Strategic Supply Chain Management with Multiple Products under Supply and Capacity Uncertainty

by

Süleyman Demirel

A dissertation submitted in partial fulfillment
of the requirements for the degree of
Doctor of Philosophy
(Business Administration)
in The University of Michigan
2012

Doctoral Committee:

Professor Izak Duenyas, Co-Chair
Professor Roman Kapuscinski, Co-Chair
Professor Yavuz A. Bozer
Professor Thomas J. Schriber
Assistant Professor Volodymyr O. Babich, Georgetown University

*"Everything about yesterday has gone with yesterday.
Today, it is needed to say new things."*

– Rumi –

For my lovely wife, Emel.

Acknowledgements

Dozens of people made this dissertation possible with their help and support, and I would like to thank them all. First, I am very fortunate to have had Roman Kapuscinski and Izak Duenyas as my advisors. Their insights, guidance, patience, and willingness to let me make mistakes on the path to discovering better answers enhanced the quality of this dissertation and were invaluable to my growth as a researcher. They were generous with their time and always available. Thank you, Roman and Izak, for your faith and confidence in me and for acting as colleagues rather than supervisors. I appreciate that you let me think independently and explore several ways to answer a question, while encouraging me to get things done without losing track. I would also like to thank the other members of my doctoral committee, Professors Yavuz Bozer, Thomas Schriber, and Volodymyr Babich. Your questions and feedback improved this dissertation and gave me new ideas to extend my research. Thank you as well to Ching-Hua Chen-Ritzo for collaborating with me, which led me to develop a portion of this dissertation. You were always eager to discuss research and provided many insightful comments during our collaboration.

The support of faculty members in the Operations and Management Science (OMS) department has also played an important role in my growth as a scholar. Doctoral seminars offered by various faculty members, including Roman Kapuscinski, Damian Beil, Hyun-Soo Ahn, Amitabh Sinha, Owen Wu, Bill Lovejoy, and Wally Hopp have been incredible sources of knowledge and ideas. I also benefited significantly from faculty members outside the OMS department, including Professors Demosthenis Teneketzis and Uday Rajan, whose classes in stochastic dynamic programming and game theory, respectively, were phenomenal.

My journey through the doctoral program would have been far less enjoyable had it not been for my fellow students in the OMS PhD program. Within this group, one person deserves special mention: Sripad Devalkar. Sripad was not merely an office-mate, but also a great friend inside the school and out. It was a pleasure to discuss ideas, papers, and conference presentations with him. Enjoying Indian food with Sripad and his wife, Sudarsana, during family visits will never be forgotten. I was also very fortunate to interact with Zhixi Wan, Man Yu, Maria Mayorga, Matt Potoff, Ravi Subramanian, Xinxin Hu, Ling Wang, Li Jiang, Wenjing Shen, and Shanshan Hu, senior PhD students, who were very friendly from the day I began my program until each graduated in turn. Whenever I needed advice, one of these scholars was always there to help. I am also thankful to Iva Iovtcheva and Eren Cetinkaya, who started the program the same year as I did. Both Iva and Eren have been great friends and colleagues. Other PhD students in the Business School have also been wonderful people with whom I have interacted. Among them are Celim Yildizhan, Deniz Anginer, Dan Gruber, Thales Teixeira, Sinem Atakan, Nilufer Aydinoglu, David Benson, David Zhu, Ali Tafti, Sun Hyun Park, John Chen, Vivek Tandon, Pranav Garg, Megan Witmer, Jim Mourey, Natalie Cotton, Robert Smith, and many others.

A very exciting period of my PhD occurred when I was given the chance to teach my own class as a graduate student instructor in the winter semester of 2008. Thanks to many sessions with Anne Harrington, director of Instructional Development, before I began teaching, a seemingly difficult job became joyful and exciting. I am also thankful for the excellent coordination I enjoyed with Professor Owen Wu, who guided me through the challenges of teaching the material. Acknowledging my teaching experience would be incomplete without mentioning (again!) Sripad Devalkar, who was teaching another section of the same course at the same time. We spent a significant amount of time together preparing for classes. At times, when I found a particular topic difficult to explain to students, Sripad offered insightful comments on how to convey the message in a simpler way. I would also like to thank the BBA Class of 2008, Sections 1 and 6 of OMS 311 for being wonderful students.

I would also like to acknowledge Kashif Rashid, my project supervisor during an internship at Schlumberger-Doll Research in 2008. My computer programming skills improved significantly as I followed his coding style during this experience. After working with Kashif, I have written computer codes more efficiently, potentially leading to better numerical results in this dissertation. Furthermore, I appreciate his ongoing friendship and willingness to explore research problems that are of common interest to us.

I am grateful to the Ross School of Business for providing me with funding during my PhD studies, without which this thesis would not have been possible. It is important for me to thank Kathy Sutcliffe, former associate dean for Faculty Development and Research, who covered my summer tuition in 2008, when I was required to register in order to complete my summer internship. I also thank Izak Duenyas, chair of the OMS department, for providing me with additional funding several times when it was most needed.

My lack of knowledge regarding administrative details and seldom being required to navigate the bureaucracy at an institution as large as the University of Michigan is thanks to the wonderful support provided by the staff in the Doctoral Studies Office at the Ross School. Thank you to Brian Jones, Roberta Perry, Kelsey Zill, Chris Gale, Martie Boron, and Linda Veltri.

I am particularly thankful as well to junior OMS PhD students Santhosh Suresh, Bam Amornpetchkul, Yao Cui, and Anyan Qi for helping me with my final presentation. Their suggestions, along with valuable input from my undergraduate classmate Ismail Civelek, significantly improved the presentation's clarity. I am also thankful to Leslie Hatch for copy-editing the dissertation within such a short time window and providing many suggestions to improve its clarity.

Beyond the Ross School, I am indebted to many other people. A small Turkish community in Ann Arbor has been a great source of friendship and support away from home. Nothing that I say here is enough to convey my gratitude to the many individuals who touched my life. I am grateful to Mucahit Bilici for picking me up from the airport when I arrived in the U.S. in 2004, welcoming me into his apartment during my first week, and helping me find a place to live and settle down. Korhan Bircan was a great first roommate, and I was extremely fortunate to have Melih Gunal as my second roommate. During the six months we were roommates, we had several discussions on science, religion, politics, philosophy, literature, music, and many other topics. I am also thankful to Burde, who made my ex-roommate the happiest person on earth by marrying him. Of course, I greatly appreciate Burde's friendship, as well.

Contrary to my initial expectations when I first arrived in the U.S., I have made many wonderful friends. Mustafa and Esra Sir, Burak and Selen Yilmaz, and Semi and Yasemin Ertan are caring individuals who offered support and friendship, not only when they were in Ann Arbor, but also when they moved after graduation. It is always fun to spend time with them whenever and wherever we can get together. Mustafa and Tuba Sanver, and Fatih and Nazik Hasoglu were also invaluable friends; indeed, they were almost like family members. Whenever we visited them in their homes in Kalamazoo, Michigan, they were welcoming and willing to generously offer their friendship. I will never forget their acts of kindness; in fact, I miss those good old days!

Kurtulus and Seyda Golcuk, Osman and Ayse Gundogdu, Adem and Berna Saglik, Husnu Kaplan, Suleyman and Ozgen Felek, Hasan and Yurdanur Korkaya, Ihsan Saracgil, Mehmet Elgin, Zahid and Muserref Samancioglu, and Hakan and Sonnur Erten were all great friends and neighbors, as well. Hasan Sert was not only a great friend, but also great teacher when he was patient enough to teach me how to drive. I have also been very fortunate to meet many friends during Turkish gatherings. I particularly enjoyed the friendship of Yavuzhan and Ebru Erdem and Emin and Asiye Kutay. Yavuzhan's stay in Ann Arbor was phenomenal, as I learned a great deal from his knowledge of and experience in Turkish classical music and his tanbur performances during all gatherings with no exception.

I cannot forget all the help I received prior to starting my PhD from people who encouraged me academically. My primary school teachers, Inci Soylu and Ergul Uzan, were first to encourage me to pursue higher goals. I enjoyed the mentorship of many friends and teachers during high school, including Ilyas Hasimov, Musa Sahin, Harun Celik, and Abdullah Demirel. In particular, I would like to thank Ulku Gurler, my undergraduate advisor, for encouraging me to pursue an academic career.

I am thankful to my parents, brothers, and sisters for the love, support, and encouragement they have given me my entire life. In particular, I thank Ali, my older brother, for providing me with financial assistance prior to my PhD to cover my initial expenses in the United States. I also thank Omer, my younger brother, for providing feedback on the frontispiece. Finally, I thank my mother-in-law, Necla Demirel, and sister-in-law, Aysegul Demirel, for visiting us and providing their help and support during my PhD studies.

Most of all, I am deeply indebted to my lovely wife, Emel, for her generous support during the entire span of my PhD studies. It is a big challenge to be the wife of a PhD student, but she never complained and always supported me with encouragement and in her prayers. We will both remember these days as a wonderful period of our life together, despite the difficulties faced. Most important of all, we stood by them together. Thank you, Emel! Without your help and support, this dissertation would not exist at all; for that very reason, I dedicate this work to you.

Table of Contents

List of Tables

List of Figures

Figure

Abstract

Effective supply chain management is essential to operational efficiency and business performance. When managing a supply chain, practitioners must carefully consider a range of issues, including inventory levels, capacity, and strategic partnerships. Although some elements of the supply chain can be planned, others cannot, including disruptions to the production process.

This dissertation examines various aspects of supply chain management taking into account the critical elements of real-world dynamics of both internal and external disruptions. Such elements include disruptions to internal resources producing multiple products on dedicated and shared resources, as well as disruptions to external suppliers of manufacturers.

In the first essay, we explore the effect of internal disruptions that manifest themselves through capacity uncertainty and study their effect on the production and inventory policies of a firm that produces multiple products in multiple stages. Focusing on a two-product case, we derive the optimal production and inventory policies and propose heuristic policies for the case of more than two products.

In the second essay, we explore the effects of external disruptions on a firm's supply chain strategy and its inventory policy. This essay carefully disentangles the active role suppliers play in defining the parameters of contracts and its effect on the firm's resulting sourcing practices. The main insight we provide is that with endogenously determined wholesale prices, the manufacturer does not necessarily benefit from flexible sourcing, whereas the suppliers may benefit.

Chapter 1

Introduction

Since the turn of millennium, managerial effort dedicated to understanding and protecting against disruptions has been unprecedented. However, most of the existing literature has omitted some of the critical elements of reality. This dissertation examines various aspects of supply chain management taking into account the critical elements of real-world dynamics of both internal and external disruptions. Such elements include disruptions to internal resources producing multiple products on dedicated and shared resources, as well as disruptions to external suppliers of manufacturers.

The existing literature has mainly focused on single-item production systems and often omitted the possibility of supply or capacity shortages, causing production disruptions. The primary motivation to analyze single-product models is that it is often possible to decompose an N-product problem into N independent single-product problems (Veinott 1966). For some real-world situations, however, applications of such models are limited. Instead, these cases serve as building blocks to analyze more complicated scenarios. Take the example of a firm producing multiple products using dedicated and shared resources, facing capacity uncertainties. In such circumstances, managing product stock is not independent across products; thus, analyzing an N-product problem is necessary.

Another element that the existing literature often omitted was decentralized decision making, even when only a single product is considered. Modern supply chains involve multiple decision makers making decisions in tandem that transform raw materials into deliverable products for customers in multiple stages. To optimize the flow of materials along a supply chain and increase economic efficiency, numerous supply contracts have been analyzed, particularly in the last two decades (Cachon, 2003). Such analyses constitute the foundations of supply chain management. Early models that addressed supply contracts, however, assumed only one supplier and one buyer, with the more powerful player determining the contract's terms. More recent studies have explored the effect of vertical competition in determining the contractual terms.

This dissertation studies various aspects of supply chain management, incorporating the elements described above. Essay 1 explores the effect of internal disruptions that manifest themselves through capacity uncertainty, and studies the effect of those internal disruptions on the production and inventory policies of a firm that produces multiple products in multiple stages. Essay 2 explores the effects of external disruptions on a firm's supply chain strategy and inventory policy, taking into account the critical elements of real-world dynamics, such as the active role suppliers play in defining the parameters of contracts. It studies the effect of contracts on the strategic behavior of suppliers and resulting sourcing practices. Table 1.1 summarizes the modeling elements that are incorporated in each of the two essays.

Table 1.1: Model Elements Included in the Essays

	Supply Disruption	**Multiple Products**	**Competitive Setting**
Essay 1 (Production Policies)	✓	✓	x
Essay 2 (Supply Contracts)	✓	x	✓

1.1. Essay 1: "Production and Inventory Control for a Make-to-Stock / Calibrate-to-Order System with Dedicated and Shared Resources"

The impact of random capacity shortages on a firm's performance is magnified when production occurs in multiple stages, with one stage requiring the output of a previous stage before it can start producing. Many internal disruptions can be interpreted as capacity uncertainty. With uncertain capacity, however, buffering and scheduling decisions often become non-obvious. Consider a company that manufactures and stocks various tools, which are later customized or calibrated based on specific orders. The firm must answer numerous questions: How much stock of the different tools should the firm produce and hold? How does the firm schedule the calibration of tools, given the uncertainties of production, calibration, and demand for the end products? To reflect the reality of the motivating situation, the tools are produced on dedicated lines, whereas calibration occurs using a shared resource. In Essay 1, we study this problem using a multiproduct inventory model for a two-stage production process with embedded capacity uncertainties at each stage.

We first analyze the case of two products. Due to capacity uncertainty of the shared resource, not only do the total quantities of planned production quantities matter, but the sequence in which the products are processed is also important. We thus characterize how the shared resource can be allocated and show that the optimal policy keeps the ending inventories as close to a so-called "target path" as possible. Due to capacity uncertainty on the dedicated lines, the firm's optimal production policy depends on the initial inventories. Notably, various structural properties of the optimal production policy can be described.

Through a numerical study, we characterize several interesting properties of the optimal production policy and demonstrate that using shared and uncertain capacity leads to counter-intuitive insights. For example, when products share limited capacity in the downstream stage, it might be optimal to stock more of the products with larger dedicated capacity or with less capacity variability. This is contrary to strategies for single-product inventory models, in which stocks are lower with larger capacity or with less variable capacity.

When a firm is making more than two products, a dynamic programming formulation is not computationally feasible. To overcome this shortcoming, we develop effective heuristic policies and subsequently test their performance. We find that the heuristic policies perform increasingly well as the number of products increases.

1.2. Essay 2: "Strategic Behavior of Suppliers in the Face of Production Disruptions"

With the globalization of operations and supply chains, firms are prone to disruption risks more than ever. Many articles in the business press, as well as our discussions with practitioners, indicate that adopting mitigation strategies to counter the negative impact of various disruptions is a top business priority. For example, Cisco and IBM are exploring and evaluating flexible sourcing through the use of backup suppliers. Others consider strategies of buffering through inventory. The operations management literature examines many of these strategies. It usually assumes, however, that the manufacturer is the sole decision maker and supplier behavior is exogenous. In such an environment, conventional wisdom suggests–and many models confirm–that a manufacturer cannot be worse off by having backup suppliers. In Essay 2, we consider suppliers as active decision makers that are aware of their strengths and weaknesses. Specifically, we develop a game-theoretical model for a multi-period sourcing setting with a manufacturer as the buyer, and two alternative suppliers who compete for manufacturer's business. In this essay's model, one of the suppliers is reliable, whereas the other is not, and the disruptions can last multiple periods. The suppliers compete by offering wholesale prices, and the manufacturer decides from which supplier (including possibly both) to buy and whether to protect itself from disruptions by carrying some inventory. The objective of Essay 2 is to evaluate the costs and benefits associated with flexible sourcing and to understand which strategy is appropriate given the manufacturer's situation. Although the initial model assumes that only one supplier is unreliable, we extend the results to a case in which both suppliers are unreliable.

To evaluate strategic supplier behavior, we consider two pricing games in which the suppliers compete. In the single-wholesale price game, each supplier announces a single (wholesale) price. In the contingent-pricing game, the reliable supplier offers wholesale prices contingent on whether it will serve as the primary or backup supplier. The unreliable supplier offers two wholesale prices: one for on-time deliveries, the other for late deliveries. A lower wholesale price for late deliveries can be viewed as equivalent to a penalty in the form of a supplier rebate or a charge-back, which are commonly observed in practice. With this model, we find that the single-wholesale price game leads to a conflict of incentives in terms of the roles suppliers want to play and the amount of business for which they contract, formally confirmed as non-existence of pure-strategy Nash equilibria in most practical situations. In contrast, the contingent-pricing game corresponds to a more intuitive relationship and has a unique pure-strategy Nash equilibrium. Except for cases that provide one of the suppliers with significant cost advantages, the manufacturer adopts flexible sourcing by using the (less expensive) unreliable supplier and the (more expensive) reliable supplier.

The economic benefits of this strategy, however, are less obvious. As noted, the conventional wisdom is that the manufacturer should never be worse off by having backup suppliers. With endogenously determined wholesale prices, however, the manufacturer does not necessarily benefit from the existence of a backup supplier; in fact, he is typically worse off. Thus, an up-front commitment to sole-sourcing and using simple wholesale price contracts may be beneficial rather than creating the opportunity for one supplier to serve as a backup through more flexible contracts. Interestingly, suppliers may benefit from flexible sourcing even though the manufacturer does not: the reliable supplier always benefits from maintaining backup capacity, whereas the unreliable supplier might benefit in some situations from the reliable supplier's backup capacity, despite reduced business volume. From a system perspective, a flexible sourcing strategy may degrade the supply chain's performance, thus requiring the need to coordinate among supply chain partners.

To extend our work, we analyze a capacity investment problem for a reliable supplier and show that the reliable supplier always prefers to offer no or full back-up capacity rather than partial availability. Furthermore, as a major extension and to complement the original model, we consider the case of two unreliable suppliers and identify the conditions under which one of the suppliers can be treated as perfectly reliable.

Chapter 2

Production and Inventory Control for a Make-to-Stock / Calibrate-to-Order System with Dedicated and Shared Resources

2.1. Introduction

Consider a firm that produces multiple products on dedicated production lines (stage 1). These products must be customized or calibrated (hereafter, "calibrated") after an order is received using a shared resource that is common to all products (stage 2). The firm stocks the products that are not yet calibrated and completes calibration after an order is received. The availability of products is crucial to the firm's performance, because it may translate into stock-outs and customers may not tolerate delays. A major challenge at both stages is that the dedicated production lines and the shared resource for calibration face capacity limitations. The capacity of calibration is finite and it is also uncertain, due to factors such as machine downtime, quality problems, and so on. Similarly, the capacity of the dedicated production lines is uncertain. To manage the production-inventory system efficiently requires that the producing and calibrating individual products is coordinated by simultaneously taking into account demand and capacity uncertainties.

The decision problem described above is common in some manufacturing and service firms. Consider the example of an oilfield services company that produces tools on dedicated production lines that are used to deliver services to oil companies. The tools are made to stock on dedicated production lines that the company maintains for tool production. Final testing, calibration, and quality inspections, however, are performed using a shared resource that calibrates each tool to satisfy specific customer requests based on the application for which the tool is intended. Tool availability, therefore, is crucial to the firm, as the firm must respond to customer requests immediately. The oilfield services company receives orders from oil companies operating rigs where each day of delay on a rig could cause million-dollar losses. It is common for an oil company to call with a special service request and provide a deadline by which the service has to be conducted on the rig. If the oilfield services provider could not provide service on time (due to delays caused by tool availability or problems with calibration, etc.), then the order would immediately be given to a competing firm. The oilfield services firm described could start producing the equipment needed for the service (stage 1 in our setting) prior to receiving the order. Final calibration and customization, however, would have to wait until the actual order was received. The non-standard nature of the tools means that the processing times for both production and calibration are inherently variable, which introduces capacity uncertainties at both stages, further complicating the production system. This specific setting motivated us to formulate this problem. To date, this issue has not been addressed in the literature.

We model the problem above as a multi-period production and inventory system. To gain insight into the optimal policy, we consider a case with two products. For the case of more than two products, we propose heuristic policies based on the results for two products. The decision maker's objective is to minimize the expected discounted inventory holding and penalty costs over a planning horizon by appropriately choosing production quantities and allocating the scarce shared resource to meet demands in every period.

For both stages, producing and calibrating the product, we model capacity uncertainties following an approach used by Ciarallo et al. (1994), Duenyas et al. (1997)

and Hu et al. (2008), in which maximum production quantity (capacity) is a random variable and its value is realized after production decisions are made. The realized capacity determines how much of the planned production can be accomplished. To the best of our knowledge, to date no published papers have dealt with capacity uncertainties when multiple products use the same resource (the calibration stage in our setting). Due to the common resource, not only the total planned production quantities are important but also the sequence in which the products are calibrated using the shared resource.

In this model, we assume that the dedicated production lines also face capacity uncertainties. We find that the optimal production quantities of stage-1 products depend on the initial inventories of both products, which differs from the commonly used independent base-stock policies. This is because the second stage of production is shared, leading to co-dependence of production quantities in the first stage. We analytically characterize the structural properties of the optimal production policy. We also conduct a numerical study to analyze how the limited capacity of the second stage influences production in the first stage. We also show how qualitative insights for finite shared capacity differ from the case of infinite shared capacity. The case of infinite shared capacity is a special case in which each product follows an independent base-stock policy. We find that the presence of finite shared capacity results in counter-intuitive qualitative insights. For example, when shared capacity is infinite, the inventory target for the product with stochastically larger dedicated capacity is lower than the other product (all other things being equal), whereas when the shared capacity is finite, this relationship may be reversed: it may be optimal to set higher inventory targets for the product with stochastically larger dedicated capacity.

