per macro-period to vary, enabling the model to reflect time-varying arrival rates. In the linear deprivation case, stochastic funding amounts can be incorporated without compromising tractability, while exponential deprivation requires full observability of the funding vector. The single-camp setting can be extended to multi-camp networks with centralized replenishment and decentralized sharing decisions—a structure conducive to scalable solution techniques such as Approximate Dynamic Programming (ADP) or Deep Reinforcement Learning (DRL). Given the structure of the optimal policy, algorithms like Proximal Policy Optimization (PPO) may offer promising approximations.

Beyond its methodological contributions, the model offers a principled decision-support tool for humanitarian logistics. By quantifying the marginal value of funding timing and size, it can inform donor strategies and support evidence-based resource allocation. In high-stakes environments—where delays can be fatal and resources are limited—the framework contributes not only to academic understanding but also to practical, equitable, and effective crisis response.

REFERENCES

[1] OCHA, "Global humanitarian overview 2025," 2024. (access date January 15, 2025).

[2] UNHCR, "Mid-year trends 2024," 2024. (access date January 15, 2025).

[3] UNHCR, "Global appeal 2025: Executive summary," 2025. (access date January 15, 2025).

[4] L. B. Landau, "Urban refugees and idps," *The Oxford Handbook of Refugee and Forced Migration Studies*, vol. 1, pp. 139–150, 2014.

[5] Sphere Association, "The sphere handbook: Humanitarian charter and minimum standards in humanitarian response," 2018. (access date January 15, 2025).

[6] J. Holguín-Veras, N. Pérez, M. Jaller, L. N. Van Wassenhove, and F. Aros-Vera, "On the appropriate objective function for post-disaster humanitarian logistics models," *Journal of Operations Management*, vol. 31, no. 5, pp. 262–280, 2013.

[7] IASC, "Donor Conditions and Their Implications for Humanitarian Response," 2016. (access date January 15, 2025).

[8] UNHCR, "UNHCR Underfunded Report 2024," 2024. (access date January 15, 2025).

[9] B. Balcik, C. D. C. Bozkir, and O. E. Kundakcioglu, "A literature review on inventory management in humanitarian supply chains," *Surveys in Operations Research and Management Science*, vol. 21, no. 2, pp. 101–116, 2016.

[10] K. V. Natarajan and J. M. Swaminathan, "Inventory management in humanitarian operations: Impact of amount, schedule, and uncertainty in funding," *Manufacturing & Service Operations Management*, vol. 16, no. 4, pp. 595–603, 2014.

[11] T. O. Kotsi, A. Aflaki, G. Aydin, and A. J. Pedraza-Martinez, "Allocation of nonprofit funds among program, fundraising, and administration," *Manufacturing & Service Operations Management*, vol. 25, no. 5, pp. 1873–1889, 2023.

[12] M. Keshvari Fard, I. Ljubić, and F. Papier, "Budgeting in international humanitarian organizations," *Manufacturing & Service Operations Management*, vol. 24, no. 3, pp. 1562–1577, 2022.

[13] C. Burkart, M. Besiou, and T. Wakolbinger, "The funding—humanitarian supply chain interface," *Surveys in Operations Research and Management Science*, vol. 21, no. 2, pp. 31–45, 2016.

[14] X. Chao, J. Chen, and S. Wang, "Dynamic inventory management with cash flow constraints," *Naval Research Logistics (NRL)*, vol. 55, no. 8, pp. 758–768, 2008.

[15] J. H. McCoy and M. L. Brandeau, "Efficient stockpiling and shipping policies for humanitarian relief: Unhcr's inventory challenge," *OR Spectrum*, vol. 33, no. 3, pp. 673–698, 2011.

[16] K. V. Natarajan and J. M. Swaminathan, "Multi-treatment inventory allocation in humanitarian health settings under funding constraints," *Production and Operations Management*, vol. 26, no. 6, pp. 1015–1034, 2017.

