Title

Towards Equity and Fairness in Data-Driven Management Systems

Author

Lee, Yoon

Publication Date

2023

Peer reviewed|Thesis/dissertation

Towards Equity and Fairness in Data-Driven Management Systems

By

Yoon Lee

A dissertation submitted in partial satisfaction of the

requirements for the degree of

Doctor of Philosophy

in

Engineering - Industrial Engineering and Operations Research

in the

Graduate Division

of the

University of California, Berkeley

Committee in charge:

Professor Anil Aswani, Chair
Professor Koushil Sreenath
Professor Rajan Udwani

Fall 2023

Towards Equity and Fairness in Data-Driven Management Systems

Abstract

Towards Equity and Fairness in Data-Driven Management Systems

by

Yoon Lee

Doctor of Philosophy in Engineering - Industrial Engineering and Operations Research

University of California, Berkeley

Professor Anil Aswani, Chair

This dissertation comprises three lines of work addressing societal challenges related to safe water access, fair incentive design, and secure data sharing, utilizing mathematical modeling to propose comprehensive solutions. These projects collectively advocate for universal resource and information access while prioritizing equity and fairness for marginalized communities. The first project centers on optimizing water storage decisions, considering factors such as cost, water quality, and wastage, with the core objective of promoting equitable access to safe drinking water, particularly for those facing severe water scarcity. In the second project, we address deficiencies in current incentive mechanisms for fairness, introducing quantitative definitions rooted in fairness principles to prevent harm to specific groups and encompass disadvantaged segments of society. Lastly, the third project delves into the intricacies of responsible scientific data sharing, presenting a novel cybersecurity insurance contract to incentivize healthcare providers while safeguarding privacy and livelihoods, especially for vulnerable populations. In sum, this research strives to develop solutions that facilitate equal and inclusive access while mitigating risks, with the goal of addressing cultural and societal issues concerning resource allocation and information sharing for underserved communities.

To my family, for their unwavering support and belief in my academic journey.

Contents

List of Figures

Acknowledgments

I am profoundly grateful for the immense support and invaluable guidance that have sustained me throughout my academic journey. My deepest appreciation goes out to all those who have played an integral role in this endeavor.

Foremost among these sources of gratitude is my dedicated advisor, Anil Aswani. His mentorship and expertise have been instrumental in my growth as a researcher, especially during challenging moments. Anil's remarkable ability to align research problems with my interests has been pivotal, and our collaborative efforts, from complex proofs to late-night paper writing, hold a special place in my heart.

My journey has been further enriched by the exceptional collaborators and committee members who have generously shared their insights and wisdom. Zuo-Jun Max Shen, Phil Kaminsky, Koushil Sreenath, and Rajan Udwani have not only provided invaluable feedback but have also been sources of inspiration throughout my research endeavors.

I must also express my heartfelt gratitude to the vibrant academic community at Berkeley, including individuals like Mahan Tajrobehkar, Yuhao Ding, and my fellow IEOR Department peers. Their camaraderie, willingness to collaborate, and intellectual contributions have added immeasurable value to my time here.

In the depths of my heart, I reserve a special place for my wife and my parents, whose unwavering love, support, and encouragement have been the cornerstone of my academic journey. Their sacrifices and belief in my abilities have been a constant source of strength and motivation.

Lastly, I fondly remember my beloved dog, Jelly, whose presence brought immeasurable joy and comfort to my life. Although Jelly is no longer with us, her memory serves as a reminder of the importance of finding happiness and solace in life's smaller moments.

To all those mentioned here and the countless others who have contributed to my academic and personal growth, I extend my deepest appreciation. Your support and faith in me have been the driving force behind my accomplishments, and I am grateful for the privilege of having you in my life.

Chapter 1

Introduction

In an era defined by the omnipresence of data and its pivotal role in decision-making across various domains, the need for equity, fairness, and privacy in data-driven management systems has never been more critical. As we increasingly rely on data to guide our choices, be it in resource allocation, incentive design, or healthcare decisions, the potential for inequities, bias, and breaches of privacy looms large. This dissertation emerges from the recognition that the power of data must be harnessed with a profound commitment to societal values. It seeks to address the pressing challenges and opportunities at the intersection of data-driven decision-making, ethics, and technology, recognizing that the decisions made within these systems can have far-reaching consequences, impacting individuals, communities, and society at large.

This study embarks on a journey to unveil innovative solutions that promote equity in inventory management, fairness in incentive design, and privacy in medical data sharing, all of which are critical pillars of responsible and ethical data utilization. At its core, this dissertation is driven by the aspiration to not only elucidate the complexities of data-driven management but also to catalyze a transformation in how we conceptualize, design, and implement these systems. By doing so, it endeavors to shape a future where data is harnessed as a force for good, serving the collective well-being of society, and advancing the cause of equity and fairness in an increasingly interconnected and data-driven world.

1.1 Contributions and Outline

In this dissertation, we present significant contributions in three key areas, each addressed in a separate chapter:

- **Chapter 2: Inventory Management with Controlled Reset**
  In Chapter 2, we address inventory management problems with periodic and controllable resets, which are commonly encountered in contexts such as managing water storage in the developing world and retailing limited-time availability products. We tackle the complexities of determining optimal policies using dynamic programming,

which becomes challenging due to the non-convex nature of the problems. Our primary contribution lies in establishing sufficient conditions that ensure an interpretable structure for the optimal policy, thus extending the well-known (s, S) policy from the operations management literature. Additionally, we demonstrate that under these mild conditions, the optimal policy exhibits a four-threshold structure. Computational experiments are conducted to illustrate these policy structures in various inventory management scenarios. This chapter is based on work from [46].

- **Chapter 3: Fairness in Incentive Design**
Chapter 3 shifts the focus to integrating fairness into incentive design, a dimension often overlooked in existing approaches. We introduce fairness into optimization problems within principal-agent models, with specific attention to scenarios involving adverse selection and moral hazard. Our contributions encompass the formulation of quantitative fairness definitions and the derivation of policy structures for fair optimal contracts. By delving into the underlying intuition behind these contracts, we emphasize the profound impact of fairness on incentive design. Furthermore, a case study is provided to illustrate the practical implications of incorporating fairness considerations into the design process.

- **Chapter 4: Cybersecurity Insurance for Medical Data Sharing**
Chapter 4 addresses the critical issue of sharing medical data, which holds immense potential for advancing healthcare but simultaneously poses risks to both patients and healthcare providers. We investigate the problem of designing optimal cybersecurity insurance contracts to incentivize responsible data sharing. Using a principal-agent model with moral hazard, we model various scenarios, derive optimal contracts, discuss their implications, and conduct numerical case studies. Two specific scenarios are considered: healthcare providers selling data to technology firms and healthcare providers sharing data for collaborative research. Our study aims to provide valuable insights into risk mitigation and the promotion of responsible scientific data sharing. This chapter draws upon work from [45].

Chapter 2

Local Water Inventory Management for the Developing World

2.1 Introduction

The study of inventory management with stochastic demand has been a key discipline in operations management since the inception of the field. In general, these problems consider a decision-maker who must choose either an order quantity or a replenishment level in order to minimize their overall expected supply chain costs or to maintain a certain level of service with high probability. Several models have been proposed in the literature to address this setting, such as the newsvendor model for single-period replenishment [3], the (r, Q) model for continuous review policies [27], and the (s, S) model for periodic review policies [75]. These models primarily aim to find an optimal inventory policy, which is a decision rule mapping the current inventory level of the facility to an order quantity. Classical models have mainly focused on balancing holding costs (i.e., the costs incurred from holding excess inventory) and shortage costs (i.e., the costs associated with not having enough units on hand to satisfy demand) to determine this policy. However, these assumptions underlying classical policies may not hold when products are perishable or subject to strict health and safety regulations.

Many settings, such as food procurement [89, 5, 22], medical supply chain management [69, 64, 78], and certain seasonally sensitive retail supply chains [13, 57], require managing perishable inventory and thus deviate from the classical inventory management models mentioned earlier. To address this constraint, strategic discarding policies are commonly employed. Essentially, decision-makers rank the products based on value and discard expired units individually, considering costs and health constraints. Although applicable to many contexts, if the managed goods are commingled or not easily separable, it becomes infeasible to discard individual units strategically, necessitating the discarding of the entire inventory. For instance, in the case of residential water storage, newly purchased water mixes with older water that is more likely to have been tainted and cannot be readily separated from the previous water in the tank.

This chapter proposes new models and techniques to address the setting of perishable or deteriorating inventory with known expiration times, where strategic discarding is not possible. Specifically, we propose modeling these systems as inventory control problems with strict reset constraints. In this setting, the decision-maker faces two key decisions at each epoch: first, whether to discard the entire inventory on hand; and second, whether to order inventory up to a certain level. Since the inventory is set to expire after a predetermined time, the decision-maker must discard the entire inventory after a fixed number of epochs. The main challenge in this problem is to find an inventory policy that balances safety constraints with holding and shortage costs by making refresh and reorder decisions.

Applications

The setting of mixing perishable inventory with reset decisions is common in many real-world applications. Here, we present two settings that exemplify the assumptions of the model. First, we consider a non-profit example of managing water storage in a residential building in the developing world. Next, we describe a profit example in the case of a retail supply chain with switching product lines.

Non-Profit: Water Storage Problem

Water availability is a significant global public health concern, with over 300 million people worldwide experiencing intermittent access to water supplies [39]. This problem is particularly severe in the developing world, where the existing distribution infrastructure fails to provide continuous water supplies. Consequently, most households resort to communal and personal water storage containers to maintain a water supply throughout the day. These containers are filled once every few days during intermittent periods of water availability [94, 92]. However, water stored in these containers is prone to contamination, and long-term storage increases the likelihood of higher concentrations of bacterial and viral pathogens [25, 95, 18, 84, 44, 38]. The problem of managing local water storage can be formulated within the framework of perishable inventory, as described previously. Each day, the decision-maker must determine the additional water quantity to purchase to satisfy the demand for that day. Since water becomes more susceptible to contamination as it sits in the tank, this can be modeled as a holding penalty to prevent excessive storage duration. Similarly, if the purchased water quantity is lower than the realized demand for the day, expediting the shipment of additional water may incur a shortage penalty. Finally, since the risk of contamination renders water perishable and contaminated water cannot be easily separated from clean water, the only available option for the decision-maker is a “reset” action, discarding the entire inventory.

Profit: Retail Management Problem

In the context of retail supply chain management, product demand is highly dependent on seasonal tastes or trends. This affects various types of products, ranging from apparel and food to consumer electronics [52, 14, 15]. Particularly, the discontinuation of products within retail outlets such as Trader Joe's and Costco warrants consideration. These retailers, in contrast to traditional supermarkets, meticulously curate their product offerings, focusing exclusively on high-demand items. This is because these retailers operate within limited shelf space, and their business models are predicated on the efficient turnover of sought-after merchandise. Consequently, when a particular product fails short of expected demand or price standards, it becomes a candidate for discontinuation. This strategic approach enables them to maintain competitive pricing structures and operational efficiency [53, 35]. In an environment where retail success hinges on precise inventory management and cost-effective supply chain practices, these discontinuations represent a concerted effort to align business models with the ever-evolving landscape of consumer preferences and supply chain complexities. Hence, the key challenge in designing an inventory policy in this setting is akin to the water storage problem, where the decision-maker must consider the optimal timing to initiate a "reset" action, corresponding to the introduction of a new product line.

Related Literature

In this chapter, we examine the problem of perishable inventory with reset control, which is a special case of multi-period inventory management with periodic review and stochastic demand. In the context of supply chain theory, a policy is a function that maps the current state of our inventory into an ordering decision [80]. For periodic-review inventory models with stochastic demands, a base-stock policy (e.g., newsvendor model) or an (s, S) policy is an example of an inventory control policy. The fundamental idea underlying the (s, S) policy is as follows: in each time period, we observe the current inventory position. If the inventory position falls below s, we place an order of sufficient size to bring the inventory position to S, where the quantity s is known as the reorder point and S is the order-up-to level [3, 75]. To this classical framework, we seek to extend the (s, S) policy by incorporating a reset control action to empty the entire inventory for the problems at hand. [54] previously studied this approach numerically, and our contributions include the theoretical analysis of the optimal policy in the presence of periodic and controllable resets. Additionally, we draw upon two main streams of literature, namely inventory management and dynamic programming, as the basis for our modeling and analysis.

Inventory and Supply Chain Management

The problem we address in this study is closely related to the management of perishable inventory [65, 58, 31, 32, 16, 36, 17], which has primarily been explored within the context of healthcare settings. In these perishable inventory models, costs arise from either not having enough inventory to meet demand (shortage costs) or holding excess inventory (holding costs) that must be discarded due to the perishable nature of the goods. Unlike our setting, this stream of literature generally assumes that expired units of inventory can be disposed of individually, allowing for salvage and disposal costs to be incorporated into the holding cost. Consequently, strategic inventory removal policies can be employed [70, 71]. In our setting, however, decision-makers cannot strategically dispose of single units of their stock. For example, in the water storage problem, fresh water may mix with still water, and they cannot be easily separated. Instead, the decision-maker must decide whether to keep or dispose of the entire inventory to reduce the risk of consuming contaminated water. Similarly, in the retail management problem, we determine whether to retain or discontinue an entire product line. If the product line is discontinued, all inventory must be discarded to free up space for the new product line.

Dynamic Programming and Optimal Control

It is well known in the operations literature that periodic-review inventory management problems can be modeled as continuous-state dynamic programs [75, 4]. While it is possible to derive structural results for the functional form of the optimal policy in most of these models, finding a closed-form solution for the parameters of the policy is often difficult. Thus, numerical algorithms are commonly used to determine the relevant parameter values. The main numerical methods developed to solve these dynamic programming problems are value and policy iteration [4]. These procedures involve finding a fixed point in either the value function or the policy space, respectively. However, these fixed points may exist in infinite-dimensional functional spaces, which can make convergence problematic. Even in finite high-dimensional state spaces, exact calculation using these approaches is numerically difficult, a phenomenon often referred to as the "curse of dimensionality." To address this challenge, several approximate dynamic programming approaches [4, 23, 67, 73, 37, 29] have been developed. These approaches perform value or policy iteration with inexact representations of the value function or policy to achieve tractability. However, the convergence rate of both exact and approximate value and policy iteration is governed by the discount factor, and convergence becomes slow for discount factors close to 1 [4, 97]. Given the specific structure of the problem we examine in our setting, we propose the use of the Binary Dynamic Search algorithm (BiDS), originally devised by [54]. In contrast to value and policy iteration, which operate in function spaces, the BiDS algorithm finds a fixed point in a vector space using binary search. Since the state space for water storage and retail management problems is small, our numerical results solve the exact dynamic program. However, in principle, the BiDS algorithm could also be used for approximate dynamic programming.

Contributions and Outline

Our study primarily focuses on the structural analysis of inventory management problems with periodic review and reset control. As part of this analysis, we theoretically characterize and prove the structure of the optimal policy for these inventory management problems. Our analysis demonstrates that the resulting policy follows a threshold structure, which can be seen as a generalization of the classic (s, S) policy. Moreover, our analysis provides new theoretical methods for analyzing inventory policies, as the non-convexity induced by our problem structure differs from that of classical inventory models.

In addition to studying the structural properties of the problem, another key contribution of ours is the implementation of a novel algorithm for solving stochastic optimal control problems with controlled resets to a single state and constraints on the maximum time between system resets. Unlike value and policy iteration, which require finding a fixed point in an infinite-dimensional functional space [4], we employ the Binary Dynamic Search (BiDS) algorithm [54]. This algorithm transforms the problem into finding a fixed point in a vector space using binary search. The BiDS algorithm significantly reduces computational cost compared to value and policy iteration. We apply BiDS to numerically solve this specific problem and experimentally validate our structural results. By doing so, we generalize the previous (s, S) inventory policy to a new threshold structure that incorporates reset control. To explore the broad applicability of this new structure, we demonstrate its interpretability and implementability in both profit and non-profit operations.

The rest of the chapter is organized as follows. In Section 2.2, we describe the stochastic optimization problem for the general reset control model and provide applications to the water storage problem and retail management problem. Section 2.3 discusses the structural properties of the optimal policy and presents the sufficient conditions that ensure the threshold structure. In Section 2.4, we introduce the Binary Dynamic Search algorithm and present the numerical results obtained for both the water and retail problems.

2.2 Reset Control Formulation and Applications

In this section, we introduce a broad class of stochastic optimal control problems with controlled resets to a single state and with constraints on the maximum time span between resets of the system. Let the subscript $n \in \mathbb{Z}_+$ denote the index of the decision epoch, and consider the discrete-time dynamical system defined as follows:

$$\begin{aligned} x_{n+1} &= h(\xi_n, \tau_n, u_n, w_n) \\ t_{n+1} &= \tau_n + 1 \\ \xi_n &= x_n \cdot (1 - r_n) + \zeta \cdot r_n \\ \tau_n &= t_n \cdot (1 - r_n) \end{aligned} \tag{2.1}$$

where $x_n \times t_n \in \mathbb{R}^{n_x} \times \mathbb{Z}_+$ are states, $\xi_n \times \tau_n \in \mathbb{R}^{n_x} \times \mathbb{Z}_+$ are pseudo-states, $u_n \times r_n \in \mathbb{R}^{n_u} \times \mathbb{B}$ are control actions, $w_n \in \mathbb{R}^{n_w}$ are i.i.d. random variables representing disturbance terms, and $h : \mathbb{R}^{n_x} \times \mathbb{Z}_+ \times \mathbb{R}^{n_u} \times \mathbb{R}^{n_w} \to \mathbb{R}$ is a deterministic function. The control $r_n = 1$ resets the system to a *known* initial state $\zeta \in \mathbb{R}^{n_x}$, the state t_n keeps track of the number of time steps since the last system reset, and the function h describes the system dynamics when no reset occurs. The pseudo-states ξ_n and τ_n capture the reset dynamics by being set to the known initial states when $r_n = 1$.

Given the discount factor $\gamma \in [0, 1)$, our goal is to solve the following stochastic control problem:

$$\begin{aligned} &\min \mathbb{E}\Big[\textstyle\sum_{n=0}^{\infty} \gamma^n \Big(g(\xi_n, \tau_n, u_n, w_n) + s(x_n, t_n, w_n) \cdot r_n \Big) \Big], \\ &\text{s.t. } (2.1),\, t_n \le k,\, u_n \in \mathcal{U}(\xi_n, \tau_n, w_n), \quad \text{for } n \ge 0. \end{aligned} \tag{2.2}$$

where $g : \mathbb{R}^{n_x} \times \mathbb{Z}_+ \times \mathbb{R}^{n_u} \times \mathbb{R}^{n_w} \to \mathbb{R}_+$ is a non-negative and continuous stage cost, and $s : \mathbb{R}^{n_x} \times \mathbb{Z}_+ \times \mathbb{R}^{n_w} \to \mathbb{R}_+$ is a non-negative and continuous reset cost with $s(\zeta, t, w) \equiv 0$. The constraint $u_n \in \mathcal{U}(\xi_n, \tau_n, w_n)$ restricts the possible control actions to lie in a set $\mathcal{U}(\xi_n, \tau_n, w_n)$, and the constraint $t_n \le k$ requires the system be reset at least once every k time steps. For notational convenience, we will simply refer to the set $\mathcal{U}_n := \mathcal{U}(\xi_n, \tau_n, w_n)$.

While this stochastic control formulation is quite general, we now present two instantiations of this model in real-world settings. First, we discuss a non-profit example of water storage control, and then we describe a for-profit example of retail management with changing product lines.

Example: Water Storage Problem

In this section, we consider the stochastic inventory control problem of water management in the developing world, where a single decision-maker must maintain the level of potable water in a residential water tank. Unlike water management in the developed world, we assume that the residence does not have access to a continuous source of water. Therefore, water must be purchased in bulk at the beginning of each day from a communal source. Additionally, the tank needs to be fully emptied every few days for cleaning to eliminate pathogen growth. In this setting, the decision-maker has two actions they can take each day: deciding whether to purge the tank and determining the amount of water to purchase at the beginning of the day. These decisions must be made in a way that optimally balances the financial cost of purchasing water or expediting purchases when the amount is short, and the implicit health costs of letting water sit longer in the tank (which can be thought of as a holding cost).

This problem can be modeled as a stochastic control problem that is a special case of the formulation presented in Section 2.2. For day n, let $x_n \in \mathbb{R}_+$ be the state variable representing the amount of water stored in the tank, and let $t_n \in \mathbb{Z}_+$ be the state variable representing the number of days since the tank was last emptied. The decision-maker's actions, the amount of water to purchase at the start of the day and whether or not to empty the tank, are represented by $u_n \in \mathbb{R}_+$ and $r_n \in \mathbb{B}$, respectively. We model the demand for each day as the i.i.d. random disturbance process $w_n \in \mathbb{R}_+$.

The dynamics of the system are described as follows:

$$x_{n+1} = (x_n \cdot (1 - r_n) + u_n - w_n)^+, \tag{2.3}$$

$$t_{n+1} = t_n \cdot (1 - r_n) + 1, \tag{2.4}$$

where $(x)^+ = \max\{x, 0\}$ and $(x)^- = \min\{x, 0\}$. Here, (2.3) states that the water level at day $n+1$ equals the current level at day n plus the amount of water purchased that day minus the amount of water demanded, but it cannot go below zero. If the tank is flushed on day n, then (2.3) states that the water level of the tank at day $n+1$ does not depend on the amount of water at day n. Likewise, (2.4) states that the number of days since the tank has been flushed increments by 1 each day until the day it is flushed, at which point the count resets to 1. Let $q : \mathbb{R}_+ \to \mathbb{R}_+$ be a non-negative and non-decreasing function that represents the increased health costs associated with letting water sit in the tank for additional days. Let $c, c_r, p \in \mathbb{R}_+$ represent the variable cost of purchasing water, the per-unit cost of purging the tank, and the shortage penalty for not having sufficient water to meet demand. Using these cost parameters, state variables, dynamics equations, and discount factor $\gamma \in [0, 1)$, the decision-maker's problem can be formulated as follows:

$$\begin{aligned}
\min \mathbb{E}\Big[\textstyle\sum_{n=0}^{\infty} \gamma^n \Big(& cu_n - p \cdot (\xi_n + u_n - w_n)^- \\
& + q(\tau_n) \cdot (\xi_n + u_n - (\xi_n + u_n - w_n)^+) + c_r x_n r_n \Big) \Big], && (2.5) \\
\text{s.t. } & x_{n+1} = (\xi_n + u_n - w_n)^+, && (2.6) \\
& t_{n+1} = \tau_n + 1, && (2.7) \\
& \xi_n = x_n \cdot (1 - r_n), && (2.8) \\
& \tau_n = t_n \cdot (1 - r_n), \qquad \text{for } n \geq 0 && (2.9) \\
& r_n \in \mathbb{B}, && (2.10) \\
& t_n \leq k, && (2.11) \\
& u_n \in \big[0, c_{\max} - \xi_n\big], && (2.12)
\end{aligned}$$

For this formulation, the state space is augmented to include pseudo-states ξ_n, τ_n that represent intermediate values of water in the tank and time since the last reset, given the reset action on day n, respectively. Using these states, the objective terms $p \cdot (\xi_n + u_n - w_n)^-$ and $q(\tau_n) \cdot (\xi_n + u_n - (\xi_n + u_n - w_n)^+)$ represent the total shortage cost and health costs at day n, respectively. The interpretation is that the decision-maker pays shortage costs only when demand exceeds water supply, and that water quality deteriorates only if a surplus remains in the tank. The quantity $(\xi_n + u_n - (\xi_n + u_n - w_n)^+)$ is the amount of water consumed on the n-th day because $\xi_n + u_n$ is the amount of water available at the beginning, and $(\xi_n + u_n - w_n)^+$ is the amount of water that is still unused at the end of the day. Since the function q is assumed to be monotonically increasing, the $q(\tau_n)$ term indicates that the quality of water deteriorates as time passes between emptying the tank and also ensures that consuming water that has been stored for longer durations of time is more heavily penalized. Likewise, the terms cu_n and $c_r x_n r_n$ represent the total costs of purchasing water and flushing the tank, respectively, where the flushing cost is only incurred if the reset action is taken. The constraint (2.11) ensures that the tank is purged at least once every k days, and the constraint (2.12) ensures that water is not purchased in excess of the tank capacity c_{max}.

A defining feature of this problem is that instead of considering the holding/shortage cost trade-off, we formulate the problem based on how long water has been stored and how much water has actually been consumed. In circumstances where people do not have a continuous supply of water and thus have no other choice but to store water in a local container that lacks disinfection capabilities, we take an optimization approach to decide when to drain the water tank and how much water to fill it when available. The objective is to reduce the risk of contamination and shortage. In Section 2.4, we will numerically solve this problem to design an efficient, interpretable policy for managing water storage systems. This policy can be easily implemented using a lookup table that can be widely distributed to the public through paper pamphlets or the internet.

Example: Retail Management Problem

Next, we consider the setting of retail inventory management in the case of changing product lines. In this setting, the decision-maker is tasked with managing the inventory of limited-time availability product lines and must decide the quantity of stock to store and the timing for changing to a new product line in order to minimize inventory costs. This involves balancing the trade-offs between ordering costs, holding costs for storing inventory, shortage costs, stock wastage from emptying the inventory, and the cost of switching the product line.

Similar to the water storage problem, this retail inventory problem can also be modeled as a special case of the reset control problem presented in Section 2.2. For each week n, let the state variables $x_n \in \mathbb{R}_+, t_n \in \mathbb{Z}_+$ represent the amount of inventory in stock for the current product line at the beginning of the week, and the number of weeks since the current product line has been offered, respectively. Let the decision-maker's weekly control actions, representing how much new inventory to add and whether or not to replace the current product line with a new one, be denoted by $u_n \in \mathbb{R}_+$ and $r_n \in \mathbb{B}$, respectively. Much like the water storage case, we model demand as an i.i.d. disturbance process $w_n \in \mathbb{R}_+$. The state dynamics for this model are identical to those in (2.3) and (2.4), with the interpretation of (2.3) reflecting that the inventory level increases with each additional product purchased, decreases by the realization of weekly demand (or the taking of a reset action), but cannot go below zero. Similarly, (2.4) now reflects the age of the current product line. Let $c, k_u > 0$ be the variable and fixed ordering costs associated with the current product line, respectively, and $c_r, k_r > 0$ be the associated variable and fixed ordering costs of purchasing units from the new product line. Furthermore, let $p, q > 0$ be the per-unit shortage and holding costs, respectively. Using these costs, states, dynamics, and discount factor $\gamma \in [0, 1)$, the decision-maker's problem can be formulated as follows:

$$\begin{aligned} \min \mathbb{E}\Big[\textstyle\sum_{n=0}^{\infty} \gamma^n \Big(&(cu_n + k_u) \cdot \mathbf{1}_{\mathbb{R}_+^*}(u_n) - p \cdot (\xi_n + u_n - w_n)^- \\ &+ q \cdot (\xi_n + u_n - w_n)^+ + (c_r x_n + k_r) \cdot r_n \Big) \Big] \end{aligned} \tag{2.13}$$

$$\begin{aligned} \text{s.t. } & x_{n+1} = (\xi_n + u_n - w_n)^+, && (2.14) \\ & t_{n+1} = \tau_n + 1, && (2.15) \\ & \xi_n = x_n \cdot (1 - r_n), && (2.16) \\ & \tau_n = t_n \cdot (1 - r_n), \qquad \text{for } n \geq 0 && (2.17) \\ & r_n \in \mathbb{B}, && (2.18) \\ & t_n \leq k, && (2.19) \\ & u_n \in \big[0, c_{\max} - \xi_n\big], && (2.20) \end{aligned}$$

where the function $\mathbf{1}_{\mathbb{R}_+^*} : \mathbb{R} \mapsto \mathbb{B}$ is defined as the indicator

$$\mathbf{1}_{\mathbb{R}_+^*}(u) = \begin{cases} 1, & \text{if } u > 0, \\ 0, & \text{otherwise.} \end{cases} \tag{2.21}$$

The total cost incurred in the n-th week comprises up to four components: (i) the purchasing cost $cu_n + k_u$, where c is the cost per unit ordered and k_u is the fixed cost associated with a positive inventory order; (ii) the shortage cost $p \cdot (\xi_n + u_n - w_n)^-$, representing the loss incurred when demand is unmet; (iii) the holding cost $q \cdot (\xi_n + u_n - w_n)^+$ for having excess inventory relative to actual demand; and (iv) the reset cost $c_r x_n + k_r$, where c_r is the waste penalty per unit discarded and k_r is the fixed cost associated with resetting the product line. The constraint (2.19) ensures that the product line is fully changed at least once every k weeks, and the constraint (2.20) ensures that the inventory capacity $c_{\max}$ is not exceeded.