Although the optimal policy can be characterized easily when the manufacturer makes only two products, the optimal policy structure is complicated for the case with more than two products. Therefore, we consider two implementable heuristics and evaluate their performance. The first heuristic follows base-stock policy for production and ignores the finite shared capacity. In the second heuristic, we take the finite shared capacity into account and observe that the heuristic performs significantly better. Using the properties of the optimal policy for two products, we extend our heuristics to multiple products and find that the performance of the heuristics improves as the number of products increases. Thus, our structural results are useful in developing heuristics that can be used in practice.

The remainder of the present paper is organized as follows. Section 2.2 reviews the related literature. Section 2.3 describes the model. Section 2.4 establishes the optimal policy structure for production and calibration. Section 2.5 presents the analytical results for the sensitivity of the optimal policy. Section 2.6 numerically examines the effect of the shared capacity on inventory levels and explores the ways various product characteristics, such as capacity, reliability of dedicated production lines, demand levels, and demand variability affect the optimal policy. Section 2.7 presents heuristic policies and explores heuristic performance under various settings. Finally, Section 2.8 concludes with a discussion of further research directions. Mathematical proofs appear in Appendix 2.9, and Appendix 2.10 provides an extension for our model that considers infinite-horizon and Markov-modulated settings.

2.2. Literature Review

Mathematical inventory theory has grown rapidly in the last century and has been extended to managing complex supply chains. Whereas the theory of inventory is considered mature at present, most of it focuses on single-item production systems; indeed, little is known about systems that produce multiple items. As Veinott (1966) points out, the primary motivation for analyzing single-product models is that it is often possible to decompose an N-product problem into N independent single-product problems. Such decomposition, however, is not possible if the products share common resources that are limited. In the present paper, we examine systems that produce multiple products that share limited capacity.

Evans (1967) first noted that independent management of inventories of multiple products may prove ineffective in the presence of a joint capacity constraint. Evans (1967) considers the classical newsvendor problem with lost sales in multiple periods, in which a production facility makes two products to stock and operates under limited capacity (henceforth referred to as Evans' model). He finds that the optimal production is a function of the existing inventories of both products and identifies switching curves that characterize the mode of operation. Thus, using an independent base-stock policy for each product is suboptimal. Nahmias and Schmidt (1984) considers the single-period version of Evans' model and examines the performance of several heuristics. A more recent paper by DeCroix and Arreola-Risa (1998) extends Evans' model to multiple products ($N \geq 2$) and provides an infinite-horizon framework and heuristic solutions. Bashyam et al. (1995) also considers Evans' model and uses perturbation analysis to compute near-optimal solutions. Shaoxiang (2004) allows for the general convex single-period cost function and finds that the structure of the optimal policy remains consistent both in the finite-horizon and infinite-horizon settings. De Vericourt et al. (2000) analyzes Evans' setting in continuous time, in which a single-machine dynamically allocates capacity to one of two queues for different products. They establish switching curves for optimal production, which are the continuous-time counterparts of the switching curves in Evans' original setting.

All the studies cited above consider production at a single facility. More often than not, however, production occurs in multiple stages, which may be operated by a single facility or by independent entities in a supply chain. Clark and Scarf (1960) pioneer models dealing with multiple stages of production. Limited production capacity in multiple stages is considered in Parker and Kapuscinski (2004) and Janakiraman and Muckstadt (2009). They, however, assume that the manufacturer is making only one product and benefits from reliable capacity. Van Mieghem and Rudi (2002) considers a class of models called *newsvendor networks*, that involve multiple products, multiple processing, and storage units. Their primary focus is on models that feature a single-period setting with reliable capacity. They also investigate dynamic settings with either a single product and single processing unit with capacity constraint or a network of multiple processing units with no capacity constraint.

The majority of papers dealing with multiple products belong to the literature on assemble-to-order production systems. Such studies investigate inventory decisions for components used to make final products or simply analyze the effect of shared resources. We refer the reader to Song and Zipkin (2003) for an excellent review of the control of assemble-to-order systems. In the most general setting, a final product may require one or more components. Although not optimal, base-stock policies are typically studied for such systems, because the optimal policy requires knowledge of the current inventories of all components, and finding the optimal policy is computationally prohibitive. Numerous papers deal with the effect of shared resources. For example, Dayanik et al. (2003) studies the effectiveness of several performance bounds for an assemble-to-order system with capacitated component production. Plambeck and Ward (2007) shows that independent control of components is optimal for a class of assemble-to-order systems with expediting.

In addition to assuming base-stock policies, papers dealing with assemble-to-order production systems typically assume that products can be made instantaneously if all the components are available, that is, shared capacity is infinite. An exception is Fu et al. (2006), who considers a single-period, single-product assemble-to-order system in which the assembler faces *finite* assembly capacity.

For systems featuring assembly processes using shared capacity on multiple products, the optimal inventory policy is complicated further. The control of such systems requires some form of capacity allocation (such as the one analyzed in the present paper) to deal with cases when demand cannot be satisfied fully. Bish et al. (2005)

analyzes a two-plant, two-product setting examining the possibility that each plant can produce either one or both of the products to order. Both the components' suppliers and the plants are capacitated. The authors study various capacity allocation schemes, such as allocating capacity to the nearest demands, to the highest-margin demands, or to a plant's primary product. Muriel et al. (2006) studies a similar problem and notes that the capacity allocation policies significantly influence the performance of flexible systems. These studies, however, assume reliable capacity and analyze the performance of commonly used heuristics rather than the optimal policy.

Aviv and Federgruen (2001) analyzes a *two-stage* system with cyclic demand in which the first stage produces a common intermediate product and operates under limited capacity, whereas the second stage produces the final product. The second stage has infinite shared capacity, and only the final products are kept in stock. The problem is solved using a lower bound approximation and the performance of the heuristic is tested. Atali and Ozer (2005) extends Aviv and Federgruen (2001) for Markov-modulated demand and production smoothing constraints.

All the studies cited above consider finite and deterministic capacity. Ciarallo et al. (1994) is the first paper to consider uncertain capacity. Unlike classical yield models, the uncertain capacity models treat capacity as a random variable that truncates planned production independently of planned production quantity. Ciarallo et al. (1994) shows that base-stock policies are optimal for the classical newsvendor problem with uncertain capacity. In the same spirit, Hu et al. (2008) consider the capacity uncertainty of two different facilities, with the possibility of inventory transshipment across locations. They find that base-stock policies are no longer optimal; indeed, the optimal inventory target in one location depends on the existing inventories of both locations.

None of the above papers study the effect of capacity uncertainties when a firm produces multiple products. When manufacturing multiple products, the actual production depends not only on the planned production of each product and the realized capacity, but also on the sequence in which the products are produced. If, for example, one of the products is prioritized (scheduled first), it is less influenced by capacity uncertainty, whereas the product with the lower priority is more exposed to capacity uncertainty. For the calibration stage in the model we propose here, we derive optimal scheduling decisions and identify conditions under which a priority-based schedule or a mixed schedule is optimal. In addition, we derive inventory policies for dedicated production lines with uncertain capacities.

2.3. The Model

Consider a firm producing two products on dedicated production lines (stage 1) that are then calibrated and customized on a shared resource (stage 2) based on customer specifications. Both the dedicated production lines and the shared resource face capacity uncertainties. The firm stocks inventory of the products after stage 1 and carries out stage 2 according to customers' order. Figure 2.1 illustrates our setting.

In our model, we adopt a periodic-review framework. The sequence of events is as follows. In the beginning of period t, the firm observes the inventories of stage-1 products, denoted by $x = (x_1, x_2)$ and raises inventories up to $y = (y_1, y_2)$, with $y \geq x$. We label y as inventory targets, as they are not necessarily achieved due to capacity uncertainties. We model the capacity uncertainty of a given production process i as a random variable K_i. Consequently, the inventory of product i is $y_i \wedge (x_i + K_i)$ after capacity is realized, where $a \wedge b = \min(a, b)$. To distinguish between the inventory targets and realized inventories after production, we denote by $\tilde{y} = y \wedge (x + K)$, where $K = (K_1, K_2)$ in vector notation. Stage 1 finishes production before customer demands are observed.

Next, the firm observes customer demands $\epsilon = (\epsilon_1, \epsilon_2)$ and receives corresponding

Figure 2.1: Model Setup

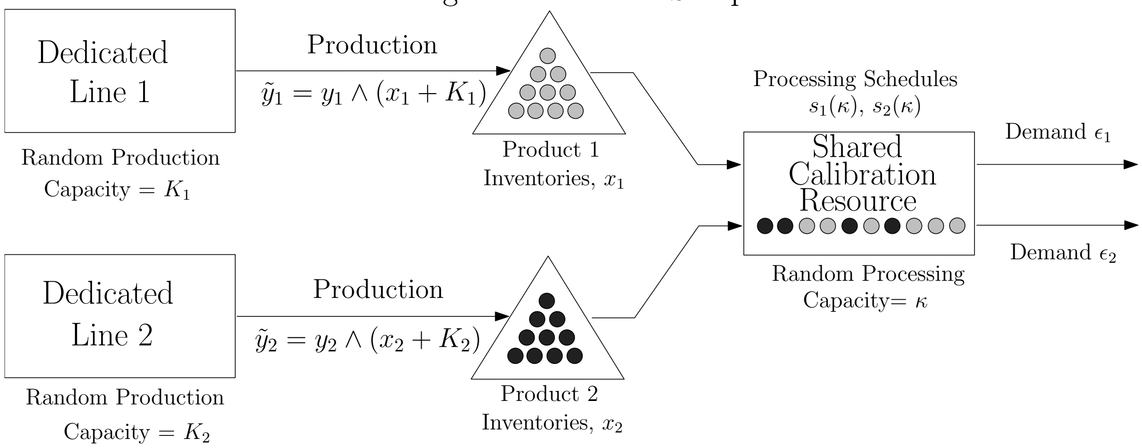

specifications for individual orders. At this point, all the uncertainties within period t are resolved except for the uncertain capacity of the flexible shared resource. In the absence of such information, the firm creates a processing schedule based on the existing inventory of products and demands that are realized. Without loss of generality, we assume that each product uses one unit of capacity of the shared resource. The more general case with different capacity requirements can be handled with slight changes in the analysis, but does not provide any additional insights.

We describe how a processing schedule is created using an example. Assume that the firm schedules calibration of 10 units of product 1 first and 10 units of product 2 thereafter. If the capacity of the shared resource turns out to be 16, then calibrating 10 units of product 1 and 6 units of product 2 can be accomplished in the current period. If, however, the firm schedules the resource evenly (i.e., sequences products as Product 1, Product 2, Product 1, etc.), then 16 units of the shared resource capacity results in processing 8 units of product 1 and 8 units of product 2. Thus, sequencing the products is an integral part of our problem.

As with almost all inventory models, we assume that the demands are continuous random variables and producing fractional quantities is allowed. With continuous variables, it is natural to view the calibration process as a purely continuous flow system with infinitely divisible products.

It is convenient to represent the scheduling decision using two functions, $s_1(\kappa)$ and $s_2(\kappa)$, where $s_i(\kappa)$ is the cumulative number of product i that is processed, when the realized capacity is κ. Consequently, $s_i(\kappa)$ of product-i customers are satisfied in the current period. To illustrate the scheduling functions, consider the continuous version of the above example. When the firm schedules calibrating 10 units of product 1 first and 10 units of product 2 thereafter, we have $s_1(\kappa) = \min(10, \kappa)$ and $s_2(\kappa) = \min(10, (\kappa - 10)^+)$. Such representation allows for processing fractional quantities. For example, if $\kappa = 11.5$, $s_1(11.5) = 10$ and $s_2(11.5) = 1.5$. The function $s_i(\kappa)$ is non-decreasing in κ and $s_1(\kappa)+s_2(\kappa) \leq \kappa$ for any realized κ. Throughout the remainder of the present paper, we refer to $s_i(\kappa)$ as the processing schedule (or simply, "schedule") of product i. To formalize our discussion, we need to present the following definition.

Definition 2.1 (Priority Scheduling for Calibration)**.** *Denote by* $\theta_i = \min(\tilde{y}_i, \epsilon_i)$ *the maximum units of product* i *that can be calibrated in the current period given the realized inventory and the demand for product* i*. We say that product* i *is prioritized if the shared resource processes* θ_i *units of product* i *first, followed by* θ_{3-i} *units of product* $3-i$*. Define* $\bar{s}_i(\kappa) := \min(\theta_i, \kappa)$ *as the processing schedule for product* i *if product* i *is prioritized. Similarly, define* $\underline{s}_i(\kappa) := \min\left\{\theta_i, (\kappa - \theta_{3-i})^+\right\}$ *as the processing schedule for product* i *if product* $3-i$ *is prioritized.*

Obviously, $\underline{s}_i(\kappa) \leq \bar{s}_i(\kappa)$ for each realization of κ and for all $\tilde{y}$ and ϵ, see Figures 2.2(a)-(b). Priority schedules are two extremes in process scheduling that bound the feasible schedules. Any feasible schedule $s_i(\kappa)$ is a continuous, non-decreasing function with $\underline{s}_i(\kappa) \leq s(\kappa) \leq \bar{s}_i(\kappa)$, where $s_1(\kappa) + s_2(\kappa)$ has a rate of increase less than or equal to 1. That is, the resource can process at most one unit using a unit of capacity, ensuring that $s_1(\kappa) + s_2(\kappa) \leq \kappa$ for all κ. We express the required monotonicity properties in terms of (right-sided) derivatives, $s_i'(\kappa) \geq 0$ for $i = 1, 2$ and $s_1'(\kappa) + s_2'(\kappa) \leq 1$. Figure 2.2(c) illustrates a feasible processing schedule for product

Figure 2.2: Scheduling Policies for Product i

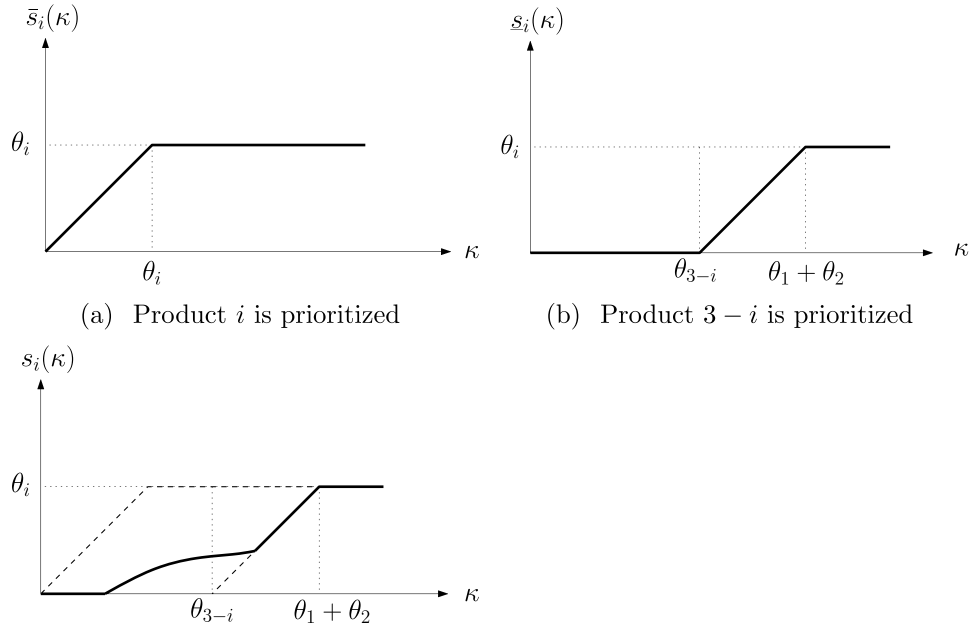

(a) Product i is prioritized

(b) Product $3 - i$ is prioritized

(c) A feasible schedule

i. Initially, no processing takes place for product i. Later, product i is processed at a rate less than 1, that is, it is processed along with product $3-i$. After all of products $3 - i$ are used or all of product $3 - i$ customers are satisfied, product i is processed at a rate of 1, that is, the shared resource processes product i only. Eventually, no more products are processed because demand is fully satisfied.

At the end of the period, the costs are accounted as follows. For any unused product i, a holding cost of h_i is incurred. We assume that unsatisfied customers are lost immediately.

We use a finite-horizon, dynamic-programming formulation to represent the firm's decision problem. Denote by $V_t(x)$ the firm's optimal cost when starting inventory in period t is $x = (x_1, x_2)$. The decision problem consists of two phases that correspond to production stages. In phase 1, inventory targets $y = (y_1, y_2)$ are set for stage 1 and then the capacities of the dedicated production lines and end-customer demands are observed. Then, in phase 2, the firm creates processing schedules for stage 2. We denote by $C_t(\tilde{y},\epsilon)$ the intermediate cost function for starting phase 2 of period t with inventory levels $\tilde{y} = (\tilde{y}_1, \tilde{y}_2)$ and demands $\epsilon = (\epsilon_1, \epsilon_2)$. The decisions in phase 2 are represented through functions $s_1(\kappa)$ and $s_2(\kappa)$, as defined previously. We denote by $s(\kappa) = (s_1(\kappa), s_2(\kappa))$ the vector of scheduling functions. The formulation is given by the following set of equations along with a zero terminal cost function, $V_{T+1}(x) = 0$.

$$\textbf{Phase One :} \qquad V_t(x) = \min_{y \geq x} \mathbb{E}_{K,\epsilon} C_t(y \wedge (x+K), \epsilon), \tag{2.1}$$

$$\begin{aligned} \textbf{Phase Two :} \quad C_t(\tilde{y},\epsilon) &= \min_{s_i(.)} \Big\{ \mathbb{E}_\kappa \{ h[\tilde{y} - s(\kappa)] + p[\epsilon - s(\kappa)] \\ &\qquad + \beta V_{t+1}(\tilde{y} - s(\kappa)) \} & (2.2) \\ s.t \quad & \underline{s}_i(\kappa) \leq s_i(\kappa) \leq \bar{s}_i(\kappa), & (2.3) \\ & s_1'(\kappa) + s_2'(\kappa) \leq 1, & (2.4) \\ & s_i'(\kappa) \geq 0 \; ; \qquad \forall \kappa \geq 0, \; i = 1,2 \Big\} & (2.5) \end{aligned}$$

The processing schedule lies between the two extremes of the priority schedules (condition (2.3)). The resource cannot process more than its capacity (condition (2.4)). Finally, $s_i(.)$ is a weakly increasing function of κ (condition (2.5)). The analysis that follows allows for non-stationary demand and capacity distributions, while assuming stationary costs. For the case of non-stationary cost parameters, our results continue to hold provided that $p_i^t + h_i^t \geq \beta p_i^{t+1}$ for each $i = 1,2$ and $t = 1, 2, \cdots, T$, which is often the case. For brevity, period index is omitted when it does not lead to confusion.

2.4. Optimal Policy

In this section, we derive by induction the optimal policy for the dynamic program formulated in Section 2.3. Assuming that the inductional assumptions hold in period $t+1$ and onward, we solve for the scheduling policy in phase 2 for a given realization of demands and inventory levels in period t. Then, we solve for the production problem in phase 1 in period t. Eventually, we prove that the inductional assumptions are preserved. We need the following definition to state the inductional assumption.

Definition 2.2. *Denote by Δ the derivative operator. A function $f : \mathbb{R}^n \to \mathbb{R}$ is said to satisfy the second-order properties in variables (x_i, x_j) if $\Delta_{x_i x_i} f(x) \geq \Delta_{x_i x_j} f(x) \geq 0$ and $\Delta_{x_j x_j} f(x) \geq \Delta_{x_i x_j} f(x) \geq 0$ for all $x \in \mathbb{R}^n$.*

Second-order properties imply both the joint convexity and supermodularity in the variables of interest. We make the following inductional assumptions for period t.

$\mathbb{A}^1_{t+1}$: The optimal cost function $V_{t+1}(x)$ satisfies *the second-order properties in* (x_1, x_2).

$\mathbb{A}^2_{t+1}$: $p_i + h_i + \beta \Delta_{x_i} V_{t+1}(x) \geq 0$ for $i = 1, 2$.

$\mathbb{A}^2_{t+1}$ implies that it is optimal to satisfy a product-i order in the current period t if there is sufficient inventory of stage-1 product and sufficient processing capacity for stage 2, as opposed to rejecting the customer (by incurring p_i) and holding the inventory for future use (by incurring h_i).[1] Clearly, the inductional assumptions $\mathbb{A}^1_{T+1}$ and $\mathbb{A}^2_{T+1}$ hold only trivially. We assume that $\mathbb{A}^1_{t+1}$ and $\mathbb{A}^2_{t+1}$ hold and will prove that $\mathbb{A}^1_t$ and $\mathbb{A}^2_t$ also hold.