[17] C. Y. Yildiz and O. E. Kundakcioglu, "Disaster relief inventory simulation: Managing resources in humanitarian camps," in *2024 Winter Simulation Conference (WSC)*, pp. 1575–1586, IEEE, 2024.

[18] M. J. Kleijn and R. Dekker, "An overview of inventory systems with several demand classes," in *New Trends in Distribution Logistics* (M. G. Speranza and P. Stähly, eds.), (Berlin, Heidelberg), pp. 253–265, Springer Berlin Heidelberg, 1999.

[19] H. Arslan, S. C. Graves, and T. A. Roemer, "A single-product inventory model for multiple demand classes," *Management Science*, vol. 53, no. 9, pp. 1486–1500, 2007.

[20] Z. Guo and H. Chen, "Joint optimization of inventory replenishment and rationing policies for an omnichannel store with both in-store and online demands," *Computers & Industrial Engineering*, vol. 191, p. 110171, 2024.

[21] K. T. Möllering and U. W. Thonemann, "An optimal critical level policy for inventory systems with two demand classes," *Naval Research Logistics (NRL)*, vol. 55, no. 7, pp. 632–642, 2008.

[22] G. Wan and Y. Cao, "A continuous cost evaluation approach for periodic review inventory systems with threshold rationing policy," *Computers & Industrial Engineering*, vol. 126, pp. 75–87, 2018.

[23] Y. Zhou and X. Zhao, "A two-demand-class inventory system with lost-sales and backorders," *Operations Research Letters*, vol. 38, no. 4, pp. 261–266, 2010.

[24] D. Wang and O. Tang, "Dynamic inventory rationing with mixed backorders and lost sales," *International Journal of Production Economics*, vol. 149, pp. 56–67, 2014.

[25] B. M. Beamon and S. A. Kotleba, "Inventory modelling for complex emergencies in humanitarian relief operations," *International Journal of Logistics: Research and Applications*, vol. 9, no. 1, pp. 1–18, 2006.

[26] B. M. Beamon and S. A. Kotleba, "Inventory management support systems for emergency humanitarian relief operations in south sudan," *The International Journal of Logistics Management*, vol. 17, no. 2, pp. 187–212, 2006.

[27] S. Azizi, C. D. C. Bozkir, A. C. Trapp, O. E. Kundakcioglu, and A. K. Kurbanzade, "Aid allocation for camp-based and urban refugees with uncertain demand and replenishments," *Production and Operations Management*, vol. 30, no. 12, pp. 4455–4474, 2021.

[28] J. Holguín-Veras, M. Jaller, L. N. Van Wassenhove, N. Pérez, and T. Wachtendorf, "On the unique features of post-disaster humanitarian logistics," *Journal of Operations Management*, vol. 30, no. 7-8, pp. 494–506, 2012.

[29] D. C. Whybark, "Issues in managing disaster relief inventories," *International Journal of Production Economics*, vol. 108, no. 1-2, pp. 228–235, 2007.

[30] B. M. Beamon and B. Balcik, "Performance measurement in humanitarian relief chains," *International Journal of Public Sector Management*, vol. 21, no. 1, pp. 4–25, 2008.

[31] B. Balcik, B. M. Beamon, C. C. Krejci, K. M. Muramatsu, and M. Ramirez, "Coordination in humanitarian relief chains: Practices, challenges and opportunities," *International Journal of Production Economics*, vol. 126, no. 1, pp. 22–34, 2010.

[32] M. Jahre, J. Kembro, A. Adjahossou, and N. Altay, "Approaches to the design of refugee camps: An empirical study in kenya, ethiopia, greece, and turkey," *Journal of Humanitarian Logistics and Supply Chain Management*, vol. 8, no. 3, pp. 323–345, 2018.

[33] O. Karsu, B. Y. Kara, and B. Selvi, "The refugee camp management: a general framework and a unifying decision-making model," *Journal of Humanitarian Logistics and Supply Chain Management*, vol. 9, no. 2, pp. 131–150, 2019.