What differentiates this inventory management problem from the classic (s, S) setting is the inclusion of reset controls. When the reset control action is taken at time n (i.e., $r_n = 1$), all products are removed, and the system is reverted to the state $(x_n, t_n) = (0, 0)$. This means that by solving this problem, the firm can change product lines earlier than planned if it is advantageous to do so. In other words, this problem determines the best business strategy tailored for limited-time products, encompassing not only the optimal inventory (i.e., how much stock to order each period) but also the optimal timing (i.e., when to switch the product line).

2.3 Structural Results

In this section, we analyze the structure of the reset control problem (2.2). We begin by describing the dynamic programming equations first. Let $J : \mathbb{R}^n \times \mathbb{Z}_+ \mapsto \mathbb{R}_+$ be the optimal cost-to-go function. That is, $J(x, t)$ is defined as the minimum value of (2.2) for the initial conditions $x_0 = x$ and $t_0 = t$. Let $J_0 = J(\zeta, 0)$ denote the cost-to-go from the reset state. Then the dynamic programming equations can be characterized using the following result from [54].

Proposition 1. *The dynamic programming equations for (2.2) are given by*

$$\begin{aligned} J(\zeta,0) &= \min_{u\in\mathcal{U}_0} \mathbb{E}\Big[g(\zeta,0,u,w) + \gamma J(h(\zeta,0,u,w),1)\Big], \\ J(x,t) &= \min\Big\{J_0 + \mathbb{E}\big(s(x,t,w)\big), \min_{u\in\mathcal{U}_t} \mathbb{E}\Big[g(x,t,u,w) + \gamma J(h(x,t,u,w),t+1)\Big]\Big\}, \\ J(x,k) &= J_0 + \mathbb{E}\big(s(x,k,w)\big), \end{aligned} \tag{2.22}$$

where the middle $J(x,t)$ holds for all x and $t = 0,\ldots,k-1$.

These are the new dynamic programming equations that result from introducing reset control. As shown here, excluding the last period ($t = k$) at which the inventory must be reset, we evaluate every period whether resetting is more favorable. The proof for the above equations can be found in [54].

Our main results in this section will prove that under a set of reasonable assumptions, the optimal fixed point policy $\pi^* : \mathbb{R}\times\mathbb{Z}_+ \mapsto \mathbb{B}\times\mathbb{R}$ to (2.22) is a threshold policy. Specifically, we show that it can be characterized by four time-dependent parameters $s_t, S_t, \sigma_t, \Sigma_t$ and a constant φ, and has the following form:

$$\pi^*(x,t) = \begin{cases} r=1, u=\varphi, & \text{for } x\in[0,\sigma_t), \\ r=0, u=S_t-x, & \text{for } x\in[\sigma_t,s_t), \\ r=0, u=0, & \text{for } x\in[s_t,\Sigma_t), \\ r=1, u=\varphi, & \text{for } x\in[\Sigma_t,\infty). \end{cases} \tag{2.23}$$

The intuition behind this policy is as follows: In the region $[0,\sigma_t)$, there is so little inventory that the cost of resetting is negligible, and thus resetting is the optimal decision. In the region $[\sigma_t, s_t)$, there is enough stock such that the reset cost is too high, and it is likely that the system will experience a shortage, so additional stock must be ordered. In the region $[s_t, \Sigma_t)$, there is enough stock present such that shortages are less likely than excess inventory, so no additional stock is ordered. Finally, in the region $[\Sigma_t, \infty)$, the stock level is so high that it will almost surely spoil, making it more beneficial to reset the system.

To prove these results, we will first describe a set of sufficient assumptions that guarantee the structure of the optimal policy. Then we will present the proof for the structure of the policy, showing that it follows from the particular structure of J by induction on the states.

Technical Assumptions

For our structural analysis of the optimal policy, we make the following technical assumptions.

Assumption 1 (State)**.** *The states are $x_n \in \mathbb{R}_+$, and the state dynamics are*

$$h(\xi_n, \tau_n, u_n, w_n) = (\xi_n + u_n - w_n)^+.$$

A1 assumes that the inventory is non-negative, meaning that unfilled demand at each stage is lost and not backlogged. This is represented by the system equation $x_{n+1} = \max\{0, \xi_n + u_n - w_n\}$, instead of $x_{n+1} = \xi_n + u_n - w_n$. The state dynamics in this assumption represent the remaining inventory from the last period, which is the positive part of the stock quantity stored in the inventory plus the amount of stock added, minus the amount of stock demanded.

Assumption 2 (Stage cost)**.** *The stage cost satisfies*

$$\mathbb{E}(g(\xi_n, \tau_n, u_n, w_n)) = cu_n + K \cdot \mathbf{1}_{\mathbb{R}_+^*}(u_n) + H(\xi_n + u_n, \tau_n) \tag{2.24}$$

for a function $H(\cdot, t)$ that satisfies $|\partial_x H(x + u, t)| \leq \kappa_t$ and $\partial_{xx} H(x + u, t) \geq m_t$. Here, $\partial_x H(\cdot, t)$ denotes the first derivative with respect to the first argument, and $\partial_{xx} H(\cdot, t)$ denotes the second derivative with respect to the first argument. The constants $c, K, \kappa_t \geq 0$ are non-negative, and the constants $m_t > 0$ are positive for $t = 0, \ldots, k-1$.

In **A2**, the stage cost is decomposed into three parts. The first term represents the ordering cost, the second term represents the fixed cost associated with a positive inventory order, and the third term represents the stage cost, which could be the holding/shortage cost. This assumption requires the function $H(\cdot, t)$ be strongly convex, which is crucial in proving the convexity of the value function with no reset and establishing the optimality of the (s, S) policy.

Assumption 3 (Reset cost)**.** *The known initial state is $\zeta = 0$. The reset cost $R(x, t) := \mathbb{E}(s(x, t, w))$ is concave and non-decreasing in x. Moreover, its derivative $\partial_x R(\cdot, t)$ is bounded as $|\partial_x R(x, t)| \leq \eta_t$, and is Lipschitz continuous as*

$$|\partial_x R(x_n, t_n) - \partial_x R(x'_n, t_n)| \leq L_t \cdot |x_n - x'_n| \tag{2.25}$$

for non-negative constants $\eta_t, L_t \geq 0$ for $t = 0, \ldots, k-1$.

A3 assumes that the entire inventory is emptied when the system is reset, and that the reset cost is concave and non-decreasing, reflecting the law of diminishing marginal utility. This assumption also requires the function $R(\cdot, t)$ be η_t-Lipshitz and L_t-smooth. These conditions are crucial in determining the threshold that ensures the optimal policy structure.

Assumption 4 (Input)**.** *The input constraints are*

$$\mathcal{U}_n = \{u : 0 \leq u \leq c_{\max} - \xi_n\},$$

where $c_{\max} \geq 0$ is a non-negative constant.

A4 enforces an upper bound $c_{\max}$ on the stock level that can be accommodated, acting as an input constraint.

Assumption 5 (Demand)**.** *The w_n are i.i.d., and the density of w_n is given by a function $f(w)$ that satisfies: $f(w) = 0$ for all $w < 0$, is Lipschitz continuous for $w \geq 0$ with a non-negative constant L, and is bounded such that $|f(w)| \leq P$ for all $w \geq 0$ with a non-negative constant P.*

A5 assumes that demand is non-negative and its density is supported on the set of non-negative real numbers. The density function $f(w)$ is required to be Lipschitz continuous and bounded from above. This assumption is fairly mild and accommodates a wide range of probability distributions, such as exponential, gamma, and truncated normal distributions.

We will use the above five assumptions for our theoretical analysis. Furthermore, we define the following function for convenience:

$$G(z,t) = cz + H(z,t) + \gamma \mathbb{E}(J((z-w)^+, t+1)).$$

Note that under the above definitions, we have the relation

$$\min_{u \in \mathcal{U}_t} \mathbb{E}\Big[g(x,t,u,w) + \gamma J(h(x,t,u,w), t+1)\Big] = \min_{z \in [x, c_{\max}]} G(z,t) + K \cdot \mathbf{1}_{\mathbb{R}_+^*}(z-x) - cx.$$

Proof Technique

To prove our main result, we will use proof by induction to show that the particular structure of the policy and value function is preserved. Our approach will be to show a series of results that hold under a temporary assumption (i.e., the induction hypothesis). These results will then be proved to hold under assumptions **A1**–**A5** in a final theorem that concludes our proof. The temporary assumption is as follows:

Assumption T (Temporary)**.** *For any fixed $t \in \{0, \ldots, k-1\}$, we can represent $J(z, t+1)$ as*

$$J(z,t+1) = J_i(z,t+1) \text{ if } z \in Z_{i,t} \text{ for } i = 1, \ldots, n_t, \tag{2.26}$$

where $Z_{i,t} = [z_{i-1,t}, z_{i,t}]$ forms a partition of the domain $[0, c_{\max}]$ with

$$0 = z_{0,t} < z_{1,t} < \cdots < z_{n,t} = c_{\max}. \tag{2.27}$$

Moreover, the derivative of each piece $\partial_z J_i(z, t+1)$ is absolutely bounded by a finite, non-negative constant M_{t+1} and is Lipschitz continuous.

AT assumes that the value function is continuous and piecewise differentiable. It can also be partitioned into multiple pieces, where the derivatives of each piece are Lipschitz and bounded. This assumption is necessary to characterize the policy structure, which is complicated by the non-convexity induced by the reset control. These pieces represent different regions of the optimal policy, which we investigate in the subsequent analysis.

Proof of the Optimal Policy Structure

To begin the proof, we first present a result on the smoothness of the cost-to-go function.

Proposition 2. *If **A1**–**A5** and **AT** hold, then $\mathbb{E}(J((z-w)^+, t+1))$ has a derivative that is Lipschitz continuous with a non-negative constant $(Lc_{\max} + P)M_{t+1}$, and the derivative is absolutely bounded by M_{t+1}.*

The detailed proof of this proposition can be found below. Here, we provide a brief overview. First, we show that $\mathbb{E}(J((z-w)^+, t+1))$ is differentiable using **AT** and the Leibniz integral rule for derivatives. Subsequently, by introducing an auxiliary function with the same expectation as the cost-to-go function, we establish the Lipschitz continuity of this derivative. Integrating over the domain and applying Lemma 2 leads to the desired result. Finally, we prove that the derivative is absolutely bounded by M_{t+1} using standard integral and absolute value inequalities.

Proof. We first show that $\mathbb{E}(J((z-w)^+, t+1))$ is differentiable. Consider any $z \in [0, c_{\max}]$, and observe that

$$\begin{aligned} J((z-w)^+, t+1) &= \begin{cases} J_1(0, t+1), & \text{if } z - w \leq 0 \\ J_i(z-w, t+1), & \text{if } z - w \in Z_{i,t} \end{cases} \\ &= \begin{cases} J_1(0, t+1), & \text{if } w \geq z \\ J_i(z-w, t+1), & \text{if } w \in [z - z_{i,t}, z - z_{i-1,t}] \end{cases} \end{aligned} \tag{2.28}$$

Now, let m be the smallest integer such that $z \in Z_{m,t}$. Then the expectation is given by

$$\begin{aligned} \mathbb{E}(J((z-w)^+, t+1)) = \int_0^\infty J((z-w)^+, t+1) f(w) dw = \int_0^{z - z_{m-1,t}} J_m(z-w, t+1) f(w) dw \\ + \sum_{i=1}^{m-1} \int_{z - z_{i,t}}^{z - z_{i-1,t}} J_i(z-w, t+1) f(w) dw + \int_z^\infty J_1(0, t+1) f(w) dw. \end{aligned} \tag{2.29}$$

Since by assumption both $J_i(z, t+1)$ and $\partial_z J_i(z, t+1)$ are continuous for $z \in Z_{i,t}$, using the

Leibniz integral rule gives

$$\begin{aligned}
&\partial_z \mathbb{E}(J((z-w)^+, t+1)) = J_m(z_{m-1,t}, t+1) f(z - z_{m-1,t}) + \int_0^{z-z_{m-1,t}} \partial_z J_m(z-w, t+1) f(w) dw \\
&+ \sum_{i=1}^{m-1} \Big(J_i(z_{i-1,t}, t+1) f(z - z_{i-1,t+1}) - J_i(z_{i,t}, t+1) f(z - z_{i,t}) + \int_{z-z_{i,t}}^{z-z_{i-1,t}} \partial_z J_i(z-w, t+1) f(w) dw \Big) \\
&\quad - J_1(0, t+1) f(z) + \int_z^\infty \partial_z J_1(0, t+1) f(w) dd = \int_0^{z-z_{m-1,t}} \partial_z J_m(z-w, t+1) f(w) dw \\
&\quad + \sum_{i=1}^{m-1} \int_{z-z_{i,t}}^{z-z_{i-1,t}} \partial_z J_i(z-w, t+1) f(w) dw + \int_z^\infty \partial_z J_1(0, t+1) f(w) dw \\
&\quad = \int_0^\infty \partial_z J((z-w)^+, t+1) f(w) dw, \qquad (2.30)
\end{aligned}$$

where the last equality follows from the fact that $f(\cdot)$ is Lipschitz continuous and $J(\cdot, t+1)$ is differentiable almost everywhere on its domain since it is piecewise differentiable by assumption. Hence, we can conclude that $\mathbb{E}(J((z-w)^+, t+1))$ is differentiable.

Next, we show that this derivative is Lipschitz continuous. Define

$$J_A(x, t+1) = \begin{cases} J(0, t+1), & x \leq 0 \\ J(x, t+1), & x \in [0, c_{\max}] \\ J(c_{\max}, t+1), & x \geq c_{\max} \end{cases} \qquad (2.31)$$

where the subscript A indicates "auxiliary". Then for $z \in [0, c_{\max}]$ and $w \in [0, \infty)$, we have $J((z-w)^+, t+1) = J_A(z-w, t+1)$. Note that $\mathbb{E}(J_A(z-w, t+1)) = \int_0^\infty J_A(z-w, t+1) f(w) dw = \int_{-\infty}^\infty J_A(z-w, t+1) f(w) dw = (J_A(\cdot, t) * f)(z)$ (i.e., the convolution of $J_A(\cdot, t+1)$ and $f(\cdot)$), where the second equality follows since $f(w) = 0$ for $w \in (-\infty, 0)$ by assumption. This implies that

$$\begin{aligned}
\partial_z (J_A(\cdot, t+1) * f)(z) &= \partial_z \mathbb{E}(J_A(z-w, t+1)) \\
&= \partial_z \mathbb{E}(J((z-w)^+, t+1)) \\
&= \int_0^\infty \partial_z J((z-w)^+, t+1) f(w) dw \\
&= \int_0^\infty \partial_z J_A(z-w, t+1) f(w) dw \\
&= (\partial_z J_A(\cdot, t+1) * f)(z).
\end{aligned} \qquad (2.32)$$

Observe that

$$\partial_z J_A(z, t+1) = \begin{cases} 0, & x < 0 \\ \partial_z J(z, t+1), & z \in (0, c_{\max}) \\ 0, & x > c_{\max} \end{cases} \qquad (2.33)$$

This means $\int_{-\infty}^{\infty} |\partial_z J_A(z,t+1)|dz = \int_0^{c_{\max}} |\partial_z J(z,t+1)|dz \leq c_{\max} M_{t+1}$ since we had assumed $|\partial_z J(z,t+1)| \leq M_{t+1}$. Thus, Lemma 2 implies $\partial_z \mathbb{E}(J((z-w)^+, t+1))$ is Lipschitz with constant $(Lc_{\max} + P)M_{t+1}$.

Lastly, we show that the derivative is absolutely bounded by M_{t+1}. From (2.32), we have

$$\begin{aligned}
|\partial_z (J_A(\cdot, t+1) * f)(z)| &= |(\partial_x J_A(x, t+1) * f)(z)| \\
&= \left| \int_0^{\infty} \partial_z J((z-w)^+, t+1) f(w) dw \right| \\
&\leq \int_0^{\infty} \left|\partial_z J((z-w)^+, t+1)\right| \cdot \left|f(w)\right| dw \\
&\leq M_{t+1} \int_0^{\infty} \left|f(w)\right| dw \\
&\leq M_{t+1},
\end{aligned} \tag{2.34}$$

where in the last line we have used the facts that $f(w)$ is a probability density, and hence non-negative, and integrates to one. $\square$

If the cost-to-go function is piecewise differentiable and the derivative of each piece is Lipschitz continuous, the expected cost-to-go has a Lipschitz derivative. This is because taking the expectation smooths the function through convolution. Next, we show a result concerning the structure of $G(z,t)$.

Proposition 3. *If **A1**–**A5** and **AT** hold, then the function $G(z,t)$ is continuous and convex on $z \in [0, c_{\max}]$ for all fixed γ such that $0 \leq \gamma \leq m_t/((Lc_{\max} + P)M_{t+1})$.*

The complete proof of this proposition is provided below, and here we present a brief outline. To establish this result, we observe that G comprises a linear function, a twice differentiable function, and, as indicated by Proposition 2, a function with a Lipschitz derivative. Since this implies that G is absolutely continuous in its first argument, we demonstrate that it has a non-decreasing derivative, meaning convexity within the specified interval.

Proof. We first recall the definition

$$G(z,t) = c \cdot z + H(z,t) + \gamma \mathbb{E}(J((z-w)^+, t+1)).$$

Note that $c \cdot z$ is linear, that $H(z,t)$ is twice differentiable by assumption, and that $\mathbb{E}(J((z-w)^+, t+1))$ has a Lipschitz derivative by Proposition 2. Hence, $G(\cdot,t)$ is absolutely continuous in its first argument. Since the domain $[0, c_{\max}]$ is closed and bounded, this means that $G(\cdot,t)$ is convex if it has a non-decreasing derivative (see page 115 of [24]). From **A2**, we have that $\partial_{zz}(c \cdot z + H(z,t)) \geq m_t$, and Proposition 2 gives that $\partial_z \mathbb{E}(J((z-w)^+, t+1))$ is Lipschitz with constant $(Lc_{\max} + P)M_{t+1}$. Thus, by Lemma 1, we have that $\partial_z G(z,t)$ is non-decreasing for γ such that $0 \leq \gamma \leq m_t/((Lc_{\max} + P)M_{t+1})$. $\square$

This proposition intuitively states that the sum of the expected single-stage cost (strongly convex) and the expected cost-to-go (non-convex but smooth) maintains convexity under suitable conditions, which can be computed using the parameters of the two functions. For example, $f_1(x) = x^2/2 + \gamma \sin(x)$ is convex for $\gamma \leq 1$, whereas $f_2(x) = x^2/2 - \gamma|x|$ will never be convex unless $\gamma = 0$. This exemplifies the critical role of the discount factor, with the smoothness of the expected cost-to-go, as proved in Proposition 2, being even more crucial in preserving the convexity of the sum.

Our next result generalizes the known result (i.e., Lemma 4.2.1 in [4]) in inventory management models to the case of a closed feasible set.

Proposition 4. *If $G(\cdot, t)$ is continuous and convex on $[0, c_{\max}]$, there exists S such that $G(S, t) \leq G(z, t)$ for all $z \in [0, c_{\max}]$. Furthermore, let $s = \inf\{z \in [0, S] \mid G(S, t) + K = G(z, t)\}$. If such s exists, then*

1. *$G(S, t) + K = G(s, t)$*

2. *$G(\cdot, t)$ is a non-increasing function on $[0, s]$*

3. *$G(S, t) + K \leq G(z, t)$ for all $z \in [0, s]$*

4. *$G(y, t) \leq G(z, t) + K$ for all y, z with $s \leq y \leq z \leq c_{\max}$*

The full proof of this proposition can be found below, but here we present a sketch of the proof. First, we note that because $G(\cdot, t)$ is continuous on the closed interval, it must have a minimizer, which shows the existence of S. Next, by continuity of $G(\cdot, t)$, we note that there may exist s such that $G(s, t) = G(S, t) + K$, thus proving **C1**. We then prove **C3** using the convexity of $G(\cdot, t)$. Then using these conditions and convexity, we prove **C2**. Finally, **C4** follows from combining these results.

Proof. Since $G(\cdot, t)$ is continuous on the closed interval $[0, c_{\max}]$, there exists a minimizer of $G(\cdot, t)$. Let $S \in \arg\min_{z \in [0, c_{\max}]} G(z, t)$. For the remainder of the proof, we consider the case where s exists. By the continuity of $G(\cdot, t)$ and the definition of s, we must have $G(S, t) + K = G(s, t)$. Next, for all $z \in [0, s]$, there exists $\lambda \in [0, 1]$ such that $s = \lambda z + (1 - \lambda)S$ and

$$G(s, t) \leq \lambda G(z, t) + (1 - \lambda)G(S, t) \leq \lambda G(z, t) + (1 - \lambda)G(z, t) = G(z, t), \tag{2.35}$$

where the inequalities follow from the convexity of $G(\cdot, t)$ and the definition of S, respectively. Next, consider any z_1 and z_2 with $0 \leq z_1 \leq z_2 \leq s$, and note that there exists $\lambda \in [0, 1]$ such that $z_2 = \lambda z_1 + (1 - \lambda)s$. Then we have

$$G(z_2, t) \leq \lambda G(z_1, t) + (1 - \lambda)G(s, t) \leq \lambda G(z_1, t) + (1 - \lambda)G(z_1, t) = G(z_1, t), \tag{2.36}$$

where the second inequality follows by (2.35). This shows $G(\cdot, t)$ is non-increasing on $[0, s]$. Since $G(\cdot, t)$ is continuous, this means by definition of s that $G(S, t) + K \leq G(s, t)$. Using the above result that $G(\cdot, t)$ is non-increasing on $[0, s]$, we have $G_t(S) + K \leq G_t(z)$ for all $z \in [0, s]$. By definition of s, for any $y \in [s, c_{\max}]$, we have $G_t(S) + K \geq G_t(y)$. However, we also know that $G_t(S) \leq G_t(z)$ for any $z \in [y, c_{\max}]$. Combining the two yields the last part of the result. $\square$

This proposition proves the optimality of the (s, S) policy when the functions $G(\cdot, t)$ are convex and have upper and lower bounds on the allowable values of the stock. We now extend this result to the functional structure present in the reset control problem.

Proposition 5. *Suppose $0 \leq \gamma \leq m_t/((Lc_{\max} + P)M_{t+1})$ for all $t = 0, \ldots, k-1$. If **A1**–**A5** and **AT** hold, then an optimal policy has the following four-stage structure with $(s_t, S_t, \sigma_t, \Sigma_t)$ thresholds:*

1. *If $x_t \in [0, \sigma_t)$, then $r_t^* = 1$ and $u_t^* = \varphi$*
2. *If $x_t \in [\sigma_t, s_t)$, then $r_t^* = 0$ and $u_t^* = S_t - x_t$*
3. *If $x_t \in [s_t, \Sigma_t)$, then $r_t^* = 0$ and $u_t^* = 0$*
4. *If $x_t \geq \Sigma_t$, then $r_t^* = 1$ and $u_t^* = \varphi$*

where $\varphi \in [0, c_{\max}]$ is a constant.

The full proof of this result can be found below, but here we present a sketch. First, we define φ as the optimal value of u from the known reset state ζ with $t = 0$. Because of this, from Proposition 1, we note that if the optimal reset action is $r^*(x, t) = 1$, that means $u^*(x, t) = \varphi$. Hence, we focus the proof on determining the reset policy. We consider the set I_R, which is the set of all states for which a reset action would not be taken. We then consider the case when this set is empty, meaning a reset action should be taken for all states. When this set is not empty, we define σ_t, σ_t' as the infimum and supremum of I_R, respectively. By Proposition 4, we note that there exists $S_t' = \arg\min_{z \in [0, c_{max}]} G(z, t)$. However, we are not guaranteed that $s_t' = \sup\{z \in [0, S_t'] \mid G(S_t', t) + K \leq G(z, t)\}$ to exist. In the case where s_t' does not exist, we demonstrate that the threshold s_t should be set to σ_t because $u^*(x, t) = 0$ will be the minimizer for all states in this case. Next, we show that if s_t' does exist, then we should set the threshold s_t to s_t', since the optimal action $u^*(x, t)$ will be to order up to level S_t' if $x < s_t'$, and order nothing otherwise. Using continuity, convexity, Proposition 3, and **A3**, we show that the region of the state space where it is optimal not to take a reset action is the interval $[\sigma_t, \sigma_t')$. This means that the optimal policy will use the previously defined σ_t as one of the thresholds, and will use $\Sigma_t = \sigma_t'$ if $\sigma_t' < c_{max}$ and $\Sigma_t = +\infty$ otherwise. Finally, we show that the threshold S_t should be set to S_t' based on its definition and the properties of s_t and σ_t.

Proof. We first define the constant

$$\varphi = \arg\min_{u \in \mathcal{U}_0} \mathbb{E}\Big[g(\zeta, 0, u, w) + \gamma J(h(\zeta, 0, u, w), 1)\Big]. \tag{2.37}$$

From Proposition 1, we have $u_t^*(x) = \varphi$ whenever $r_t^*(x) = 1$. So, our proof will be structured around determining an optimal reset policy $r_t^*(x)$. We define the functions

$$\begin{aligned} J_R(x,t) &= J_0 + \mathbb{E}(s(x,t,w)), \\ J_N(x,t) &= \min_{u\in[0,c_{\max}-x]} G(x+u,t) + K \cdot \mathbf{1}_{\mathbb{R}_+^*}(u) - c \cdot x, \end{aligned} \tag{2.38}$$

where the subscript R indicates "reset", and the subscript N indicates "no reset". Observe that these are the value functions corresponding to $r_t = 1$ and $r_t = 0$, respectively. Next, define the set $I_R = \{x \in [0, c_{\max}] \mid J_R(x,t) \geq J_N(x,t)\}$, and observe that I_R is bounded by construction. Now, we consider two cases:

The first case is when $I_R = \emptyset$. Then by definition of I_R, we have $J_R(x,t) < J_N(x,t)$ for all $x \in [0, c_{\max}]$. For this case, an optimal policy is to choose $\sigma_t = s_t = \Sigma_t = +\infty$ since it is optimal to choose $r_t^*(x) = 1$ for all $x \in [0, c_{\max}]$.