2.4.1 Scheduling Policy – Stage 2

In this section, we analyze the optimal policy in stage 2, that is, the scheduling problem, $(s_1(\kappa), s_2(\kappa))$. As it turns out, the optimal scheduling policy is best described in terms of the remaining inventories $(z_1(\kappa), z_2(\kappa))$ as a function of realized shared capacity. When the realized capacity is κ, $s_i(\kappa)$ units will be withdrawn and the remaining inventory of stage-1 product will become $z_i(\kappa) := \tilde{y}_i - s_i(\kappa)$. A trajectory is defined as the path that the remaining inventories of stage-1 products,

[1] This condition may be violated with non-stationary costs. For example, when the penalty cost in the next period is significantly higher, it may, in fact, be optimal to ration inventory for the next period.

(z_1, z_2), follow as calibration takes place until (ϵ_1, ϵ_2) units of stage-2 products are produced. The optimal scheduling problem can be fully expressed as finding the optimal trajectory to follow. Figure 2.3 illustrates three trajectories (solid, arrowed curves). In Figure 2.3(a), calibration of product 1 is prioritized. Similarly, in Figure 2.3(b), product 2 is prioritized. Figure 2.3(c) illustrates a feasible trajectory in which neither product is prioritized.

Figure 2.3: Illustration of Trajectories

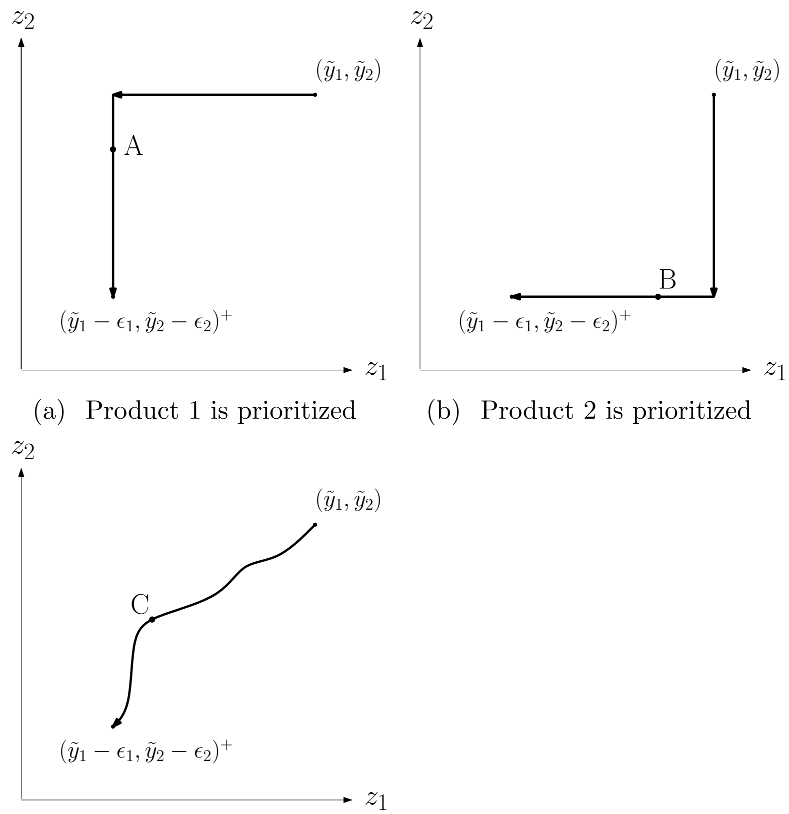

(a) Product 1 is prioritized

(b) Product 2 is prioritized

(c) A feasible trajectory

All the trajectories illustrated in Figure 2.3 terminate at $(\tilde{y} - \epsilon)^+$. Since the choice of the trajectory is made before the actual capacity is known, and the actual capacity may be smaller than needed, some customers may not be satisfied, and some stage-1 products may be carried to the next period. For example, for the same realization of capacity level, trajectory in Figure 2.3(a) terminates at point A, while trajectories in Figures 2.3(b)-(c) terminate at points B and C, respectively. The trajectory determines which customers will be satisfied and which of the products will be carried to the next period.

As illustrated in Figure 2.3, the same capacity realization leads to three different outcomes, A, B, and C, with the same sum of the remaining inventories expressed as $z_1(\kappa) + z_2(\kappa) = \tilde{y}_1 + \tilde{y}_2 - \kappa$. This is an important observation for determining the optimal trajectory. The structure of the policy is characterized in relation to a *target path*, which we describe next. Define $\gamma = p + h$. The firm's cost-to-go as a function of the realized capacity is

$$\begin{aligned} J_t(\kappa|\tilde{y}, \epsilon, s(.)) &:= h(\tilde{y} - s(\kappa)) + p(\epsilon - s(\kappa)) + \beta V_{t+1}(\tilde{y} - s(\kappa)) \\ &= \underbrace{p(\epsilon - \tilde{y})}_{\substack{\text{Treated as constant} \\ \text{in stage 2}}} + \underbrace{\gamma z(\kappa) + \beta V_{t+1}(z(\kappa))}_{\text{Function of trajectory}} \end{aligned} \tag{2.6}$$

Definition 2.3 (Target Path). *For all* $\omega \geq 0$, *we define* $\zeta(\omega) = (\zeta_1(\omega), \zeta_2(\omega)) = \arg\min_{z \geq 0} \left\{ \gamma z + \beta V_{t+1}(z) \mid z_1 + z_2 = \omega \right\}$. *The* "target path" *is defined as* $\mathcal{Z} = \bigcup_{\omega \geq 0} \zeta(\omega)$.

To obtain the *target path*, we minimize the firm's cost for a given sum, ω, of the ending inventory levels to carry to the next period, while ignoring feasibility constraints, and repeat the same procedure for each value of ω. Our objective is to choose a trajectory before observing the exact value of the capacity level. As it turns out, the optimal trajectory is one that moves along a path of ending inventories that minimizes the costs for every realization of capacity. Theorem 2.1 shows that such a path is one that comes as close to the *target path* as possible. We note that the *target path* is independent of the initial state (i.e., realized inventories and demands), as well as the capacity distribution of the shared resource in the current period. Denote by $x|_{[a,b]} = \arg\min_{x' \in [a,b]} |x - x'|$. In other words, $x|_{[a,b]}$ is the closest point to x in the interval $[a, b]$.

Theorem 2.1 (Optimal Scheduling Policy)**.** (a) *The functions* $\zeta_i(\omega)$ *satisfy* $\zeta_1(\omega) + \zeta_2(\omega) = \omega$ *and* $0 \leq \zeta_i'(\omega) \leq 1$. (b) *Let* $z_i^*(\kappa) = \zeta_i(\tilde{y}_1 + \tilde{y}_2 - \kappa)|_{[\tilde{y}_i - \bar{s}_i(\kappa), \tilde{y}_i - \underline{s}_i(\kappa)]}$ *for* $i = 1, 2$. *The optimal trajectory is given by* $\bigcup_{\kappa \geq 0} z^*(\kappa)$*, leading to the optimal scheduling functions,* $s_i^*(\kappa) = [\tilde{y}_i - \zeta_i(\tilde{y}_1 + \tilde{y}_2 - \kappa)]|_{[\underline{s}_i(\kappa), \bar{s}_i(\kappa)]}$ *for* $i = 1, 2$. (c) *The optimal trajectory and the scheduling functions are independent of the current-period capacity distribution of the shared stage 2 resource.* (d) $C_t(\tilde{y}, \epsilon) = \mathbb{E}_\kappa J_t(\kappa | \tilde{y}, \epsilon, s^*(.))$.

For Theorem 2.1, we can also provide a more intuitive description. Figure 2.4 illustrates the optimal trajectory (solid, arrowed curves) for three different starting configurations. The *target path* $\mathcal{Z}$ is shown as a dotted curve. In Figure 2.4(a), the starting inventories are initially off the *target path*; hence, we move toward the *target path* in the most direct manner. Then, we move along the *target path* provided it is feasible to do so. Finally, we move toward the final destination, $(\tilde{y} - \epsilon)^+$. In Figures 2.4(b)-(c), prioritizing the calibration of one of the products turns out to be optimal. In Figure 2.4(b), getting closer to the *target path* is possible by calibrating product 1 only. However, all product 1 demand is satisfied before we reach the *target path*; therefore, we move toward the final destination by calibrating product 2 only. A similar interpretation applies for Figure 2.4(c). The optimal trajectory moves along a path of ending inventories that minimizes the costs for every realization of capacity, leading to part (c) of Theorem 2.1.

As a special case, we consider a symmetric setting with two products that have the same holding and penalty costs, as well as identical demand and dedicated capacity distributions.

Corollary 2.1 (Identical Products)**.** *Due to symmetry, we have* $\zeta_1(\omega) = \zeta_2(\omega) = \frac{\omega}{2}$. *The* target path *corresponds to the 45-degree line* ($z_1 = z_2$)*, implying that the shared resource will process both products at an equal rate along the* target path.

Figure 2.4: Optimal Trajectory of the Remaining Inventories

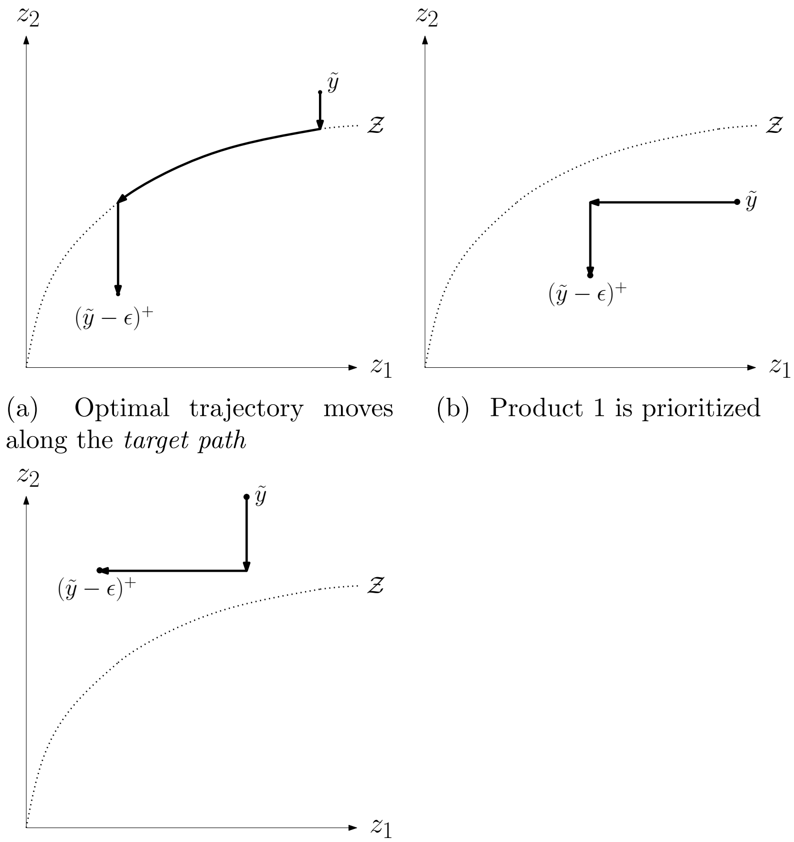

(a) Optimal trajectory moves along the *target path*

(b) Product 1 is prioritized

(c) Product 2 is prioritized

The intuition is straightforward. Although we are indifferent between satisfying demand 1 and 2, we are not indifferent in terms of the remaining inventory of stage-1 products for the next period. To illustrate the policy stated in Corollary 2.1, consider $(\tilde{y}_1, \tilde{y}_2) = (6, 8)$ and $(\epsilon_1, \epsilon_2) = (4, 8)$. Figure 2.5 shows the optimal trajectory. If the products were processed in discrete units, the sequence of calibration would be given by [2-2]-[1-2-1-2-1-2-1-2]-[2-2].

Considering our motivating example, the oilfield services company for which we worked was processing orders in a FIFO manner, without accounting for inventories and capacity. What we find here is that the sequence in which the service requests should be fulfilled depends on the inventory of the tools and the number of requests in each service category.

Figure 2.5: Identical Products Case, $(\tilde{y}_1, \tilde{y}_2) = (6, 8)$ and $(\epsilon_1, \epsilon_2) = (4, 8)$

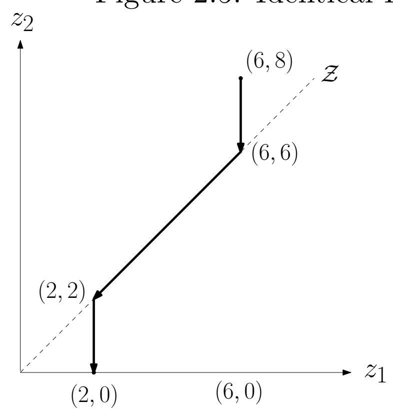

The next proposition characterizes the properties of the intermediate cost function and shows that the second-order properties are preserved, which is an intermediate step for our inductional procedure.

Proposition 2.1. *The intermediate cost function $C_t(\tilde{y}, \epsilon)$ satisfies the second-order properties in $(\tilde{y}_1, \tilde{y}_2)$ (Definition 2.2). Also, $\Delta_{\tilde{y}_i} C_t(\tilde{y}, \epsilon) \geq -p_i$ for $i = 1, 2$.*

2.4.2 Production Policy – Stage 1

We now analyze the optimal production policy. At the beginning of the period, the firm observes the starting inventories for stage-1 products, $x = (x_1, x_2)$ and sets production targets for each product, $y = (y_1, y_2)$. After making the production decisions, the capacities $K = (K_1, K_2)$ for stage-1 products are realized and the inventories of stage 1 products, $\tilde{y} = y \wedge (x + K)$, are observed. Denote by $G_t(\tilde{y}) := \mathbb{E}_\epsilon C_t(\tilde{y}, \epsilon)$ the expected cost-to-go, assuming the inventories of products are known, but the demands are not yet realized. It is convenient to express the optimal cost function as

$$V_t(x) \quad = \quad \min_{y \geq x} \mathbb{E}_K G_t(y \wedge (x + K)) \tag{2.7}$$

The optimal production policy is obtained by assuming that the scheduling policy of stage 2 will be created optimally, as captured by the function $G_t(\tilde{y})$. Let us define functions $\bar{y}_1(y_2)$ and $\bar{y}_2(y_1)$ as $\bar{y}_1(y_2) = \arg\min_{y_1 \geq 0} G_t(y)$ and $\bar{y}_2(y_1) = \arg\min_{y_2 \geq 0} G_t(y)$. If there are multiple minima, we take the smallest one. Let (y_1^0, y_2^0) be the global minimizer of $G_t(y)$, where $\bar{y}_1(y_2^0) = y_1^0$ and $\bar{y}_2(y_1^0) = y_2^0$. In the following theorem, we characterize the optimal production policy.

Theorem 2.2 (Optimal Production Policy)**.** *(a) $\bar{y}_i(x_{3-i})$ is a non-increasing function of x_{3-i} with $\frac{d\bar{y}_i(x_{3-i})}{dx_{3-i}} \geq -1$ $(i = 1, 2)$. (b) Denote by $y^*(x)$ the optimal inventory targets for starting inventories $x = (x_1, x_2)$. An optimal production policy exists that satisfies the following.*

i) If $x_1 \geq \bar{y}_1(x_2)$ and $x_2 \geq \bar{y}_2(x_1)$ (Region 1), then, it is not optimal to produce, $y^*(x) = x$.

ii) If $x_i < \bar{y}_i(x_{3-i})$ and $x_{3-i} \geq \bar{y}_{3-i}(x_i)$ (Regions 2 and 3), $y_i^*(x) = \bar{y}_i(x_{3-i})$ and $y_{3-i}^*(x) = x_{3-i}$.

iii) Let $x_1 \leq \bar{y}_1(x_2)$ and $x_2 \leq \bar{y}_2(x_1)$ (Region 4). Define $\hat{y}(x) = \{y \mid \mathbb{E}_{K_2}\Delta_{y_1} G_t(y_1, y_2 \wedge (x_2 + K_2)) = 0,\ \mathbb{E}_{K_1}\Delta_{y_2} G_t(y_1 \wedge (x_1 + K_1), y_2) = 0\}$. If $\hat{y}(x)$ is a feasible target (that is, $\hat{y}(x) \geq x$), then the optimal inventory target is $y^*(x) = \hat{y}(x)$. Furthermore, $\bar{y}_i(y_{3-i}^*(x)) \leq y_i^*(x) \leq \bar{y}_i(x_{3-i})$ for each $i = 1, 2$. If, however, $\hat{y}_i(x) \geq x_i$ and $\hat{y}_{3-i}(x) < x_{3-i}$ for some i, then it is optimal to set targets as $y_i^*(x) = \bar{y}_i(x_{3-i})$ and $y_{3-i}^*(x) = x_{3-i}$.

Figure 2.6 illustrates the properties of the optimal production policy described in Theorem 2.2. Figure 2.6(a) shows the optimal production policies in Regions 1 through 3. The arrows represent the optimal targets. For example, in Region 2, product 2 is not produced, whereas the product 1 target is set so that inventory reaches the curve $\bar{y}_1(x_2)$. Figures 2.6(b)-(c) illustrate the area in which optimal targets are located for different initial states in Region 4.

Figure 2.6: Optimal Production Policies

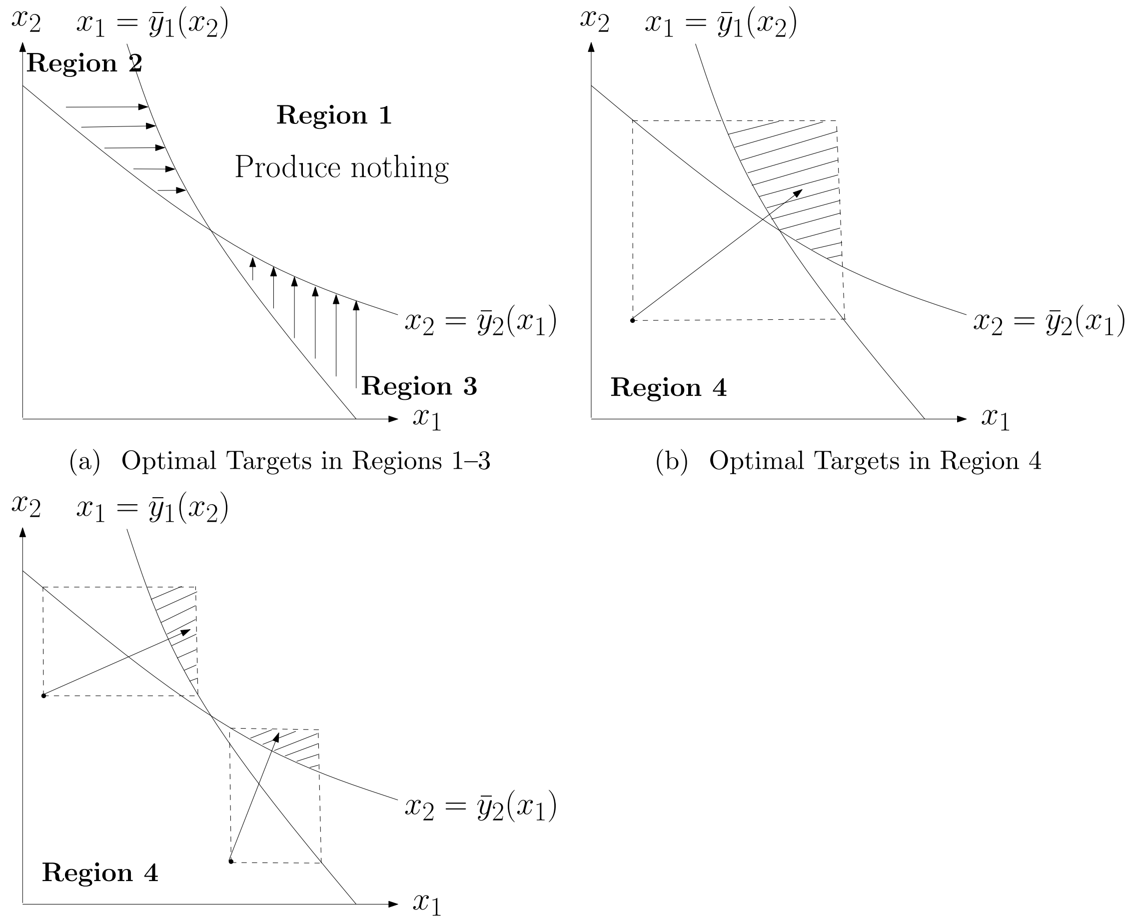

(a) Optimal Targets in Regions 1–3

(b) Optimal Targets in Region 4

(c) Optimal Targets in Region 4

Table 2.1 includes an example problem and shows the optimal targets for various starting inventory levels. Figure 2.7 shows the curves and the regions that characterize the optimal policy for this example. In this example, $(y_1^0, y_2^0) = (11, 11)$ is the minimizer of the function $G_t(y)$. If the starting inventories exceed 11, we do not produce any of the products. When the starting inventories are close enough to $(11, 11)$, the optimal target is $(11, 11)$. If, however, both starting inventories are small, for example, $(x_1, x_2) = (0, 0)$, we set a target of $(12, 12)$. The objective is to achieve a higher production on one of the dedicated production lines in case the other dedicated production line realizes low capacity. With this policy, we are more likely to carry sufficient inventories so that the capacity of the shared resource is not wasted in stage 2. When one of the products has sufficient inventories, we do not produce that product, and the target of the other product does not exceed 11. For example, when $x_2 = 12$, the product 1 target is 11 regardless of x_1. If $x_2 = 13$, we set a smaller target for product 1, equal to 10. Note that although the products are symmetric, the optimal targets are not necessarily equal (i.e., $y_1^*(x) \neq y_2^*(x)$). For example, when $(x_1, x_2) = (6, 0)$, we have $y^*(x) = (12, 11)$. Interestingly, a higher target is set for the product that has a higher inventory. This is the case because having high initial inventory makes it more likely to achieve higher targets.