[34] N. Altay and A. Narayanan, "Forecasting in humanitarian operations: Literature review and research needs," *International Journal of Forecasting*, vol. 38, no. 3, pp. 1234–1244, 2022.

[35] UNHCR, "Unhcr supply strategy 2024-2030," 2024. (access date January 15, 2025).

[36] OCHA, "Disbursement," 2025. (access date January 15, 2025).

[37] UNHCR, "Unhcr global report 2023: Income," 2023. (access date January 15, 2025).

[38] W. B. Powell, *Approximate Dynamic Programming: Solving the Curses of Dimensionality*. John Wiley & Sons, 2011.

APPENDICES

APPENDIX A. SUPPLEMENTARY MODEL DERIVATIONS

To maintain clarity in the main text, we provide selected technical derivations and supplementary calculations in the appendix. These include detailed formulations that support the development of the model and ensure the completeness and reproducibility of the results.

A.1 Earliest Funding Installment

Deterministic Funding When internal demand arrives and cannot be satisfied due to insufficient inventory, we determine the macro period at which the system becomes feasible—i.e., able to fulfill the new arrived internal demand. We define the function $d(x_t, r_t)$, which returns the arrival macro period of earliest installment that enables satisfaction of the backlog.

In the deterministic case, the timing and value of each installment are known in advance. Let $\eta(t)$ be the macro period associated with micro period t. Then, the macro period is given by:

$$d(x_t, r_t) = \begin{cases} \eta(t) + 1, & \text{if } -x_t \leq r_t, \\ \eta(t) + \min\left\{ k \in \{1, \ldots, H - \eta(t)\} \;\middle|\; x_t + r_t + \sum_{i=1}^{k} z_{\eta(t)+i} \geq 0 \right\}, & \text{if } \sum_{i=1}^{H-\eta(t)} z_{\eta(t)+i} \geq -x_t - r_t, \\ H, & \text{otherwise.} \end{cases} \tag{A.1}$$

If there is no deprived population, no further investigation is necessary. For that reason, the condition is not provided. If the available budget is sufficient to satisfy the entire deprived population, then the system will fulfill the unmet demand in the first macro period. However, if the deprivation exceeds the current budget with new arrival, it becomes necessary to identify the earliest funding installment that enables the system to fully satisfy the deprived population. In the deterministic setting, funding is assumed to arrive in every macro period. Therefore, the earliest macro period in which the cumulative funding is sufficient to increase x_t to zero determines the fulfillment time. If this is not possible within the current macro period, the internal demand is penalized accordingly with end of the horizon.

Stochastic Funding In the stochastic case, the number of remaining installments m_t is part of the state, but their realization is uncertain. The total funding vector is given by $\mathbf{z} = (z_1, \ldots, z_K)$, and m_t indicates the number of installments yet to arrive. Unlike the deterministic case, the exact timing of future funding is unknown. Nevertheless, the following function is used to check not to exact satisfaction period, but which installment ahead needed to satisfy whether a newly received internal demand changes the earliest installment in which the deprivation can be fully satisfied, i.e., $k(x_t - 1, r_t, m_t) > k(x_t, r_t, m_t)$ or not:

$$k(x_t, r_t, m_t) = \begin{cases} 0, & \text{if } -x_t \leq r_t, \\ \min\left\{k \in \{k_0, \cdots, K\} : \sum_{k'=k}^{K} z_{k'} \geq -x_t + r_t\right\}, & \text{if } \sum_{k'=k_0}^{K} z_{k'} \geq -x_t + r_t, \\ H + 1, & \text{otherwise,} \end{cases} \tag{A.2}$$

where $k_0 = K - m_t + 1$ excludes installments that have already arrived.