The second case is when $I_R \neq \emptyset$. Let $\sigma_t = \inf\{x \in [0, c_{\max}] \mid J_R(x,t) \geq J_N(x,t)\}$ and $\sigma_t' = \sup\{x \in [0, c_{\max}] \mid J_R(x,t) \geq J_N(x,t)\}$, and note that σ_t and σ_t' are finite since I_R is bounded and non-empty. Let $S_t' = \arg\min_{z\in[0,c_{\max}]} G(z,t)$ and define $s_t' = \sup\{z \in [0, S_t'] \mid G(S_t',t) + K \leq G(z,t)\}$. Note that by Proposition 4, S_t' is guaranteed to exist whereas s_t' may or may not exist. We consider two subcases based on the existence of s_t':

The first subcase is when s_t' does not exist. Then an optimal policy chooses $s_t = \sigma_t$ because in this subcase, a minimizer to the optimization problem defining $J_N(x,t)$ is $u_t^*(x) = 0$ for all $x \in [0, c_{\max}]$. Note that by setting s_t equal to σ_t, the second policy region vanishes, which ensures that $u_t^*(x) = 0$ for all $x \in [0, c_{\max}]$ with our policy.

The rest of the proof considers the second subcase in which s_t' exists. We set $s_t = s_t'$, since

$$u_t^*(x) = \begin{cases} S_t' - x, & \text{if } 0 \leq x < s_t' \\ 0, & \text{if } s_t' \leq x \leq c_{\max} \end{cases} \tag{2.39}$$

is optimal for the optimization problem defining $J_N(x,t)$. We now observe that for $x \in [0, s_t)$, $J_N(x,t) = G(S_t',t) + K - c \cdot x$, which is non-increasing in x. For $x \in [s_t, c_{\max}]$, $J_N(x,t) = G(x,t) - c \cdot x$, which is convex in x since $G(x,t)$ is convex by Proposition 3. Clearly, $J_N(x,t)$ is continuous for $x \in [0, s_t)$ since it is linear in this region, and $J_N(x,t)$ is continuous for $x \in (s_t, c_{\max}]$ since $G(\cdot,t)$ is continuous by Proposition 3. The only question about continuity occurs at $x = s_t$. Since $G(S_t',t) + K = G(s_t,t)$ by definition, we have that $G(S_t',t) + K - c \cdot s_t = G(s_t,t) - c \cdot s_t$. This proves that $J_N(x,t)$ must be continuous at $x = s_t$ since the left and right side of the last equality with $G(\cdot,t)$ are the limits of $J_N(\cdot)$ in the two respective regions. This means $J_N(x,t)$ is continuous on $x \in [0, c_{\max}]$. Because $J_R(\cdot,t)$ is continuous by **A3**, this means that $J_R(\sigma_t,t) \geq J_N(\sigma_t,t)$ and $J_R(\sigma_t',t) \geq J_N(\sigma_t',t)$ by the definitions of σ_t and σ_t'.

Next, consider any $x \in [\sigma_t, s_t]$ (note that our argument still holds even if this set is empty). Since $J_N(x,t)$ is non-increasing for $x \in [0, s_t]$, and since $J_R(x,t)$ is non-decreasing in x by **A3**, we have

$$J_R(x,t) \geq J_R(\sigma_t,t) \geq J_N(\sigma_t,t) \geq J_N(x,t) \tag{2.40}$$

for $x \in [\sigma_t, s_t]$. This argument also implies that $J_R(s_t,t) \geq J_N(s_t,t)$. Next, consider any $x \in [s_t, \sigma'_t]$, and observe that there exists $\mu \in [0,1]$ such that $x = \mu s_t + (1-\mu)\sigma'_t$. Since we showed above that $J_N(\cdot,t)$ is convex, this means that for any $x \in [s_t, \sigma'_t]$, we have

$$\begin{aligned} J_N(x,t) = J_N(\mu s_t + (1-\mu)\sigma'_t, t) &\leq \mu J_N(s_t,t) + (1-\mu) J_N(\sigma'_t, t) \\ &\leq \mu J_R(s_t,t) + (1-\mu) J_R(\sigma'_t, t) \\ &\leq J_R(\mu s_t + (1-\mu)\sigma'_t, t) = J_R(x,t), \end{aligned} \tag{2.41}$$

where the last inequality follows because $J_R(\cdot,t)$ is concave by **A3**. Combining the above shows that $J_R(x,t) \geq J_N(x,t)$ for all $x \in [\sigma_t, \sigma'_t]$.

Now, observe that $J_R(x,t) < J_N(x,t)$ for $x \in [0,\sigma_t)$ and for $x \in (\sigma'_t, c_{\max}]$. (If this last statement were not true, then we could choose an $x' \in [0,\sigma_t)$ or an $x' \in (\sigma'_t, c_{\max}]$ such that $J_R(x',t) \geq J_N(x',t)$. Hence, we would have $x' \in I_R$, which reaches a contradiction since, by the definition of σ_t and σ'_t, we would have $\sigma_t \leq x'$ or $\sigma'_t \geq x'$.) Consequently, an optimal policy uses this value of σ_t. If $\sigma'_t < c_{\max}$, then an optimal policy uses $\Sigma_t = \sigma'_t$, and if $\sigma'_t = c_{\max}$, then an optimal policy chooses $\Sigma_t = +\infty$.

Finally, we must choose a correct value for S_t. We must consider three sub-subcases. The first sub-subcase is when $S'_t < \sigma_t$. Since $s'_t \leq S'_t$ by definition of s'_t, this means the second policy region is empty when $s_t = s'_t$, which ensures that $u^*_t(x) = 0$ for all $x \in [\sigma_t, \sigma'_t]$ with our policy. Thus, we can choose $S_t = S'_t$. (Any arbitrary choice of S_t would give an optimal policy because the corresponding region is empty.) The second sub-subcase is when $\sigma_t \leq S'_t \leq \sigma'_t$. Then an optimal policy chooses $S_t = S'_t$. The third sub-subcase is when $S'_t > \sigma'_t$. By definition of σ_t and σ'_t, this means $J_R(S'_t,t) < J_N(S'_t,t)$ and $\sigma'_t \geq \sigma_t$. Thus, in this sub-subcase, we have

$$J_N(\sigma'_t,t) = G(\sigma'_t,t) - c\cdot\sigma'_t \geq G(S'_t,t) - c\cdot S'_t = J_N(S'_t,t), \tag{2.42}$$

where we have used the definition of S'_t in the last inequality. Recalling that $J_R(x,t) \geq J_N(x,t)$ for all $x \in [\sigma_t, \sigma'_t]$, we have

$$J_R(S'_t,t) < J_N(S'_t,t) \leq J_N(\sigma'_t,t) \leq J_R(\sigma'_t,t). \tag{2.43}$$

However, this last statement is a contradiction since $\sigma'_t < S'_t$ and $J_R(\cdot,t)$ is non-decreasing by **A3**. Therefore, this sub-subcase is not possible. □

The regions in Proposition 5 represent that (i) the inventory is sufficiently small such that reset cost is negligible, and thus resetting the system and ordering a one-step quantity is optimal; (ii) the stock is sufficient such that ordering up to a given quantity and consuming the majority prior to the stock becoming too obsolete or contaminated is optimal; (iii) the stock in the system is sufficient to satisfy possible future demands, and thus resetting the system or ordering additional stock is not optimal (a.k.a., do-nothing region); and (iv) there is so much stock in the inventory that it will almost certainly never be consumed, hence making it optimal to reset the system and reorder up to a baseline amount. Next, we show how this policy influences the structure of the cost-to-go function.

I_R	s'_t	S'_t	σ_t	s_t	S_t	Σ_t	Interpretation
$I_R = \emptyset$			$+\infty$	$+\infty$	-	$+\infty$	Always reset
$I_R \neq \emptyset$	$\nexists s'_t$		σ_t	σ_t	-	σ'_t ($\sigma'_t < c_{\max}$), $+\infty$ ($\sigma'_t = c_{\max}$)	Missing $[\sigma_t, s_t)$
	$\exists s'_t$	$S'_t < \sigma_t$	σ_t	s'_t	-	σ'_t or $+\infty$	Missing $[\sigma_t, s_t)$
		$\sigma_t \leq S'_t \leq \sigma'_t$	σ_t	s'_t	S'_t	σ'_t or $+\infty$	All four stages
		$S'_t > \sigma'_t$	-	-	-	-	Not possible

Figure 2.1: Proof sketch for Proposition 5

Proposition 6. *Suppose* $0 \leq \gamma \leq m_t/((Lc_{\max} + P)M_{t+1})$ *for all* $t = 0, \ldots, k-1$. *If* ***A1***–***A5*** *and* ***AT*** *hold, then the function* $J(x,t)$ *for* $t = 0, \ldots, k-1$ *has the following form with* $(s_t, S_t, \sigma_t, \Sigma_t)$ *thresholds:*

$$J(x,t) = \begin{cases} J_0 + \mathbb{E}(s(x,t,w)), & x \in [0, \sigma_t) \\ c(S_t - x) + K + H(S_t, t) + \gamma \mathbb{E}(J((S_t - w)^+, t+1)), & x \in [\sigma_t, s_t) \\ H(x,t) + \gamma \mathbb{E}(J((x-w)^+, t+1)), & x \in [s_t, \Sigma_t) \\ J_0 + \mathbb{E}(s(x,t,w)), & x \geq \Sigma_t \end{cases} \tag{2.44}$$

Furthermore, each piece has a Lipschitz derivative, and the derivative of each piece is absolutely bounded by $M_t = \eta_t + c + \kappa_t + m_t/(Lc_{\max} + P)$ *for* $t = 0, \ldots, k-1$.

The full proof of this proposition can be found below. Here, we present a proof sketch. The functional structure of the policy follows directly as a consequence of Proposition 5, so we focus on showing the properties of the derivative of the cost-to-go function. We examine each piece of the function individually. The result is proven for the first and fourth pieces as a consequence of **A3**, and we note that in the second region, $J(x,t)$ is linear and thus satisfies the result. For the third piece, the result follows as a consequence of Proposition 2 and **A2**.

Proof. The policy structure in Proposition 5 implies that the value function $J(x,t)$ takes the form (28). This means we have to analyze the derivative of at most four pieces. In the first and fourth regions, by **A3**, we have that $J(x,t)$ has a derivative that is Lipschitz with constant L_t, and that the derivative is absolutely bounded by η_t. In the second region, $J(x,t)$ is linear in x and thus has a Lipschitz derivative that is absolutely bounded by c. In the third region, we know that the first term has a Lipschitz derivative because it is twice differentiable by **A2**, and that the second term has a Lipschitz derivative by Proposition 2. Thus, the third region has a Lipschitz derivative because it is the sum of two functions with Lipschitz derivatives. To bound its derivative, we note that $\partial_x H(x,t)$ is absolutely bounded by κ_t by **A2**. From Proposition 2, we know that the derivative of the second term is absolutely bounded by γM_{t+1}. Thus, the third piece has a derivative that is absolutely bounded by $\kappa_t + \gamma M_{t+1}$. This means the derivative of each piece is absolutely bounded by $\eta_t + c + \kappa_t + \gamma M_{t+1}$. As $0 \leq \gamma \leq m_t/((Lc_{\max} + P)M_{t+1})$, the derivative of each piece is absolutely bounded by $M_t = \eta_t + c + \kappa_t + m_t/(Lc_{\max} + P)$. □

The policy structure from Proposition 5 implies that the value function has up to four thresholds. The first two thresholds, namely s and S, are obtained from the value function when there is no reset, which is linear and then convex. The points of intersection that result from juxtaposing the value functions with and without reset give us the last two thresholds, σ and Σ, thereby yielding an optimal cost-to-go function that is continuous and piecewise differentiable, as shown in Figure 2.2. We now complete the proof by induction and demonstrate that the optimal policy and cost-to-go functions indeed follow the forms described above.

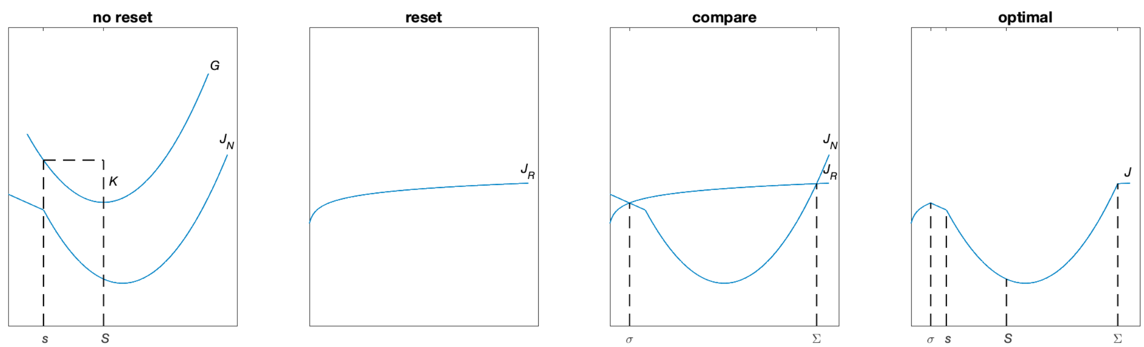

linear then convex (s, S) thresholds	concave	points of intersection (σ, Σ) thresholds	continuous, piecewise differentiable

Figure 2.2: Proof sketch for Proposition 6

Theorem 1. *Suppose γ is such that $0 \leq \gamma \leq \min\{\gamma_t \mid t \in \{0, \dots, k-1\}\}$ for*

$$
\begin{aligned}
\gamma_{k-1} &= \frac{m_{k-1}}{(Lc_{\max} + P) \cdot \eta_k}, \\
\gamma_t &= \frac{m_t}{(Lc_{max} + P) \cdot (\eta_{t+1} + c + \kappa_{t+1}) + m_{t+1}} \quad \text{for } t = 0, \dots, k-2.
\end{aligned}
\tag{2.45}
$$

*If **A1**–**A5** hold, then the policy described in Proposition 5 is optimal, and the value function has the structure described in Proposition 6.*

The complete proof of the theorem can be found below, but here we present a sketch. First, we note that when $t = k$, the optimal policy will choose to reset the system for all $x \in [0, c_{max}]$, which means $J(x, k) = J_0 + \mathbb{E}(s(x, k, w))$. The remaining proof follows from Propositions 5 and 6, along with induction on t.

Proof. First, note that we can choose $M_k = \eta_k$ using **A3**, since an optimal policy at $t = k$ is to choose $r_k^*(x) = 1$ for all $x \in [0, c_{\max}]$, which means $J(x, k) = J_0 + \mathbb{E}(s(x, k, w))$. Inductively applying Propositions 5 and 6 implies the desired result. □

This concludes our proof by induction, from which we obtain the threshold for the discount factor that guarantees the structural properties of the optimal policy. We have shown that assumptions **A1**-**A5** are sufficient to establish the policy structure, and the temporary assumption **AT** can be safely removed. Proposition 6 demonstrates that the value function can be partitioned into at most four pieces, each of which is Lipschitz continuous and absolutely bounded.

Policy Structure for Water Storage and Retail Management Problems

In this section, we provide sufficient conditions for both water storage and retail problems to exhibit the four-threshold structure of the reset problem. We outline the conditions under which each model satisfies assumptions **A1**-**A5**. First, we discuss the conditions for the water storage problem.

Proposition 7. *Suppose **A5** (the assumption about the distribution of w_n) holds, and that $f(w) > 0$ for all $w \in [0, c_{\max}]$. If $p - q(t) > 0$ for all $t = 0, \dots, k-1$, then the water storage problem described in Section 2.2 satisfies **A1**–**A5**.*

The complete proof of this proposition can be found below, but here we present a sketch. First, we establish that **A1**, **A4**, and **A5** hold based on the problem definition and assumptions. We then demonstrate that by reformulating the stage cost and setting the fixed cost to zero, the problem satisfies **A2**. Finally, we show that **A3** follows from the reset conditions in the water storage problem.

Proof. **A1** and **A4** hold by the definition of the water problem, and **A5** holds by assumption. To show **A2**, we first note that

$$\begin{aligned} \xi_n + u_n - (\xi_n + u_n - w_n)^+ &= \xi_n + u_n - (\xi_n + u_n - w_n) + (\xi_n + u_n - w_n)^- \\ &= w_n + (\xi_n + u_n - w_n)^-. \end{aligned} \tag{2.46}$$

This means the single stage cost can be rewritten as

$$\begin{aligned} g(\xi_n, \tau_n, u_n, w_n) &= cu_n - p \cdot (\xi_n + u_n - w_n)^- + q(\tau_n) \cdot (\xi_n + u_n - (\xi_n + u_n - w_n)^+) \\ &= cu_n - (p - q(\tau_n)) \cdot (\xi_n + u_n - w_n)^- + q(\tau_n) \cdot w_n. \end{aligned} \tag{2.47}$$

Hence, (2.24) is satisfied by setting $K = 0$ and

$$H(z, t) = -(p - q(t)) \int_z^\infty (z - w) f(w) dw + q(t) \cdot \mathbb{E}(w). \tag{2.48}$$

Note that using the Leibniz integral rule twice gives

$$\begin{aligned}\partial_z H(z,t) &= -(p-q(t))\int_z^\infty f(w)dw,\\ \partial_{zz} H(z,t) &= (p-q(t))\cdot f(z).\end{aligned} \tag{2.49}$$

Thus, we have $|\partial_z H(z,t)| = |-(p-q(t))\int_z^\infty f(w)dw| \le p-q(t)$ since $f(w)$ is a density. Since $H(\cdot,t)$ is twice differentiable, this means it is absolutely continuous. Thus, $H(\cdot,t)$ is strongly convex on $[0,c_{\max}]$ since $\partial_{zz}H(z,t) = (p-q(t))\cdot f(z) \ge (p-q(t))\cdot \inf_{z\in[0,c_{\max}]} f(z) = (p-q(t))\cdot \min_{z\in[0,c_{\max}]} f(z) > 0$ for $z \in [0,c_{\max}]$, where the second equality holds because a Lipschitz continuous function attains its minimum in a compact domain. This shows that **A2** holds. To show that **A3** holds, we first note that the water storage problem as described in Section 2.1 implies $\zeta = 0$. Next, note that the reset cost $R(x,t) = c_r x$ is linear. Hence, the conditions of **A3** follow immediately. □

The following proposition presents sufficient conditions for the retail problem to have the same threshold policy structure.

Proposition 8. *Suppose* ***A5*** *holds, and that $f(w) > 0$ for all $w \in [0,c_{\max}]$. Then the retail management problem described in Section 2.2 satisfies* ***A1–A5****.*

The full proof of this proposition can be found below, but here we provide a brief outline. First, we note that **A1**, **A4**, and **A5** hold based on the problem definition and assumptions. Next, letting the fixed cost $K = k_u$ and using the Leibniz integral rule, we demonstrate that **A2** is satisfied. Finally, we show that **A3** follows from the reset cost structure in the retail management problem.

Proof. **A1** and **A4** hold by the definition of the retail problem, and **A5** holds by assumption. To show **A2**, we first note that the single stage cost can be written as

$$g(\xi_n,\tau_n,u_n,w_n) = cu_n + k_u\cdot \mathbf{1}_{\mathbb{R}_{++}}(u_n) - p\cdot(\xi_n+u_n-w_n)^- + q\cdot(\xi_n+u_n-w_n)^+ \tag{2.50}$$

Hence, (2.24) is satisfied by setting $K = k_u$ and

$$H(z,t) = -p\int_z^\infty (z-w)f(w)dw + q\int_{-\infty}^z (z-w)f(w)dw. \tag{2.51}$$

Note that using the Leibniz integral rule twice gives

$$\begin{aligned}\partial_z H(z,t) &= -p\int_z^\infty f(w)dw + q\int_{-\infty}^z f(w)dw,\\ \partial_{zz} H(z,t) &= (p+q)\cdot f(z).\end{aligned} \tag{2.52}$$

Thus, we have $|\partial_z H(z,t)| = |-p\int_z^\infty f(w)dw + q\int_{-\infty}^z f(w)dw| = (p+q)\cdot F(z) - p \le q$, since $F(z) \le 1$ for $z \in [0,c_{\max}]$. Since $H(\cdot,t)$ is twice differentiable, this means it is absolutely continuous. Thus, $H(\cdot,t)$ is strongly convex on $[0,c_{\max}]$ since $\partial_{zz}H(z,t) = (p+q)\cdot f(z) \geq (p+q)\cdot \min_{z\in[0,c_{\max}]} f(z) > 0$ for $z \in [0,c_{\max}]$. This shows that **A2** holds. To show **A3** holds, we first note that the retail management problem as described in Section 2.2 implies $\zeta = 0$. Next, note that the reset cost $R(x,t) = c_r x + k_r$ is linear. Hence, the conditions of **A3** follow immediately. $\square$

2.4 Numerical Results

In this section, we present numerical studies to validate the theoretical structure derived in Section 2.3. The computations are performed in MATLAB 2018b on a laptop computer with a 2.6GHz processor and 16GB of RAM. We begin by introducing the numerical dynamic programming algorithm used to compute the optimal value functions and policies. We then explore two case studies, the water storage problem and the retail management problem, to numerically verify that their value functions and optimal policies align with the theoretical structure.

Binary Dynamic Search Algorithm

Algorithm 1 Binary Dynamic Search (BiDS) Algorithm [54]

initialize $\underline{v} \leftarrow 0$ and $\overline{v} \leftarrow (\frac{1}{1-\gamma}) \min_u \mathbb{E}[g(\zeta,0,u,w_0) + \gamma\cdot s(h(\zeta,0,u,w_0),1,w_1)]$
repeat
    set $v \leftarrow (\overline{v}+\underline{v})/2$
    set $V(x,k,v) = v + \mathbb{E}[s(x,w,k)]$
    for $t = (k-1),(k-2),\ldots,0$ **do**
        set $V(x,t,v) = \min\left\{v + \mathbb{E}[s(x,w,t)], \min_{u\in\mathcal{U}_t} \mathbb{E}[g(x,t,u,w) + \gamma V(h(x,t,u,w),t+1,v)]\right\}$
    end for
    set $\Upsilon(v) = \min_{u\in\mathcal{U}_0} \mathbb{E}[g(\zeta,0,u,w) + \gamma V(h(\zeta,0,u,w),1,v)]$
    if $v > \Upsilon(v)$ **then**
        set $\overline{v} \leftarrow v$
    else
        set $\underline{v} \leftarrow v$
    end if
until $(\overline{v} - \underline{v}) \leq \epsilon$
set $v^* = (\overline{v}+\underline{v})/2$

The Binary Dynamic Search (BiDS) algorithm, initially developed by [54], is employed to solve dynamic programming equations. In classical discounted reward settings, value functions and optimal policies are typically computed using value iteration (VI) and policy iteration (PI) algorithms. However, when dealing with an infinite state space, this computation becomes equivalent to finding a fixed point in an infinite-dimensional functional space. While strong convergence guarantees exist in a discounted setting, the rate of convergence is highly dependent on the discount factor and can be prohibitively slow [4]. As a result, approximate dynamic programming methods are commonly used in practice for such problems. However, the BiDS algorithm leverages the specific structure of reset control problems to calculate the optimal policy and value function with arbitrary precision. BiDS converts the problem into finding a fixed point in a vector space using binary search, with the vector representing the optimal value function evaluated at the reset state and time 0. The algorithm initializes the search space using upper and lower bounds on the value function at the reset state, and then uses backward induction to compute the implied reset state cost. By comparing this value with the candidate for the iteration, a new search interval is selected using the same procedures as a binary search. The algorithm terminates either when a true fixed point is found or when the numerical tolerance ϵ is reached. The theoretical convergence and computational guarantees of BiDS are outlined in [54].

Example: Water Storage Problem

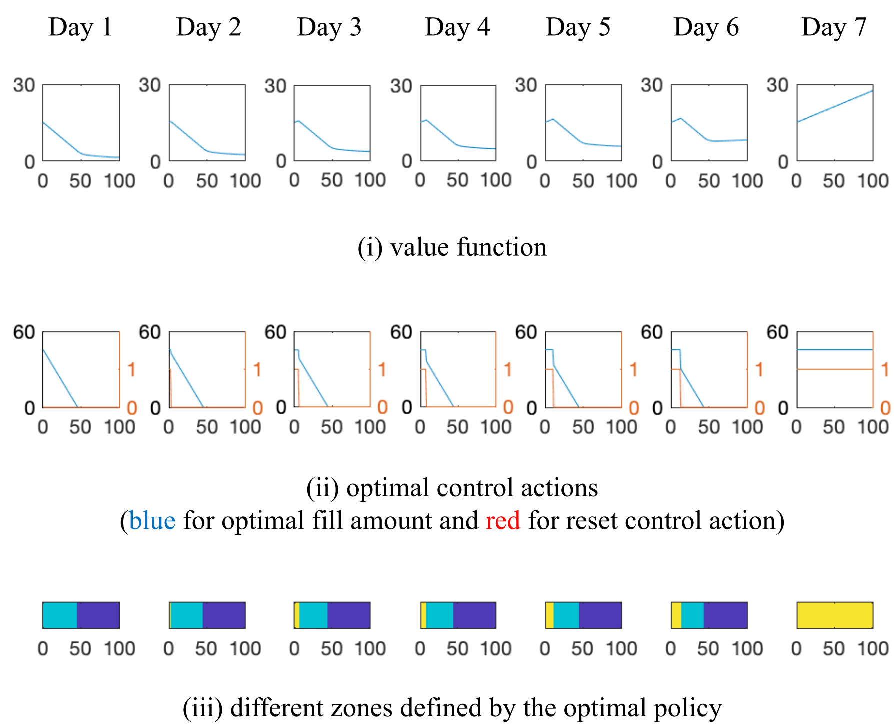

Figure 2.3: Example of an optimal policy for the water storage problem

We use BiDS to numerically solve the water storage problem introduced in Section 2.2. Random demand is generated from a truncated normal distribution bounded from below by zero [11]. To ensure flushing occurs at least once a week, we set $k = 7$. The cost parameters are carefully chosen to place more weight on shortage cost compared to purchasing and flushing costs (i.e., $p > c > c_r$). Furthermore, the cost of consuming water that has been stored long-term exceeds the cost of flushing water (i.e., $q(t) < p$ for all t, but $q(t) \geq c_r$ for t close to k). This design favors resetting the tank more frequently to prevent contamination. In addition, while exponential microbial growth is normally assumed [50], we present a simplified model with linear growth to account for fluctuations in the microbial load that result from water usage and refilling. It should be noted that any other health penalty could be used as long as it satisfies the monotonicity condition.

In Figure 2.3, the top row shows the value function, the second row displays the optimal control actions (blue for the optimal fill amount and red for the reset control action), and the bottom row presents the different zones for the optimal policy. In the policy visualization, yellow corresponds to flushing the tank and reordering water, cyan represents ordering water without flushing, and dark blue indicates not taking any action. The x-axis represents the state x (amount of water in the tank), and the subplots from left to right correspond to states $t = 1, \ldots, 7$ (number of days since the tank was last emptied).