Figure 2.7: Illustration of the Regions

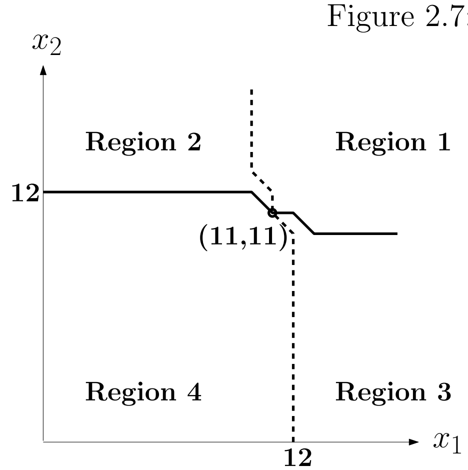

The optimal production policy depends on the initial inventories of the stage-1 products. In addition, unlike stage 2, in which the optimal policy is decoupled from the current-period capacity distribution of the shared resource, the optimal production policy in stage 1 depends on the current-period capacity distributions for the two products. This is the case because the production of one product (hence, its capacity) influences the production of the other product. The sensitivity of the optimal policy to the current-period capacities will be explored in more detail in Section 2.5. The next proposition explores the dependence of the targets on the starting inventories and shows that the optimal targets satisfy certain monotonicity properties in the starting inventories.

Table 2.1: Example – Optimal Targets

13	(10,∼)	(10,∼)	(10,∼)	(10,∼)	(10,∼)	(10,∼)	(10,∼)
12	(11,∼)	(11,∼)	(11,∼)	(11,∼)	(11,∼)	(11,∼)	(11,∼)
11	(11,12)	(11,12)	(11,12)	(11,12)	(11,12)	(11,12)	(11,12)
10	(11,12)	(11,12)	(11,12)	(11,12)	(11,12)	(11,12)	(11,12)
9	(11,12)	(11,12)	(11,12)	(11,12)	(11,12)	(11,12)	(11,12)
8	(11,12)	(11,12)	(11,12)	(11,12)	(11,12)	(11,12)	(11,12)
7	(11,12)	(11,12)	(11,12)	(11,12)	(11,12)	(11,12)	(11,12)
6	(11,12)	(11,12)	(11,12)	(11,12)	(11,12)	(11,12)	(12,11)
5	(11,12)	(11,12)	(11,12)	(11,12)	(11,12)	(11,12)	(12,11)
4	(12,12)	(12,12)	(12,12)	(12,12)	(12,12)	(12,11)	(12,11)
3	(12,12)	(12,12)	(12,12)	(12,12)	(12,12)	(12,11)	(12,11)
2	(12,12)	(12,12)	(12,12)	(12,12)	(12,12)	(12,11)	(12,11)
1	(12,12)	(12,12)	(12,12)	(12,12)	(12,12)	(12,11)	(12,11)
0	(12,12)	(12,12)	(12,12)	(12,12)	(12,12)	(12,11)	(12,11)
(x_1, x_2)	0	1	2	3	4	5	6

13	(10,∼)	(10,∼)	(10,∼)	(∼,∼)	(∼,∼)	(∼,∼)	(∼,∼)
12	(11,∼)	(11,∼)	(11,∼)	(11,∼)	(∼,∼)	(∼,∼)	(∼,∼)
11	(11,12)	(11,∼)	(11,∼)	(11,∼)	(∼,∼)	(∼,∼)	(∼,∼)
10	(11,12)	(11,11)	(11,11)	(11,11)	(∼,11)	(∼,11)	(∼,∼)
9	(11,12)	(11,11)	(11,11)	(11,11)	(∼,11)	(∼,11)	(∼,10)
8	(11,12)	(11,11)	(11,11)	(11,11)	(∼,11)	(∼,11)	(∼,10)
7	(11,12)	(12,11)	(12,11)	(12,11)	(12,11)	(∼,11)	(∼,10)
6	(12,11)	(12,11)	(12,11)	(12,11)	(12,11)	(∼,11)	(∼,10)
5	(12,11)	(12,11)	(12,11)	(12,11)	(12,11)	(∼,11)	(∼,10)
4	(12,11)	(12,11)	(12,11)	(12,11)	(12,11)	(∼,11)	(∼,10)
3	(12,11)	(12,11)	(12,11)	(12,11)	(12,11)	(∼,11)	(∼,10)
2	(12,11)	(12,11)	(12,11)	(12,11)	(12,11)	(∼,11)	(∼,10)
1	(12,11)	(12,11)	(12,11)	(12,11)	(12,11)	(∼,11)	(∼,10)
0	(12,11)	(12,11)	(12,11)	(12,11)	(12,11)	(∼,11)	(∼,10)
(x_1, x_2)	7	8	9	10	11	12	13

Note. Optimal targets are shown for a 50-period problem. The notation "∼" indicates that the product is not produced. The products are symmetric with costs $p_i/(p_i + h_i) = 0.95$, the average demand is 6, the average dedicated capacity for each product is 7, and the average capacity of the shared resource is 13. All distributions are triangular. Demand variability is 0.4, capacity variabilities are all 0.8.

Proposition 2.2. *The optimal inventory targets satisfy the following monotonicity properties in the starting inventories: for each* $i = 1, 2$, (i) $0 \leq \frac{\partial y_i^*(x)}{\partial x_i} \leq 1$, (ii) $-1 \leq \frac{\partial y_i^*(x)}{\partial x_{3-i}} \leq 0$, (iii) $\frac{\partial y_i^*(x)}{\partial x_i} - \frac{\partial y_i^*(x)}{\partial x_{3-i}} \leq 1$.

The optimal targets are increasing in their own inventory and decreasing in the other product's inventory. Figure 2.8 illustrates the monotonicity properties of the targets in x_1 for a given value of x_2. If x_2 is increased slightly, then the curve representing the product 1 target would shift downward, whereas the curve representing the product 2 target would shift upward.

Figure 2.8: Monotonicity of Targets

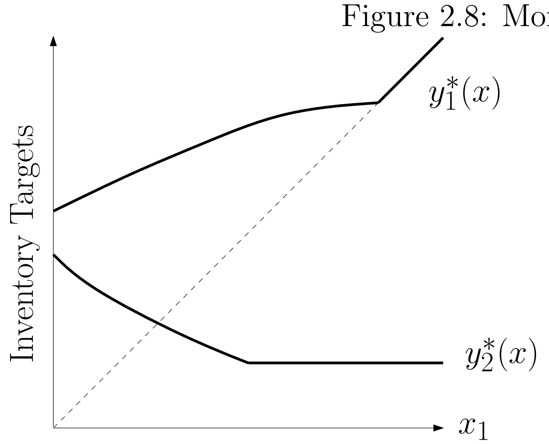

The monotonicity properties we derived in Proposition 2.2 are also vital in completing the inductional step in the next proposition.

Proposition 2.3. *The optimal cost function* $V_t(x)$ *satisfies the second-order properties in* (x_1, x_2) *(Definition 2.2). In addition,* $p_i + h_i + \beta \Delta_{x_i} V_t(x) \geq 0$ *for* $i = 1, 2$.

With Proposition 2.3, the optimal policy structure is established for the most general case in which the capacities of the dedicated production lines and the shared resource are uncertain. Theorems 2.1 and 2.2 characterize the optimal production policy and the optimal scheduling policy fully. Although we established the optimal policy assuming a finite-horizon setting and independent distributions across periods, we can relax these assumptions. In Appendix 2.10, we address infinite-horizon and Markov-modulated settings. Furthermore, we describe the optimal policy structure for special cases of the original model.

Only One Dedicated Production Line with Uncertain Capacity. Assume that dedicated production line 1 has infinite capacity, and dedicated production line 2 faces uncertain capacity. Due to infinite capacity assumption, Theorem 2.2 simplifies and first-order conditions no longer depend on x_1. In this case, the production targets for each product are only functions of the initial inventory x_2 in Region 4.

Capacitated Dedicated Production Lines with No Uncertainty. Assume that dedicated production line i has a certain finite capacity, K_i. Recall that (y_1^0, y_2^0) is the global minimizer of $G_t(y)$. Define $\tau_i(x_i) = \Big((x_i + K_i) \wedge y_i^0\Big) \vee x_i$. We describe the optimal policy in relation to $\tau_i(x_i)$.

Proposition 2.4. *If the dedicated production line i has a deterministic capacity K_i for $i = 1, 2$, then the inventory target for product i depends on the initial inventory of the other product only and is given by $\bar{y}_i(\tau_{3-i}(x_{3-i}))$. The optimal targets satisfy $-1 \leq \frac{\bar{y}_i(\tau_{3-i}(x_{3-i}))}{x_{3-i}} \leq 0$ for each i.*

Proposition 2.4 provides a simpler characterization for the optimal production policy when the dedicated production lines face deterministic capacity. The optimal target of a product does not depend on its own inventory, while it is decreasing in the inventory of the other product. In addition, if each of the dedicated production lines has infinite capacity, then a base-stock policy (y_1^0, y_2^0) is optimal, provided as well that the initial state is below the base-stocks.

Uncapacitated Stage 2. Assume that the shared capacity has infinite capacity. The setting with uncapacitated stage 2 reduces to two independent, single-product settings, which Ciarallo et al. (1994) has studied. The inventory policy for each product is a base-stock policy. The scheduling policy is thus irrelevant, because any schedule can be completed due to infinite capacity.

2.5. Sensitivity of Optimal Policy

In Section 2.4, we analyzed the optimal policy and found that the optimal scheduling policy is independent of the shared capacity distribution in the current-period, whereas the production policy depends on the distribution of the dedicated capacities. In this section, we explore the sensitivity of the optimal policy (including the production policy and the scheduling policy) with respect to the current-period parameters, such as dedicated capacities and shared capacity, as well as demand levels and cost coefficients. Below, we use "increase" and "decrease" in the non-strict sense. For random variables, such as capacity and demand, "increased" means stochastic dominance.

2.5.1 Sensitivity to Capacity

We first examine the sensitivity of the optimal policy to the dedicated capacities. Because the scheduling decisions are based on realizing the dedicated capacities and demands, changing the capacity uncertainty of the dedicated production lines in the current period does not influence the scheduling policy. It does, however, impact the production policies for both products.

Proposition 2.5. *If the capacity of dedicated production line 1 (2) is stochastically increased in the current period t, then the inventory target for product 1 (2) increases, while the inventory target for product 2 (1) decreases.*

Note that this is different from inventory target for a one-product, one-stage case. It is known (Ciarallo et al., 1994, Hu et al., 2008) that for one stage system, the current period target is independent of variability of capacity. The difference results from the shared capacity. When the capacity of dedicated production line 1 is increased, achieving the product 1 target becomes more likely, leading to an increase in the product 1 target and a decrease in the product 2 target. This is similar in spirit to the effect of increasing the starting inventory of product 1, which increases the product 1 target and decreases the product 2 target (see Table 2.1 and Proposition 2.2). Next, we explore how the optimal policy changes as the capacity of the shared resource changes.

Proposition 2.6. *If the shared capacity is stochastically increased in the current period t, the sum of the inventory targets increases. If the capacities of the dedicated production lines and the shared resource are deterministic and the shared capacity is increased, inventory targets increase for both products.*

While Proposition 2.6 shows that individual targets increase when the deterministic shared capacity is increased, similar behavior is observed with uncertain shared capacity. Numerical evidence suggests that in the case of uncertain shared capacity, not only the sum of targets increases, but both targets also increase. The interesting result here is that an increase in capacity can lead to increase in inventory, which initially appears to be counter-intuitive. The intuition, however, can be explained as follows. If the capacity of the shared resource is increased, the firm has a higher potential to satisfy demand; hence, the inventory targets increase for products. Consider two extreme cases. If shared capacity is close to zero, little opportunity exists to satisfy demand, even if inventories are sufficient to meet all the demand. Thus, the shared resource is the bottleneck and maintaining significant inventories is not a useful strategy. On the other hand, if the shared capacity is plentiful, the potential exists to satisfy demand; hence, maintaining more inventories is beneficial.

2.5.2 Sensitivity to Demand and Cost Parameters

Next, we examine the sensitivity of the optimal policy to the demand levels and the cost parameters (holding and penalty costs). The scheduling policy is not affected by a change in current-period demand distribution. We characterize how the inventory targets change as the demand distribution changes.

Proposition 2.7. *Assume that the dedicated capacities are finite and certain. If the* *demand for product 1 (2) is stochastically increased in the current period t, then the product 1 (2) target increases, while the product 2 (1) target decreases.*

The fact that increasing product 1 demand increases the product 1 target and decreases the product 2 target is expected, because we do not want to starve the shared resource. Having more inventory of the product with the higher demand achieves this objective. Next, we analyze the sensitivity to the cost parameters. The cost parameters impact both the production policy and the scheduling policy. The results are shown for the most general case when both the dedicated capacities and the shared capacity are uncertain.

Consider two production schedules, $s(\kappa)$ and $\hat{s}(\kappa)$. We say that the schedule $s(\kappa)$ produces (calibrates) more of product i and less of product $3-i$ than does schedule $\hat{s}(\kappa)$, if $s_i(\kappa) \leq \hat{s}_i(\kappa)$ for each κ.

Proposition 2.8. *(i) If the penalty cost for product 1 (2) is increased in the current period t, then the product 1 (2) target increases, whereas the product 2 (1) target decreases. In addition, we produce (calibrate) more of product 1 (2) and less of product 2 (1) using the shared resource.*

(ii) If the holding cost for product 1 (2) is increased in the current period t, then the product 1 (2) target decreases, whereas the product 2 (1) target increases. In addition, we calibrate less of product 1 (2) and more of product 2 (1) using the shared resource.

2.6. Effect of Product Asymmetries on Inventory Policy

Assume that one of the dedicated production lines has lower capacity. Everything else being equal, do we hold more inventory of the product with higher capacity or the product with lower capacity? Similarly, assume that one of the dedicated production lines has more reliable (less uncertain) capacity. Everything else being equal, do we hold more inventory of the product with more reliable capacity or the product with less reliable capacity? These questions can be repeated for demand levels and demand variability. Based on the conventional wisdom of the single-product, single-stage inventory models, we may expect that we should hold more inventory of the product that has lower or less reliable capacity. Similarly, we might expect that we should hold more inventory of the product with larger demand or more demand variability. Does our intuition hold when we have two stages of production, in which stage 2 is shared?

The short answer for these questions is that it depends on the level of shared capacity. Our intuition holds, for example, when stage 2 has infinite capacity, in which case our problem reduces to two independent, single-product settings. Surprisingly, we find that our intuition may not hold with finite shared capacity for stage 2, and that the actual result may be the opposite of what simple intuition might suggest.

To address the above questions, we consider symmetric settings across products (i.e., identical demand and capacity distributions, penalty and holding costs) and assume an asymmetric setting only in the dimension of interest (i.e., dedicated capacity levels, reliability of the dedicated production lines, demand levels, and demand variability). In all of the numerical studies that follow, we solve a 100-period problem with stationary parameters.

Base Case: Consider a base case that assumes identical configurations across products with $p_i = 10$, $h_i = 1.5$, $\mathbb{E}[\epsilon_i] = 20$, $cv(\epsilon_i) = 0.4$ (i.e., *coefficient of variation*), $\mathbb{E}[K_i] = 20$, $cv(K_i) = 0.4$ for $i = 1, 2$. Stage 2 has no capacity uncertainty, $cv(\kappa) = 0$. The random variables (capacity and demand distributions) are modeled as a two-point mass with equal probabilities.

The following cases are defined relative to the base case.

- **Effect of Asymmetric Dedicated Capacity:** We first explore how the level of dedicated capacity and its reliability affect the optimal inventory policy. Figure 2.9 shows two numerical examples. In the first example (a), product 1 has less dedicated capacity than product 2. In the second example (b), product 1 has less reliable dedicated capacity than product 2. The inventory targets for each product are shown for initial inventories $(0, 0)$ as a function of the shared capacity level.

Figure 2.9: Effect of Asymmetric Dedicated Capacity on Inventory Targets

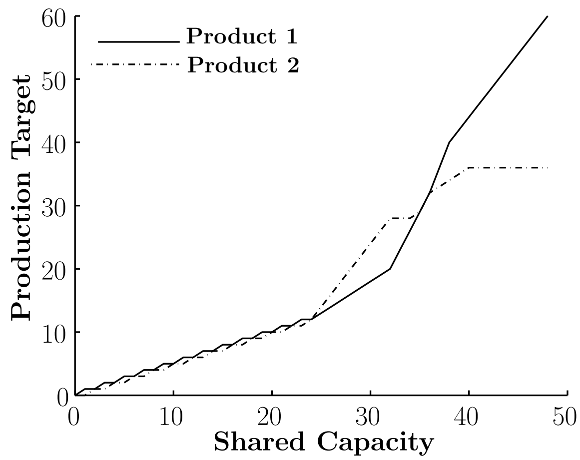

(a) Effect of Capacity Levels
($K_1 = 12$, $K_2 = 20$, w.p.1)

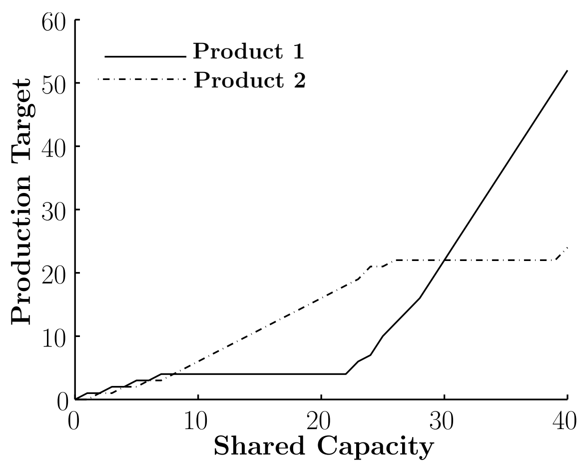

(b) Effect of Capacity Reliability
($cv(K_1) = 0.8$, $cv(K_2) = 0.1$)

Considering first the effect of capacity level, we clearly observe three behaviors corresponding to the tightness of the shared capacity. Dedicated capacity is irrelevant for very tight shared capacity: target levels are equal and add up to the shared capacity level. With low shared capacity, demand cannot be fully satisfied anyway; hence, targets are low and the likelihood of achieving those targets is high for both products. Thus, no need exists for an extra inventory buffer. When shared capacity is increased, product 2 has a higher target. A higher target for the product with a higher dedicated capacity allows us to avoid starving the scarce shared capacity. This is the case because producing product 2 can be ramped up more quickly in the face of large demands, resulting in better utilization of the shared resource. Eventually, if shared capacity is sufficiently large, it becomes possible to better satisfy demands. Stocking more units of the product with a lower dedicated capacity is needed, because inventory can be depleted when demand is large, and it takes several periods to recover. Similar patterns of (a) equal targets, (b) more inventory for the "more favorable" product, and (c) more inventory for the "less favorable" product are observed for asymmetric capacity reliability.

- **Effect of Asymmetric Demand:** We next explore how the demand level and its variability affect the optimal inventory policy. Figure 2.10 shows two numerical examples. In the first example (a), product 1 has larger demand than product 2. In the second example (b), product 1 has more demand variability than product 2.

With higher product 1 demand and high shared capacity, the target for product 1 is significantly higher in order to buffer it from uncertain dedicated capacity and to satisfy the demand. When shared capacity is low, however, we do not buffer too much more of product 1, because the likelihood of starving the shared resource will be low in any case. Considering the effect of demand variability, we observe a similar pattern as above. Demand variability is irrelevant for very tight shared capacity. In the middle range of the shared capacity, the product with less variable demand has a higher inventory target. Because shared capacity is still scarce, stocking the product with a highly uncertain demand may result in under-utilization of the shared resource. This risk is avoided by stocking product with more predictable demand. Finally, if shared capacity is sufficiently large, stocking more units of the product with higher demand variability allows the company to satisfy more demand.