APPENDIX B. STRUCTURAL PROPERTY PROOFS

B.1 Proof of Lemma 3.1

Proof. We proceed by backward induction, starting from the terminal period. At the end of the planning horizon, the terminal cost function is defined as:

$$V_T(x_T, r_T) = \begin{cases} c^H x_T, & \text{if } x_T \geq 0, \\ \tilde{c}^I(-x_T - r_T)^+, & \text{otherwise,} \end{cases} \tag{B.1}$$

which is convex in x_T because it is a sum of the positive part of a linear function and a convex function.
For micro to micro period transitions, the Bellman recursion is:

$$\begin{aligned} V_t(x_t, r_t) =& (1-p)\left(\tilde{c}^H(x_t)^+ + \tilde{c}^I(-x_t)^+ + V_{t+1}(x_t, r_t)\right) \\ &+ p^I\left(\tilde{c}^H(x_t-1)^+ + \tilde{c}^I(-x_t+1)^+ + V_{t+1}(x_t-1, r_t)\right) \\ &+ p^E \min_{a_t \in \mathcal{A}(x_t)} \Big(\tilde{c}^H(x_t-a_t)^+ + \tilde{c}^I(-x_t)^+ \\ &\quad + c^E(1-a_t) + V_{t+1}(x_t-a_t, r_t)\Big). \end{aligned} \tag{B.2}$$

Holding and deprivation costs are convex in x_t by construction. By the induction hypothesis, V_{t+1} is convex in inventory. For each fixed a_t, the inner expression for external demand is convex in x_t; since the minimum of convex functions is convex, and convexity is preserved under expectation, $V_t(x_t, r_t)$ is convex in x_t. At macro periods, the planner selects a replenishment amount $y_t \in [\min\{r_t, (-x_t)^+\}, r_t]$, and the state update follows the same structure as micro periods. Because V_{t+1} is convex in x and both $x_t + y_t$ and $r_t - y_t$ are affine in y_t, convexity is preserved. At the last micro period of a macro period, funding realization occurs, but no replenishment decision is made:

$$V_t(x_t, r_t) = [\text{stage cost}] + \mathbb{E}_q\left[V_{t+1}(x_t, r_t + qz, m_t - q)\right], \tag{B.3}$$

where q is the random variable for funding installment arrival. Since expectation preserves convexity, $V_t(x_t, r_t)$ is convex in x_t for all t. □

B.2 Proof of Theorem 3.2

Proof. The sharing decision compares the marginal benefit of keeping one unit of inventory with the penalty of referring an external demand and holding for one period. Sharing is optimal when

$$V_t(x_t - 1, r_t) - V_t(x_t, r_t) \leq c^E + \tilde{c}^H. \tag{B.4}$$

Define the marginal cost:

$$\Delta V_t(x_t) := V_t(x_t - 1, r_t) - V_t(x_t, r_t). \tag{B.5}$$

If $\Delta V_t(x_t)$ is non-increasing in x_t, then there exists a threshold $\Omega_t(r_t)$ such that sharing is optimal if and only if $x_t \geq \Omega_t(r_t)$. Convexity of $V_t(x_t, r_t)$ implies:

$$V_t(x_t + 1, r_t) - 2V_t(x_t, r_t) + V_t(x_t - 1, r_t) \geq 0, \tag{B.6}$$

so the forward differences $V_t(x_t, r_t) - V_t(x_t - 1, r_t)$ are non-decreasing. Hence, $\Delta V_t(x_t)$ is non-increasing, and the threshold $\Omega_t(r_t)$ exists. □

B.3 Proof of Theorem 3.3

Proof. Given the convexity of $V_{t+1}(x_{t+1}, r_{t+1})$ in x_{t+1}, the post-replenishment cost function

$$C(y_t) = V_{t+1}(x_t + y_t, r_t - y_t)$$

is convex in y_t because it is the composition of a convex function with an affine mapping. Therefore, the minimizer of $C(y_t)$ defines the *base-stock level* $\mathcal{X}_t(r_t)$. □