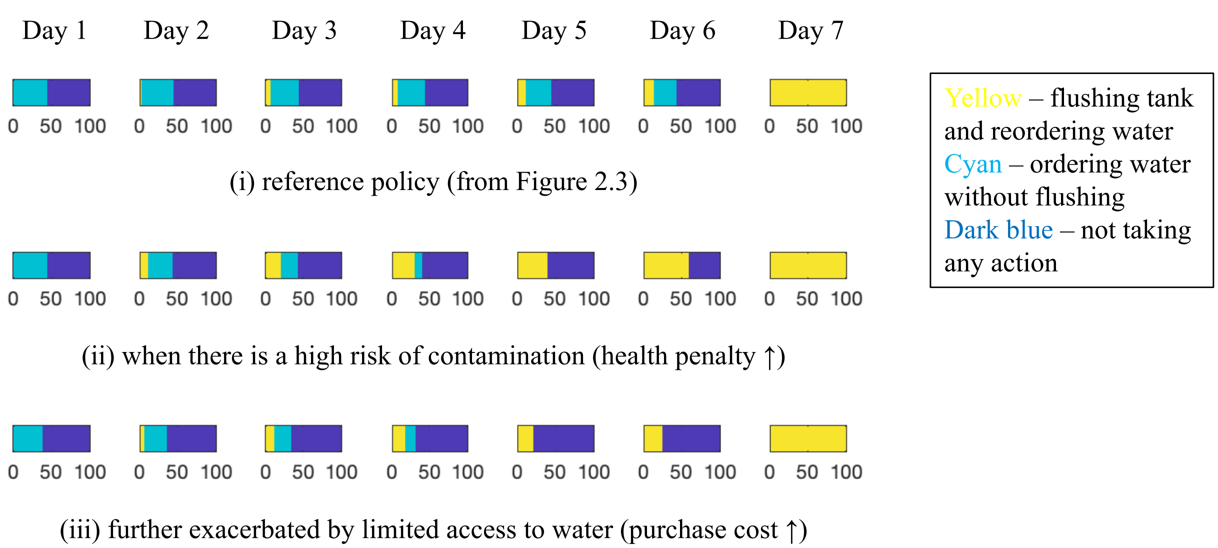

Figure 2.4: Optimal policy for different cases of the water storage problem

With all other parameters fixed from Figure 2.3 (top), we analyze the following two scenarios in Figure 2.4. When there is a high risk of contamination (middle), we impose a larger penalty on water consumption to reflect this condition. It can be observed that the optimal policy suggests flushing the tank more frequently, as indicated by the thicker yellow and thinner cyan bands, in order to reduce the risk of consuming contaminated water.

Further exacerbation due to limited access to water (bottom), which is often the case in the developing world [88, 91], can be modeled by assigning a higher cost to purchasing water. In this case, the control policy is primarily composed of the reset (yellow) and do-nothing (dark blue) regions, with their thresholds slightly shifted to the left compared to the previous policy. As illustrated by wider blue and narrower yellow regions, this regrettably implies that resetting cannot be afforded as often. Nevertheless, our model provides robust and reliable guideline, especially when faced with difficult choices, whether it suggest pursuing a breakthrough or adopting a compromise strategy.

Example: Retail Management Problem

In this section, we solve the retail management problem described in Section 2.2 using the BiDS algorithm under two different cases, as depicted in Figure 2.5. In the first case, where a firm lacks the required flexibility to swiftly change the product line (top), we model this by assigning a higher reset cost. It can be observed that the optimal policy essentially reduces to an (s, S) policy, similar to classical inventory management problems where firms do not have the infrastructure and resources to frequently switch the product line. Conversely, in the second case, where a company is situated in a rapidly evolving industry and thus equipped with the speed and adaptability demanded by the market (bottom), we assign a lower reset cost but a higher holding cost to take the ephemeral nature of our merchandise into consideration. In this case, the inventory policy features a yellow region on the right, indicating the need to change inventory when there is too much stock. This epitomizes the appropriate response to the transience of a trend.

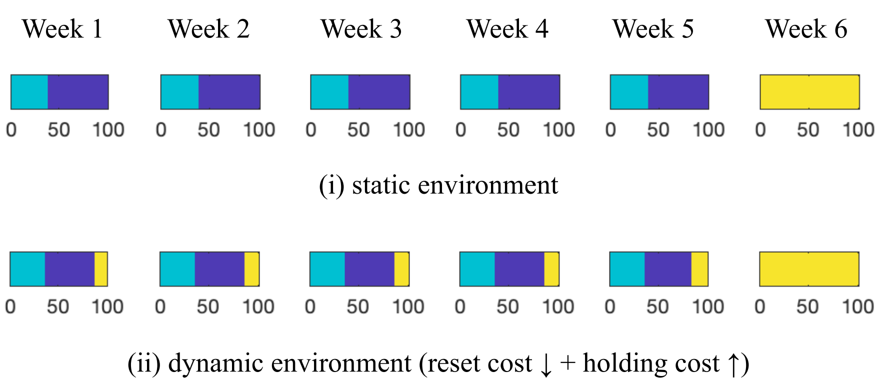

Figure 2.5: Optimal policy for different cases of the retail management problem

At each t, there are at most three thresholds that separate the control policy into four regions. In the first region (yellow), the trivial amount of stock renders the inventory reset cost negligible. Hence, emptying the inventory and ordering an optimal one-step quantity is optimal. In the next region (cyan), the sufficiency of stock in the inventory makes it optimal to order up to some quantity and still consume most of the stock before it becomes obsolete. In the third 'do-nothing region' (blue), it is not optimal to reset the inventory or order further stock as potential future demand can already be satisfied. In the final region (yellow), the inventory stock is almost certain to never be consumed, and it is thus optimal to empty the inventory and reorder up to some optimal quantity, changing the product line if appropriate. Also, note that while these regions cover all possible stages suggested by the optimal policy, the actual policy may not include all of them, as seen in the case of Figure 2.5.

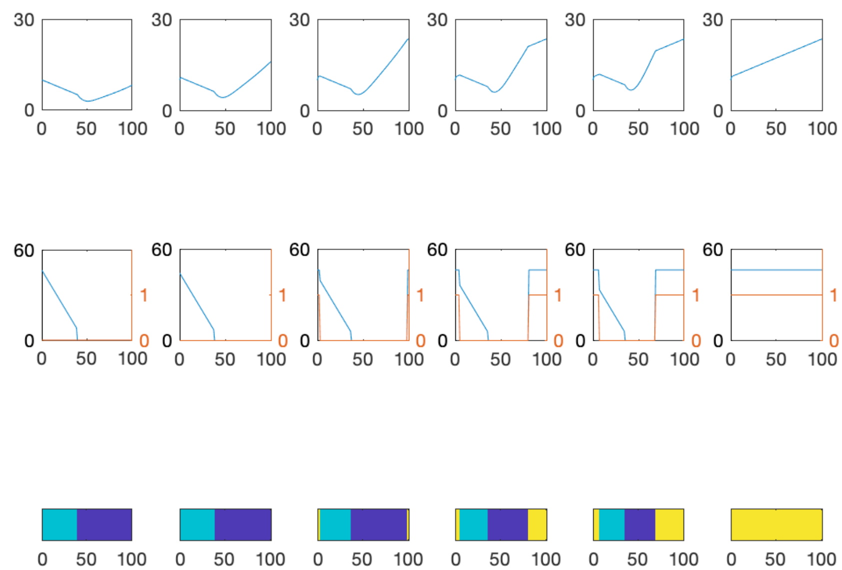

Figure 2.6: Example of an optimal policy for deteriorating item with linear time-dependent holding cost

2.5 Conclusion

Inventory management that empowers us to flexibly adapt to change is necessary to survive in a market and society that continues to evolve with growing acceleration. Previous inventory policies were contingent upon the trade-off balance between holding and shortage costs, and informed us only of the optimal inventory position, or the quantity that must be ordered, for each period. By introducing the option to reset inventory into such classical frameworks, we are pioneering the design of more dynamic policies that overcome the limitations of pre-existing static policies. Furthermore, this enables us to continuously assess whether resets are more profitable at the preset periodic intervals or at earlier time points. In this chapter, we (i) conducted a theoretical investigation of the structural properties underlying the optimal policy for this problem, and (ii) implemented the appropriate algorithm to numerically solve and empirically validate the policy structure. Finally, we demonstrated the broad utility of this problem by providing examples in both non-profit and for-profit contexts.

The original motivation for this study stemmed from the management of water inventory in developing nations [54] where clean water is not consistently available throughout the day due to the absence of comprehensive water distribution networks. As a result, residential buildings rely on water storage systems that are filled during limited time windows when the distribution network is active. However, these systems are not equipped with disinfection capabilities. This leads to water being stored for long durations, thereby exposing people to substantially increased bacterial and viral contamination [18, 44, 39]. Furthermore, obtaining even the minimum quantity of water required for survival poses significant challenges for individuals without a continuous water supply at home, forcing them to compromise their health by using less water or collecting water from unsafe sources [91]. Our model aspires to promote universal and equitable access to safe water by providing decision-making guidance in such healthcare dilemmas.

Our formulation and proof are also sufficiently flexible to account for general reset control problems in supply chain management. For instance, we can incorporate modifications, such as including a time-varying holding cost that considers the depreciation rate and penalizes based on the duration of storage, to effectively manage aging inventories (see Figure 2.6). Finally, our results suggest that the sufficient conditions for ensuring the threshold structure of the optimal policy are marginally conservative as they are not necessary conditions, and that there may exist a relaxation of these conditions, which we leave for future studies.

Chapter 3

Incorporating Fairness into Principal-Agent Models with Adverse Selection and Moral Hazard

3.1 Introduction

Principal-agent models form a foundational pillar in the study of incentive design, with particular emphasis on navigating the complexities of adverse selection and moral hazard. Historically, the objective has been the efficient alignment of interests and the mitigation of associated risks. Yet, an evident lacuna in this methodology has been the marginalization of fairness as a vital consideration, which can inadvertently perpetuate disparities for specific demographic cohorts, notably delineated by race, gender, or other salient attributes. Such an omission not only threatens the integrity of these models but also raises pressing ethical quandaries about their broader societal implications. Against this backdrop, this chapter endeavors to reconcile these concerns by addressing two cardinal principles: (i) the theoretical pursuit of what is fair and equitable; and (ii) the numerical considerations of what is feasible and optimal, in the context of contract design. By incorporating fairness into the optimization framework, this study seeks to augment the traditional paradigm, fostering a more holistic and ethically robust approach to incentive design.

Principal-Agent Model

Let us consider the following hypothetical situation: a company is organizing a picnic and intends to purchase sandwiches. The company is capable of utilizing every sandwich produced, but there are diminishing returns associated with the value of each additional sandwich. Additionally, the restaurant from which the company sources the sandwiches can be either efficient or inefficient, although the company is unaware of the restaurant's type. In this context, efficiency is assessed based on the cost of producing each sandwich, which encompasses factors such as raw ingredients, energy inputs, and equipment purchases. Given this scenario, an intriguing problem arises: how should the company structure the contract for hiring the restaurant? It is worth noting that this example is just one instance within a broader range of similar situations.

Principal-agent models represent situations in which a principal (e.g., manager) wishes to delegate a task to an agent (e.g., employee). The example of a company (principal) hiring a restaurant (agent) aligns with this pattern. There is a well-established literature on principal-agent models in economics [26, 28] and in operations management [41]. These models are classified based on the type of the information asymmetry arising between the principal and the agent. This asymmetry typically appears in the form of either *adverse selection* (hidden information) or *moral hazard* (hidden actions) [10]. The company-restaurant context above is an example of an adverse selection setting, wherein the agent is aware of their efficiency level, while the principal lacks this knowledge. On the other hand, moral hazard arises when the agent possesses knowledge about their effort level, which remains unknown to the principal [30]. These information asymmetries complicate the contract design problem, as an effective contract is one that encourages the agent to reveal their private knowledge, even though doing so may require offering some information rent to the agent.

Contributions and Outline

The motivation behind this study stems from the observation that existing approaches to incentive design often overlook the incorporation of fairness. Fairness is a crucial aspect of incentive systems since inadequate design can result in harm to individuals from specific groups (e.g., based on race or gender). For instance, consider financial incentives for demand-responsive electricity usage that impose varying fees based on current transmission costs and supply availability. Improperly designed incentives can lead to disadvantages for individuals residing in certain locations, thus introducing biases related to race, ethnicity, and socioeconomic factors, given the unequal distribution of the population across demographics. Similar concerns regarding fairness also arise in the realm of healthcare, spanning aspects such as insurance contract composition, procedure pricing, and the allocation of scarce medical resources.

Our contribution lies in integrating the concept of fairness into the framework of incentive design, particularly within the context of principal-agent models involving either adverse selection (Section 3.2) or moral hazard (Section 3.3). In both cases, we will formulate the problem mathematically, derive the optimal contract that satisfies fairness criteria, and explore its implications. By incorporating fairness considerations into incentive design, we aspire to address the potential biases and disparities that may arise from traditional approaches. Our analysis will provide insights into how to construct contracts that not only optimize economic objectives but also ensure fairness among different groups or individuals.

3.2 Adverse Selection

Suppose the principal contracts the agent to produce q units of a good. The value to the principal of the q units of the good is given by the function $S(q)$, where $S(0) = 0$, $S' > 0$, and $S'' < 0$. This means that the marginal value of the good is positive (i.e., producing more goods is better) but exhibits diminishing utility. In this model, the production cost of the agent is unobservable to the principal. However, there are certain facts known to both parties. Both the principal and the agent are aware of the fixed costs F, the cost function of the agent $C(q, \theta) = \theta q + F$, where θ represents the marginal cost. Additionally, it is known that the agent can be either inefficient ($\theta = \theta^I$) or efficient ($\theta = \theta^E$), with $\theta^I > \theta^E$, indicating higher marginal costs for the inefficient type. While the agent is aware of its own type, the principal does not possess this information. Instead, the principal has knowledge of, or can estimate, the probability that an agent is efficient, denoted as $\mathbb{P}(\theta = \theta^E) = \alpha$, and the probability that an agent is inefficient, denoted as $\mathbb{P}(\theta = \theta^I) = 1 - \alpha = \beta$. The decision-making process in this model revolves around the design of a menu of contracts. The principal has the authority to specify multiple production levels along with the associated payments or transfers t for each production level. Each contract in the menu is represented by a tuple (q, t), and the menu consists of multiple such tuples. In this simple case, there should be two contract tuples, one for the efficient agent and the other for the inefficient agent.

Consider a scenario where the principal has full knowledge of the agent's type. In this case, the principal should design the contract in such a way that the principal's marginal utility of a good equals the agent's marginal cost, resulting in the principal achieving maximum utility. This implies selecting production levels q_1^I and q_1^E that satisfy $S'(q_1^I) = \theta^I$ and $S'(q_1^E) = \theta^E$, respectively. However, it is important to note that agents will only agree to participate in a contract if the offered transfer t ensures they do not incur a financial loss. Thus, we must consider *participation constraints* that guarantee the transfers in the contract are at least as large as the costs incurred by the agents:

$$
\begin{aligned}
t^E - \theta^E q^E - F &\geq 0, \\
t^I - \theta^I q^I - F &\geq 0.
\end{aligned}
$$

Taking everything into account, our menu of contracts is as follows: If the agent is of type θ^I, the principal offers the contract $(q_1^I, \theta^I q_1^I + F)$. Similarly, if the agent is of type θ^E, the principal offers the contract $(q_1^E, \theta^E q_1^E + F)$. These production levels are referred to as the first-best production levels. What is interesting about these contracts is that the agents make no profit. This means that their transfer is equal to their cost regardless of their type, indicating that the agents receive no additional financial benefit from participating in the contract.

Now, return to the general setting where the principal does not have knowledge of the agent's type. The initial naïve approach might be for the principal to offer a menu of contracts $(q_1^I, t_1^I), (q_1^E, t_1^E)$ and hope that the inefficient agent chooses the contract designed for inefficiency (q_1^I, t_1^I), while the efficient agent chooses the contract designed for efficiency (q_1^E, t_1^E). However, this is not what occurs in practice. It is in the best interest of the efficient agent, who has a lower marginal cost, to pretend to be inefficient and accept the contract designed for inefficiency (q_1^I, t_1^I). This behavior is known as adverse selection. Adverse selection arises because for a fixed production level, say q_1^I, the efficient agent has lower production costs compared to the inefficient agent $(\theta^E q_1^I + F < \theta^I q_1^I + F)$. Since the contract is designed such that $t_1^I - \theta^I q_1^I + F = 0$, it implies that $t_1^I - \theta^E q_1^I + F > 0$, resulting in a profit for the efficient agent who pretends to be inefficient.

Ideally, our goal is to design a menu of contracts in such a way that each agent selects the contract that aligns with their true type. This means we want the agents to be truthful about their type. The property of being truthful is also known as being strategy-proof or incentive compatible, and it is captured by the following *incentive compatibility constraints*:

$$t^E - \theta^E q^E - F \geq t^I - \theta^E q^I - F,$$
$$t^I - \theta^I q^I - F \geq t^E - \theta^I q^E - F.$$

The intuition behind these incentive compatibility constraints is that the profit an agent earns when selecting a contract originally designed for their own type should be greater than or equal to the profit they would earn if they were to select a contract designed for the opposite type. By enforcing these constraints, we ensure that agents have no incentive to misrepresent their type and that it is in their best interest to be truthful about their type when selecting a contract.

The profit earned by an agent is referred to as information rent. In the case of efficient and inefficient agents, their information rents are defined as

$$U^E = t^E - \theta^E q^E - F,$$
$$U^I = t^I - \theta^I q^I - F.$$

Even when a contract is designed to satisfy the incentive compatibility constraints, it is common for the efficient agent to receive some amount of information rent. To illustrate this, consider the scenario where the efficient agent pretends to be inefficient. In this case, its payment would be

$$U^E = t^E - \theta^E q^E - F \geq t^I - \theta^E q^I - F = t^I - \theta^I q^I - F + (\theta^I - \theta^E)q^I = U^I + (\theta^I - \theta^E)q^I.$$

Hence, even if the contract is designed such that $U^I = 0$, there would still exist a nonzero information rent of $(\theta^I - \theta^E)q^I$. However, this rent can be reduced by decreasing the value of q^I, thereby posing an important problem of how to achieve this reduction.

Considering that the efficient agent is entitled to receive some information rent, an intriguing question arises: How can the contract be designed to minimize the rent provided and enhance overall efficiency for the principal? This can be formulated as an optimization problem:

$$\begin{aligned} \max_{(q^E,t^E),(q^I,t^I)} \quad & \alpha(S(q^E)-t^E)+\beta(S(q^I)-t^I) \\ \text{s.t.} \quad & t^E-\theta^E q^E-F\geq 0 \\ & t^I-\theta^I q^I-F\geq 0 \\ & t^E-\theta^E q^E-F\geq t^I-\theta^E q^I-F \\ & t^I-\theta^I q^I-F\geq t^E-\theta^I q^E-F \end{aligned}$$

Recall that the first two constraints represent the participation constraints, ensuring that the transfers are sufficient to cover the costs for both types of agents. The last two constraints are the incentive compatibility constraints, guaranteeing that each agent is better off choosing their own contract rather than pretending to be the other type.

This optimization problem can be solved by analyzing its optimality conditions, leading to solutions known as the second-best production levels. Here, we provide a summary of the obtained results. The production level for the efficient agent remains unchanged and is denoted as $q_2^E = q_1^E$. However, the production level for the inefficient agent is reduced to q_2^I, satisfying $S'(q_2^I) = \theta^I + \frac{\alpha}{\beta}(\theta^I - \theta^E)$. The corresponding second-best transfers are given by $t_2^E = \theta^E q_2^E + (\theta^I - \theta^E)q_2^I + F$ and $t_2^I = \theta^I q_2^I + F$. Notably, only the efficient agent receives a strictly positive information rent, denoted as $U_2^E = (\theta^I - \theta^E)q_2^I$. Furthermore, this rent is lower than the information rent U_1^E extracted from the menu of contracts designed for the first-best production levels.

Using this framework, we will formulate a principal-agent model with adverse selection that incorporates fairness as follows:

$$\begin{aligned} \text{maximize} \quad & \text{principal's expected utility} \\ \text{subject to} \quad & \text{agent's participation} \\ & \text{agent's incentive compatibility} \\ & \text{(+ contract fairness)} \end{aligned}$$

Fair Principal-Agent Model with Adverse Selection

In this section, we will address the following questions:

- What if we introduce a protected group alongside the efficiency types?
- If we have a protected group, how can we design a contract that ensures fairness among different groups of agents?
- Most importantly, how should we precisely define fairness in this context?

Consider a scenario with two distinct groups and two different types of agents, resulting in a total of four possible combinations. To clarify, agents are classified into two groups that we aim to ensure fairness with respect to, and they also possess efficiency types based on their production costs. An additional assumption is made that one group, as a whole, exhibits higher efficiency compared to the other. This implies that each group comprises a different proportion of efficient and inefficient agents, leading to one group having a higher conditional probability of being efficient than the other.

Our objective is to explore the possibility of formulating a contract that provides the same expected profit for agents from both groups. This quantitative definition of fairness involves equating the expected profit of agents across all groups, thereby making profit independent of group membership. In the following sections, we will examine two different menus of contracts to incorporate fairness into the design of incentives.

Model 1

Assume that there are two groups of agents, A and B. Their probabilities are given by the matrix $P = \begin{pmatrix} \alpha & \beta \\ \gamma & \delta \end{pmatrix}$, where $P_{ij} = \mathbb{P}(\text{group} = i,\ \text{efficiency} = j)$, with $0 < \alpha, \beta, \gamma, \delta < 1$ and $\alpha + \beta + \gamma + \delta = 1$. If we choose to offer a menu of contracts $\{(q^E, t^E), (q^I, t^I)\}$, then the optimization problem is given by

$$
\begin{aligned}
\max_{(q^E,t^E),(q^I,t^I)} \quad & \alpha(S(q^E) - t^E) + \beta(S(q^I) - t^I) + \delta(S(q^E) - t^E) + \gamma(S(q^I) - t^I) && (3.1)\\
\text{s.t.} \quad & t^E - \theta^E q^E - F \geq 0 && (3.2)\\
& t^I - \theta^I q^I - F \geq 0 && (3.3)\\
& t^E - \theta^E q^E - F \geq t^I - \theta^E q^I - F && (3.4)\\
& t^I - \theta^I q^I - F \geq t^E - \theta^I q^E - F && (3.5)\\
& \frac{\alpha}{\alpha+\beta}(t^E - \theta^E q^E - F) + \frac{\beta}{\alpha+\beta}(t^I - \theta^I q^I - F) && (3.6)\\
& = \frac{\gamma}{\gamma+\delta}(t^E - \theta^E q^E - F) + \frac{\delta}{\gamma+\delta}(t^I - \theta^I q^I - F)
\end{aligned}
$$

The following theorem presents an important result that proves the infeasibility of the aforementioned optimization problem, thereby highlighting the impossibility of achieving an optimal fair contract when contracts are offered solely based on agent type.

Theorem 2. *If $\frac{\alpha}{\alpha+\beta} \neq \frac{\gamma}{\gamma+\delta}$ and $q^I > 0$, then there exists no feasible solution that simultaneously satisfies both incentive compatibility and fairness constraints.*

Proof. Fairness constraint (3.6) implies that

$$\begin{aligned}\left(\frac{\alpha}{\alpha+\beta}-\frac{\gamma}{\gamma+\delta}\right)(t^E-\theta^E q^E-F) &= \left(\frac{\delta}{\gamma+\delta}-\frac{\beta}{\alpha+\beta}\right)(t^I-\theta^I q^I-F)\\ &= \left(\left(1-\frac{\gamma}{\gamma+\delta}\right)-\left(1-\frac{\alpha}{\alpha+\beta}\right)\right)(t^I-\theta^I q^I-F)\\ &= \left(\frac{\alpha}{\alpha+\beta}-\frac{\gamma}{\gamma+\delta}\right)(t^I-\theta^I q^I-F),\end{aligned}$$

so we have $t^E-\theta^E q^E-F=t^I-\theta^I q^I-F$. However, combined with incentive compatibility constraint (3.4), this implies that $t^I-\theta^E q^I-F \le t^E-\theta^E q^E-F=t^I-\theta^I q^I-F$, which means $\theta^I q^I \le \theta^E q^I$ and thus $\theta^I \le \theta^E$. This is a contradiction since we had assumed $\theta^I > \theta^E$. Hence, fairness and incentive compatibility are not jointly feasible. □

The conditions outlined in Theorem 2 are justifiable. The first condition indicates that there must be a difference in overall efficiency levels between the two groups. Otherwise, the groups would be identical, and there would be no need to incorporate fairness into contract design. The second condition ensures that the production level must be strictly positive to prevent the less efficient type from being shut down, which may result in excessive screening of types.

The key takeaway is that if we offer two contracts, one for each agent type, we discover that there is no feasible solution. In other words, if we only consider the agent's type when designing contracts, it is not possible to achieve our definition of fairness across different groups. It is intriguing to observe that we cannot formulate a fair optimal contract by solely considering the types of agents when there is a discrepancy between groups, such as a gap in their group-wise efficiency. This epitomizes the importance of factoring in protected groups when designing incentives to ensure fairness.

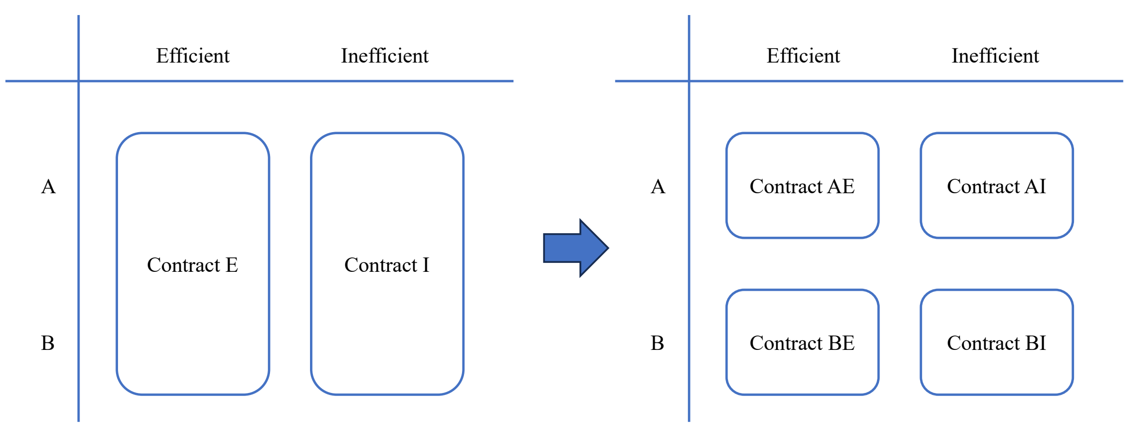

Figure 3.1: Visualization of the transition from Model 1 to Model 2

Model 2

In addition to our previous assumptions, we make another assumption that agents can pretend to be of the opposite type, but they must truthfully report which group they belong to. This assumption is crucial in our analysis because we are primarily concerned with designing contracts that do not discriminate against agents based on their group membership. By allowing agents to potentially misrepresent their type only, we are able to focus our attention on designing contracts that are fair with respect to group membership while still incentivizing truthful reporting of individual efficiency characteristics.

With a menu of contracts $\{(q_A^E, t_A^E), (q_A^I, t_A^I), (q_B^E, t_B^E), (q_B^I, t_B^I)\}$, the optimization problem is given by

$$\max_{\{(q,t)\}} \quad \alpha(S(q_A^E) - t_A^E) + \beta(S(q_A^I) - t_A^I) + \gamma(S(q_B^E) - t_B^E) + \delta(S(q_B^I) - t_B^I) \tag{3.7}$$

$$\text{s.t.} \quad t_A^E - \theta^E q_A^E - F \geq 0 \tag{3.8}$$

$$t_A^I - \theta^I q_A^I - F \geq 0 \tag{3.9}$$

$$t_B^E - \theta^E q_B^E - F \geq 0 \tag{3.10}$$

$$t_B^I - \theta^I q_B^I - F \geq 0 \tag{3.11}$$

$$t_A^E - \theta^E q_A^E - F \geq t_A^I - \theta^E q_A^I - F \tag{3.12}$$

$$t_A^I - \theta^I q_A^I - F \geq t_A^E - \theta^I q_A^E - F \tag{3.13}$$

$$t_B^E - \theta^E q_B^E - F \geq t_B^I - \theta^E q_B^I - F \tag{3.14}$$

$$t_B^I - \theta^I q_B^I - F \geq t_B^E - \theta^I q_B^E - F \tag{3.15}$$

$$\begin{aligned} &\frac{\alpha}{\alpha+\beta}(t_A^E - \theta^E q_A^E - F) + \frac{\beta}{\alpha+\beta}(t_A^I - \theta^I q_A^I - F) \\ &= \frac{\gamma}{\gamma+\delta}(t_B^E - \theta^E q_B^E - F) + \frac{\delta}{\gamma+\delta}(t_B^I - \theta^I q_B^I - F) \end{aligned} \tag{3.16}$$

Define information rents for each group-type combination:

$$\begin{aligned} U_A^E &= t_A^E - \theta^E q_A^E - F, \\ U_A^I &= t_A^I - \theta^I q_A^I - F, \\ U_B^E &= t_B^E - \theta^E q_B^E - F, \\ U_B^I &= t_B^I - \theta^I q_B^I - F. \end{aligned}$$

By substituting transfers in the principal's objective with functions of information rents and outputs, we reformulate the decision variables as $\{(q, U)\}$. This modification enhances the economic insights derived from the optimal contract and enables us to evaluate the distributional effects of asymmetric information, with a particular emphasis on information rents.