Figure 2.10: Effect of Asymmetric Demand on Inventory Targets

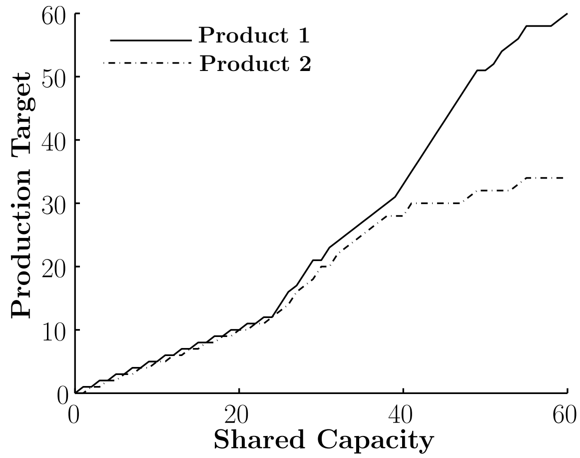

(a) Effect of Demand Levels
($\mathbb{E}(\epsilon_1) = 25$, $\mathbb{E}(\epsilon_2) = 20$, $\sigma(\epsilon_1) = \sigma(\epsilon_2) = 8$)

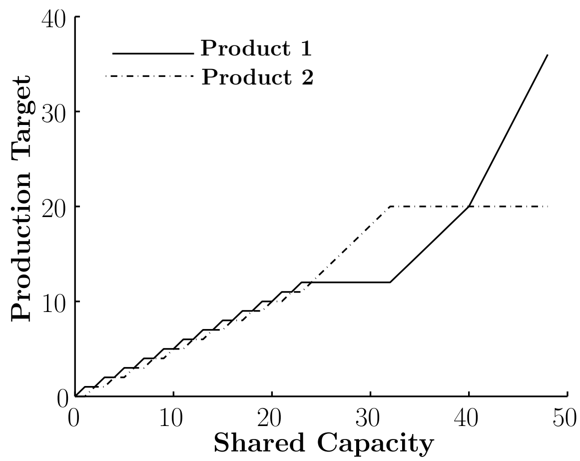

(b) Effect of Demand Variability
($cv(\epsilon_1) = 0.4$, $cv(\epsilon_2) = 0$, $cv(K_1) = cv(K_2) = 0$)

2.7. Heuristic Policy

While the optimal policy is easy to explain, the case of multiple products suffers from the curse of dimensionality. As the state space or the sample space of the random variables grows, the dynamic-programming formulation becomes computationally challenging. In this section, therefore, we propose two implementable heuristic policies, a *straw policy* and a *fixed-rate scheduling policy* and report on their performance.

Straw Policy (Heuristic S): A *straw policy* follows a base-stock policy for production, in which base stocks are obtained by ignoring the shared capacity. With the shared capacity ignored, the problem is reduced to solving two independent, single-product problems (the single-product version of the problem is analyzed in Ciarallo et al. (1994)). After the base stocks are obtained, calibrating the product with a higher penalty cost is prioritized.

Fixed-Rate Scheduling Policy (Heuristic F): This heuristic also follows a base-stock policy for production, but does not ignore the shared capacity in setting the base-stock levels for the calibration stage. We thus propose an easy-to-implement scheduling policy. We approximate the *target path* by a linear trajectory and schedule the processing of products at a fixed rate that is proportional to base stocks along the *target path*. When the starting inventories are off the approximated *target path*, we move toward it in the most direct manner, as is the case in the optimal policy. Note that our *fixed-rate scheduling policy* is optimal for the identical products case, as stated in Corollary 2.1. For two products, let y_1^* and y_2^* be the base stocks for products 1 and 2. Let $\alpha_i := \frac{y_i^*}{y_1^*+y_2^*}$ and $\zeta_i(k) := \alpha_i k$ for $i = 1, 2$. Heuristic F is then given by $s_i(\kappa) = [y_i - \alpha_i(y_1 + y_2 - \kappa)]|_{[\underline{s}_i(\kappa),\bar{s}_i(\kappa)]}$. To obtain the base stocks under this policy, we use a simulation-based optimization model. The heuristic for stage 2 is optimal in the symmetric settings, and the base stocks for all products are equal. The only difference is that the inventory targets in the optimal policy depend on the initial inventories; this is not captured by Heuristic F.

To evaluate the heuristic's performance, we start with the case of two products for which the optimal policy is available. Then, we consider the case of more than two products, each having dedicated resources and all using a common shared resource. Given the base-stock levels for each product, implementing the two heuristics is straightforward. Obtaining the base-stock levels, however, involves optimization in n-dimensions. To obtain the base-stock level for a given product, we aggregate all remaining products and search for the best base-stocks for a two-product model.[2] We compare the heuristic's performance to a lower bound on the optimal costs, obtained by relaxing the shared-capacity constraint. The gap we report is an upper bound on the actual optimality gap, except for a three-product case analyzed separately, for which the optimal costs are obtained computationally. We evaluate the effect of the number of products on the performance of the heuristics and show that the heuristics perform better as the number of products is increased.

2.7.1 Performance of the Heuristics

We first illustrate how the shared capacity influences the heuristic performance.[3] Using the same numerical example used in Figure 2.9(b), which considered asymmetric products with different dedicated capacity reliability, Figure 2.11 shows the performance of the heuristics.

Figure 2.11: Heuristic Performance for Two Products

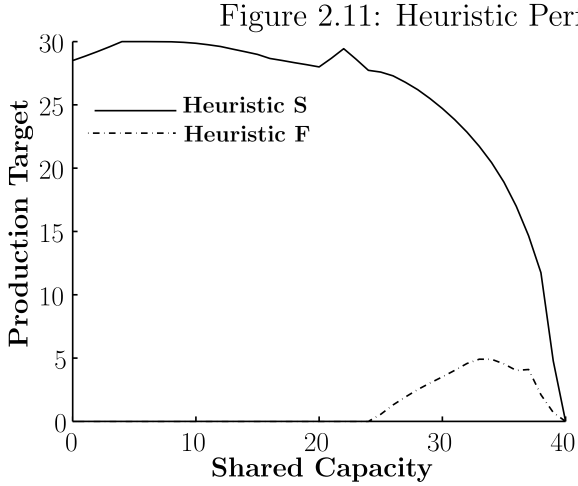

[2]To aggregate the products, we pool the demands and dedicated capacities of the products and use holding and penalty costs obtained by a weighted average of the costs of individual products. The weight for a product is the ratio of the average demand for that product and the average demand for all products being aggregated.

[3]To evaluate the heuristic performance for a problem instance, we first derive the optimal policy using the dynamic programming formulation for a horizon of $T = 100$ periods. The straw base stocks also depend on time horizon, and we find them for $T = 100$. To obtain the fixed-rate scheduling policy, we simulate the system for $N = 10,000$ periods for a given set of base stocks. We then search for the best base stocks using the same sample path for all simulations. Next, we simulate the system using the optimal policy and each of the heuristics for 10,000 periods along 10 sample paths.

For Heuristic S, the optimality gap is high (around 30%) when the shared capacity is low. This is expected, because the *straw policy* ignores the shared capacity. Heuristic F results in much lower optimality gaps. For moderately high shared capacity, the optimality gap is up to 5%. For lower capacity, Heuristic F is near-optimal. We observe a similar pattern with the other numerical examples in Section 2.6.

To evaluate the performance of the heuristics further, we ran several test studies. We considered both symmetric settings (identical costs, demands, and dedicated capacities) and asymmetric settings. To create a test bed, we varied the parameters in multiple dimensions. Table 2.2 summarizes the parameters used for the symmetric settings,[4] and Table 2.3 summarizes the parameters used for asymmetric settings.[5]

Table 2.2: Parameters for the Test Cases (Symmetric Settings)

Capacity Utilization	80%, 90%, 96%, 100%
Capacity Variability	0.25, 0.50, 0.75, 1
Uncapacitated Service Levels	80%, 90%, 95%
Demand Variability	0.25, 0.50, 0.75, 1

Note. We fixed the average demand to 24 and chose the average capacities to control for the utilization of each resource (dedicated production line i's utilization is $\mathbb{E}[\epsilon_i]/\mathbb{E}[K_i]$, and shared resource utilization is $\mathbb{E}[\epsilon_1 + \epsilon_2]/\mathbb{E}[\kappa]$). The holding and penalty costs are chosen to control for uncapacitated optimal service levels, $p_i/(p_i + h_i)$.

Table 2.3: Parameters for the Test Cases (Asymmetric Settings)

Capacity Utilization	75%, 80%, 85%, 90%, 95%
Penalty Cost	40, 60, 80, 100, 120
Holding Cost	0.4, 0.6, 0.8, 1.0, 1.2
Average Demand	100, 120, 140, 160, 180

Note. Demand and dedicated capacities are exponentially distributed, whereas the shared capacity is Erlang$-n$.

Table 2.4 reports the performance of both heuristics for symmetric systems. Clearly, Heuristic F performs extremely well in symmetric settings, and it outperforms Heuris-

[4]Capacity utilization and variability are the same for production, but could differ for the calibration stage. A total of 3,072 test cases are considered (3 configurations for costs, 4 for demands, 16 for stage 1, and 16 for stage 2).

[5]To evaluate the heuristic performance, we randomly generate 1,000 problem instances and report the average gap across all instances.

tic S. The numerical evidence also suggests that the performance of the *straw policy* is similar for the two-product case and the three-product case, whereas the *fixed-rate scheduling policy* results in lower optimality gaps for the three-product case compared to two-product case.

Table 2.4: Performance of the Heuristics (Symmetric Settings)

	Two Products		**Three Products**	
Optimality Gap	**Heuristic S**	**Heuristic F**	**Heuristic S**	**Heuristic F**
Average	7.28%	0.05%	6.94%	0.01%
99^{th} **Percentile**	16.89%	0.59%	15.77%	0.18%

Next, we consider the case of more than three products and investigate the effect of the number of products on the heuristics' performance. When we compare performance for different numbers of products, we keep the average shared capacity per product constant; therefore, utilization of the shared resource remains the same when we modify the number of products. Figure 2.12 illustrates the heuristic performance for the test cases we consider. Recall that we are comparing our heuristics to a lower bound on costs. In all cases, the *fixed-rate scheduling policy* performs better than the *straw policy.* As the number of products is increased along with proportional capacity increase, however, the performance of both heuristics improves and the optimality gap rapidly converges to zero. Although the gap is initially very high, we expect that the actual gap between the heuristics and the optimal cost is much smaller based on our findings from the two-product case. Thus, the bound on the gap may be significantly inflated when the number of products is small.

When the number of products is increased and the shared capacity per product remains fixed, more opportunity emerges to reap the benefits of capacity pooling for the calibration stage. Given that the total shared capacity is large relative to the demand for a single product, our conclusion is that the shared capacity is less relevant when the number of products is high. In such situations, the shared capacity can be

Figure 2.12: Comparison of Heuristic Performance to a Lower Bound on Optimal Cost

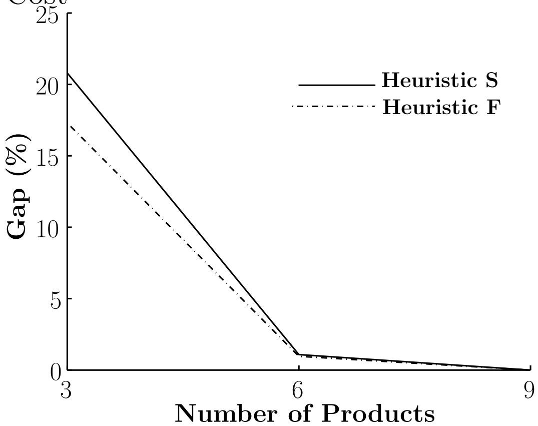

treated as infinite and can be ignored. If the number of products is small and the capacity is tight, however, ignoring the shared capacity will result in poor outcomes.

2.8. Conclusions and Further Research

In this essay, we have considered optimal production and inventory control for a make-to-stock/calibrate-to-order system with dedicated resources for each product in stage 1 and a common resource that all products share in stage 2. We fully characterized the optimal policy for the case of two products and proposed heuristic policies for the case of multiple products based on the optimal policy structure for two products. We numerically explored the effect of product asymmetries on the optimal policy, showing that depending on shared capacity level, three different modes of behavior are present. We also showed numerically that the performance of the heuristics is near-optimal when the number of products is sufficiently large and the shared capacity is large relative to the demand for individual products.

Our model and assumptions are driven by our work with an oilfield services company that made tools in advance but had to calibrate/customize them right before use in a short time or faced lost sales. An interesting extension of our problem would be to consider backlogging. Incorporating demand backlogging, however, causes explosion of the state space. Not only do we need to track inventories in our state description, but we also we need to keep track of the number of backlogged customers for each product. This departs from the classical inventory models in which inventory position, defined as inventory-on-hand minus backlogs, is sufficient as a state variable. Due to shared capacity in our model, positive inventories and positive backlogs may coexist. Hence, inventory position alone is insufficient to describe the state of the system. This is an interesting case for future research.

Furthermore, in our setting, the calibrated tools could not be held in stock first because the calibration depended on the actual use (e.g., a tool used at a different depth needed to be calibrated differently) and the type of use information would only be revealed when the demand arrived. Another interesting extension for other environments might be to consider a case in which the firm may be able to keep inventory of products in stage 2 as well.

2.9. Appendix: Mathematical Proofs

Proof of Theorem 2.1. Consider first the case when the shared capacity in period t is deterministic and equals κ. Clearly, only production quantities $s_1(\kappa)$ and $s_2(\kappa)$ matter – the trajectory of production (i.e., scheduling) is irrelevant. To reflect this, we refer to the scheduling problem with deterministic capacity in the current period as *capacity allocation problem.* We denote by $J_t(\tilde{y}, \epsilon, \kappa)$ the intermediate cost function. Defining $z_i = \tilde{y}_i - s_i(\kappa)$ as the stage-1 product i inventory carried to the next period, the capacity allocation problem reduces to the following:

$$J_t(\tilde{y}, \epsilon, \kappa) = \min_{(z_1, z_2)} \Big\{ hz + p[\epsilon - \tilde{y} + z] + \beta V_{t+1}(z) \tag{2.8}$$

$$s.t \quad \underline{s}_i(\kappa) \leq \tilde{y}_i - z_i \leq \bar{s}_i(\kappa), \; i = 1, 2 \tag{2.9}$$

$$(\tilde{y}_1 - z_1) + (\tilde{y}_2 - z_2) \leq \kappa \Big\} \tag{2.10}$$

The decision problem is expressed in terms of the inventory to carry over to the next period. (2.8) replaces (2.2), (2.9) replaces (2.3), and (2.10) replaces (2.4) and (2.5) in the original formulation (the trajectory of production is ignored). The optimal solution should satisfy the following condition.

Optimality Condition: If the shared capacity is not binding, $\min(\tilde{y}_1, \epsilon_1)+\min(\tilde{y}_2, \epsilon_2) \leq \kappa$, then $s_i(\kappa) = \min(\tilde{y}_i, \epsilon_i)$, implying $s_i(\kappa) = \underline{s}_i(\kappa) = \bar{s}_i(\kappa)$ and $z_i = (\tilde{y}_i - \epsilon_i)^+$. Otherwise, all the shared capacity must be utilized, $s_1(\kappa) + s_2(\kappa) = \kappa$, implying that $z_1 + z_2 = \tilde{y}_1 + \tilde{y}_2 - \kappa$.

Due to $\mathbb{A}^2_{t+1}$, i.e., $p_i + h_i + \beta\Delta_{x_i}V_{t+1}(z) \geq 0$, the function minimized in (2.8) is non-decreasing in z_1 and z_2, hence the optimality condition is justified. Recall that $\gamma = p + h$. When the shared capacity is binding, we replace $z_2 = \tilde{y}_1 + \tilde{y}_2 - \kappa - z_1$, leading to optimization in single variable.

$$\begin{aligned} J_t(\tilde{y}, \epsilon, \kappa) &= p(\epsilon - \tilde{y}) + \gamma_2(\tilde{y}_1 + \tilde{y}_2 - \kappa) \\ &\quad + \min_{z_1}\Big\{(\gamma_1 - \gamma_2)z_1 + \beta V_{t+1}(z_1, \tilde{y}_1 + \tilde{y}_2 - \kappa - z_1) \\ &\qquad\qquad s.t \quad \tilde{y}_1 - \bar{s}_1(\kappa) \leq z_1 \leq \tilde{y}_1 - \underline{s}_1(\kappa)\Big\} \end{aligned} \tag{2.11}$$

Define $f_t(z_1, \omega) := (\gamma_1 - \gamma_2)z_1 + \beta V_{t+1}(z_1, \omega - z_1)$ as the function to be minimized in (2.11). Since $V_{t+1}(x)$ is jointly convex, $f_t(z_1, \omega)$ is convex in z_1 for $z_1 \in [0, \omega]$ and for any $\omega \geq 0$. To obtain the optimal capacity allocation, let $\zeta_1(\omega) = \arg\min_{0\leq z_1\leq\omega} f_t(z_1, \omega)$. In case of multiple optimal solutions, we choose the smallest one. As a consequence, the optimal capacity allocation satisfies $z_1 = \zeta_1(\tilde{y}_1 + \tilde{y}_2 - \kappa)|_{[\tilde{y}_1-\bar{s}_1(\kappa),\tilde{y}_1-\underline{s}_1(\kappa)]}$. Defining $\zeta_2(\omega) = \omega - \zeta_1(\omega)$ and by a symmetric argument, we also have $z_2 = \zeta_2(\tilde{y}_1 + \tilde{y}_2 - \kappa)|_{[\tilde{y}_2-\bar{s}_2(\kappa),\tilde{y}_2-\underline{s}_2(\kappa)]}$, which implies $s_i(\kappa) = \{\tilde{y}_i - \zeta_i(\tilde{y}_1 + \tilde{y}_2 - \kappa)\}|_{[\underline{s}_i(\kappa),\bar{s}_i(\kappa)]}$ for each $i = 1, 2$. Before proceeding to the case with stochastic shared capacity, we show that $0 \leq \zeta_i'(\omega) \leq 1$. First, note that due to the optimality we have:

$$\left.\frac{\partial f_t(z_1,\omega)}{\partial z_1}\right|_{z_1=\zeta_1(\omega)} = 0 = (\gamma_1-\gamma_2)+\beta\Delta_{x_1}V_{t+1}(\zeta_1(\omega),\zeta_2(\omega)) -\beta\Delta_{x_2}V_{t+1}(\zeta_1(\omega),\zeta_2(\omega)) \quad (2.12)$$

To evaluate $\zeta_1'(\omega)$, we take the derivative of both sides of (2.12) and obtain

$$\zeta_1'(\omega) = \frac{(\Delta_{x_2x_2}-\Delta_{x_1x_2})V_{t+1}(\zeta_1(\omega),\zeta_2(\omega))}{(\Delta_{x_1x_1}+\Delta_{x_2x_2}-2\Delta_{x_1x_2})V_{t+1}(\zeta_1(\omega),\zeta_2(\omega))} \quad (2.13)$$

$0 \leq \zeta_1'(\omega) \leq 1$ follows due to the inductional assumptions on $V_{t+1}(x)$. Since $\zeta_2(\omega)=\omega-\zeta_1(\omega)$, we immediately have $0\leq\zeta_2'(\omega)\leq 1$ and $\zeta_1'(\omega)+\zeta_2'(\omega)=1$.

Next, consider stochastic shared capacity. First, we relax (2.3) and (2.4), that is, we ignore $s_1'(\kappa)+s_2'(\kappa)\leq 1$ and $s_i'(\kappa)\geq 0$. Since choice of $s_i(\kappa)$ for a given κ is now decoupled from any other capacity $\kappa'\neq\kappa$, the problem reduces to independently minimizing the objective function for each realization of κ, which becomes the deterministic capacity case. Thus, the solution to the relaxed problem is given by $s_i(\kappa)=\left\{\tilde{y}_i-\zeta_i(\tilde{y}_1+\tilde{y}_2-\kappa)\right\}|_{[\underline{s}_i(\kappa),\bar{s}_i(\kappa)]}$, which satisfies ignored conditions (2.3) and (2.4), hence, it is feasible for the original problem with stochastic shared capacity. (2.3) follows from $\zeta_1'(\omega)+\zeta_2'(\omega)=1$ and (2.4) follows from $0\leq\zeta_i'(\omega)\leq 1$. Hence, $s_i(\kappa)$ is indeed optimal for the original problem with stochastic shared capacity. Clearly, the optimal schedule is independent of the distribution of the current-period shared capacity. As a consequence of the optimal policy, we have $C_t(\tilde{y},\epsilon)=\mathbb{E}_\kappa J_t(\tilde{y},\epsilon,\kappa)$: the firm incurs $J_t(\tilde{y},\epsilon,\kappa)$ for each realization of κ. □

Proof of Proposition 2.1. From the proof of Theorem 2.1, we have $C_t(\tilde{y},\epsilon)=\mathbb{E}_\kappa J_t(\tilde{y},\epsilon,\kappa)$. It is sufficient to show that the derivatives of $J_t(\tilde{y},\epsilon,\kappa)$ satisfy the desired properties, as expectation operator preserves them. It is useful to first establish the convexity of $J_t(\tilde{y},\epsilon,\kappa)$. An equivalent formulation for the *capacity allocation problem* expresses all constraints as linear in the state and decision variables.

$$J_t(\tilde{y},\epsilon,\kappa) = \min_{(z_1,z_2)} \Big\{ hz + p[\epsilon - \tilde{y} + z] + \beta V_{t+1}(z) \tag{2.14}$$

$$s.t \quad 0 \le z_i \le \tilde{y}_i \ i = 1,2 \tag{2.15}$$

$$\tilde{y}_i - z_i \le \epsilon_i \ i = 1,2 \tag{2.16}$$

$$z_1 + z_2 \ge \tilde{y}_1 + \tilde{y}_2 - \kappa \Big\} \tag{2.17}$$

Function minimized in (2.14) is convex due to inductional assumptions, the set of constraints, (2.15)–(2.17) is convex due to linearity. Hence, convexity of $J_t(\tilde{y},\epsilon,\kappa)$ follows from Heyman and Sobel (1984), Property B–4, page 525. Next, we investigate the properties of $J_t(\tilde{y},\epsilon,\kappa)$. Due to stationarity of holding and penalty costs, it is always beneficial to satisfy demand, if there exists sufficient capacity. This simplifies the analysis below.