The revised optimization problem can be expressed as follows:

$$
\begin{aligned}
\max_{\{(q,U)\}} \quad & \alpha(S(q_A^E) - \theta^E q_A^E) + \beta(S(q_A^I) - \theta^I q_A^I) + \gamma(S(q_B^E) - \theta^E q_B^E) + \delta(S(q_B^I) - \theta^I q_B^I) && (3.17)\\
& - (\alpha U_A^E + \beta U_A^I + \gamma U_B^E + \delta U_B^I) \\
\text{s.t.} \quad & U_A^E \geq 0 && (3.18)\\
& U_A^I \geq 0 && (3.19)\\
& U_B^E \geq 0 && (3.20)\\
& U_B^I \geq 0 && (3.21)\\
& U_A^E \geq U_A^I + \Delta\theta q_A^I && (3.22)\\
& U_A^I \geq U_A^E - \Delta\theta q_A^E && (3.23)\\
& U_B^E \geq U_B^I + \Delta\theta q_B^I && (3.24)\\
& U_B^I \geq U_B^E - \Delta\theta q_B^E && (3.25)\\
& \frac{\alpha}{\alpha+\beta} U_A^E + \frac{\beta}{\alpha+\beta} U_A^I = \frac{\gamma}{\gamma+\delta} U_B^E + \frac{\delta}{\gamma+\delta} U_B^I && (3.26)
\end{aligned}
$$

where $\Delta\theta = \theta^I - \theta^E > 0$.

A primary hurdle in tackling this problem is discerning the constraints imposed by participation and incentive compatibility that are crucial, in other words, those that are binding at the optimal solution. One might initially consider employing Lagrangian techniques. However, given the number of constraints, a more pragmatic tactic would be to conjecture the active constraints first, and subsequently verify that the disregarded constraints are indeed strictly satisfied, as in [41]. In the context of our model, this method proves more beneficial in both solving the optimization problem and gaining a clearer understanding of the underlying economic rationale within this model.

Firstly, let us examine contracts featuring positive production levels or outputs, denoted as $q_A^I, q_B^I > 0$. The capacity of the efficient agent to impersonate the inefficient agent ensures that the efficient agent's participation constraint is consistently strictly satisfied. In fact, this can be readily understood, as the conjunction of the inefficient agent's participation constraint (3.19, 3.21) and the efficient agent's incentive constraint (3.22, 3.24) directly substantiates the efficient agent's participation constraint (3.18, 3.20). Moreover, the inefficient agent's incentive constraint (3.23, 3.25) is immaterial given that the challenge arises from an efficient agent potentially feigning inefficiency, rather than vice versa. This reduction in the number of pertinent constraints results in only three remaining constraints: the efficient agent's incentive compatibility constraint (3.22, 3.24), the inefficient agent's participation constraint (3.19, 3.21), and the fairness constraint (3.26).

Special Case

To gain a better understanding of the impact of adding a fairness constraint to the optimization problem, we will first consider a special case where one group is composed entirely of a single type of agents, for instance, inefficient agents. Specifically, we will set the probabilities $\gamma = 0$ and $\delta = 1 - \alpha - \beta$. With these assumptions, the optimization problem takes the following form:

$$\begin{aligned}
\max_{\{(q,U)\}} \quad & \alpha(S(q_A^E) - \theta^E q_A^E) + \beta(S(q_A^I) - \theta^I q_A^I) + \delta(S(q_B^I) - \theta^I q_B^I) && (3.27)\\
& - (\alpha U_A^E + \beta U_A^I + \delta U_B^I) && \\
\text{s.t.} \quad & U_A^I \geq 0 && (3.28)\\
& U_B^I \geq 0 && (3.29)\\
& U_A^E \geq U_A^I + \Delta\theta q_A^I && (3.30)\\
& \frac{\alpha}{\alpha+\beta} U_A^E + \frac{\beta}{\alpha+\beta} U_A^I = U_B^I && (3.31)
\end{aligned}$$

In the following theorem, we will present a characterization of the optimal fair contract for this problem.

Theorem 3. *If $\gamma = 0$ and $\delta = 1 - \alpha - \beta$, then an optimal solution exists and follows a specific structure outlined as follows:*

$$\begin{aligned}
q_A^E &: S'(q_A^E) - \theta^E = 0, \quad U_A^E = \Delta\theta q_A^I\\
q_B^I &: S'(q_B^I) - \theta^I = 0, \quad U_B^I = \frac{\alpha}{\alpha+\beta}\Delta\theta q_A^I\\
q_A^I &: S'(q_A^I) - \theta^I = \frac{\alpha/\beta}{\alpha+\beta}\Delta\theta, \quad U_A^I = 0
\end{aligned}$$

Proof. We can easily verify that the participation constraint for group A (3.28) and the fairness constraint (3.31), when combined, imply the participation constraint for group B (3.29). We then incorporate the fairness constraint (3.31) into the objective (3.27) as follows:

$$\alpha U_A^E + \beta U_A^I + \delta U_B^I = \left(1 + \frac{\delta}{\alpha+\beta}\right)(\alpha U_A^E + \beta U_A^I) = \frac{1}{\alpha+\beta}(\alpha U_A^E + \beta U_A^I).$$

To see that the two remaining constraints (3.28, 3.30) are binding, we use a proof by contradiction. If they were not binding, we could lower U_A^I and U_A^E to increase the objective while satisfying all constraints and keeping all outputs the same. Therefore, we must have $U_A^I = 0$ and $U_A^E = \Delta\theta q_A^I$. With these values, we have an unconstrained optimization problem as follows:

$$\max_{q_A^E, q_A^I} \alpha(S(q_A^E) - \theta^E q_A^E) + \beta(S(q_A^I) - \theta^I q_A^I) + \delta(S(q_B^I) - \theta^I q_B^I) - \frac{\alpha}{\alpha+\beta}\Delta\theta q_A^I$$

The first-order optimality conditions yield:

$$\begin{aligned}
q_A^E &: S'(q_A^E) - \theta^E = 0\\
q_B^I &: S'(q_B^I) - \theta^I = 0\\
q_A^I &: S'(q_A^I) - \theta^I = \frac{\alpha/\beta}{\alpha+\beta}\Delta\theta
\end{aligned}$$

The information rents and corresponding transfers are given by:

$$
\begin{aligned}
&U_A^E = \Delta\theta q_A^I, \quad t_A^E = \theta^E q_A^E + F + \Delta\theta q_A^I \\
&U_B^I = \frac{\alpha}{\alpha+\beta}\Delta\theta q_A^I, \quad t_B^I = \theta^I q_B^I + F + \frac{\alpha}{\alpha+\beta}\Delta\theta q_A^I \\
&U_A^I = 0, \quad t_A^I = \theta^I q_A^I + F
\end{aligned}
$$

□

The interpretation of this result is that when there are no efficient agents in group B, inefficient agents in group B will receive a strictly positive information rent. Interestingly, the amount of this information rent is dependent upon the production level assigned to the inefficient agents in group A. This is due to the fairness constraint that ensures expected profit equivalence across the two groups. Here, the objective can be seen as the expected utility minus the expected information rent. Thus, the principal can balance the trade-off and accept some distortions away from efficiency to decrease the agent's information rent. Consequently, the inefficient agents in group A, who determine the actual amount of information rents paid to all stakeholders, face a downward output distortion. We can verify this by noting that $S'(q_A^I) - \theta^I = \frac{\alpha/\beta}{\alpha+\beta}\Delta\theta$, which is greater than the second-best output level of $\frac{\alpha}{\beta}\Delta\theta$. Another interesting aspect of this contract is that the inefficient agent of group B will be assigned to their first-best production level. This is because we assumed that efficient agents are missing from group B, and thus there is no need to consider incentive compatibility for them.

This is the first step towards understanding how to design a fair contract in accordance with our quantitative definition when we have two groups that differ in their group-wise efficiency levels. This special case will be particularly useful in contexts where one group is significantly more disadvantaged and underprivileged relative to the other group. We aim to prepare a menu of contracts that does not discriminate against certain groups and, hopefully, can even narrow the differential between them.

General Case

Let us now return to the general case where we assume $\alpha, \beta, \gamma, \delta \in (0,1)$ and $\alpha+\beta+\gamma+\delta = 1$. It is important to recall that the optimization problem can be written concisely as follows:

$$\begin{aligned}
\max_{\{(q,U)\}} \quad & \alpha(S(q_A^E) - \theta^E q_A^E) + \beta(S(q_A^I) - \theta^I q_A^I) + \gamma(S(q_B^E) - \theta^E q_B^E) + \delta(S(q_B^I) - \theta^I q_B^I) \\
& - (\alpha U_A^E + \beta U_A^I + \gamma U_B^E + \delta U_B^I) \\
\text{s.t.} \quad & U_A^I \geq 0 \quad \textbf{(PC-A)} \\
& U_B^I \geq 0 \quad \textbf{(PC-B)} \\
& U_A^E \geq U_A^I + \Delta\theta q_A^I \quad \textbf{(IC-A)} \\
& U_B^E \geq U_B^I + \Delta\theta q_B^I \quad \textbf{(IC-B)} \\
& \frac{\alpha}{\alpha+\beta} U_A^E + \frac{\beta}{\alpha+\beta} U_A^I = \frac{\gamma}{\gamma+\delta} U_B^E + \frac{\delta}{\gamma+\delta} U_B^I \quad \textbf{(FC)}
\end{aligned}$$

Here, the acronyms stand for the following constraints: **PC** represents participation constraints, **IC** denotes incentive compatibility constraints for each group, and **FC** stands for the fairness constraint.

The main challenge in addressing this problem is identifying the constraints that are relevant in the optimal solution. Our approach is to initially hypothesize the active constraints and then verify afterwards that the neglected constraints are indeed strictly satisfied.

Proposition 9. *At least one of the participation constraints must be binding, implying that either $U_A^I = 0$ or $U_B^I = 0$. This indicates that an inefficient agent from either group receives zero information rent.*

Proof. Suppose there exists an optimal solution with $U_A^I > 0$ and $U_B^I > 0$ (this is a proof by contradiction). Then we can construct another solution with a better objective. Without loss of generality, assume $U_B^I \geq U_A^I = \epsilon > 0$. Now, subtract ϵ from U_A^I and U_B^I, while keeping the outputs q_A^I and q_B^I unchanged, to obtain such a solution. We can verify that the constraints still hold:

- **(PC-A)** $U_A^I - \epsilon = 0$,
- **(PC-B)** $U_B^I - \epsilon \geq 0$,
- **(IC-A)** $U_A^E - \epsilon \geq U_A^I - \epsilon + \Delta\theta q_A^I$,
- **(IC-B)** $U_B^E - \epsilon \geq U_B^I - \epsilon + \Delta\theta q_B^I$,
- **(FC)** $\frac{\alpha}{\alpha+\beta}(U_A^E - \epsilon) + \frac{\beta}{\alpha+\beta}(U_A^I - \epsilon) = \frac{\alpha}{\alpha+\beta}U_A^E + \frac{\beta}{\alpha+\beta}U_A^I - \epsilon = \frac{\gamma}{\gamma+\delta}U_B^E + \frac{\delta}{\gamma+\delta}U_B^I - \epsilon = \frac{\gamma}{\gamma+\delta}(U_B^E - \epsilon) + \frac{\delta}{\gamma+\delta}(U_B^I - \epsilon)$.

The objective strictly increases by ϵ since $-(\alpha(U_A^E-\epsilon)+\beta(U_A^I-\epsilon)+\gamma(U_B^E-\epsilon)+\delta(U_B^I-\epsilon)) = -(\alpha U_A^E+\beta U_A^I+\gamma U_B^E+\delta U_B^I)+(\alpha+\beta+\gamma+\delta)\cdot\epsilon = -(\alpha U_A^E+\beta U_A^I+\gamma U_B^E+\delta U_B^I)+\epsilon$. Thus, this solution is not optimal, and it must be the case that $U_A^I = 0$ or $U_B^I = 0$. □

This proposition states that in the optimal solution, at least one of the participation constraints must be active, which implies that either the inefficient agent from group A (U_A^I) or the inefficient agent from group B (U_B^I) receives zero information rent. This result provides a weaker guarantee compared to the contract solely based on the incentive compatibility constraints. Nevertheless, it underscores the necessity of efficiency in the incentive design, wherein at least one group's inefficient agents must forgo any information rent.

We will now eliminate the equality constraint by incorporating it into the objective function and the inequality constraints, expressing them in terms of U_A^E and U_A^I. Using **FC**, the last term in the objective can be rewritten as:

$$\alpha U_A^E+\beta U_A^I+\gamma U_B^E+\delta U_B^I = \left(\frac{\alpha+\beta+\gamma+\delta}{\alpha+\beta}\right)(\alpha U_A^E+\beta U_A^I) = \frac{1}{\alpha+\beta}(\alpha U_A^E+\beta U_A^I).$$

Next, we proceed to the constraints. **FC** implies that $U_B^I = \frac{\gamma+\delta}{\delta}(\frac{\alpha}{\alpha+\beta}U_A^E+\frac{\beta}{\alpha+\beta}U_A^I-\frac{\gamma}{\gamma+\delta}U_B^E)$. Substituting this into **IC-B** and **PC-B** yields:

$$U_B^E \geq \frac{\alpha}{\alpha+\beta}U_A^E+\frac{\beta}{\alpha+\beta}U_A^I+\frac{\delta}{\gamma+\delta}\Delta\theta q_B^I \quad \textbf{(IC-B*)}$$
$$U_B^E \leq \frac{\gamma+\delta}{\gamma}\cdot\left(\frac{\alpha}{\alpha+\beta}U_A^E+\frac{\beta}{\alpha+\beta}U_A^I\right) \quad \textbf{(PC-B*)}$$

Based on Proposition 9, assume $U_A^I = 0$. Later, we will show that it does not matter which we assume to be binding, as an optimal contract will induce inefficient agents from both groups to receive zero information rent. The optimization problem now becomes:

$$\begin{aligned}
\max_{\{(q,U)\}} \quad & \alpha(S(q_A^E)-\theta^E q_A^E)+\beta(S(q_A^I)-\theta^I q_A^I) \\
& +\gamma(S(q_B^E)-\theta^E q_B^E)+\delta(S(q_B^I)-\theta^I q_B^I)-\frac{\alpha}{\alpha+\beta}U_A^E \\
\text{s.t.} \quad & U_A^E \geq \Delta\theta q_A^I \quad \textbf{(IC-A)} \\
& U_B^E \geq \frac{\alpha}{\alpha+\beta}U_A^E+\frac{\delta}{\gamma+\delta}\Delta\theta q_B^I \quad \textbf{(IC-B*)} \\
& U_B^E \leq \frac{\alpha}{\alpha+\beta}\cdot\frac{\gamma+\delta}{\gamma}U_A^E \quad \textbf{(PC-B*)}
\end{aligned}$$

An arbitrary feasible U_B^E exists as long as it satisfies both **IC-B*** and **PC-B***. Since the problem is now independent of U_B^E (but still a function of q_B^E), we can combine these constraints as $\frac{\alpha}{\alpha+\beta}U_A^E+\frac{\delta}{\gamma+\delta}\Delta\theta q_B^I \leq \frac{\alpha}{\alpha+\beta}\cdot\frac{\gamma+\delta}{\gamma}U_A^E$ to obtain a new constraint:

$$U_A^E \geq \frac{\alpha+\beta}{\alpha}\cdot\frac{\gamma}{\gamma+\delta}\Delta\theta q_B^I \quad \textbf{(IPC-B)}$$

Finally, we have the following optimization problem:

$$\begin{aligned}
\max_{\{(q,U)\}} \quad & \alpha(S(q_A^E) - \theta^E q_A^E) + \beta(S(q_A^I) - \theta^I q_A^I) \\
& + \gamma(S(q_B^E) - \theta^E q_B^E) + \delta(S(q_B^I) - \theta^I q_B^I) - \frac{\alpha}{\alpha+\beta} U_A^E \\
\text{s.t.} \quad & U_A^E \geq \Delta\theta q_A^I \quad \textbf{(IC-A)} \\
& U_A^E \geq \frac{\alpha+\beta}{\alpha} \cdot \frac{\gamma}{\gamma+\delta} \Delta\theta q_B^I \quad \textbf{(IPC-B)}
\end{aligned}$$

We will now solve the problem above and present the following theorem characterizing the optimal fair contract.

Theorem 4. *If we choose to offer four contracts, one for each combination of type and group, an optimal solution exists and follows a specific structure outlined as follows:*

$$\begin{aligned}
& q_A^E : S'(q_A^E) - \theta^E = 0, \quad U_A^E = \rho\Delta\theta q_B^I \\
& q_A^I : q_A^I = \rho q_B^I, \quad U_A^I = 0 \\
& q_B^E : S'(q_B^E) - \theta^E = 0, \quad U_B^E = \Delta\theta q_B^I \\
& q_B^I : \rho \cdot \beta(S'(\rho q_B^I) - \theta^I) + \delta(S'(q_B^I) - \theta^I) - \frac{\gamma}{\gamma+\delta}\Delta\theta = 0, \quad U_B^I = 0
\end{aligned}$$

where $\rho = \frac{\alpha+\beta}{\alpha} \cdot \frac{\gamma}{\gamma+\delta}$.

Proof. Since U_A^E is included as a negative term in the objective function that we wish to maximize, it is necessary for either **IC-A** or **IPC-B** to be binding, as they impose lower bounds on the value of U_A^E. In this proof, we will consider both cases and solve the resulting problems.

First, consider the case where **IC-A** is binding. This implies that $U_A^E = \Delta\theta q_A^I \geq \frac{\alpha+\beta}{\alpha} \cdot \frac{\gamma}{\gamma+\delta}\Delta\theta q_B^I$. Consequently, the optimization problem transforms into:

$$\begin{aligned}
\max_{q} \quad & \alpha(S(q_A^E) - \theta^E q_A^E) + \beta(S(q_A^I) - \theta^I q_A^I) \\
& + \gamma(S(q_B^E) - \theta^E q_B^E) + \delta(S(q_B^I) - \theta^I q_B^I) - \frac{\alpha}{\alpha+\beta}\Delta\theta q_A^I \\
\text{s.t.} \quad & q_B^I \leq \frac{\alpha}{\alpha+\beta} \cdot \frac{\gamma+\delta}{\gamma} q_A^I
\end{aligned}$$

By observing the objective, we can deduce that we should increase q_B^I until the marginal utility equals the marginal cost (i.e., $S'(q_B^I) = \theta^I$), which corresponds to the first-best production level discussed earlier. However, due to the presence of the incentive compatibility constraint, which is implicit by the existence of a feasible U_B^E as assumed, this condition will not be met. Therefore, we can safely assume that the remaining constraint is also binding, as it places an upper bound on q_B^I. Hence, the resulting problem is to maximize:

$$\alpha(S(q_A^E) - \theta^E q_A^E) + \beta(S(q_A^I) - \theta^I q_A^I) + \gamma(S(q_B^E) - \theta^E q_B^E) + \delta(S(q_B^I) - \theta^I q_B^I) - \frac{\alpha}{\alpha+\beta}\Delta\theta q_A^I,$$

where $q_B^I = \frac{\alpha}{\alpha+\beta} \cdot \frac{\gamma+\delta}{\gamma} q_A^I$.

Now, consider the case where **IPC-B** is binding. This implies that $U_A^E = \frac{\alpha+\beta}{\alpha} \cdot \frac{\gamma}{\gamma+\delta} \Delta\theta q_B^I \geq \Delta\theta q_A^I$. Consequently, the optimization problem transforms into:

$$\begin{aligned}
\max_q \quad & \alpha(S(q_A^E) - \theta^E q_A^E) + \beta(S(q_A^I) - \theta^I q_A^I) \\
& + \gamma(S(q_B^E) - \theta^E q_B^E) + \delta(S(q_B^I) - \theta^I q_B^I) - \frac{\gamma}{\gamma+\delta} \Delta\theta q_B^I \\
\text{s.t.} \quad & q_A^I \leq \frac{\alpha+\beta}{\alpha} \cdot \frac{\gamma}{\gamma+\delta} q_B^I
\end{aligned}$$

Using a similar reasoning as in the previous case, we can assume that the constraint on q_A^I is also binding. Hence, the resulting problem is to maximize:

$$\alpha(S(q_A^E) - \theta^E q_A^E) + \beta(S(q_A^I) - \theta^I q_A^I) + \gamma(S(q_B^E) - \theta^E q_B^E) + \delta(S(q_B^I) - \theta^I q_B^I) - \frac{\gamma}{\gamma+\delta} \Delta\theta q_B^I,$$

where $q_A^I = \frac{\alpha+\beta}{\alpha} \cdot \frac{\gamma}{\gamma+\delta} q_B^I$. The important point to note is that this problem is equivalent to the one derived when assuming **IC-A** is binding. Hence, irrespective of the assumed binding constraint, we ultimately arrive at the same optimization problem.

Next, we assess the feasibility of U_B^E by leveraging **IPC-B**. This leads to the following implications:

$$\begin{aligned}
U_B^E &\geq \frac{\alpha}{\alpha+\beta} U_A^E + \frac{\delta}{\gamma+\delta} \Delta\theta q_B^I = \Delta\theta q_B^I \quad \textbf{(IC-B*)} \\
U_B^E &\leq \frac{\alpha}{\alpha+\beta} \cdot \frac{\gamma+\delta}{\gamma} U_A^E = \Delta\theta q_B^I \quad \textbf{(PC-B*)}
\end{aligned}$$

Hence, the feasible region defined collectively by **IC-B*** and **PC-B*** implies $U_B^E = \Delta\theta q_B^I$. Since $U_A^E = \frac{\alpha+\beta}{\alpha} \cdot \frac{\gamma}{\gamma+\delta} \Delta\theta q_B^I$ and $U_A^I = 0$, applying **FC** reveals that $U_B^I = 0$. This explains why the choice of assuming U_A^I or U_B^I to be zero, as stated in Proposition 9, is inconsequential. Regardless of which one we select, it becomes evident that the optimal contract induces both U_A^I and U_B^I to be zero.

Finally, the first-order optimality conditions yield:

$$\begin{aligned}
& q_A^E : S'(q_A^E) - \theta^E = 0 \\
& q_A^I : q_A^I = \rho q_B^I \\
& q_B^E : S'(q_B^E) - \theta^E = 0 \\
& q_B^I : \rho \cdot \beta(S'(\rho q_B^I) - \theta^I) + \delta(S'(q_B^I) - \theta^I) - \frac{\gamma}{\gamma+\delta} \Delta\theta = 0
\end{aligned}$$

where $\rho = \frac{\alpha+\beta}{\alpha} \cdot \frac{\gamma}{\gamma+\delta}$. $\square$

This theorem indicates that efficient agents from both groups will continue to produce at their first-best levels (i.e., $q_A^E = q_B^E = q_1^E$), while inefficient agents will experience a downward (or possibly upward) distortion based on their group affiliation. The reason behind this can be understood by recalling that the incentive compatible contract involves a trade-off between minimizing information rent and sacrificing overall efficiency. Introducing a fairness constraint further complicates the problem. In the absence of fairness, if one group is more efficient than the other, the more efficient group would receive a higher amount of information rent as a whole. Therefore, in addition to the factors already considered, achieving optimal balance in expected profit across different groups while minimizing efficiency loss (i.e., decline in total production levels) becomes crucial. This relationship will be further explored in the following corollary.

Corollary 1. *Depending on the value of ρ, the optimal fair contract exhibits the following characteristics:*

- *If $\rho > 1$, then $q_A^I < q_B^I$ and $U_A^E < U_B^E$.*
- *If $\rho < 1$, then $q_A^I > q_B^I$ and $U_A^E > U_B^E$.*
- *If $\rho = 1$, then $q_A^I = q_B^I = q_2^I$ and $U_A^E = U_B^E = U_2^E$.*

where $\rho = \frac{\alpha+\beta}{\alpha} \cdot \frac{\gamma}{\gamma+\delta}$.

Proof. For $\rho \neq 1$, the relationships $q_A^I = \rho q_B^I$ and $U_A^E = \rho U_B^E$ are straightforward. However, when $\rho = 1$, further attention is required. The condition $\rho = 1$ implies $\frac{\alpha}{\alpha+\beta} = \frac{\gamma}{\gamma+\delta}$, or equivalently, $\frac{\beta}{\alpha+\beta} = \frac{\delta}{\gamma+\delta}$. The value of q_B^I is determined by the equation $(\beta+\delta)(S'(q_B^I) - \theta^I) - \frac{\gamma}{\gamma+\delta}\Delta\theta = 0$. Simplifying, we have:

$$\begin{aligned}
S'(q_B^I) - \theta^I &= \frac{\gamma}{(\beta+\delta)(\gamma+\delta)}\Delta\theta \\
&= \frac{\gamma}{\beta(\gamma+\delta) + \delta(\gamma+\delta)}\Delta\theta \\
&= \frac{\gamma}{\delta(\alpha+\beta) + \delta(\gamma+\delta)}\Delta\theta \\
&= \frac{\gamma}{\delta(\alpha+\beta+\gamma+\delta)}\Delta\theta \\
&= \frac{\gamma}{\delta}\Delta\theta,
\end{aligned}$$

which is identical to the equation for the second-best production level in the incentive compatible contract. □

The corollary provides insights into the impact of the coefficient ρ, which quantifies the disparity between the efficiency levels across the groups, on the optimal fair contract. In cases where $\rho \neq 1$, indicating the presence of a gap in group-wise efficiency levels, the optimal fair contract takes this into account. It achieves fairness by redistributing the profit through assigning more production to the less efficient group. This allocation of resources empowers them to obtain a greater share of the information rent.

When $\rho = 1$, it signifies that the groups are indistinguishable in terms of their collective efficiency. In this case, the optimal fair contract coincides with the incentive compatible contract. Consequently, inefficient agents from both groups will produce at their second-best levels (i.e., $q_A^I = q_B^I = q_2^I$). This results in efficient agents from both groups receiving the same information rents as in the incentive compatible contract (i.e., $U_A^E = U_B^E = U_2^E$). This contract ensures equal treatment for both groups, acknowledging their similar efficiency levels and providing equitable outcomes.

In conclusion, to design contracts that align with our definition of fairness, which mandates equal expected profit for both groups, an optimal menu of contracts includes additional assignments and payments to the less efficient group in order to compensate for the disparity resulting from the implementation of the fairness constraint. Our interpretation is that by using this particular definition of fairness to derive an optimal contract, we are supporting the less efficient group, which may have limited resources or inadequate physical infrastructure compared to the more efficient group. This approach aims to subsidize the less efficient group, leveling the playing field for all participants involved.

3.3 Moral Hazard

Now, let us consider another form of information asymmetry known as moral hazard. Returning to the example of the sandwich restaurant, we previously discussed adverse selection where the company lacks knowledge about the restaurant's efficiency type. In moral hazard, the focus shifts from efficiency types to unobservable effort levels chosen by the agents. In this case, the restaurant has the ability to select either a high or low level of effort when making sandwiches. The chosen effort level influences the production outcome in a probabilistic manner. If the restaurant exerts a high level of effort, it increases the likelihood of achieving a high production level. Conversely, if they exert low effort, the probability of obtaining a low production level is higher. Importantly, the company's payment to the restaurant is based on the production level achieved, rather than the level of effort exerted. This is because only the restaurant is privy to the information regarding their actual effort exertion.

Mathematically, the agent decides on one of two effort levels: low ($e = 0$) or high ($e = 1$). These efforts incur costs, $\psi(e)$, where $\psi(e = 0) = 0$ and $\psi(e = 1) = \psi$. We model an agent's utility function, given payment t, as $U = u(t) - \psi(e)$, where $u(0) = 0$, $u' > 0$, and $u'' < 0$. Similarly, the principal's valuation of goods produced q is represented by $S(q)$, with $S(0) = 0$, $S' > 0$, and $S'' < 0$. To account for information asymmetry, we introduce a random production function, where an agent's production level can be either low ($q = q^L$) or high ($q = q^H$). The agent's choice to exert high effort strictly increases the probability of high production, with the probabilities defined as:

$$\mathbb{P}(q = q^H \mid e = 0) = \pi_0,$$
$$\mathbb{P}(q = q^H \mid e = 1) = \pi_1 > \pi_0.$$

Within this framework, the contract design problem entails the principal offering a menu of contracts to the agent, each represented as (q, t), where q denotes the production level and t the associated transfer payment. This menu of contracts allows the agent to choose the most suitable option. The principal's objective is to design an optimal menu of contracts that maximizes their utility.

When the agent exerts low effort, the principal's expected utility is expressed as:

$$\pi_0(S(q^H) - t^H) + (1 - \pi_0) \cdot (S(q^L) - t^L).$$

Similarly, when the agent exerts high effort, the principal's expected utility is given by:

$$\pi_1(S(q^H) - t^H) + (1 - \pi_1) \cdot (S(q^L) - t^L).$$

To optimize this situation, we consider the concept of the first-best level of effort, assuming that the principal can observe effort. To do this, the payments must ensure that the agent does not anticipate negative utility, leading to the *participation constraint*:

$$\pi_1 u(t^H) + (1 - \pi_1) u(t^L) - \psi \geq 0.$$

In this case, the principal can enforce zero expected utility with $t^L = t^H = t_1$, where $u(t_1) = \psi$. Next, we delve into the question of whether the principal can induce effort:

$$\pi_1(S(q^H) - t^H) + (1 - \pi_1) \cdot (S(q^L) - t^L) \geq \pi_0 \cdot S(q^H) + (1 - \pi_0) \cdot S(q^L).$$

This equation compares the expected utility with payment and high effort (left-hand side) against the expected utility with no payment and low effort (right-hand side). Rearranging the terms reveals:

$$(\pi_1 - \pi_0)(S(q^H) - S(q^L)) \geq \pi_1 t^H + (1 - \pi_1) t^L.$$

Here, the left-hand side reflects the expected gain from exerting effort, while the right-hand side indicates the first-best cost of inducing effort, that is, $t_1 = u^{-1}(\psi)$.

Now, returning to the situation where the principal cannot verify effort, and with a *risk-neutral* agent, we have $U = u(t) - \psi(e) = t - \psi(e)$. To induce high effort, the principal seeks to maximize:

$$\max_{t^H, t^L} \quad \pi_1(S(q^H) - t^H) + (1 - \pi_1)(S(q^L) - t^L) \tag{3.32}$$

$$\text{s.t.} \quad \pi_1 t^H + (1 - \pi_1) t^L - \psi \geq 0 \tag{3.33}$$

$$\pi_1 t^H + (1 - \pi_1) t^L - \psi \geq \pi_0 t^H + (1 - \pi_0) t^L \tag{3.34}$$

where the second constraint is the *incentive compatibility* constraint, ensuring that the agent prefers to exert high effort. Upon solving this, if the principal intends to induce high effort, the optimal payments can be determined as follows:

$$t^H = \frac{1-\pi_0}{\pi_1-\pi_0}\psi, \quad t^L = -\frac{\pi_0}{\pi_1-\pi_0}\psi.$$

The expected payment equals $\pi_1 t^H + (1-\pi_1)t^L = \psi$, aligning with the first-best cost of inducing effort, making moral hazard less problematic for a risk-neutral agent.

We can further impose a limited liability constraint, $t^L \geq -\ell$, where $\ell \geq 0$, to ensure that transfers remain above a certain threshold. If $\ell > \frac{\pi_0}{\pi_1-\pi_0}\psi$, the optimal payments remain unchanged, as previously discussed. However, if $\ell \leq \frac{\pi_0}{\pi_1-\pi_0}\psi$, then the optimal payments are revised to:

$$t^H = -\ell + \frac{1}{\pi_1-\pi_0}\psi, \quad t^L = -\ell.$$

Consequently, the agent earns a limited liability rent of $\pi_1 t^H + (1-\pi_1)t^L - \psi = -\ell + \frac{\pi_0}{\pi_1-\pi_0}\psi > 0$, resulting in a higher cost of inducing effort.

In the context of this framework, we will formulate the principal-agent model with moral hazard as follows:

$$\begin{array}{rl} \text{maximize} & \text{principal's expected profit} \\ \text{subject to} & \text{agent's participation} \\ & \text{agent's incentive compatibility} \\ & \text{agent's limited liability} \\ & \text{(+ contract fairness)} \end{array}$$

Similar to the previous model, our objective in this context is to design a contract that achieves agent participation and motivates them to exert a high level of effort. However, we also wish to address the fairness aspect to bridge the gap between the protected groups in terms of their abilities, which may have originated from disparities in resources, infrastructure, or socioeconomic barriers.

In this particular model, we consider the existence of two groups of agents, where one group has a higher probability of success, indicating a greater likelihood of attaining a high production level compared to the other group.

Fair Principal-Agent Model with Moral Hazard

Model 1

With a menu of contracts $\{(q^H, t^H), (q^L, t^L)\}$, the optimization problem is given by

$$
\begin{aligned}
\max_{t^H, t^L} \quad & \pi_{1,A}(S(q^H) - t^H) + (1 - \pi_{1,A})(S(q^L) - t^L) \\
& + \pi_{1,B}(S(q^H) - t^H) + (1 - \pi_{1,B})(S(q^L) - t^L) && (3.35) \\
\text{s.t.} \quad & \pi_{1,A}t^H + (1 - \pi_{1,A})t^L - \psi \geq 0 && (3.36) \\
& \pi_{1,B}t^H + (1 - \pi_{1,B})t^L - \psi \geq 0 && (3.37) \\
& \pi_{1,A}t^H + (1 - \pi_{1,A})t^L - \psi \geq \pi_{0,A}t^H + (1 - \pi_{0,A})t^L && (3.38) \\
& \pi_{1,B}t^H + (1 - \pi_{1,B})t^L - \psi \geq \pi_{0,B}t^H + (1 - \pi_{0,B})t^L && (3.39) \\
& t^L \geq -\ell && (3.40) \\
& \pi_{1,A}t^H + (1 - \pi_{1,A})t^L - \psi = \pi_{1,B}t^H + (1 - \pi_{1,B})t^L - \psi && (3.41)
\end{aligned}
$$

where the probabilities are constrained within the range of 0 to 1, specifically $0 < \pi_{1,A}, \pi_{1,B}, \pi_{0,A}, \pi_{0,B} < 1$, and with the conditions that $\pi_{1,A} > \pi_{0,A}$ and $\pi_{1,B} > \pi_{0,B}$ (indicating that an agent choosing to exert high effort strictly increases the probability of high production). Additionally, we assume $\pi_{1,A} \neq \pi_{1,B}$ to introduce a distinction in the overall group-wise productivity levels between the two groups, and $\ell \geq 0$ to guarantee that the agent's transfer always exceed a certain pre-determined level.

The following theorem presents a significant result demonstrating the infeasibility of the above optimization problem, thus underscoring the impossibility of attaining an optimal fair contract when contracts are solely based on production levels.

Theorem 5. *If $\pi_{1,A} \neq \pi_{1,B}$ and $\psi > 0$, then there exists no feasible solution that satisfies both incentive compatibility and fairness constraints.*

Proof. Fairness constraint (3.41) implies that $(\pi_{1,A} - \pi_{1,B})t^H = (\pi_{1,A} - \pi_{1,B})t^L$, so we have $t^H = t^L$. However, combined with incentive compatibility constraint (3.38), this implies that $t^L - \psi \geq t^L$, which means $\psi \leq 0$. This is a contradiction since we had assumed $\psi > 0$. Hence, fairness and incentive compatibility are not jointly feasible. □

The conditions outlined in Theorem 5 are well-founded. The first condition posits that there must be a discernible distinction in overall productivity levels between the two groups. Without such differentiation, the groups would be indistinguishable, and there would be no imperative to incorporate fairness into the contract design. The second condition ensures that the cost associated with exerting high effort must be unequivocally positive. This factor accounts for the reality that agents bear a cost commensurate with their effort level, which is at the heart of the moral hazard issue stemming from the information asymmetry.

Similar to our findings in the context of adverse selection, a key observation arises: when offering two contracts, each tailored to specific production levels, we encounter a fundamental challenge. In essence, if our contract design focuses solely on production levels, achieving our definition of fairness across distinct groups becomes infeasible. An intriguing aspect to note is that we cannot craft a fair optimal contract by considering production alone, especially when disparities exist between groups, such as variations in their group-wise productivity. This embodies the vital importance of inclusively considering protected groups when fashioning incentive structures to uphold principles of fairness in our analysis.

Model 2

With $\{(q^H, t_A^H), (q^L, t_A^L)\}$ for A and $\{(q^H, t_B^H), (q^L, t_B^L)\}$ for B, the optimization problem is given by

$$
\begin{aligned}
\max_{t^H, t^L} \quad & \pi_{1,A}(S(q^H) - t_A^H) + (1 - \pi_{1,A})(S(q^L) - t_A^L) \\
& + \pi_{1,B}(S(q^H) - t_B^H) + (1 - \pi_{1,B})(S(q^L) - t_B^L) && (3.42) \\
\text{s.t.} \quad & \pi_{1,A} t_A^H + (1 - \pi_{1,A}) t_A^L - \psi \geq 0 \quad \textbf{(PC-A)} && (3.43) \\
& \pi_{1,B} t_B^H + (1 - \pi_{1,B}) t_B^L - \psi \geq 0 \quad \textbf{(PC-B)} && (3.44) \\
& \pi_{1,A} t_A^H + (1 - \pi_{1,A}) t_A^L - \psi \geq \pi_{0,A} t_A^H + (1 - \pi_{0,A}) t_A^L \quad \textbf{(IC-A)} && (3.45) \\
& \pi_{1,B} t_B^H + (1 - \pi_{1,B}) t_B^L - \psi \geq \pi_{0,B} t_B^H + (1 - \pi_{0,B}) t_B^L \quad \textbf{(IC-B)} && (3.46) \\
& t_A^L \geq -\ell \quad \textbf{(LC-A)} && (3.47) \\
& t_B^L \geq -\ell \quad \textbf{(LC-B)} && (3.48) \\
& \pi_{1,A} t_A^H + (1 - \pi_{1,A}) t_A^L - \psi = \pi_{1,B} t_B^H + (1 - \pi_{1,B}) t_B^L - \psi \quad \textbf{(FC)} && (3.49)
\end{aligned}
$$

Here, the acronyms represent the following constraints: **PC** for participation constraints, **IC** for moral hazard incentive compatibility constraints, **LC** for limited liability constraints for each group, and **FC** for the fairness constraint.

Similar to our approach in addressing adverse selection, a primary challenge in solving this problem lies in identifying the constraints imposed by participation and limited liability that are binding at the optimal solution. Due to the multitude of constraints, our strategy involves first hypothesizing the active constraints and subsequently verifying that the disregarded constraints indeed hold strictly. This method proves advantageous in both solving the optimization problem and enhancing our comprehension of the underlying economic rationale within the context of our model.

To simplify the problem, we identify redundant constraints and eliminate the equality constraint by integrating it into the objective function and the inequality constraints, expressing them in terms of t_A^H and t_A^L. This approach allows us to initially postulate the active constraints, with subsequent verification that the omitted constraints remain satisfied.

We can easily observe that the participation constraint for group A (3.43) and the fairness constraint (3.49), when combined, inherently imply the participation constraint for group B

(3.44). We then incorporate the fairness constraint (3.49) into the objective (3.42) as follows:

$$\begin{aligned}\max\ (3.42) &= \max\ -(\pi_{1,A}t_A^H + (1-\pi_{1,A})t_A^L + \pi_{1,B}t_B^H + (1-\pi_{1,B})t_B^L) \\ &= \min\ \pi_{1,A}t_A^H + (1-\pi_{1,A})t_A^L + \pi_{1,B}t_B^H + (1-\pi_{1,B})t_B^L \\ &= \min\ 2(\pi_{1,A}t_A^H + (1-\pi_{1,A})t_A^L) \\ &= \min\ \pi_{1,A}t_A^H + (1-\pi_{1,A})t_A^L.\end{aligned}$$

Likewise, we can incorporate the fairness constraint (3.49) into the incentive compatibility constraint for group B (3.46). This reformulates the optimization problem using only t_A^H and t_A^L, making it independent of t_B^H and t_B^L, and can be expressed concisely as follows:

$$\min_{t_A^H, t_A^L} \quad \pi_{1,A}t_A^H + (1-\pi_{1,A})t_A^L \tag{3.50}$$

$$\text{s.t.} \quad \pi_{1,A}t_A^H + (1-\pi_{1,A})t_A^L - \psi \geq 0 \quad \textbf{(PC-A)} \tag{3.51}$$

$$\pi_{1,A}t_A^H + (1-\pi_{1,A})t_A^L - \psi \geq \pi_{0,A}t_A^H + (1-\pi_{0,A})t_A^L \quad \textbf{(IC-A)} \tag{3.52}$$

$$t_A^L \geq -\ell \quad \textbf{(LC-A)} \tag{3.53}$$

Now, we can proceed to solve the simplified problem for group A. However, it is essential to verify that the omitted constraints, namely **(PC-B)**, **(IC-B)**, and **(LC-B)**, are satisfied when identifying a feasible solution for group B. We will present the following theorem characterizing the optimal fair contract.

Theorem 6. *If we choose to offer four contracts, one for each combination of production and group, an optimal solution exists and follows a specific structure outlined as follows:*

- *If $\ell > \tau_A$, then*

$$t_A^H = \frac{1-\pi_{0,A}}{\pi_{1,A}-\pi_{0,A}}\psi, \quad t_A^L = -\tau_A$$

$$t_B^H = \frac{1-\pi_{0,B}}{\pi_{1,B}-\pi_{0,B}}\psi, \quad t_B^L = -\tau_B$$

- *If $\tau_B \leq \ell \leq \tau_A$, then*

$$t_A^H = -\ell + \frac{1}{\pi_{1,A}-\pi_{0,A}}\psi, \quad t_A^L = -\ell$$

$$t_B^H = -\ell + \frac{\pi_{0,A}/\pi_{0,B}}{\pi_{1,A}-\pi_{0,A}}\psi, \quad t_B^L = -\ell$$

- *If $0 \leq \ell < \tau_B$, then*

$$t_A^H = -\ell + \frac{1}{\pi_{1,A}-\pi_{0,A}}\psi, \quad t_A^L = -\ell$$

$$t_B^H = -\ell + \left(\frac{1}{\pi_{1,A}-\pi_{0,A}} + \frac{\pi_{1,A}}{\pi_{1,B}}(\tau_A - \tau_B)\right)\psi, \quad t_B^L = -\ell + (\tau_A - \tau_B)\psi$$

where $\tau_A = \frac{\pi_{0,A}}{\pi_{1,A}-\pi_{0,A}}\psi$ *and* $\tau_B = \frac{\pi_{0,B}}{\pi_{1,B}-\pi_{0,B}}\psi$.

Proof. To solve this problem, we will begin by introducing two key constants: τ_A and τ_B, defined as follows: $\tau_A = \frac{\pi_{0,A}}{\pi_{1,A}-\pi_{0,A}}\psi$ and $\tau_B = \frac{\pi_{0,B}}{\pi_{1,B}-\pi_{0,B}}\psi$. These constants represent threshold values that determine when agents become protected by the limited liability clause. They are essential for considering scenarios where either group A, group B, or both, are subject to limited liability constraints. For the purpose of this proof, we will make the assumption that $\tau_A \geq \tau_B$. It is important to note that a parallel proof can be formulated for the opposite direction. We will now analyze the problem by exploring three distinct cases: (i) $\ell > \tau_A$, (ii) $\tau_B \leq \ell \leq \tau_A$, and (iii) $0 \leq \ell < \tau_B$.

Firstly, if $\ell > \tau_A$, the problem effectively reduces to solving it in the absence of limited liability constraints, and the optimal solution aligns with the first-best contract. This outcome arises because the value of ℓ exceeds the thresholds at which limited liability constraints provide meaningful protection. In simpler terms, since $\ell > \tau_A$ and $\ell > \tau_B$, this implies that agents fare better by receiving either τ_A or τ_B instead of ℓ. Furthermore, the fairness constraint is automatically met in this specific scenario, as agents from both groups anticipate zero profit.

Secondly, when $\tau_B \leq \ell \leq \tau_A$, the limited liability constraint becomes active. In this scenario, the optimal solution for group A is equivalent to the second-best contract. Now, the challenge lies in identifying a feasible solution for group B that adheres to the omitted constraints. An optimal contract for group B is jointly determined by

$$\pi_{1,B}t_B^H + (1-\pi_{1,B})t_B^L - \psi \geq \pi_{0,B}t_B^H + (1-\pi_{0,B})t_B^L \quad \textbf{(IC-B)}$$
$$t_B^L \geq -\ell \quad \textbf{(LC-B)}$$

$$\pi_{1,B}t_B^H + (1-\pi_{1,B})t_B^L - \psi = \pi_{1,A}t_A^H + (1-\pi_{1,A})t_A^L - \psi = -\ell + \frac{\pi_{0,A}}{\pi_{1,A}-\pi_{0,A}}\psi \quad \textbf{(FC)}$$

If we let both **(IC-B)** and **(LC-B)** be binding, then $t_B^L = -\ell$, and the problem becomes

$$\pi_{0,B}t_B^H - (1-\pi_{0,B})\ell = -\ell + \frac{\pi_{0,A}}{\pi_{1,A}-\pi_{0,A}}\psi.$$

Solving for t_B^H yields

$$t_B^H = -\ell + \frac{\pi_{0,A}/\pi_{0,B}}{\pi_{1,A}-\pi_{0,A}}\psi.$$

Finally, if $0 \leq \ell < \tau_B$, the limited liability constraint becomes applicable to both groups. As a result, the optimal solution for group A aligns with the second-best contract. Now, the task at hand is to identify a feasible solution for group B that fulfills the unconsidered constraints. In this scenario, it also involves finding an optimal balance between group A and group B, as both are governed by limited liability constraints, and the fairness constraint dictates that they should have the same expected profit. An optimal contract for group B is jointly established by

$$\pi_{1,B}t_B^H + (1-\pi_{1,B})t_B^L - \psi \geq \pi_{0,B}t_B^H + (1-\pi_{0,B})t_B^L \quad \textbf{(IC-B)}$$
$$t_B^L \geq -\ell \quad \textbf{(LC-B)}$$
$$\pi_{1,B}t_B^H + (1-\pi_{1,B})t_B^L - \psi = -\ell + \frac{\pi_{0,A}}{\pi_{1,A}-\pi_{0,A}}\psi \quad \textbf{(FC)}$$

If we assume **(IC-B)** to be binding, then we have the following system of equations:

$$\pi_{1,B}t_B^H + (1-\pi_{1,B})t_B^L - \psi = -\ell + \frac{\pi_{0,A}}{\pi_{1,A}-\pi_{0,A}}\psi$$
$$\pi_{0,B}t_B^H + (1-\pi_{0,B})t_B^L = -\ell + \frac{\pi_{0,A}}{\pi_{1,A}-\pi_{0,A}}\psi$$

Solving for t_B^H and t_B^L yields

$$t_B^H = -\ell + \left(\frac{1}{\pi_{1,A}-\pi_{0,A}} + \frac{\pi_{1,A}}{\pi_{1,B}}(\tau_A - \tau_B)\right)\psi, \quad t_B^L = -\ell + (\tau_A - \tau_B)\psi.$$

where **(LC-B)** is satisfied because $t_B^L = -\ell + (\tau_A - \tau_B)\psi \geq -\ell$ since we had assumed $\tau_A \geq \tau_B$ and $\psi > 0$. □

In an intuitive sense, these regions can be understood as follows:

- $\ell > \tau_A$: In this scenario, the limited liability constraints are essentially irrelevant, and the principal can enforce a contract that results in zero expected profit, irrespective of production levels and group membership.

- $\tau_B \leq \ell \leq \tau_A$: This case is particularly interesting. Without the fairness constraint, it would lead to an outcome where group A is covered by limited liability, while group B is not. In other words, an optimal contract would leave zero profit for group B, whereas group A expects strictly positive profit due to their coverage by the limited liability clause. However, the presence of the fairness constraint complicates matters. To achieve a fair contract, a delicate balance must be struck between the information rent the principal would originally have paid and the allocation of this rent to the group that would not have received it had fairness not been a consideration. Moreover, the value of t_B^H is determined in such a way that it offers higher compensation to groups with a lower probability of success under this contract. It becomes evident as follows: $t_B^H > t_A^H$ when $\pi_{0,A} > \pi_{0,B}$, and conversely, $t_A^H > t_B^H$ when $\pi_{0,A} < \pi_{0,B}$. This essentially translates into providing support or a subsidy to the group characterized by lower overall productivity and, consequently, a reduced chance of success.

- $0 \leq \ell < \tau_B$: This represents a situation where the limited liability constraints apply to both group A and group B. Under these constraints, both groups are protected, ensuring that agents receive strictly positive expected profits. This scenario holds a particular interest, as it demands the simultaneous fulfillment of both fairness and limited liability constraints, necessitating the precise determination of supplementary compensation for agents within group B. Specifically, for t_B^H, this additional payment amounts to $\frac{\pi_{1,A}}{\pi_{1,B}}(\tau_A - \tau_B)\psi$. This term hinges on the interplay between $\pi_{1,A}$ and $\pi_{1,B}$, reflecting that the extra compensation is intricately calibrated based on the genuine disparity between these two groups. In essence, the greater the gap in their probabilities of achieving high production levels, the larger this supplementary payment becomes. This can also be construed as a form of financial support extended to the group with a lower probability of attaining high production.

We would like to briefly discuss the implications of the assumption made in this proof. It was assumed that $\tau_A \geq \tau_B$, but it can be demonstrated in a similar manner for the reverse direction. In practical applications, we can always designate the group with the higher threshold as group A and the other as group B. Upon closer examination of these thresholds, the numerator represents the probability of success when low effort is exerted, while the denominator reflects the difference between the probabilities of success when high and low effort is applied. It is numerically possible for these threshold values to be equal, even when one group, say group A, has a significantly higher chance of success than the other, group B. For instance, we might have $\tau_A = \tau_B = 1$ when $\pi_{1,A} = 2\pi_{0,A} = 2\pi_{1,B} = 4\pi_{0,B}$. However, it is more reasonable to assert that $\tau_A \geq \tau_B$ for two key reasons. First, the numerator will be larger for group A because group A has a higher probability of success when they choose not to exert high effort, i.e., $\pi_{0,A} \geq \pi_{0,B}$. Furthermore, the denominator will be smaller for group A. In real-world scenarios, the difference between the probabilities for different levels of effort is often relatively minimal for group A, whereas it may be substantial for group B. This could be attributed to factors such as group B's relative lack of resources or infrastructure, which could disproportionately affect their probabilities of success.

To sum up, by adopting a menu of contracts specifically tailored to each group and potential outcome, we can paradoxically create a contract that is agnostic to group membership in terms of the expected profit agents derive from participating in the incentive scheme. If we disregard the fairness constraint and optimize the problem separately for each group, we would generate a set of contracts that favor the group with a higher likelihood of achieving a high production level. This approach could perpetuate and potentially exacerbate the existing disparities between the two groups. Therefore, our objective is to devise a contract that is independent of group membership, ensuring that agents from both groups attain the same expected profit regardless of their group affiliation. To achieve this, we propose the idea of an additional payment, which can be seen as aid or subsidy provided to the less privileged group with a lower probability of success. This additional payment serves as an investment to narrow the gap between the groups in the future, promoting fairness as defined in our model.

Example: Fair Incentives for Green Retrofitting

The buildings sector accounts for approximately 76% of electricity and 40% of total energy consumption, along with associated greenhouse gas emissions, in the U.S. [85]. According to the UN Environment Programme [87], over 80% of a building's energy consumption occurs during the occupancy operation stage, rather than the construction stage. To mitigate the environmental and economic costs of existing buildings, governments increasingly promote green (i.e., low-carbon) retrofitting projects through financial incentive mechanisms [12, 48].

A crucial aspect of green retrofitting involves the installation of solar panels to harness renewable energy technologies. The design of solar retrofitting projects comprises a principal-agent problem between the government (i.e., principal) and building owners (i.e., agents). Various alternatives for solar panels (e.g., monocrystalline, polycrystalline, etc.) come with differing energy efficiency levels and price points. Building owners decide on the type of solar panels for their roofs primarily based on government incentives. However, a significant challenge arises because many private building owners lack enthusiasm for adopting energy-conscious behaviors [77, 81, 48, 68]. As a result, governments often encounter a moral hazard problem when designing incentive policies, as they remain uninformed about the owners' efforts in selecting more efficient panels. The level of effort exerted by owners directly influences the cost savings and energy efficiency achieved through solar retrofitting.

In addition to the building owner's effort level, the effectiveness of government incentive policies is significantly influenced by the socioeconomic barriers in the property market in which the building is situated—whether it falls within a low-income or high-income region. Research indicates that current governmental incentives have deterred building owners in low-income regions from engaging in retrofitting [1, 74, 93, 42]. Recognizing these disparities, our analytical results highlight the need to devise new subsidy structures. The fair incentive policies proposed in this study aim to bridge the gap in retrofitting subsidies between the two income groups. These policies not only tackle the socioeconomic barriers but also encourage greater participation in green retrofitting projects, thereby advancing the broader goals of environmental sustainability and equitable development.

3.4 Conclusion

In the realm of incentive design, principal-agent models have long been a cornerstone for ensuring efficient outcomes. Traditionally, the primary focus has been on aligning interests, mitigating risks associated with adverse selection, and overcoming challenges related to moral hazard. However, a glaring oversight in this framework has been the omission of fairness considerations, potentially leading to detrimental consequences for certain demographic groups, especially when viewed through the lenses of race, gender, or other defining characteristics. The repercussions of such oversight are manifold, ranging from societal inequalities to a pervasive sense of inequity. It prompted us to question not only the efficiency of these incentive structures but also their ethical underpinnings. Through this work, we delved into the core of these concerns by addressing two fundamental dimensions: (i) discerning what is feasible and optimal, and (ii) unraveling what is fair and equitable, particularly in the milieu of contract design. In doing so, we sought to illuminate the path towards a more inclusive and just framework for incentive design by introducing fairness into optimization problems and elucidating the profound implications of this integration.