Case 1: $\underline{\epsilon_1 + \epsilon_2 \le \kappa}$. If the total demand does not exceed the available shared capacity, then the demands are satisfied subject to the availability of the inventories of stage-1 products, and the cost that is incurred for any $(\tilde{y}_1,\tilde{y}_2)$ is $J_t(\tilde{y},\epsilon,\kappa) = h(\tilde{y}-\epsilon)^+ + p(\epsilon-\tilde{y})^+ + \beta V_{t+1}(\tilde{y}-\epsilon)^+$.

First, we analyze the first-order derivatives. By Theorem 23.1 of Rockafellar (1997), the one-sided directional partial derivatives of the convex function $J_t(\tilde{y},\epsilon,\kappa)$ exist. Therefore, whenever the first-order derivatives are not continuous, we use the right-hand derivatives below.

$$\begin{aligned}\Delta_{\tilde{y}_i} J_t(\tilde{y},\epsilon,\kappa) &= [h_i + \beta\Delta_{x_i} V_{t+1}(\tilde{y}-\epsilon)^+] 1_{\tilde{y}_i \ge \epsilon_i} - p_i 1_{\tilde{y}_i < \epsilon_i} \\ &\ge -p_i 1_{\tilde{y}_i \ge \epsilon_i} - p_i 1_{\tilde{y}_i < \epsilon_i} \ = \ -p_i\end{aligned}$$

Second, we analyze the second-order derivatives. The function $J_t(\tilde{y},\epsilon,\kappa)$ obviously preserves all the second-order properties that the function $V_{t+1}(x)$ has. This can be verified in a straightforward way through considering four subregions, $\tilde{y}_i \geq (<)\epsilon_i$. For example, for $\tilde{y}_i \geq \epsilon_i$ for $i = 1, 2$, second-order properties hold as $\Delta_{\tilde{y}_i\tilde{y}_i} J_t(\tilde{y}, \epsilon, \kappa) = \beta\Delta_{x_ix_i} V_{t+1}(\tilde{y} - \epsilon) \geq \beta\Delta_{x_1x_2} V_{t+1}(\tilde{y} - \epsilon) = \Delta_{\tilde{y}_1\tilde{y}_2} J_t(\tilde{y}, \epsilon, \kappa) \geq 0$. Other subregions can be analyzed similarly. We also need to verify these properties on the boundaries, $\tilde{y}_1 = \epsilon_1$ and $\tilde{y}_2 = \epsilon_2$. Observe that the first-order derivatives are not continuous on the boundaries. Denoting by Δ^+ the right-hand derivative operator, we instead show that the following properties hold on the boundaries.

(i) $\Delta^+_{\tilde{y}_i} J_t(\tilde{y}, \epsilon, \kappa)$ is (weakly) increasing in $\tilde{y}_1$ and $\tilde{y}_2$.

(ii) $\Delta^+_{\tilde{y}_1} J_t(\tilde{y}, \epsilon, \kappa) - \Delta^+_{\tilde{y}_2} J_t(\tilde{y}, \epsilon, \kappa)$ is (weakly) increasing in $\tilde{y}_1$ and (weakly) decreasing in $\tilde{y}_2$.

Property (i) replicates $\Delta_{\tilde{y}_i\tilde{y}_i} J_t(\tilde{y}, \epsilon, \kappa) \geq 0$ and $\Delta_{\tilde{y}_1\tilde{y}_2} J_t(\tilde{y}, \epsilon, \kappa) \geq 0$, and Property (ii) replicates $\Delta_{\tilde{y}_i\tilde{y}_i} J_t(\tilde{y}, \epsilon, \kappa) \geq \Delta_{\tilde{y}_1\tilde{y}_2} J_t(\tilde{y}, \epsilon, \kappa)$. These properties are sufficient for the preservation of the second-order properties over the boundaries. (For further details, see Shaoxiang, 2004, Proposition 1.)

For conciseness, denote by $J(\tilde{y}_1, \tilde{y}_2) = J_t(\tilde{y}, \epsilon, \kappa)$ and $z = (\tilde{y} - \epsilon)^+$. We have $\Delta^+_{\tilde{y}_i} J(\tilde{y}_1, \tilde{y}_2) = [h_i + \beta\Delta_{x_i} V_{t+1}(z)]1_{\tilde{y}_i \geq \epsilon_i} - p_i 1_{\tilde{y}_i < \epsilon_i}$. $\Delta^+_{\tilde{y}_1} J(\tilde{y}_1, \tilde{y}_2)$ is increasing $\tilde{y}_1$, since $-p_1 \leq h_1 + \beta\Delta_{x_1}(0, z_2)$ on the boundary $\tilde{y}_1 = \epsilon_1$, due to $\mathbb{A}^2_{t+1}$. $\Delta^+_{\tilde{y}_1} J(\tilde{y}_1, \tilde{y}_2)$ is increasing in $\tilde{y}_2$ for $\tilde{y}_1 < \epsilon_1$, as $\Delta^+_{\tilde{y}_1} J(\tilde{y}_1, \tilde{y}_2) = -p_1$ is constant. Let $\tilde{y}_1 \geq \epsilon_1$. In this case, $\Delta^+_{\tilde{y}_1} J(\tilde{y}_1, \tilde{y}_2) = h_1 + \beta\Delta_{x_1}(\tilde{y}_1 - \epsilon_1, 0)$ for $\tilde{y}_2 < \epsilon_2$ and $\Delta^+_{\tilde{y}_1} J(\tilde{y}_1, \tilde{y}_2) = h_1 + \beta\Delta_{x_1}(\tilde{y}_1 - \epsilon_1, \tilde{y}_2 - \epsilon)$ for $\tilde{y}_2 \geq \epsilon_2$ is increasing in $\tilde{y}_2$ since $\Delta_{x_1x_2} V_{t+1}(x) \geq 0$. Similarly, we can show that $\Delta^+_{\tilde{y}_2} J(\tilde{y}_1, \tilde{y}_2)$ is increasing in $\tilde{y}_1$ and $\tilde{y}_2$, which completes property (i).

Next, we prove property (ii).

$$\begin{aligned}
\Delta^+_{\tilde{y}_1} J(\tilde{y}_1, \tilde{y}_2) - \Delta^+_{\tilde{y}_2} J(\tilde{y}_1, \tilde{y}_2) &= [h_1 + \beta\Delta_{x_1} V_{t+1}(z)]1_{\tilde{y}_1 \geq \epsilon_1} - p_1 1_{\tilde{y}_1 < \epsilon_1} \\
&\quad -[h_2 + \beta\Delta_{x_2} V_{t+1}(z)]1_{\tilde{y}_2 \geq \epsilon_2} + p_2 1_{\tilde{y}_2 < \epsilon_2}
\end{aligned}$$

First, consider the boundary $\tilde{y}_1 = \epsilon_1$ with $\tilde{y}_2 < \epsilon_2$. $\Delta^+_{\tilde{y}_1} J(\tilde{y}_1, \tilde{y}_2) - \Delta^+_{\tilde{y}_2} J(\tilde{y}_1, \tilde{y}_2)$ is increasing in $\tilde{y}_1$, since over the boundary of $\tilde{y}_1 = \epsilon_1$, we have $-p_1 \leq h_1 + \beta\Delta_{x_1}(0, 0)$, and decreasing in $\tilde{y}_2$, since $-p_2 \leq h_2 + \beta\Delta_{x_2}(0, 0)$ over the boundary of $\tilde{y}_2 = \epsilon_2$. Next, consider the boundary of $\tilde{y}_1 = \epsilon_1$ with $\tilde{y}_2 \geq \epsilon_2$. In this case, $\Delta^+_{\tilde{y}_1} J(\tilde{y}_1, \tilde{y}_2) - \Delta^+_{\tilde{y}_2} J(\tilde{y}_1, \tilde{y}_2)$ is increasing in $\tilde{y}_1$ and decreasing in $\tilde{y}_2$, since $\Delta_{x_1} V_{t+1}(\tilde{y} - \epsilon) - \Delta_{x_2} V_{t+1}(\tilde{y} - \epsilon)$ is increasing in $\tilde{y}_1$ and decreasing in $\tilde{y}_2$ due to the inductional assumptions. The other case that considers the boundary of $\tilde{y}_2 = \epsilon_2$ can be handled similarly.

Case 2: $\underline{\epsilon_1 + \epsilon_2 > \kappa}$. The shared capacity is a binding resource, provided also that sufficient inventories exist. Figure 2.13 illustrates various regions induced by the optimal policy.

Figure 2.13: Regions Induced by Optimal Policy When $\epsilon_1 + \epsilon_2 > K$

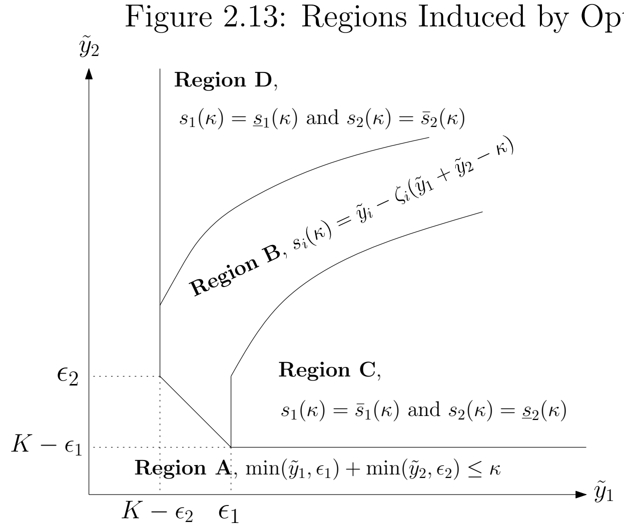

Region A: $\underline{\min(\tilde{y}_1, \epsilon_1) + \min(\tilde{y}_2, \epsilon_2) \leq \kappa}$. In this region, the realized capacity is not a binding resource. Although the total demand exceeds capacity, the amount of inventory does not allow for full utilization of the shared capacity. As a result, the cost that is incurred is $J_t(\tilde{y}, \epsilon, \kappa) = h(\tilde{y} - \epsilon)^+ + p(\epsilon - \tilde{y})^+ + \beta V_{t+1}(\tilde{y} - \epsilon)^+$, satisfying the desired properties, as discussed in Case 1.

Region B: $s_i(\kappa) = \tilde{y}_i - \zeta_i(\tilde{y}_1 + \tilde{y}_2 - \kappa)$. Denote by $\zeta_i = \zeta_i(\tilde{y}_1 + \tilde{y}_2 - \kappa)$. Note that $\zeta_2 = \tilde{y}_1 + \tilde{y}_2 - \kappa - \zeta_1$. As a result, the following holds.

$$\begin{aligned} J_t(\tilde{y}, \epsilon, \kappa) &= h_1\zeta_1 + h_2\zeta_2 + p_1(\epsilon_1 - \tilde{y}_1 + \zeta_1) + p_2(\epsilon_2 - \tilde{y}_2 + \zeta_2) + \beta V_{t+1}(\zeta_1, \zeta_2) \\ &= h_1\zeta_1 + h_2(\tilde{y}_1 + \tilde{y}_2 - \kappa - \zeta_1) + p_1(\epsilon_1 - \tilde{y}_1 + \zeta_1) \\ &\quad + p_2(\epsilon_2 + \tilde{y}_1 + \kappa - \zeta_1) + \beta V_{t+1}(\zeta_1, \tilde{y}_1 + \tilde{y}_2 - \kappa - \zeta_1) \end{aligned}$$

Using the envelope theorem, we have $\Delta_{\tilde{y}_1} J_t(\tilde{y}, \epsilon, \kappa) = -p_1 + p_2 + h_2 + \beta\Delta_{x_2} V_{t+1}(\zeta_1, \zeta_2) \geq -p_1$. By a similar argument, we can show $\Delta_{\tilde{y}_2} J_t(\tilde{y}, \epsilon, \kappa) \geq -p_2$. Second-order conditions are derived as follows.

$$\begin{aligned} \Delta_{\tilde{y}_1\tilde{y}_1} J_t(\tilde{y}, \epsilon, \kappa) &= \beta\Delta_{x_1x_2} V_{t+1}(\zeta_1, \zeta_2)\zeta_1' + \beta\Delta_{x_2x_2} V_{t+1}(\zeta_1, \zeta_2)\zeta_2' \\ \Delta_{\tilde{y}_1\tilde{y}_2} J_t(\tilde{y}, \epsilon, \kappa) &= \beta\Delta_{x_1x_2} V_{t+1}(\zeta_1, \zeta_2)\zeta_1' + \beta\Delta_{x_2x_2} V_{t+1}(\zeta_1, \zeta_2)\zeta_2' \end{aligned}$$

Obviously, $\Delta_{\tilde{y}_1\tilde{y}_1} J_t(\tilde{y}, \epsilon, \kappa) = \Delta_{\tilde{y}_1\tilde{y}_2} J_t(\tilde{y}, \epsilon, \kappa) \geq 0$ due to inductional assumptions on $V_{t+1}(x)$ and the fact that the functions $\zeta_i'(k) \geq 0$. By symmetry, we have $\Delta_{\tilde{y}_2\tilde{y}_2} J_t(\tilde{y}, \epsilon, \kappa) = \Delta_{\tilde{y}_1\tilde{y}_2} J_t(\tilde{y}, \epsilon, \kappa) \geq 0$.

Region C: $s_1(\kappa) = \bar{s}_1(\kappa)$ and $s_2(\kappa) = \underline{s}_2(\kappa)$. First, we prove that $\tilde{y}_1 \geq \min(\epsilon_1, \kappa)$ in this region. The optimal solution satisfies $s_1(\kappa) = \bar{s}_1(\kappa)$ and $s_2(\kappa) = \underline{s}_2(\kappa)$ if and only if $\bar{s}_1(\kappa) < \tilde{y}_1 - \zeta_1(\tilde{y}_1 + \tilde{y}_2 - \kappa)$, since the optimal $s_i(\kappa)$ satisfies $s_i(\kappa) = [\tilde{y}_i - \zeta_i(\tilde{y}_1 + \tilde{y}_2 - \kappa)]|_{[\underline{s}_i(\kappa), \bar{s}_i(\kappa)]}$ for $i = 1, 2$, by Theorem 2.1. To derive a contradiction, assume that $\tilde{y}_1 < \min(\epsilon_1, \kappa)$, in which case $\bar{s}_1(\kappa) = \min(\tilde{y}_1, \epsilon_1, \kappa) = \tilde{y}_1$. Thus, either $\bar{s}_1(\kappa) = \tilde{y}_1 < \tilde{y}_1 - \zeta_1(\tilde{y}_1 + \tilde{y}_2 - \kappa)$ or $\zeta_1(\tilde{y}_1 + \tilde{y}_2 - \kappa) < 0$. This is, however, a contradiction, since $\zeta_1(\omega) \geq 0$ by Theorem 2.1. As a consequence, $s_1(\kappa) = \min(\epsilon_1, \kappa)$ and $s_2(\kappa) = \kappa - \min(\epsilon_1, \kappa)$. Using $s_i = s_i(\kappa)$, we express the cost function as: $J_t(\tilde{y}, \epsilon, \kappa) = h_1(\tilde{y}_1 - s_1) + p_1(\epsilon_1 - s_1) + h_2(\tilde{y}_2 - s_2) + p_2(\epsilon_2 - s_2) + \beta V_{t+1}(\tilde{y} - s)$. It is easy to check that $\Delta_{\tilde{y}_i} J_t(\tilde{y}, \epsilon, \kappa) = h_i + \beta\Delta_{x_i} V_{t+1}(\tilde{y} - s) \geq -p_i$. Obviously, the second-order properties of $J_t(\tilde{y}, \epsilon, \kappa)$ in this region are the same as these of $V_{t+1}(x)$.

Hence, the second-order properties are preserved.

Region D: $\underline{s_1(\kappa) = \underline{s}_1(\kappa) \text{ and } s_2(\kappa) = \bar{s}_2(\kappa)}$. The proof follows the same reasoning as in Region C, except that product indices are switched. Hence, we omit the details.

We have shown that $J_t(\tilde{y}, \epsilon, \kappa)$ is convex, $\Delta_{\tilde{y}_i} J_t(\tilde{y}, \epsilon, \kappa) \geq -p_i$ and satisfies the second-order properties in $(\tilde{y}_1, \tilde{y}_2)$. Since $C_t(\tilde{y}, \epsilon) = \mathbb{E}_\kappa J_t(\tilde{y}, \epsilon, \kappa)$, we have that $C_t(\tilde{y}, \epsilon)$ is also convex, $\Delta_{\tilde{y}_i} C_t(\tilde{y}, \epsilon) \geq -p_i$ and satisfies the second-order properties in $(\tilde{y}_1, \tilde{y}_2)$.

Crossing from Region A to D, A to C, B to A and B to C does not cause problems, since the first derivatives are continuous. We only need to pay attention to the boundary between Regions A and B, however, the proof follows the same logic as in Case 1, by use of assumption $\mathbb{A}^2_{t+1}$. □

Proof of Theorem 2.2. Since $G_t(\tilde{y}) = \mathbb{E}_\epsilon C_t(\tilde{y}, \epsilon)$ and $C_t(\tilde{y}, \epsilon)$ is convex and satisfies the second order properties in $(\tilde{y}_1, \tilde{y}_2)$ due to Proposition 2.1, $G_t(\tilde{y})$ is convex in $(\tilde{y}_1, \tilde{y}_2)$ and it satisfies the second-order properties, as the expectation operator preserves them.

(a) Recall that $\bar{y}_1(y_2) = \arg\min_{y_1 \geq 0} G_t(y)$, hence $\Delta_{y_1} G_t(\bar{y}_1(y_2), y_2) = 0$. Taking the derivative of both sides with respect to y_2, we obtain $\frac{d\bar{y}_1(y_2)}{dy_2} = -\frac{\Delta_{y_1 y_2} G_t(\bar{y}_1(y_2), y_2)}{\Delta_{y_1 y_1} G_t(\bar{y}_1(y_2), y_2)}$. It follows that $-1 \leq \frac{d\bar{y}_1(y_2)}{dy_2} \leq 0$ since $G_t(.)$ satisfies the second-order properties. By a similar argument, $-1 \leq \frac{d\bar{y}_2(y_1)}{dy_1} \leq 0$. *(b)* Recall that the feasible set of targets is $y \geq x$. Also, $\tilde{y} = y \wedge (x + K)$ denotes the inventories for realized capacity levels $K = (K_1, K_2)$. Obviously, $x \leq \tilde{y} \leq y$.

Region 1: $\underline{x_1 \geq \bar{y}_1(x_2) \text{ and } x_2 \geq \bar{y}_2(x_1)}$. For any capacity realization, we have $\tilde{y} \geq x$. As a consequence, $G_t(x_1, x_2) \leq G_t(\tilde{y}_1, x_2) \leq G_t(\tilde{y}_1, \tilde{y}_2)$ holds due to convexity, and the definitions of the curves $\bar{y}_1(.)$ and $\bar{y}_2(.)$. Consequently, $G_t(x_1, x_2) \leq \mathbb{E}_{K_1, K_2} G_t(y_1 \wedge (x_1 + K_1), y_2 \wedge (x_2 + K_2))$ for any $y \geq x$, hence, it is optimal not to produce any of the products.

Region 2: $x_1 < \bar{y}_1(x_2)$ and $x_2 \geq \bar{y}_2(x_1)$. We first show that the optimal inventory targets would be $(\bar{y}_1(x_2), x_2)$ if the dedicated capacity realizations, K_1 and K_2, were known beforehand. First, we show that it is not optimal to produce product 2. Consider any $y_2 > x_2$ with $y_1 \geq x_1$. Since $\bar{y}_2(y_1)$ is non-increasing, it implies that $\bar{y}_2(x_1) \geq \bar{y}_2(\tilde{y}_1)$. And, since $x_2 \geq \bar{y}_2(x_1)$ by assumption, we have $x_2 \geq \bar{y}_2(\tilde{y}_1)$. Since $\tilde{y}_2 \geq x_2 \geq \bar{y}_2(\tilde{y}_1)$, we immediately have $G_t(\tilde{y}_1, x_2) \leq G_t\tilde{y}_1, \tilde{y}_2)$ by the definition of $\bar{y}_2(\tilde{y}_1)$ and due to convexity. Hence, not producing product 2 is optimal. Recall that $\bar{y}_1(x_2)$ is a minimizer of $G_t(y_1, x_2)$ in y_1. Due to convexity of $G_t(.)$, it follows that $\bar{y}_1(x_2)$ is also a minimizer of $G_t(y_1 \wedge (x_1 + K_1), x_2)$. Therefore, $(\bar{y}_1(x_2), x_2)$ is optimal. Note that the optimal inventory targets are independent of the capacity realizations and knowing the capacity levels in advance does not influence the optimal decisions. Therefore, inventory targets $(\bar{y}_1(x_2), x_2)$ are also optimal when capacities are not known *a priori*.

Region 3: $x_1 \geq \bar{y}_1(x_2)$ and $x_2 < \bar{y}_2(x_1)$. We follow the same reasoning as for Region 2.