Chapter 4

Optimally Designing Cybersecurity Insurance Contracts to Encourage the Sharing of Medical Data

4.1 Introduction

The rapid development of new artificial intelligence algorithms for health care has the potential to lead to an era of computational precision health [72, 43, 55, 96, 60, 2, 33]. The development of these algorithms requires access to large sets of medical data. Nonetheless, the sharing of such medical data poses risks to patients due to the possible loss in privacy or livelihood that can occur when medical data is stolen or used in non-permitted ways. New ideas for the cybersecurity of medical data are needed to ensure that these new advances can continue to be developed.

Privacy Risks from Sharing Medical Data

A unique aspect of medical data is that even when it has been anonymized/deidentified [82, 49, 47, 20] (in accordance with legislation like HIPAA [86] or GDPR [21]) prior to sharing, the data can often be deanonymized/reidentified [59, 62, 19, 63]. Examples include deanonymization of a Massachusetts hospital database by joining it with a public voter database [83] and reidentification of a physical activity data set from the National Health and Nutrition Examination Survey (NHANES) using standard machine learning [56].

In addition, a recent study has revealed that more than two-thirds of hospital data breaches include sensitive demographic and financial information that could lead not only to fraud and identity theft but also to discrimination and violation of fundamental rights [34]. This highlights the necessity of developing approaches to safeguard patients and health care providers against cybersecurity threats.

Cybersecurity of Medical Data

The above described privacy risks deter health care providers from sharing their data [90, 66]. One possible approach to mitigating some of the risks with sharing health data is through the design of cybersecurity insurance contracts. For instance, cybersecurity insurance can be used to partially compensate for the costs involved with recovery from a cyber-related security breach or similar incidents [51].

A growing literature studies cybersecurity insurance. For instance, [40, 61, 9] focus on the interdependent security problem to verify whether firms have adequate incentives to invest in protection against a risk whose magnitude depends on the action of others. The work in [6, 7] introduces new models and measures for the correlation of cyber-risks within and across independent firms, while [79] investigates the issue of information asymmetries, namely in the form of moral hazard, when cyber-insurers cannot observe individual user security levels. The studies [8, 76] provide a unifying framework to address the aforementioned hurdles that complicate risk management via cyber-insurance.

Contributions and Outline

In this chapter, we study the problem of incorporating cybersecurity insurance, which has been mainly explored in the setting of interdependent and correlated networks, into the design of contracts that govern the sharing of medical data. Such contracts would not only protect health care providers against losses resulting from a cyber-attack, but have the potential to foster the sharing of medical data.

In Section 4.2, we analyze the scenario in which a health care provider sells medical data to a technology firm that uses the data to develop new artificial intelligence algorithms. We provide mathematical models for both parties, formulate a contract design problem in the setting of a principal-agent model with moral hazard [41], derive the optimal contract, and discuss insights gained from the optimal contract. In Section 4.3, we analyze a second scenario in which a group of health care providers forms a consortium to share medical data with each other for the purpose of conducting scientific research and improving patient care. Again, we provide a mathematical model for the health care providers, formulate a contract design problem, derive the optimal contract, and discuss insights gained from the optimal contract.

4.2 Scenario A: Health Care Provider Selling Medical Data to Technology Firm

The first scenario we study is that of a health care provider selling medical data to a technology firm that is developing artificial inteligence algorithms using the shared data. Here, an important consideration to the health care provider is the quality of cybersecurity that the technology firm uses to protect any medical data they receive. If the firm suffers from a data breach, then the health care provider itself will face liability from those patients whose medical data has been breached. Thus, the health care provider will want to structure their contract with the firm in such a way that the firm is incentivized to invest in the cybersecurity of the medical data.

In this scenario that we consider, the health care provider has two options available to mitigate the liability risks associated with a data breach. The first is that the health care provider is able to impose a fine or penalty on the firm if the technology firm suffers from a data breach. If this fine or penalty is sufficiently large, it can incentivize the firm to invest in cybersecurity that protects the data. This fine or penalty is in addition to the fee that the firm is charged in return for access to the medical data. The second is that the health care provider is able to purchase cybersecurity insurance from an external insurance agency.

Technology Firm Model

The financial value to the technology firm of the shared medical data is V. This financial value is derived from the firm's ability to use the data to develop new artificial intelligence algorithms for health care, which can be sold to various health care providers. To get this data, the firm must pay the quantity ϕ to the health care provider. The technology firm is responsible for securing the data they receive. If the firm suffers from a data breach, they are required by the contract to pay a fine or penalty t to the health care provider.

The technology firm chooses an investment level i in cybersecurity that protects the data. The firm chooses between a high $(i = 1)$ or low $(i = 0)$ level of investment. If the firm chooses high investment, then the probability of a breach is $\alpha \in (0, 1)$, and the firm spends ψ for this investment level. If the firm chooses low investment, then the probability of a breach is $\gamma \in (0, 1)$, and (without loss of generality) the firm has zero expenditure for this investment level. We assume that $\alpha < \gamma$, meaning that a high level of investment strictly *lowers* the probability of a data breach. We assume that the technology firm chooses their investment level by maximizing their expected profit:

$$i^*(\phi, t) = \arg\max_{i \in \{0,1\}} \ (1 - i) \cdot F^l(\phi, t) + i \cdot F^h(\phi, t), \tag{4.1}$$

where expected profit under a low level of investment is

$$F^l(\phi, t) = V - \phi - \gamma \cdot t, \tag{4.2}$$

and the expected profit under a high level of investment is

$$F^h(\phi, t) = V - \phi - \psi - \alpha \cdot t. \tag{4.3}$$

Note that the technology firm is *risk neutral* in this model.

Health Care Provider Model

The financial value to the health care provider of their own medical data is W. This financial value is derived from the provider's ability to use the data to self-improve the quality of its health care services through improved patient treatment and care delivery processes, as well as through medical research. If the technology firm suffers a data breach, then the health care provider has to spend L to address its various liabilities to the affected patients. Since t is a fine or penalty on the firm in the event of a breach, we assume $t \leq L$. Having $t > L$ is unrealistic because it would mean the healthcare provider profits from a data breach at the firm.

Furthermore, the health care provider can choose to purchase a policy to insure against their liabilities in the event of a breach. Under the assumption of an actuarially fair policy [76], which would be expected to occur when there are a large number of insurers in the insurance marketplace, the health care provider can purchase an insurance policy that pays out L_c under the event of a data breach at the cost of pL_c, where p is the probability of a data breach.

Finally, we assume the health care provider is *risk averse*. This means that if the health care provider earns a financial revenue of x, then their utility for that revenue is $U(x)$ for a function $U(\cdot)$ that is strictly increasing and concave. Under the additional assumption that $U(\cdot)$ is differentiable, this risk aversion assumption means that $U'(\cdot) > 0$ and $U''(\cdot) < 0$.

Contract Design Problem

In this scenario, the health care provider faces a contract design problem in which their goal is to pick the purchase price ϕ, the value of the fine or penalty t, and the insurance policy payout L_c so as to maximize their own expected utility. This contract design problem can be written as the following bilevel program:

$$\begin{aligned} \max_{\phi,t,L_c} \ & (1 - i^*(\phi, t)) \cdot H^l(\phi, t, L_c) + i^*(\phi, t) \cdot H^h(\phi, t, L_c) \\ \text{s.t. } & i^*(\phi, t) = \arg \max_{i \in \{0,1\}} (1 - i) \cdot F^l(\phi, t) + i \cdot F^h(\phi, t) \\ & (1 - i^*(\phi, t)) \cdot F^l(\phi, t) + i^*(\phi, t) \cdot F^h(\phi, t) \geq 0 \\ & \phi \geq 0,\ t \in [0, L],\ L_c \geq 0 \end{aligned} \tag{4.4}$$

where we note that the health care provider's expected utility when the technology firm has a low level of investment in cybersecurity of the health data is given by

$$H^l(\phi, t, L_c) = \gamma \cdot U(W + \phi - \gamma \cdot L_c - L + t + L_c) + (1 - \gamma) \cdot U(W + \phi - \gamma \cdot L_c), \tag{4.5}$$

the health care provider's expected utility when the technology firm has a high level of investment is

$$H^h(\phi, t, L_c) = \alpha \cdot U(W + \phi - \alpha \cdot L_c - L + t + L_c) + (1 - \alpha) \cdot U(W + \phi - \alpha \cdot L_c), \tag{4.6}$$

and the second constraint in (4.4) is a *participation constraint* that ensures the purchase cost ϕ and fine or penalty t are such that the technology firm does not expect to lose money.

Optimal Contract

We next proceed to solve the contract design problem (4.4) through a series of steps. Let ϕ^*, t^*, L_c^* denote the optimal contract, meaning they maximize the objective function of (4.4). We first characterize the optimal insurance coverage pay out.

Proposition 10. *We have that $L_c^* = L - t^*$.*

Proof. We first consider the case where $i^*(\phi^*, t^*) = 0$. In this case, the objective function of (4.4) is $H^l(\phi, t, L_c)$, and the second constraint in (4.4) is $F^l(\phi, t) \geq 0$. Next, note that the first-order stationarity condition is

$$\begin{aligned} 0 = \partial_{L_c} H^l(\phi, t, L_c) = \gamma \cdot (1-\gamma) \cdot U'(W + \phi - \gamma \cdot L_c - L + t + L_c) \\ - \gamma \cdot (1-\gamma) \cdot U'(W + \phi - \gamma \cdot L_c). \end{aligned} \tag{4.7}$$

Since we assumed that $U''(\cdot) < 0$, this means the above is satisfied when $W + \phi - \gamma \cdot L_c - L + t + L_c = W + \phi - \gamma \cdot L_c$. Hence at optimality we have $L_c^* = L - t^*$, which is feasible since $t^* < L$ implies $L_c^* \geq 0$. The proof for the case $i^*(\phi^*, t^*) = 1$ proceeds almost identically. □

The implication of the above result is that we can rewrite the contract design problem as

$$\begin{aligned} \max_{\phi,t}\ & (1 - i^*(\phi, t)) \cdot H^l(\phi, t) + i^*(\phi, t) \cdot H^h(\phi, t) \\ \text{s.t. }\ & i^*(\phi, t) = \arg\max_{i \in \{0,1\}}\ (1 - i) \cdot F^l(\phi, t) + i \cdot F^h(\phi, t) \\ & (1 - i^*(\phi, t)) \cdot F^l(\phi, t) + i^*(\phi, t) \cdot F^h(\phi, t) \geq 0 \\ & \phi \geq 0,\ t \in [0, L] \end{aligned} \tag{4.8}$$

where

$$H^l(\phi, t) := H^l(\phi, t, L - t) = U(W + \phi - \gamma \cdot (L - t)), \tag{4.9}$$

$$H^h(\phi, t) := H^h(\phi, t, L - t) = U(W + \phi - \alpha \cdot (L - t)). \tag{4.10}$$

We next use the above reformulation to characterize the optimal purchase price and fine or penalty amount.

Proposition 11. *If $i^*(\phi^*, t^*) = 0$, then an optimal choice solves the optimization problem*

$$\begin{aligned} \max_{\phi,t}\ & \phi + \gamma \cdot t \\ \text{s.t. }\ & \phi + \gamma \cdot t \leq V \\ & (\gamma - \alpha) \cdot t < \psi \\ & \phi \geq 0,\ t \in [0, L] \end{aligned} \tag{4.11}$$

If $i^*(\phi^*, t^*) = 1$, then an optimal choice solves the optimization problem

$$\begin{aligned} \max_{\phi,t} \ & \phi + \alpha \cdot t \\ \text{s.t. } & \phi + \alpha \cdot t \le V - \psi \\ & (\gamma - \alpha) \cdot t \ge \psi \\ & \phi \ge 0, \ t \in [0, L] \end{aligned} \tag{4.12}$$

Proof. We first consider the case where $i^*(\phi^*, t^*) = 0$. In this case, we have that $F^l(\phi, t) > F^h(\phi, t)$ (which is equivalent to $\psi > (\gamma - \alpha) \cdot t$), that the second constraint of (4.8) is

$$F^l(\phi, t) = V - \phi - \gamma \cdot t \ge 0, \tag{4.13}$$

(which is equivalent to $\phi + \gamma \cdot t \le V$), and that the objective function of (4.8) is $H^l(\phi, t)$. Since $U'(\cdot) > 0$, this means $H^l(\phi, t)$ is strictly increasing in $\phi + \gamma \cdot t$. The above observations imply that (4.11) provides an optimal choice. The proof for the case $i^*(\phi^*, t^*) = 1$ is nearly identical. □

We conclude by using the above characterization to finish deriving an optimal contract for this scenario.

Theorem 7. *If $\psi > (\gamma - \alpha) \cdot L$ or $\psi > (\gamma - \alpha) \cdot V/\gamma$, then an optimal contract is $(\phi^*, t^*, L_c^*) = (V, 0, L)$. If $\psi \le (\gamma - \alpha) \cdot L$ and $\psi \le (\gamma - \alpha) \cdot V/\gamma$, then an optimal contract is $(\phi^*, t^*, L_c^*) = (V - \gamma/(\gamma - \alpha)\psi, \psi/(\gamma - \alpha), L - \psi/(\gamma - \alpha))$.*

Proof. If $\psi > (\gamma - \alpha) \cdot L$, then this means $W + V - \gamma \cdot L > W + V - \psi - \alpha \cdot L$, which implies $H^l > H^h$ since $U'(\cdot) > 0$. Hence, the optimal choice is $i^*(\phi, t) = 0$. This means the choice $t^* = 0$ and $\phi^* = V$ is optimal by Proposition 11. If $\psi > (\gamma - \alpha) \cdot V/\gamma$, then (4.12) is infeasible. This means applying Proposition 11 tells us that the optimal choice is $i^*(\phi, t) = 0$, and that we can again choose $t^* = 0$ and $\phi^* = V$. Finally, if $\psi \le (\gamma - \alpha) \cdot L$ and $\psi \le (\gamma - \alpha) \cdot V/\gamma$, then this means $W + V - \gamma \cdot L \le W + V - \psi - \alpha \cdot L$, which implies $H^l \le H^h$ since $U'(\cdot) > 0$. Moreover, the condition $\psi \le (\gamma - \alpha) \cdot V/\gamma$ implies that (4.12) is feasible. This means the choice $t^* = \psi/(\gamma - \alpha)$ (and note $t^* \le L$ since $\psi/(\gamma - \alpha) \le L$ in this case) and $\phi^* = V - \psi - \alpha/(\gamma - \alpha)\psi = V - \gamma/(\gamma - \alpha)\psi$ provides an optimal contract for this case. □

Insights from the Optimal Contract

Several insights can be gained from the optimal contract in Theorem 7. The most interesting insights are related to the conditions that result in a contract where the technology firm makes a high or low investment in cybersecurity:

When $\psi > (\gamma - \alpha) \cdot L$ or $\psi > (\gamma - \alpha) \cdot V/\gamma$, the optimal contract leads to the technology firm making a low investment in cybersecurity. If $\psi > (\gamma - \alpha) \cdot V/\gamma$, then a high investment ψ by the technology firm in cybersecurity is relatively costly compared to the financial value V

to the *technology firm* of the medical data, and the technology firm will not want to make a high investment in cybersecurity. If $\psi > (\gamma - \alpha) \cdot L$, then this means that a high investment ψ by the technology firm in cybersecurity is relatively costly compared to the liability L of the *health care provider* in the event of a data breach. It is surprising that the optimal contract when $\psi > (\gamma - \alpha) \cdot L$ holds also leads to the technology firm making a low investment in cybersecurity.

Another interesting aspect of the optimal contract when $\psi > (\gamma - \alpha) \cdot L$ or $\psi > (\gamma - \alpha) \cdot V/\gamma$ is that the optimal contract is such that there is no penalty or fine $t^* = 0$ to the technology firm in the event of a data breach. In this case, it is instead optimal to charge as much as possible for the data, meaning it is optimal to charge $\phi^* = V$.

On the other hand, the optimal contract induces the technology firm to invest in cybersecurity only when $\psi \leq (\gamma - \alpha) \cdot L$ and $\psi \leq (\gamma - \alpha) \cdot V/\gamma$. A numerical example of these thresholds are shown in Fig. 4.1. Here, a high investment ψ by the technology firm in cybersecurity is relatively cheap compared to the financial value V to the *technology firm* of the medical data and relatively small compared to the liability L of the *health care provider* in the event of a data breach.

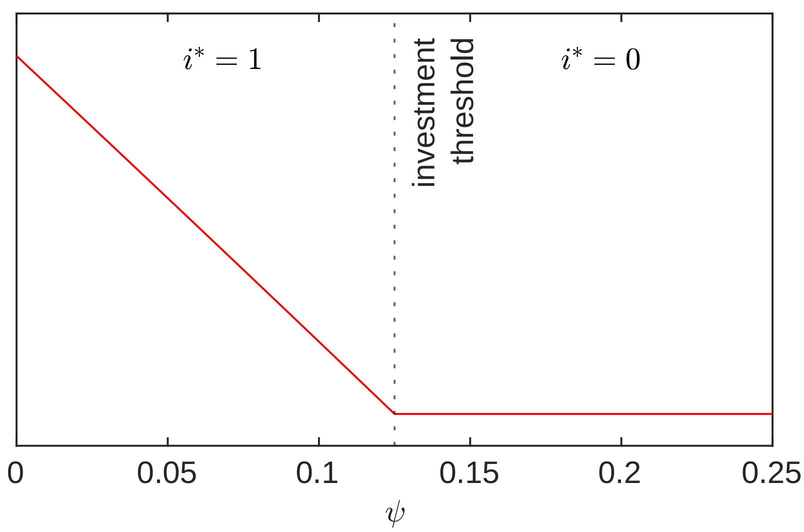

Figure 4.1: This plot shows the expected utility from an optimal contract as a function of ψ. We can observe that there exists a threshold beyond which the cost of investment is relatively high compared to either the financial value of the medical data or the liability in the event of a data breach, and thus it is optimal to induce $i^* = 0$ if ψ is above this threshold value.

4.3 Scenario B: Consortium of Health Care Providers Sharing Data

The second scenario we study is that of a group of health care providers sharing medical data amongst themselves. We assume that the cybersecurity level of each health care provider is fixed, and that all health care providers have identical models. Here, the primary consideration is each health care provider's decision of whether or not to join the consortium. If any single health care provider in the consortium suffers from a data breach, then all the health care providers in the consortium will face liability from those patients whose medical data has been breached. Thus, the health care providers must decide whether any benefits accrued from being in the consortium outweigh the increased risks of data breach due to sharing medical data.

In this scenario that we consider, each health care provider has two options available to mitigate the liability risks associated with a data breach. The first is that consortium can impose a fine or penalty on the health care provider that suffers from a data breach, which is equally shared by the remaining health care providers. The second is that each health care provider is able to purchase cybersecurity insurance from an external insurance agency.

Health Care Provider Model

The financial value to a health care provider of their own medical data is W, and the financial value to a health care provider of the medical data from a consortium of k health care providers is $v(k) \cdot W$, where the function $v(\cdot) > 0$ is strictly increasing and concave with $v(1) = 1$. This financial value is derived from the provider's ability to use the data to self-improve its health care services through improved patient treatment and care delivery processes, as well as through medical research. This model says that more quantity of data gives more financial value, but that there are diminishing financial returns to increasing quantities of data.

If any consortium member suffers a data breach, then each health care provider in the consortium has to spend L to address its various liabilities to the affected patients. The probability of a data breach among k health care providers is given by $p(k) \in (0, 1)$. We assume that this function $p(\cdot)$ is concave and increases with a sublinear growth rate such that $p(k) < k \cdot p(1)$. Furthermore, the health care provider responsible for the data breach is required to pay a fine or penalty of $(k-1) \cdot t$ that is equally divided among the $(k-1)$ remaining health care providers, where we assume $t \leq L$. Having $t > L$ is unrealistic because it would mean a healthcare provider profits from a data breach elsewhere.

Each health care provider can choose to purchase a policy to insure against their liabilities in the event of a breach. Under the assumption of an actuarially fair policy [76], each health care provider can purchase an insurance policy that pays out L_c under the event of a data breach at the cost of pL_c, where p is the probability of a data breach.

Finally, we assume each health care provider is *risk averse*. This means that if a health care provider earns a financial revenue of x, then their utility for that revenue is $U(x)$ for a function $U(\cdot)$ that is strictly increasing and concave. Under the additional assumption that $U(\cdot)$ is differentiable, this assumption of risk aversion means that $U'(\cdot) > 0$ and $U''(\cdot) < 0$.

Contract Design Problem

In this scenario, the consortium faces a contract design problem in which their goal is to pick the value of the fine or penalty t and the insurance policy payout L_c so as to encourage participation in the consortium and thus motivate data sharing. This contract design problem can be written as the following:

$$\begin{aligned} \max_{s,t,L_c}\ & (1-s)\cdot H^1(L_c) + s\cdot H^k(t,L_c) \\ \text{s.t. } & s\in\{0,1\},\ t\in[0,L],\ L_c\geq 0 \end{aligned} \tag{4.14}$$

where we note that a health care provider's expected utility when they do not participate in the consortium is given by

$$H^1(L_c) = p(1)\cdot U(W - p(1)\cdot L_c - L + L_c) + (1-p(1))\cdot U(W - p(1)\cdot L_c), \tag{4.15}$$

and a health care provider's expected utility when they do participate in the consortium is given by

$$\begin{aligned} H^k(t,L_c) = p(1)\cdot U(v(k)\cdot W - p(k)\cdot L_c - L - (k-1)\cdot t + L_c) \\ + (p(k)-p(1))\cdot U(v(k)\cdot W - p(k)\cdot L_c - L + t + L_c) \\ + (1-p(k))\cdot U(v(k)\cdot W - p(k)\cdot L_c). \end{aligned} \tag{4.16}$$

Optimal Contract

We next proceed to solve the contract design problem (4.14) through a series of steps. Let s^*, t^*, L_c^* denote the optimal contract, meaning they maximize (4.14). We first characterize the optimal fine or penalty amount.

Proposition 12. *We have that $t^* = 0$.*

Proof. If $s^* = 0$, then the objective function of (4.14) is $H^1(L_c)$. This means the objective function value does not depend on t, and so any feasible t is optimal. Hence, we can pick $t^* = 0$ in this case. If $s^* = 1$, then the objective function of (4.14) is $H^k(t, L_c)$. Now, consider the partial derivative with respect to t:

$$\begin{aligned} \partial_t H^k(t,L_c) = -(k-1)\cdot p(1)\cdot U'(v(k)\cdot W - p(k)\cdot L_c - L - (k-1)\cdot t + L_c) \\ + (p(k)-p(1))\cdot U'(v(k)\cdot W - p(k)\cdot L_c - L + t + L_c). \end{aligned} \tag{4.17}$$

Since we assumed that $U''(\cdot) < 0$, this means $U'(v(k)\cdot W - p(k)\cdot L_c - L - (k-1)\cdot t + L_c) >$ $U'(v(k)\cdot W - p(k)\cdot L_c - L + t + L_c)$ since $v(k)\cdot W - p(k)\cdot L_c - L - (k-1)\cdot t + L_c <$

$v(k) \cdot W - p(k) \cdot L_c - L + t + L_c$. Recalling that $p(k) > p(1)$ and $U'(\cdot) > 0$ by assumption, we thus have

$$\begin{aligned}\partial_t H^k(t, L_c) \leq \big[- (k-1) \cdot p(1) + (p(k) - p(1)) \big] \\ \times U'(v(k) \cdot W - p(k) \cdot L_c - L - (k-1) \cdot t + L_c). \quad (4.18)\end{aligned}$$

Since we assumed that $p(k) < k \cdot p(1)$ and $U'(\cdot) > 0$, this means $\partial_t H^k(t, L_c) < 0$. Thus choosing $t^* = 0$ is optimal because we are constrained in (4.14) to choose $t \in [0, 1]$. □

The implication of the above result is that we can rewrite the contract design problem as

$$\begin{aligned}\max_{s, L_c} \ & (1-s) \cdot H^1(L_c) + s \cdot H^k(L_c) \\ \text{s.t. } & s \in \{0, 1\}, \ L_c \geq 0\end{aligned} \quad (4.19)$$

where $H^1(\cdot)$ is as defined in (4.15), and

$$\begin{aligned}H^k(L_c) := H^k(0, L_c) = p(k) \cdot U(v(k) \cdot W - p(k) \cdot L_c - L + L_c) \\ + (1 - p(k)) \cdot U(v(k) \cdot W - p(k) \cdot L_c). \quad (4.20)\end{aligned}$$

We next use the above reformulation to characterize the optimal insurance coverage payout.

Proposition 13. *We have that $L_c^* = L$.*

Proof. If $s^* = 0$, then the objective function of (4.19) is $H^1(L_c)$. Next, note the first-order stationarity condition is

$$\begin{aligned}0 = \partial_{L_c} H^1(L_c) = p(1) \cdot (1 - p(1)) \cdot U'(W - p(1) \cdot L_c - L + L_c) \\ - p(1) \cdot (1 - p(1)) \cdot U'(W - p(1) \cdot L_c). \quad (4.21)\end{aligned}$$

Since we assumed that $U''(\cdot) < 0$, this means the above is satisfied when $W - p(1) \cdot L_c - L + L_c = W - p(1) \cdot L_c$. Hence at optimality we have $L_c^* = L$, which is feasible since $L > 0$ implies $L_c^* \geq 0$. The proof for the case $s^* = 1$ proceeds almost identically. □

The implication of the above result is that we can rewrite the contract design problem as

$$\max_{s \in \{0,1\}} (1-s) \cdot H^1 + s \cdot H^k, \quad (4.22)$$

where

$$H^1 := H^1(L) = U(W - p(1) \cdot L), \quad (4.23)$$
$$H^k := H^k(L) = U(v(k) \cdot W - p(k) \cdot L). \quad (4.24)$$

We conclude by using the above characterization to finish deriving an optimal contract for this scenario.

Theorem 8. *If $W - p(1) \cdot L > v(k) \cdot W - p(k) \cdot L$, then an optimal contract is given by $(s^*, t^*, L_c^*) = (0, 0, L)$. If $W - p(1) \cdot L \leq v(k) \cdot W - p(k) \cdot L$, then an optimal contract is given by $(s^*, t^*, L_c^*) = (1, 0, L)$.*

Proof. If $W - p(1) \cdot L > v(k) \cdot W - p(k) \cdot L$, then this means $H^1 = U(W - p(1) \cdot L) > H^k = U(v(k) \cdot W - p(k) \cdot L)$ since $U'(\cdot) > 0$. Hence, the optimal choice is $s^* = 0$. This means the choice $t^* = 0$ and $L_c^* = L$ is optimal by Propositions 12 and 13. If $W - p(1) \cdot L \leq v(k) \cdot W - p(k) \cdot L$, then this means $H^1 = U(W - p(1) \cdot L) \leq H^k = U(v(k) \cdot W - p(k) \cdot L)$ since $U'(\cdot) > 0$. Hence, the optimal choice is $s^* = 1$. This means the choice $t^* = 0$ and $L_c^* = L$ is optimal by Propositions 12 and 13. □

Insights from the Optimal Contract

Several insights can be gained from the optimal contract in Theorem 8. One interesting insight is that the optimal contract has $t^* = 0$, meaning that there is no penalty or fine in the event of a data breach, even when the health care providers participate in the consortium. This has an important practical implication, which is that participation in the data sharing consortium can only be encouraged by the ability of a health care provider to purchase an insurance policy from an external insurance company. Specifically, the optimal contract has $L_c^* = L$. This means $H^k = H^k(L) > H^k(0)$, or in words that purchasing insurance gives each participating health care provider a strictly higher utility than *not* purchasing insurance. Restated, the ability to purchase insurance for the event of a data breach makes it more likely for a health care provider to be willing to share data. Furthermore, Fig. 4.2 shows that depending on the particular functional forms, there is often a maximum consortium size beyond which costs associated with the increased likelihood of data breaches exceeds the value of sharing more data.

4.4 Conclusion

In this chapter, we designed contracts that help to mitigate the risks associated with data sharing so as to encourage health care providers to share their medical data. We first studied a scenario where a single health care provider sells medical data to a technology firm that is interested in using the data to develop new artificial intelligence algorithms. We next studied a scenario where multiple health care providers share data with each other for the purpose of conducting scientific research and improving patient care. Both cases required managing a trade-off between the value of sharing data with the liabilities associated with data breaches. The key concepts towards managing the risks associated with data breaches were the ideas of imposing a fine or penalty and purchasing external insurance to mitigate liabilities in the event of a data breach. Our results suggest that it is possible to devise contracts that promote the sharing of medical data while preserving the integrity and privacy of the data. By implementing the correct incentives, it may be possible to overcome the barriers to data sharing and facilitate the use of health information for science, technology, and policy.

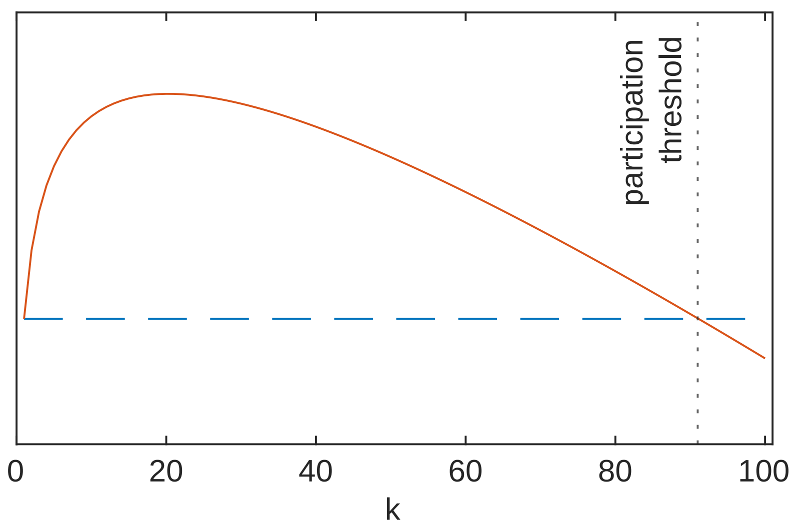

Figure 4.2: The red solid line is a health care provider's expected utility when they participate in the consortium, and the blue dashed line is the expected utility when they do not participate, where k is the number of health care providers in the consortium. The black dotted line represents the participation threshold beyond which the risks outweigh the benefits of sharing data by participating in the consortium.

Chapter 5

Conclusion

This dissertation has combined three lines of work that investigate societal problems related to safe water access, secure data sharing, and fair mechanism design. These projects employed mathematical modeling to propose solutions that advocate for universal and equitable access to information and resources while taking into account issues of equity and fairness, especially for underserved populations.

In Chapter 2, we focused on safe water access, a critical public health crisis in many developing countries. We proposed a model that helps communities determine the optimal quantity of water to store in local water containers. Our model considered the financial costs of purchasing water, the potential degradation of water quality, and the possibility of water wastage. Our goal was to empower communities to make informed decisions that promote equitable access to safe drinking water, especially for those who face difficulties obtaining even the minimal amount of water necessary for survival. A potential future direction would be to explore improving visualization for better interpretability and relaxing the set of conditions that ensure the threshold structure of the optimal policy.

In Chapter 3, we delved into the shortcomings of current approaches to incentive design with respect to fairness considerations. We introduced quantitative fairness criteria that integrated principles of equity and justice into the mechanism design framework. Our model took into account different forms of inequalities to ensure that the incentive system did not adversely affect individuals from specific socioeconomic backgrounds. Our objective was to establish several quantitative definitions of fairness that comprehensively captured the full spectrum of qualitative attributes, thereby encompassing marginalized segments of society. A possible future direction would be to generalize the results to more than two types of efficiency (adverse selection) and levels of effort (moral hazard), and to consider risk-averse agents instead of augmenting risk-neutrality with limited liability.

Chapter 4 explored the challenges of medical data sharing, which has the potential to advance healthcare but also exposes patients and healthcare providers to various risks. We proposed an optimal cybersecurity insurance contract that incentivized healthcare providers to responsibly share medical data while effectively mitigating associated risks. Our goal was to encourage the sharing of medical data for furthering scientific research, all while ensuring the protection of privacy and livelihood for all stakeholders, including the most vulnerable populations. A possible future direction would be to relax the actuarily fair assumption that allows us to simplify the insurance policy structure, and also to consider multiple levels of cybersecurity as a function of investment.

Overall, this dissertation aimed to contribute to the development of innovative solutions that facilitated equal and inclusive access to resources while mitigating risks and ensuring fairness. This work intended to support disadvantaged and underrepresented populations by addressing cultural and societal problems pertinent to resource allocation and information sharing.

Bibliography

[1] Aimee R. Ambrose. "Improving energy efficiency in private rented housing: Why don't landlords act?" In: *Indoor and Built Environment* 24.7 (2015), pp. 913–924.

[2] American Hospital Association. *Sharing Data, Saving Lives: The Hospital Agenda for Interoperability*. Tech. rep. 2019.

[3] Kenneth J. Arrow, Theodore Harris, and Jacob Marschak. "Optimal Inventory Policy". In: *Econometrica* 19.3 (1951), pp. 250–272.

[4] Dimitri P. Bertsekas. *Dynamic programming and optimal control*. Vol. 1. Athena Scientific, 1995.

[5] Joseph Blackburn and Gary Scudder. "Supply chain strategies for perishable products: the case of fresh produce". In: *Production and Operations Management* 18.2 (2009), pp. 129–37.

[6] Rainer Böhme. "Cyber-Insurance Revisited". In: *WEIS*. 2005.

[7] Rainer Böhme and Gaurav Kataria. "Models and Measures for Correlation in Cyber-Insurance". In: *WEIS*. 2006.

[8] Rainer Böhme and Galina Schwartz. "Modeling Cyber-Insurance: Towards a Unifying Framework". In: *WEIS*. 2010.

[9] Jean C. Bolot and Marc Lelarge. "A New Perspective on Internet Security using Insurance". In: *IEEE INFOCOM 2008*. 2008, pp. 1948–1956.

[10] Patrick Bolton and Mathias Dewatripont. *Contract theory*. MIT press, 2004.

[11] Zdravko I. Botev. "The normal law under linear restrictions: simulation and estimation via minimax tilting". In: *Journal of the Royal Statistical Society: Series B (Statistical Methodology)* 79.1 (2017), pp. 125–148.

[12] Marta Bottero, Chiara D'Alpaos, and Federico Dell'Anna. "Boosting Investments in Buildings Energy Retrofit: The Role of Incentives". In: *New Metropolitan Perspectives*. Springer, 2019, pp. 593–600.

[13] Felipe Caro and Jérémie Gallien. "Dynamic Assortment with Demand Learning for Seasonal Consumer Goods". In: *Management Science* 53.2 (2007), pp. 276–292.

[14] Felipe Caro and Jérémie Gallien. "Inventory Management of a Fast-Fashion Retail Network". In: *Operations Research* 58.2 (2010), pp. 257–273.

[15] Felipe Caro and Victor Martínez-de-Albéniz. "Fast Fashion: Business Model Overview and Research Opportunities". In: *Retail Supply Chain Management*. Ed. by Narendra Agrawal and Stephen A. Smith. International Series in Operations Research & Management Science. Springer, 2015, pp. 237–264.

[16] Leon Yang Chu, Vernon Ning Hsu, and Zuo-Jun Max Shen. "An economic lot-sizing problem with perishable inventory and economies of scale costs: Approximation solutions and worst case analysis". In: *Naval Research Logistics* 52.6 (2005), pp. 536–548.

[17] Leandro C Coelho and Gilbert Laporte. "Optimal joint replenishment, delivery and inventory management policies for perishable products". In: *Computers & Operations Research* 47 (2014), pp. 42–52.

[18] S. T. Coelho et al. "Controlling water quality in intermittent supply systems". In: *Water Supply* 3.1-2 (2003), pp. 119–125.

[19] Anupam Datta, Divya Sharma, and Arunesh Sinha. "Provable de-anonymization of large datasets with sparse dimensions". In: *International Conference on Principles of Security and Trust*. Springer. 2012, pp. 229–248.

[20] Cynthia Dwork. "Differential Privacy". In: *Automata, Languages and Programming*. Ed. by Michele Bugliesi et al. Berlin, Heidelberg: Springer, 2006, pp. 1–12. ISBN: 978-3-540-35908-1.

[21] European Union. *Regulation (EU) 2016/679 of the European Parliament and of the Council of 27 April 2016 on the protection of natural persons with regard to the processing of personal data and on the free movement of such data, and repealing Directive 95/46/EC (General Data Protection Regulation)*. (Accessed: Mar. 31, 2022). URL: `https://eur-lex.europa.eu/eli/reg/2016/679/oj`.

[22] Poorya Farahani, Martin Grunow, and H-O Günther. "Integrated production and distribution planning for perishable food products". In: *Flexible Services and Manufacturing Journal* 24 (2012), pp. 28–51.

[23] Daniela Pucci de Farias and Benjamin Van Roy. "The linear programming approach to approximate dynamic programming". In: *Operations Research* 51.6 (2003), pp. 850–865.

[24] David H. Fremlin. *Measure Theory*. Vol. 2. Torres Fremlin, 2001.

[25] Ashok Gadgil. "Drinking water in developing countries". In: *Annual Review of Energy and the Environment* 23.1 (1998), pp. 253–286.

[26] Sanford Grossman and Oliver Hart. "An Analysis of the Principal-Agent Problem". In: *Econometrica* 51.1 (1983), pp. 7–45.

[27] George Hadley and Thomson M. Whitin. *Analysis of inventory systems*. Prentice-Hall, Inc., 1963.

[28] Oliver Hart and Bengt Holmström. "The theory of contracts". In: *Advances in economic theory: Fifth world congress*. Vol. 71. Cambridge. 1987, p. 155.

[29] William B. Haskell, Rahul Jain, and Dileep Kalathil. "Empirical dynamic programming". In: *Mathematics of Operations Research* 41.2 (2016), pp. 402–429.

[30] Bengt Holmström. "Moral hazard and observability". In: *The Bell journal of economics* (1979), pp. 74–91.

[31] Vernon Ning Hsu. "Dynamic economic lot size model with perishable inventory". In: *Management Science* 46.8 (2000), pp. 1159–1169.

[32] Vernon Ning Hsu and Timothy J. Lowe. "Dynamic economic lot size models with period-pair-dependent backorder and inventory costs". In: *Operations Research* 49.2 (2001), pp. 316–321.

[33] Tim Hulsen. "Sharing Is Caring—Data Sharing Initiatives in Healthcare". In: *International Journal of Environmental Research and Public Health* 17.9 (2020). ISSN: 1660-4601.

[34] John Jiang and Ge Bai. "Types of Information Compromised in Breaches of Protected Health Information". In: *Annals of Internal Medicine* 172.2 (2020), pp. 159–160.

[35] Jaewon Kang. "How Trader Joe's Beat Out Conventional Supermarkets by Offering Less". In: *Wall Street Journal* (2023).

[36] Itir Z. Karaesmen, Alan Scheller-Wolf, and Borga Deniz. "Managing Perishable and Aging Inventories: Review and Future Research Directions". In: *Planning Production and Inventories in the Extended Enterprise: A State of the Art Handbook*. Ed. by Karl G. Kempf, Pınar Keskinocak, and Reha Uzsoy. Springer, 2011, pp. 393–436.

[37] Nikolaos Kariotoglou et al. "Approximate dynamic programming for stochastic reachability". In: *2013 European Control Conference (ECC)*. 2013, pp. 584–589.

[38] Philipp Klingel. "Technical causes and impacts of intermittent water distribution". In: *Water Supply* 12.4 (2012), pp. 504–512.

[39] Emily Kumpel and Kara L. Nelson. "Intermittent Water Supply: Prevalence, Practice, and Microbial Water Quality". In: *Environmental Science & Technology* 50.2 (2016), pp. 542–553.

[40] Howard Kunreuther and Geoffrey Heal. "Interdependent security". In: *Journal of Risk and Uncertainty* 26.2 (2003), pp. 231–249.

[41] Jean-Jacques Laffont and David Martimort. *The Theory of Incentives: The Principal-Agent Model*. Princeton University Press, 2002.

[42] Michaela Lang et al. "Systematic review: Landlords' willingness to retrofit energy efficiency improvements". In: *Journal of cleaner production* 303 (2021), p. 127041.

[43] Choong Ho Lee and Hyung-Jin Yoon. "Medical big data: promise and challenges". In: *Kidney Research and Clinical Practice* 36.1 (2017), pp. 3–11.

[44] Ellen J. Lee and Kellogg J. Schwab. “Deficiencies in drinking water distribution systems in developing countries”. In: *Journal of Water and Health* 3.2 (2005), pp. 109–127.

[45] Yoon Lee and Anil Aswani. “Optimally Designing Cybersecurity Insurance Contracts to Encourage the Sharing of Medical Data”. In: *2022 IEEE 61st Conference on Decision and Control (CDC)*. 2022, pp. 6782–6787. DOI: 10.1109/CDC51059.2022.9992544.

[46] Yoon Lee et al. *Optimal Policy for Inventory Management with Periodic and Controlled Resets*. 2022. arXiv: 2209.03571 [math.OC].

[47] Ninghui Li, Tiancheng Li, and Suresh Venkatasubramanian. “t-Closeness: Privacy Beyond k-Anonymity and l-Diversity”. In: *2007 IEEE 23rd International Conference on Data Engineering*. 2007, pp. 106–115.

[48] Xin Liang et al. “Making incentive policies more effective: An agent-based model for energy-efficiency retrofit in China”. In: *Energy Policy* 126 (2019), pp. 177–189.

[49] Ashwin Machanavajjhala et al. “L-diversity: privacy beyond k-anonymity”. In: *2006 IEEE 22nd International Conference on Data Engineering*. 2006, pp. 24–24.

[50] Raina M. Maier, Ian L. Pepper, and Charles P. Gerba. *Environmental Microbiology*. Elsevier, 2009.

[51] Angelica Marotta et al. “Cyber-insurance survey”. In: *Computer Science Review* 24 (2017), pp. 35–61.

[52] Hirofumi Matsuo and Susumu Ogawa. “Innovating innovation: The case of seven-eleven Japan.” In: *International Commerce Review* 7.2 (2007).

[53] Nathaniel Meyersohn. “Why Costco and Trader Joe’s stopped selling your favorite food”. In: *CNN Business* (2022).

[54] Yonatan Mintz, Zuo-Jun Max Shen, and Anil Aswani. “Local Water Storage Control for the Developing World”. In: *2017 IEEE 56th Annual Conference on Decision and Control (CDC)*. 2017, pp. 5074–5079.

[55] Yonatan Mintz et al. “Behavioral analytics for myopic agents”. In: *arXiv preprint arXiv:1702.05496* (2017).

[56] Liangyuan Na et al. “Feasibility of Reidentifying Individuals in Large National Physical Activity Data Sets From Which Protected Health Information Has Been Removed With Use of Machine Learning”. In: *JAMA Network Open* 1.8 (2018). ISSN: 2574-3805.

[57] Anna Nagurney et al. *Networks against time: Supply chain analytics for perishable products*. Springer, 2013.

[58] Steven Nahmias. “Perishable Inventory Theory: A Review”. In: *Operations Research* 30.4 (1982), pp. 680–708.

[59] Arvind Narayanan and Vitaly Shmatikov. “Robust De-anonymization of Large Sparse Datasets”. In: *2008 IEEE Symposium on Security and Privacy*. 2008, pp. 111–125.

[60] Office of Public Health Scientific Services. *Public Health Surveillance: Preparing for the Future*. Tech. rep. Centers for Disease Control and Prevention, 2018.

[61] Hulisi Ogut, Nirup M Menon, and Srinivasan Raghunathan. "Cyber Insurance and IT Security Investment: Impact of Interdependence Risk". In: *WEIS*. 2005.

[62] Paul Ohm. "Broken Promises of Privacy: Responding to the Surprising Failure of Anonymization". In: *UCLA Law Review* 57 (2010), pp. 1701–1777.

[63] Esther Oxford. "Hundreds of patient data breaches are left unpunished". In: *BMJ* 377 (2022).

[64] William P Pierskalla. "Supply chain management of blood banks". In: *Operations research and health care*. Springer, 2005, pp. 103–145.

[65] William P. Pierskalla and Chris D. Roach. "Optimal issuing policies for perishable inventory". In: *Management Science* 18.11 (1972), pp. 603–614.

[66] Elizabeth Pisani and Carla AbouZahr. "Sharing health data: good intentions are not enough". In: *Bulletin of the World Health Organization* 88.6 (2010), pp. 462–466.

[67] Warren B. Powell. *Approximate Dynamic Programming: Solving the curses of dimensionality*. John Wiley & Sons, 2007.

[68] Tharindu Prabatha, Kasun Hewage, and Rehan Sadiq. "An Incentives Planning Framework for Residential Energy Retrofits: A Life Cycle Thinking-Based Analysis under Uncertainty". In: *Sustainability* 15.6 (2023), p. 5479.

[69] Gregory P. Prastacos. "Blood inventory management: an overview of theory and practice". In: *Management Science* 30.7 (1984), pp. 777–800.

[70] Donald B. Rosenfield. "Disposal of excess inventory". In: *Operations Research* 37.3 (1989), pp. 404–409.

[71] Donald B. Rosenfield. "Optimality of myopic policies in disposing excess inventory". In: *Operations Research* 40.4 (1992), pp. 800–803.

[72] Joachim Roski, George W. Bo-Linn, and Timothy A. Andrews. "Creating Value In Health Care Through Big Data: Opportunities And Policy Implications". In: *Health Affairs* 33.7 (2014), pp. 1115–1122.

[73] Ilya O. Ryzhov, Warren B. Powell, and Peter I. Frazier. "The knowledge gradient algorithm for a general class of online learning problems". In: *Operations Research* 60.1 (2012), pp. 180–195.

[74] Celso Santos. "Energy-efficient technology retrofit investment behaviors of households in lower and higher income regions". PhD thesis. Iowa State University, 2020.

[75] Herbert Scarf. *The optimality of (S, s) policies in the dynamic inventory problem*. Tech. rep. Stanford University, 1959.

[76] Galina A. Schwartz and S. Shankar Sastry. "Cyber-Insurance Framework for Large Scale Interdependent Networks". In: *HiCoNS*. 2014, pp. 145–154. ISBN: 9781450326520.

[77] Azlan Shah Ali, Ismail Rahmat, and Hasnanywati Hassan. "Involvement of key design participants in refurbishment design process". In: *Facilities* 26.9/10 (2008), pp. 389–400.

[78] Zuo-Jun Max Shen, M. Dessouky, and F. Ordóñez. "Perishable inventory management system with a minimum volume constraint". In: *Journal of the Operational Research Society* 62.12 (2011), pp. 2063–2082.

[79] Nikhil Shetty et al. "Competitive Cyber-Insurance and Internet Security". In: *Economics of Information Security and Privacy*. Ed. by Tyler Moore, David Pym, and Christos Ioannidis. Boston, MA: Springer, 2010, pp. 229–247. ISBN: 978-1-4419-6967-5.

[80] Lawrence V. Snyder and Zuo-Jun Max Shen. *Fundamentals of Supply Chain Theory*. Wiley, 2019.

[81] Immanuel Stieß and Elisa Dunkelberg. "Objectives, barriers and occasions for energy efficient refurbishment by private homeowners". In: *Journal of Cleaner Production* 48 (2013), pp. 250–259.

[82] Latanya Sweeney. "K-Anonymity: A Model for Protecting Privacy". In: *International Journal of Uncertainty, Fuzziness and Knowledge-Based Systems* 10.5 (2002), pp. 557–570. ISSN: 0218-4885.

[83] Latanya Sweeney. "Weaving Technology and Policy Together to Maintain Confidentiality". In: *The Journal of Law, Medicine & Ethics* 25.2-3 (1997), pp. 98–110.

[84] S. Tokajian and F. Hashwa. "Water quality problems associated with intermittent water supply". In: *Water Science & Technology* 47.3 (2003), pp. 229–234.

[85] U.S. Department of Energy. *Quadrennial Technology Review: An Assessment of Energy Technologies and Research Opportunities*. Tech. rep. 2015.

[86] U.S. Department of Health and Human Services. *The Health Insurance Portability and Accountability Act Privacy Rule*. (Accessed: Mar. 31, 2022). URL: `https://www.hhs.gov/hipaa/for-professionals/privacy/index.html`.

[87] UNEP. "Buildings and Climate Change: Status, Challenges and Opportunities". In: *Sustainable Construction and Building Initiative (SBCI)* (2007).

[88] United Nations. *The Right to Water*. Tech. rep. Office of the United Nations High Commissioner for Human Rights, 2015.

[89] Karel van Donselaar et al. "Inventory control of perishables in supermarkets". In: *International Journal of Production Economics* 104.2 (2006), pp. 462–472.

[90] Willem G. Van Panhuis et al. "A systematic review of barriers to data sharing in public health". English (US). In: *BMC Public Health* 14.1 (2014). ISSN: 1471-2458.

[91] WaterAid. *Water: At What Cost? The State of the World's Water*. Tech. rep. 2016.

[92] WHO and UNICEF. *Progress on drinking water, sanitation and hygiene: 2017 update and SDG baselines*. Tech. rep. WHO/UNICEF Joint Monitoring Programme for Water Supply, Sanitation and Hygiene, 2017.

[93] Nicola Willand et al. “Retrofit poverty: socioeconomic spatial disparities in retrofit subsidies uptake”. In: *Buildings and Cities* 1.1 (2020).

[94] World Bank. *The Karnataka Urban Water Sector Improvement Project: 24 x 7 Water Supply is Achievable*. Tech. rep. Water and Sanitation Program, 2010.

[95] Marie-Hélène Zérah. “How to assess the quality dimension of urban infrastructure: The case of water supply in Delhi”. In: *Cities* 15.4 (1998), pp. 285–290.

[96] Mo Zhou et al. “Evaluating machine learning–based automated personalized daily step goals delivered through a mobile phone app: Randomized controlled trial”. In: *JMIR mHealth and uHealth* 6.1 (2018), e9117.

[97] Robert A. E. Zidek and Ilya V. Kolmanovsky. “Stochastic drift counteraction optimal control and enhancing convergence of value iteration”. In: *2016 IEEE 55th Conference on Decision and Control (CDC)*. 2016, pp. 1119–1124.

Appendix A

Statement and Proof of Lemmas

Lemma 1. *Consider two functions $f(x)$ and $g(x)$ defined on a compact interval $[a, b]$. Suppose that $f(x)$ is differentiable, that its derivative is bounded from below by a positive constant $D_x f(x) \geq M > 0$, and that $g(x)$ is Lipschitz continuous with constant $L \geq 0$. Then for any fixed λ such that $0 \leq \lambda \leq M/L$, the function $f(x) + \lambda g(x)$ is non-decreasing and continuous.*

Proof. The continuity of $f(x) + \lambda g(x)$ is immediate by the differentiability of $f(x)$ and the Lipschitz continuity of $g(x)$. So we focus on showing $f(x)+\lambda g(x)$ is non-decreasing. Consider any $x_1, x_2 \in [a, b]$ with $x_2 \geq x_1$. Since $D_x f(x) \geq M$, we have $f(x_2) - f(x_1) \geq M(x_2 - x_1)$. Since $g(x)$ is Lipschitz continuous with constant L, we have $|g(x_2) - g(x_1)| \leq L(x_2 - x_1)$. Combining these two inequalities implies

$$\begin{aligned} f(x_2) - f(x_1) + \lambda(g(x_2) - g(x_1)) &\geq M(x_2 - x_1) - \lambda L(x_2 - x_1) \\ &\geq (M - \lambda L) \cdot (x_2 - x_1) \geq 0. \end{aligned} \tag{A.1}$$

where the last inequality follows because $x_2 \geq x_1$ and because $0 \leq \lambda \leq M/L$, which means $M - \lambda L \geq 0$. Since we have shown $f(x_2) - f(x_1) + \lambda(g(x_2) - g(x_1)) \geq 0$ for any $x_2 \geq x_1$, this implies $f(x) + \lambda g(x)$ is non-decreasing. $\square$

Lemma 2. *If $\int_{-\infty}^{\infty} |f(x)|dx$ is finite, $\int_{-\infty}^{\infty} |g(x)|dx$ is finite, $f(x) = 0$ for $x < 0$, $|f(x)| \leq M$, $|g(x)| \leq P$, and on $x \geq 0$, we have that $f(x)$ is Lipschitz with constant L, then the convolution $(f*g)(x) = \int_{-\infty}^{\infty} f(x-z)g(z)dz$ is Lipschitz with constant $(L\cdot\int_{-\infty}^{\infty} |g(z)|dz+MP)$.*

Proof. Without loss of generality, we assume that $x \le y$. Observe that

$$\begin{aligned}
|(f*g)(x)-(f*g)(y)| &= |\textstyle\int_{-\infty}^{\infty} f(x-z)g(z)dz - \int_{-\infty}^{\infty} f(y-z)g(z)dz| \\
&= |\textstyle\int_{-\infty}^{\infty}(f(x-z)-f(y-z))\cdot g(z)dz| \\
&= |\textstyle\int_{-\infty}^{x}(f(x-z)-f(y-z))\cdot g(z)dz - \int_{x}^{y} f(y-z)\cdot g(z)dz| \\
&\le |\textstyle\int_{-\infty}^{x}(f(x-z)-f(y-z))\cdot g(z)dz| + |\int_{x}^{y} f(y-z)\cdot g(z)dz| \\
&\le \textstyle\int_{-\infty}^{x} |f(x-z)-f(y-z)|\cdot|g(z)|dz + \int_{x}^{y} |f(y-z)|\cdot|g(z)|dz \\
&\le \textstyle\int_{-\infty}^{x} L|(x-z)-(y-z)|\cdot|g(z)|dz + \int_{x}^{y} MP\cdot dz \\
&\le \textstyle\int_{-\infty}^{x} L|x-y|\cdot|g(z)|dz + MP\cdot|x-y| \\
&\le (L\cdot \textstyle\int_{-\infty}^{\infty}|g(z)|dz + MP)\cdot|x-y|.
\end{aligned} \tag{A.2}$$

This shows the convolution is Lipschitz continuous. □

Lemma 3. *Consider a function $f(x)$ that is finite at the point y (i.e., $|f(y)|$ is bounded). If $f(x)$ is Lipschitz with constant L on a compact domain $[a,b]$ with $y \in [a,b]$, then it is finitely bounded $|f(x)| \le L(b-a)+|f(y)|$ for $x \in [a,b]$.*

Proof. For any $x \in [a,b]$, we have

$$\begin{aligned}
|f(x)| &= |f(x)-f(y)+f(y)| \\
&\le |f(x)-f(y)|+|f(y)| \\
&\le L|x-y|+|f(y)| \\
&\le L(b-a)+|f(y)|.
\end{aligned} \tag{A.3}$$

Since the interval $[a,b]$ is compact, this means a,b are finite. This gives the desired bound. □