Region 4: $x_1 \leq \bar{y}_1(x_2)$ and $x_2 \leq \bar{y}_2(x_1)$. We prove that the optimal inventory targets (y_1, y_2) satisfy $\bar{y}_i(y_{3-i}) \leq y_i \leq \bar{y}_i(x_{3-i})$ for each i:

- $y_i \leq \bar{y}_i(x_{3-i})$ for each i: Assume that one of the products, say product 1, violates the condition. Consider any $y_1 > \bar{y}_1(x_2)$ and $y_2 \geq x_2$. Since $\bar{y}_1(y_2)$ is non-increasing, we have $\bar{y}_1(x_2) \geq \bar{y}_1(y_2)$ for all $y_2 \geq x_2$, hence, $y_1 > \bar{y}_1(y_2)$. By the definition of $\bar{y}_1(y_2)$ and due to convexity, revising the production decision by keeping y_2 the same and setting $y_1 = \bar{y}_1(y_2)$, the firm could have incurred a lower cost for any realization of K_1 and K_2. (The details are identical to the one for Region 2.) Therefore, $y_i \leq \bar{y}_i(x_{3-i})$ for each i.

- $\bar{y}_i(y_{3-i}) \leq y_i$ for each i: Assume that and one of the products, say product 1, violates the condition. Consider any $y_1 < \bar{y}_1(y_2)$ and $y_2 \geq x_2$. Since $\bar{y}_1(y_2)$ is non-increasing and $\tilde{y}_2 \leq y_2$, we have $\bar{y}_1(y_2) \leq \bar{y}_1(\tilde{y}_2)$. By the definition of $\bar{y}_1(y_2)$ and due to convexity, revising the production decision by keeping y_2 the same and setting $y_1 = \bar{y}_1(y_2)$, the firm could have incurred a lower cost for any realization of K_1 and K_2 (as argued for Region 2). Therefore, $\bar{y}_i(y_{3-i}) \leq y_i$ for each i.

The property above, derived for Region 4, has the following two consequences.

Consequence 1: Assume that producing product j is not optimal, $y_j = x_j$. Since $\bar{y}_i(y_j) \leq y_i \leq \bar{y}_i(x_j)$, the expression reduces to $y_i = \bar{y}_i(x_i)$.

Consequence 2: If $x_i < y_i^0$ for each i, then it is optimal to produce both products: $y_i > x_i$. To reach a contradiction, assume that it is not optimal to produce product 1. Then, by *Consequence 1*, inventory target for product 2 is $\bar{y}_2(x_1)$. However, a better solution can be obtained by setting $y_1 = \bar{y}_1(\bar{y}_2(x_1))$ and keeping $y_2 = \bar{y}_2(x_1)$, as shown for Region 4. Therefore, both products are produced in an optimal policy.

Consequence 2 always holds when $G_t(y)$ is strictly convex. If $G_t(y)$ is not strictly convex, there could be multiple optimal solutions, however, we focus on the one that satisfies *Consequence 2*. We next present an exact characterization of the optimal policy in Region 4. Define $EG(y;x) = \mathbb{E}_{K_1,K_2} G_t(y \wedge (x + K))$. The objective is to minimize $EG(y;x)$ in y for a given x with $y \geq x$. $EG(y;x)$ is not necessarily convex in y, however, as we will show, the first-order conditions are sufficient for optimality. Denote by $f_i(.)$ and $F_i(.)$ the probability density function and the cumulative distribution function of the product i's capacity. We derive the first-order conditions as follows.

$$
\begin{aligned}
\Delta_{y_1} EG(y;x) &= [1 - F_1(y_1 - x_1)]\mathbb{E}_{K_2} \Delta_{y_1} G_t(y_1, y_2 \wedge (x_2 + K_2)) \\
\Delta_{y_2} EG(y;x) &= [1 - F_2(y_2 - x_2)]\mathbb{E}_{K_1} \Delta_{y_2} G_t(y_1 \wedge (x_1 + K_1), y_2)
\end{aligned}
$$

The following two steps show that $EG(y;x)$ is unimodal in y_1 for a given y_2 (and vice versa).

Step A1: The function $G_t(y_1, y_2 \wedge (x_2 + K_2))$ is convex in y_1 for any K_2 due to the convexity of $G_t(.)$. Hence, $\mathbb{E}_{K_2} G_t(y_1, y_2 \wedge (x_2 + K_2))$ is convex in y_1. Define $\check{y}_1(y_2; x) = \arg\min_{y_1 \geq 0} \mathbb{E}_{K_2} G_t(y_1, y_2 \wedge (x_2 + K_2))$. For convenience, we treat x as a constant and write $\check{y}_2(y_1)$ instead of $\check{y}_2(y_1; x)$.

Step A2: For any K_1, $\check{y}_1(y_2)$ is also a minimizer of $\mathbb{E}_{K_2} G_t(y_1 \wedge (x_1 + K_1), y_2 \wedge (x_2 + K_2))$. Specifically, $\mathbb{E}_{K_2} G_t(y_1 \wedge (x_1 + K_1), y_2 \wedge (x_2 + K_2))$ is unimodal in y_1 and it is minimized at $y_1 = \check{y}_1(y_2)$ for any K_1. For $y_1 \leq \check{y}_1(y_2)$, the function $\mathbb{E}_{K_2} G_t(y_1 \wedge (x_1 + K_1), y_2 \wedge (x_2 + K_2))$ is convex and decreasing in y_1. And for $y_1 \geq \check{y}_1(y_2)$, it is increasing (not necessarily convex). Thus, $EG(y; x) = \mathbb{E}_{K_1, K_2} G_t(y_1 \wedge (x_1 + K_1), y_2 \wedge (x_2 + K_2))$ is also unimodal in y_1 and it is minimized at $y_1 = \check{y}_1(y_2)$.

Steps A1–A2 show that $EG(y; x)$ is unimodal in y_i for given y_{3-i} and x. To establish the optimal solution, we show that $\check{y}_i(y_{3-i}) \geq \bar{y}_i(y_{3-i})$ for $i = 1, 2$. Without loss of generality, consider $i = 1$.

$$
\begin{aligned}
\Delta_{y_1} EG(\bar{y}_1(y_2), y_2; x) &= [1 - F_1(\bar{y}_1(y_2) - x_1)] \mathbb{E}_{K_2} \Delta_{y_1} G_t(\bar{y}_1(y_2), y_2 \wedge (x_2 + K_2)) \\
&\leq [1 - F_1(\bar{y}_1(y_2) - x_1)] \mathbb{E}_{K_2} \Delta_{y_1} G_t(\bar{y}_1(y_2), y_2) \\
&= 0 = \Delta_{y_1} EG(\check{y}_1(y_2), y_2; x),
\end{aligned}
$$

where the inequality is due to second-order properties of G. Hence, $\check{y}_1(y_2) \geq \bar{y}_1(y_2)$. For $y_2 \leq x_2$, the inequality holds with equality. Recall that x satisfies $x_i \leq \bar{y}_i(x_{3-i})$ for each i (in Region 4). Consequently, the solution to the first-order conditions, denoted by $\hat{y} = \hat{y}(x)$ (the intersection of the curves $y_i = \check{y}_i(y_{3-i})$) is such that it is either feasible (i.e. $\hat{y} \geq x$), or $\hat{y}_i \geq x_i$ and $\hat{y}_{3-i} < x_{3-i}$ for some i. In other words, the first-order conditions never imply $\hat{y} < x$. Otherwise, a contradiction is reached, in which x would not belong to Region 4 using the fact that $\check{y}_i(y_{3-i}) \geq \bar{y}_i(y_{3-i})$. Consider three subcases.

Case 1, $\hat{y} \geq x$: First, consider the case where the solution to the first-order conditions is feasible. We will show that $\Delta_{y_iy_i}EG(\hat{y};x) \geq \Delta_{y_1y_2}EG(\hat{y};x) \geq 0$ for $i = 1, 2$. In other words, whenever the first-order conditions are satisfied, then second-order properties are also satisfied. Obviously, second-order properties imply that the Hessian is positive, further implying the sufficiency of the first-order conditions for optimality. We obtain the second-order conditions as follows.

$$\begin{aligned}
\Delta_{y_1y_1}EG(y;x) &= [1 - F_1(y_1 - x_1)]\mathbb{E}_{K_2}\Delta_{y_1y_1}G_t(y_1, y_2 \wedge (x_2 + K_2)) \\
&\quad -f_1(y_1 - x_1)\mathbb{E}_{K_2}\Delta_{y_1}G_t(y_1, y_2 \wedge (x_2 + K_2)) \\
\Delta_{y_1y_2}EG(y;x) &= [1 - F_1(y_1 - x_1)][1 - F_2(y_2 - x_2)]\Delta_{y_1y_2}G_t(y_1, y_2)
\end{aligned}$$

$\Delta_{y_1y_2}EG(y;x) \geq 0$ since $\Delta_{y_1y_2}G_t(y_1, y_2) \geq 0$. When the first-order conditions are satisfied, $\Delta_{y_1y_1}EG(\hat{y};x)$ reduces to $\Delta_{y_1y_1}EG(\hat{y};x) = [1-F_1(\hat{y}_1-x_1)]\mathbb{E}_{K_2}\Delta_{y_1y_1}G_t(\hat{y}_1, \hat{y}_2\wedge (x_2 + K_2))$. Thus,

$$\begin{aligned}
\mathbb{E}_{K_2}\Delta_{y_1y_1}G_t(\hat{y}_1, \hat{y}_2 \wedge (x_2 + K_2)) &= \int_0^{\hat{y}_2-x_2} \Delta_{y_1y_1}G_t(\hat{y}_1, x_2 + K_2)f_2(k_2)dk_2 \\
&\quad + \int_{\hat{y}_2-x_2}^{\infty} \Delta_{y_1y_1}G_t(\hat{y}_1, \hat{y}_2)f_2(k_2)dk_2 \\
&\geq \int_{\hat{y}_2-x_2}^{\infty} \Delta_{y_1y_1}G_t(\hat{y}_1, \hat{y}_2)f_2(k_2)dk_2 \\
&= [1 - F_2(\hat{y}_2 - x_2)]\Delta_{y_1y_1}G_t(\hat{y}_1, \hat{y}_2)
\end{aligned}$$

Multiplying both sides of the inequality by $1-F_1(\hat{y}_1-x_1)$, we obtain $\Delta_{y_1y_1}EG(\hat{y};x) \geq \Delta_{y_1y_2}EG(\hat{y};x)$.Similarly, $\Delta_{y_2y_2}EG(\hat{y};x) \geq \Delta_{y_1y_2}EG(\hat{y};x)$. As a result, the Hessian is positive when the first-order conditions are satisfied, hence, $y^*(x) = \hat{y}(x)$.

Case 2, $\hat{y}_1 \geq x_1$, $\hat{y}_2 < x_2$: We have either $x_1 > y_1^0$ or $x_2 > y_2^0$, since otherwise, producing both products would be optimal, and the first-order conditions would be feasible (*Consequence 2*). Since $\hat{y}_2 < x_2$, we have $x_2 > y_2^0$ (it can be shown by contradiction). From Step A1, $\mathbb{E}_{K_1}\Delta_{y_2}G_t(y_1\wedge(x_1+K_1), \check{y}_2) = 0$. Taking the derivative of both sides with respect to y_1, we obtain the following.

$$\frac{d\check{y}_2}{dy_1} = -\frac{[1-F_1(y_1-x_1)]\Delta_{y_1y_2}G_t(y_1,\check{y}_2)}{[1-F_1(y_1-x_1)]\Delta_{y_1y_2}G_t(y_1,\check{y}_2)+\int_0^{y_1-x_1}\Delta_{y_2y_2}G_t(x_1+K_1,\check{y}_2)f_1(k_1)dk_1}$$

Due to the second-order properties of $G_t(y)$, $-1 \le \frac{d\check{y}_2}{dy_1} \le 0$. Similarly, $-1 \le \frac{d\check{y}_1}{dy_2} \le 0$. Figure 2.14 illustrates the shaded region, where the optimal inventory targets lie. The curves reflect the relationship $\check{y}_i(y_{3-i}) \ge \bar{y}_i(y_{3-i})$, and that for $y_{3-i} \le x_{3-i}$, the inequality becomes an equality.

Figure 2.14: Optimal Solution in Region 4, $\hat{y}_1 \ge x_1$, $\hat{y}_2 < x_2$

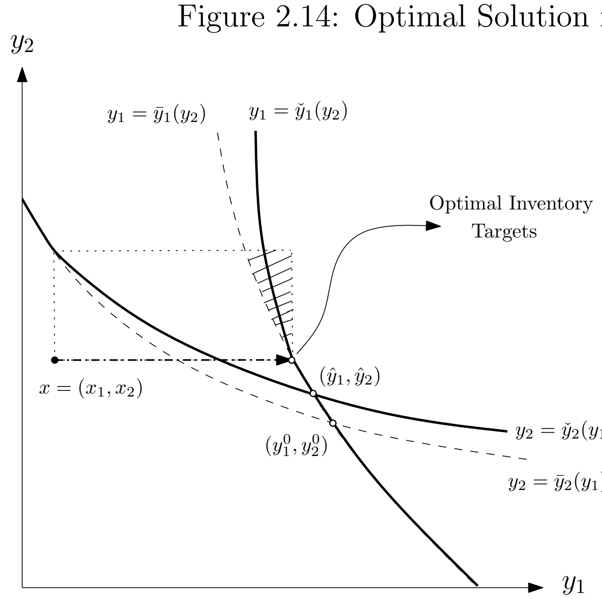

We show that it is optimal not to produce product 2. To reach a contradiction, consider any (y_1, y_2) (from the shaded region), satisfying $y_2 > x_2$ with $\bar{y}_2(y_1) \le y_2 \le \bar{y}_2(x_1)$, and assume that it is optimal. Due to unimodality of $EG(y;x)$ in y_1, we must have $y_1 = \check{y}_1(y_2)$ for optimality. In this case, revising the target for product 2 as $\max(\check{y}_2(y_1), x_2)$ improves the objective function due to unimodality of $EG(y;x)$ in y_2. Thus, we are able to construct a better solution, which is in contradiction with the optimality of (y_1, y_2). Hence, not ordering product 2 is optimal. As a result, the product 1 target is set to be $\bar{y}_1(x_2)$ by *Consequence 1.*

Case 3, $\hat{y}_1 < x_1$, $\hat{y}_2 \geq x_2$: Similar analysis for Case 2 applies. □

Proof of Proposition 2.2. Assume first that (x_1, x_2) belongs to the region, where it is optimal to produce both products. In this case, the optimal targets (y_1^*, y_2^*) satisfy the first-order conditions: $\mathbb{E}_{K_2}\Delta_{y_1}G_t(y_1^*, y_2^* \wedge (x_2 + K_2)) = 0$ and $\mathbb{E}_{K_1}\Delta_{y_2}G_t(y_1^* \wedge (x_1 + K_1), y_2^*) = 0$. Taking the derivatives of both sides of the two equalities with respect to x_1 and x_2, we obtain

$$\begin{aligned} 0 &= \mathbb{E}_{K_2}\Delta_{y_1y_1}G_t(y_1^*, y_2^* \wedge (x_2 + K_2))\frac{\partial y_1^*}{\partial x_1} + [1 - F_2(y_2^* - x_2)]\Delta_{y_1y_2}G_t(y_1^*, y_2^*)\frac{\partial y_2^*}{\partial x_1} \\ &:= a_{11}\frac{\partial y_1^*}{\partial x_1} + a_{12}\frac{\partial y_2^*}{\partial x_1} \end{aligned}$$

$$\begin{aligned} 0 &= [1 - F_1(y_1^* - x_1)]\Delta_{y_1y_2}G_t(y_1^*, y_2^*)\frac{\partial y_1^*}{\partial x_1} + \mathbb{E}_{K_1}\Delta_{y_2y_2}G_t(y_1^* \wedge (x_1 + K_1), y_2^*)\frac{\partial y_2^*}{\partial x_1} \\ &\quad + \int_0^{y_1^* - x_1}\Delta_{y_1y_2}G_t(x_1 + k_1, y_2^*)f_1(k_1)dk_1 \;:=\; a_{21}\frac{\partial y_1^*}{\partial x_1} + a_{22}\frac{\partial y_2^*}{\partial x_1} + b_1 \end{aligned}$$

$$\begin{aligned} 0 &= \mathbb{E}_{K_2}\Delta_{y_1y_1}G_t(y_1^*, y_2^* \wedge (x_2 + K_2))\frac{\partial y_1^*}{\partial x_2} + [1 - F_2(y_2^* - x_2)]\Delta_{y_1y_2}G_t(y_1^*, y_2^*)\frac{\partial y_2^*}{\partial x_2} \\ &\quad + \int_0^{y_2^* - x_2}\Delta_{y_1y_2}G_t(y_1^*, x_2 + k_2)f_2(k_2)dk_2 \;:=\; a_{11}\frac{\partial y_1^*}{\partial x_2} + a_{12}\frac{\partial y_2^*}{\partial x_2} + b_2 \end{aligned}$$

$$\begin{aligned} 0 &= [1 - F_1(y_1^* - x_1)]\Delta_{y_1y_2}G_t(y_1^*, y_2^*)\frac{\partial y_1^*}{\partial x_2} + \mathbb{E}_{K_1}\Delta_{y_2y_2}G_t(y_1^* \wedge (x_1 + K_1), y_2^*)\frac{\partial y_2^*}{\partial x_2} \\ &:= a_{21}\frac{\partial y_1^*}{\partial x_2} + a_{22}\frac{\partial y_2^*}{\partial x_2} \end{aligned}$$

We have four equations with four unknowns. The solution of the above set of equations is straightforward, and given by $\frac{\partial y_1^*}{\partial x_1} = \frac{a_{12}b_1}{a_{11}a_{22} - a_{12}a_{21}}$, $\frac{\partial y_2^*}{\partial x_1} = -\frac{a_{11}b_1}{a_{11}a_{22} - a_{12}a_{21}}$, $\frac{\partial y_1^*}{\partial x_2} = -\frac{a_{22}b_2}{a_{11}a_{22} - a_{12}a_{21}}$, and $\frac{\partial y_2^*}{\partial x_2} = \frac{a_{21}b_2}{a_{11}a_{22} - a_{12}a_{21}}$. To establish the monotonicity properties, we prove $a_{22} \geq a_{21} + b_1$.

$$\begin{aligned} a_{22} &= \mathbb{E}_{K_1} \Delta_{y_2 y_2} G_t(y_1^* \wedge (x_1 + K_1), y_2^*) \\ &= \int_{y_1^*-x_1}^{\infty} \Delta_{y_2 y_2} G_t(y_1^*, y_2^*) f_1(k_1) dk_1 + \int_0^{y_1^*-x_1} \Delta_{y_2 y_2} G_t(x_1 + k_1, y_2^*) f_1(k_1) dk_1 \\ &= [1 - F_1(y_1^* - x_1)] \Delta_{y_2 y_2} G_t(y_1^*, y_2^*) + b_1 \\ &\geq [1 - F_1(y_1^* - x_1)] \Delta_{y_1 y_2} G_t(y_1^*, y_2^*) + b_1 \; = \; a_{21} + b_1 \end{aligned}$$

The inequality follows from the second-order properties. Thus, we have $0 \leq b_1 \leq a_{22} - a_{21}$. In the same manner, we have $0 \leq b_2 \leq a_{11} - a_{12}$ by symmetry of the argument. Obviously, $a_{11} \geq a_{12} \geq 0$ and $a_{22} \geq a_{21} \geq 0$. Therefore, $\frac{\partial y_1^*}{\partial x_1} \geq 0$, $\frac{\partial y_2^*}{\partial x_1} \leq 0$, $\frac{\partial y_1^*}{\partial x_2} \leq 0$, and $\frac{\partial y_2^*}{\partial x_2} \geq 0$ follow immediately. Thus,

$$\begin{aligned} \frac{\partial y_1^*}{\partial x_1} &= \frac{a_{12} b_1}{a_{11} a_{22} - a_{12} a_{21}} \leq \frac{a_{12}(a_{22} - a_{21})}{a_{11} a_{22} - a_{12} a_{21}} \leq \frac{a_{11}(a_{22} - a_{21})}{a_{11} a_{22} - a_{12} a_{21}} \leq 1 \\ \frac{\partial y_2^*}{\partial x_1} &= -\frac{a_{11} b_1}{a_{11} a_{22} - a_{12} a_{21}} \geq -\frac{a_{11}(a_{22} - a_{21})}{a_{11} a_{22} - a_{12} a_{21}} \geq -1 \end{aligned}$$

$$\frac{\partial y_1^*}{\partial x_1} - \frac{\partial y_1^*}{\partial x_2} - 1 = \frac{a_{22}(b_2 - a_{11}) + a_{12}(b_1 + a_{21})}{a_{11} a_{22} - a_{12} a_{21}} \leq 0$$

The last inequality is due to the fact that $b_2 - a_{11} \leq -a_{12}$ and $b_1 + a_{21} \leq a_{22}$. Using a symmetric argument, we have $\frac{\partial y_1^*}{\partial x_2} \geq -1$, $\frac{\partial y_2^*}{\partial x_2} \leq 1$ and $\frac{\partial y_2^*}{\partial x_2} - \frac{\partial y_2^*}{\partial x_1} \leq 1$.

Assume now that (x_1, x_2) belongs to the region, where it is optimal to order only one product, say product 1. In this case, $y_1^* = \bar{y}_1(x_2)$ and $y_2^* = x_2$. From Theorem 2.2, part (a), we have $-1 \leq \frac{d\bar{y}_1(y_2)}{dy_2} \leq 0$. Hence, $-1 \leq \frac{\partial y_1^*}{\partial x_2} \leq 0$ and $\frac{\partial y_1^*}{\partial x_1} = 0$. In addition, $\frac{\partial y_2^*}{\partial x_1} = 0$ and $\frac{\partial y_2^*}{\partial x_2} = 1$. Hence, the desired monotonicity results hold. In the region, where it is not optimal to order, $y_i^* = x_i$, the desired monotonicity results hold trivially: $\frac{\partial y_i^*}{\partial x_i} = 1$ and $\frac{\partial y_i^*}{\partial x_{3-i}} = 0$. □

Proof of Proposition 2.3. Since $G_t(\tilde{y}) = \mathbb{E}_\epsilon C_t(\tilde{y}, \epsilon)$, we have that $\Delta_{\tilde{y}_i} G_t(\tilde{y}) \geq -p_i$ and $G_t(\tilde{y})$ satisfies the second order properties in $(\tilde{y}_1, \tilde{y}_2)$. In all of the cases considered below, it is easy to show that $\Delta_{x_i} V_t(x) \geq -p_i$, and as a consequence, $p_i + h_i + \beta\Delta_{x_i} V_t(x) \geq p_i + h_i - \beta p_i = (1-\beta)p_i + h_i \geq 0$, as desired. Therefore, we focus on the second-order properties. Consider (x_1, x_2) in each of the regions.

• Region 1: In this case, $V_t(x_1, x_2) = G_t(x_1, x_2)$, hence it satisfies the second-order properties.

• Region 2: In this case, $V_t(x_1, x_2) = \mathbb{E}_{K_1} G_t(\bar{y}_1(x_2) \wedge (x_1 + K_1), x_2)$. To show that $V_t(x_1, x_2)$ satisfies the second-order properties in this region, it suffices to show that the function $G_t(\bar{y}_1(x_2) \wedge (x_1 + K_1), x_2)$ satisfies the second-order properties for any given K_1. When that is the case, expectation with respect to K_1 preserves the second-order properties.

When $x_1 + K_1 \leq \bar{y}_1(x_2)$, we have $G_t(\bar{y}_1(x_2) \wedge (x_1 + K_1), x_2) = G_t(x_1 + K_1, x_2)$, which preserves the second-order properties. When $x_1 + K_1 > \bar{y}_1(x_2)$, we have $G_t(\bar{y}_1(x_2) \wedge (x_1 + K_1), x_2) = G_t(\bar{y}_1(x_2), x_2)$. Below, we show that the function $G_t(\bar{y}_1(x_2), x_2)$ satisfies the second-order properties in (x_1, x_2). Obviously, $\frac{\partial G_t(\bar{y}_1(x_2), x_2)}{\partial x_1} = 0$, hence, $\frac{\partial^2 G_t(\bar{y}_1(x_2), x_2)}{\partial x_1^2} = \frac{\partial^2 G_t(\bar{y}_1(x_2), x_2)}{\partial x_1 \partial x_2} = 0$. The first-order derivative with respect to x_2 is obtained using the envelope theorem. The second-order derivative is obtained using the expression for $\bar{y}_1'(x_2)$ from Theorem 2.2, part(a).

$$\begin{aligned}
\frac{\partial G_t(\bar{y}_1(x_2), x_2)}{\partial x_2} &= \Delta_{y_2} G_t(\bar{y}_1(x_2), x_2) \\
\frac{\partial^2 G_t(\bar{y}_1(x_2), x_2)}{\partial x_2^2} &= \Delta_{y_1 y_2} G_t(\bar{y}_1(x_2), x_2)\bar{y}_1'(x_2) + \Delta_{y_2 y_2} G_t(\bar{y}_1(x_2), x_2) \\
&= \frac{(\Delta_{y_1 y_1}\Delta_{y_2 y_2} - \Delta^2_{y_1 y_2}) G_t(\bar{y}_1(x_2), x_2)}{\Delta_{y_1 y_1} G_t(\bar{y}_1(x_2), x_2)} \geq 0
\end{aligned}$$

Therefore, $V_t(x_1, x_2)$ satisfies the second-order properties in Region 2 (hence in Region 3).

• Region 4: We first derive the first-order derivative of the optimal cost function. The envelope theorem is useful in obtaining the following result.

$$\begin{aligned} V_t(x_1,x_2) &= \mathbb{E}_{K_1,K_2} G_t(y_1^* \wedge (x_1+K_1), y_2^* \wedge (x_2+K_2)) \\ \Delta_{x_1} V_t(x_1,x_2) &= \int_0^{y_1^*-x_1} \mathbb{E}_{K_2} \Delta_{y_1} G_t(x_1+k_1, y_2^* \wedge (x_2+K_2)) f_1(k_1) dk_1 \end{aligned}$$

Next, we derive the second-order derivatives.

$$\begin{aligned} \Delta_{x_1x_1} V_t(x_1,x_2) &= \int_0^{y_1^*-x_1} \int_0^{y_2^*-x_2} \Delta_{y_1y_1} G_t(x_1+k_1, x_2+k_2) f_1(k_1) f_2(k_2) dk_2 dk_1 \\ &\quad + \int_0^{y_1^*-x_1} \int_{y_2^*-x_2}^{\infty} \Delta_{y_1y_1} G_t(x_1+k_1, y_2^*) f_1(k_1) f_2(k_2) dk_2 dk_1 \\ &\quad + \left[\int_0^{y_1^*-x_1} \int_{y_2^*-x_2}^{\infty} \Delta_{y_1y_2} G_t(x_1+k_1, y_2^*) f_1(k_1) f_2(k_2) dk_2 dk_1 \right] \frac{\partial y_2^*}{\partial x_1} \end{aligned}$$

$$\begin{aligned} \Delta_{x_1x_2} V_t(x_1,x_2) &= \int_0^{y_1^*-x_1} \int_0^{y_2^*-x_2} \Delta_{y_1y_2} G_t(x_1+k_1, x_2+k_2) f_1(k_1) f_2(k_2) dk_2 dk_1 \\ &\quad + \left[\int_0^{y_1^*-x_1} \int_{y_2^*-x_2}^{\infty} \Delta_{y_1y_2} G_t(x_1+k_1, y_2^*) f_1(k_1) f_2(k_2) dk_2 dk_1 \right] \frac{\partial y_2^*}{\partial x_2} \end{aligned}$$

$\Delta_{x_1x_2} V_t(x_1,x_2) \geq 0$ due to the second-order properties of $G_t(.)$ and $\frac{\partial y_2^*}{\partial x_2} \geq 0$, as shown in Proposition 2.2. Next, we prove that $\Delta_{x_1x_1} V_t(x_1,x_2) \geq \Delta_{x_1x_2} V_t(x_1,x_2)$. From Property (iii) of, Proposition 2.2 we have $\frac{\partial y_2^*}{\partial x_1} - \frac{\partial y_2^*}{\partial x_2} \geq -1$. Thus,

$$\begin{aligned} & (\Delta_{x_1x_1} - \Delta_{x_1x_2}) V_t(x_1,x_2) \\ = & \int_0^{y_1^*-x_1} \int_0^{y_2^*-x_2} (\Delta_{y_1y_1} - \Delta_{y_1y_2}) G_t(x_1+k_1, x_2+k_2) f_1(k_1) f_2(k_2) dk_2 dk_1 \\ & + \int_0^{y_1^*-x_1} \int_{y_2^*-x_2}^{\infty} \Delta_{y_1y_1} G_t(x_1+k_1, y_2^*) f_1(k_1) f_2(k_2) dk_2 dk_1 \\ & + \left[\int_0^{y_1^*-x_1} \int_{y_2^*-x_2}^{\infty} \Delta_{y_1y_2} G_t(x_1+k_1, y_2^*) f_1(k_1) f_2(k_2) dk_2 dk_1 \right] \left(\frac{\partial y_2^*}{\partial x_1} - \frac{\partial y_2^*}{\partial x_2} \right) \geq 0 \end{aligned}$$

By a similar argument, we have $\Delta_{x_2x_2}V_t(x_1,x_2) \geq \Delta_{x_1x_2}V_t(x_1,x_2) \geq 0$, hence, $V_t(x_1,x_2)$ satisfies the second-order properties in this region. Finally, crossing from one region to another does not cause any problem since the first derivatives are continuous. □

Proof of Proposition 2.4. We use Theorem 2.2 and explicitly derive the optimal targets in each region. We show that the optimal targets (y_1^*, y_2^*) from Theorem 2.2 and the targets $(y_1(\tau_2(x_2)), y_2(\tau_1(x_1)))$, claimed in this proposition, either coincide, or lead to the same production decisions. That is, if it is optimal to order product i, $y_i^* > x_i$, then $y_i^* \wedge (x_i + K_i) = \bar{y}_i(\tau_{3-i}(x_{3-i})) \wedge (x_i + K_i)$. Otherwise, if ordering product i is not optimal, $y_i^* = x_i$, then $\bar{y}_i(\tau_{3-i}(x_{3-i})) \leq x_i$. Recall that the functions $\bar{y}_i(.)$ are non-increasing with $-1 \leq \bar{y}_i'(.) \leq 0$ from Theorem 2.2, which is used below.

Region 1: $x_1 \geq \bar{y}_1(x_2)$ and $x_2 \geq \bar{y}_2(x_1)$. From Theorem 2.2, it is not optimal to produce any of the products. Below, we show for product 1 that the initial inventory x_1 exceeds the target defined by $\bar{y}_1(\tau_2(x_2))$ in Region 1. (Similar argument applies to product 2, hence omitted.)

- If $\tau_2(x_2) = x_2 + K_2$, then, $\bar{y}_1(\tau_2(x_2)) = \bar{y}_1(x_2 + K_2) \leq \bar{y}_1(x_2) \leq x_1$.
- If $\tau_2(x_2) = y_2^0 (< x_2)$, then $x_1 > y_1^0$ in Region 1, because otherwise $x_2 > y_2^0$, and $\tau_2(x_2) = x_2$. Thus, $\bar{y}_1(\tau_2(x_2)) = \bar{y}_1(y_2^0) = y_1^0 < x_1$.
- Finally, if $\tau_2(x_2) = x_2$, then, $\bar{y}_1(\tau_2(x_2)) = \bar{y}_1(x_2) \leq x_1$.

Region 2: $x_1 < \bar{y}_1(x_2)$ and $x_2 \geq \bar{y}_2(x_1)$. From Theorem 2.2, it is not optimal to produce product 2, while product 1 target equals $\bar{y}_1(x_2)$. In Region 2, we have $x_1 \leq y_1^0$ and $x_2 \geq y_2^0$. So, $\tau_1(x_1)$ equals either y_1^0 or $x_1 + K_1$, while $\tau_2(x_2) = x_2$. As a result, $\bar{y}_1(\tau_2(x_2)) = \bar{y}_1(x_2)$. Product 1 target is indeed as given by Theorem 2.2. For product 2, consider the following two cases.

- If $\tau_1(x_1) = x_1 + K_1$, then, $\bar{y}_2(\tau_1(x_1)) = \bar{y}_2(x_1 + K_1) \leq \bar{y}_2(x_1) \leq x_2$.
- If $\tau_1(x_1) = y_1^0$, then, $\bar{y}_2(\tau_1(x_1)) = \bar{y}_2(y_1^0) = y_2^0 \leq x_2$, likewise.

Region 3: $\underline{x_1 \geq \bar{y}_1(x_2) \text{ and } x_2 < \bar{y}_2(x_1)}$. Same as Region 2.

Region 4: $\underline{x_1 \leq \bar{y}_1(x_2) \text{ and } x_2 \leq \bar{y}_2(x_1)}$. Recall that $\Delta_{y_1} G_t(\bar{y}_1(y_2), y_2) = 0$ and $\Delta_{y_2} G_t(y_1, \bar{y}_2(y_1)) = 0$ by definition. Also, $\bar{y}_1(y_2^0) = y_1^0$ and $\bar{y}_2(y_1^0) = y_2^0$. From Theorem 2.2, targets satisfy: $\Delta_{y_1} G_t(y_1, y_2 \wedge (x_2 + K_2)) = 0$ and $\Delta_{y_2} G_t(y_1 \wedge (x_1 + K_1), y_2) = 0$.

• Let $\tau_i(x_i) = y_i^0$ for each i, i.e., $y_i^0 - K_i \leq x_i \leq y_i^0$. Then, $\bar{y}_i(\tau_{3-i}(x_{3-i})) = \bar{y}_i(y_{3-i}^0) = y_i^0$. Clearly, $(y_1, y_2) = (y_1^0, y_2^0)$, satisfies the first order conditions: $\Delta_{y_1} G_t(y_1^0, y_2^0 \wedge (x_2 + K_2)) = \Delta_{y_1} G_t(y_1^0, y_2^0) = 0$ and $\Delta_{y_2} G_t(y_1^0 \wedge (x_1 + K_1), y_2^0) = \Delta_{y_2} G_t(y_1^0, y_2^0) = 0$, hence it must be optimal.

• Let $\tau_i(x_i) = x_i + K_i$ for each i, i.e., $x_i + K_i < y_i^0$. We have $\bar{y}_1(x_2 + K_2) \geq \bar{y}_1(y_2^0) = y_1^0 \geq x_1 + K_1$, since $\bar{y}_1(.)$ is non-increasing. Similarly, $\bar{y}_2(x_1 + K_1) \geq x_2 + K_2$. For $(y_1, y_2) = (\bar{y}_1(x_2 + K_2), \bar{y}_2(x_1 + K_1))$, first-order conditions are satisfied: $\Delta_{y_1} G_t(\bar{y}_1(x_2 + K_2), \bar{y}_2(x_1 + K_1) \wedge (x_2 + K_2)) = \Delta_{y_1} G_t(\bar{y}_1(x_2 + K_2), x_2 + K_2) = 0$ and $\Delta_{y_2} G_t(\bar{y}_1(x_2 + K_2) \wedge (x_1 + K_1), \bar{y}_2(x_1 + K_1)) = \Delta_{y_2} G_t(x_1 + K_1, \bar{y}_2(x_1 + K_1)) = 0$.

• Let $\tau_1(x_1) = x_1 + K_1$ and $\tau_2(x_2) = y_2^0$. Then, $\bar{y}_1(\tau_2(x_2)) = \bar{y}_1(y_2^0) = y_1^0$ and $\bar{y}_2(\tau_1(x_1)) = \bar{y}_2(x_1 + K_1)$. Define $z_2 = \bar{y}_2(x_1 + K_1) \wedge (x_2 + K_2)$ and $\acute{y}_1 = \bar{y}_1(z_2)$. Below, we show that $(y_1, y_2) = (\acute{y}_1, \bar{y}_2(x_1 + K_1))$ satisfies first-order conditions, and $\acute{y}_1 \geq x_1 + K_1$ (i.e., $\acute{y}_1$ is not reachable), establishing the optimality of targets $\bar{y}_i(\tau_{3-i}(x_{3-i}))$. Indeed, since $\acute{y}_1$ is not reachable, it can be replaced by $\bar{y}_1(\tau_2(x_2))$, which is also not reachable ($\bar{y}_1(\tau_2(x_2)) = y_1^0 \geq x_1 + K_1$).

First, we prove that $\acute{y}_1 \geq x_1 + K_1$. Since $x_1 + K_1 \leq y_1^0$, and $z_2 \leq \bar{y}_2(x_1 + K_1)$ by definition, the point $(x_1 + K_1, z_2)$ belongs to Region 4. From the definition of Region 4, $\acute{y}_1 = \bar{y}_1(z_2) \geq x_1 + K_1$. The first-order conditions are satisfied: $\Delta_{y_1} G_t(\acute{y}_1, \bar{y}_2(x_1 + K_1) \wedge (x_2 + K_2)) = \Delta_{y_1} G_t(\bar{y}_1(z_2), z_2) = 0$, and $\Delta_{y_2} G_t(\acute{y}_1 \wedge (x_1 + K_1), \bar{y}_2(x_1 + K_1)) = \Delta_{y_2} G_t(x_1 + K_1, \bar{y}_2(x_1 + K_1)) = 0$.

Since $\acute{y}_1 \geq x_1 + K_1$, product 1 target is not reachable. And, since $y_1^0 \geq x_1 + K_1$, replacing product 1 target with y_1^0 does not alter the ending inventories, hence it is optimal. In addition, we have $\bar{y}_2(\tau_1(x_1)) = \bar{y}_2(x_1 + K_1)$ as desired.

• The case when $\tau_1(x_2) = x_2 + K_2$ and $\tau_1(x_1) = y_1^0$ is similar as above.

• Let $\tau_i(x_i) = x_i$ for some i, the only case not analyzed for Region 4. In Region 4, we cannot have $\tau_i(x_i) = x_i$ for both i. Without loss of generality, we analyze the case when $\tau_2(x_2) = x_2$, i.e., $x_2 > y_2^0$, in which case, we must have $x_1 \leq y_1^0$ (otherwise, (x_1, x_2) would fall into Region 1). Hence, $\tau_1(x_1)$ equals either $x_1 + K_1$ or y_1^0.

∘ Start with $\tau_1(x_1) = y_1^0$. First, we derive the optimal targets based on Theorem 2.2. We can show that (y_1^0, y_2^0) is a solution to the first-order conditions, but it is not a feasible production decision, since $x_2 > y_2^0$. Using Theorem 2.2, product 2 is not produced, while product 1 target is $\bar{y}_1(x_2)$. Comparing to the targets given by $\bar{y}_i(\tau_{3-i}(x_{3-i}))$, we have $\bar{y}_1(\tau_2(x_2)) = \bar{y}_1(x_2)$ as desired, and $\bar{y}_2(\tau_1(x_1)) = \bar{y}_2(y_1^0) = y_2^0 < x_2$, which is not feasible, as desired.

∘ Next, let $\tau_1(x_1) = x_1 + K_1$. Define $z_2 = \bar{y}_2(x_1 + K_1) \wedge (x_2 + K_2)$ and $\acute{y}_1 = \bar{y}_1(z_2)$. We show that $(y_1, y_2) = (\acute{y}_1, \bar{y}_2(x_1 + K_1))$ satisfies the first-order conditions. Since $x_1 + K_1 \leq y_1^0$ and $z_2 \leq \bar{y}_2(x_1 + K_1)$ by definition, the point $(x_1 + K_1, z_2)$ belongs to Region 4. Thus, it must satisfy $\acute{y}_1 = \bar{y}_1(z_2) \geq x_1 + K_1$. The first-order conditions are satisfied: $\Delta_{y_1} G_t(\acute{y}_1, \bar{y}_2(x_1+K_1) \wedge (x_2+K_2)) = \Delta_{y_1} G_t(\bar{y}_1(z_2), z_2) = 0$ and $\Delta_{y_2} G_t(\acute{y}_1 \wedge (x_1+K_1), \bar{y}_2(x_1+K_1)) = \Delta_{y_2} G_t(x_1+K_1, \bar{y}_2(x_1+K_1)) = 0$. Using Theorem 2.2, the optimal targets are as follows. If $\bar{y}_2(x_1 + K_1) \geq x_2$, then the optimal target is $(\acute{y}_1, \bar{y}_2(x_1 + K_1))$. Otherwise, the product 2 target is not feasible, hence, product 1 target must be $\bar{y}_1(x_2)$, while product 2 should not be produced.

We now compare the optimal targets from Theorem 2.2 to the targets given by $\bar{y}_i(\tau_{3-i}(x_{3-i}))$. Consider first the case that product 2 is not produced, $\bar{y}_2(x_1 + K_1) \leq x_2$. Since $\tau_2(x_2) = x_2$, product 1 target is $\bar{y}_1(\tau_2(x_2)) = \bar{y}_1(x_2)$ as desired. Because $\tau_1(x_1) = x_1 + K_1$, the product 2 target satisfies $\bar{y}_2(\tau_1(x_1)) = \bar{y}_2(x_1 + K_1) \leq x_2$, as desired. Consider next the case of $\bar{y}_2(x_1 + K_1) \geq x_2$, that is, the optimal targets are $(\acute{y}_1, \bar{y}_2(x_1 + K_1))$. In this case, $\bar{y}_2(\tau_1(x_1)) = \bar{y}_2(x_1 + K_1)$, as desired. For product 1, since $\acute{y}_1 \geq x_1 + K_1$, it suffices to show that $\bar{y}_1(\tau_2(x_2)) = \bar{y}_1(x_2) \geq x_1 + K_1$. First, notice that $\bar{y}_2(\bar{y}_1(x_2)) \leq x_2$ when $x_2 \geq y_2^0$ due to the monotonic behavior of the curves $\bar{y}_i(.)$ (follows from a simple tatonnement argument). Since $\bar{y}_2(x_1+K_1) \geq x_2$, we have $\bar{y}_2(\bar{y}_1(x_2)) \leq \bar{y}_2(x_1+K_1)$. Since $\bar{y}_2(.)$ is non-increasing, $\bar{y}_1(x_2) \geq x_1 + K_1$, as desired. □

Proof of Proposition 2.5. The proof is for product 1. Refer to Theorem 2.2 for the description of the inventory policy. The optimal targets do not depend on dedicated capacities in Regions 1, 2, and 3. Hence, the optimal policy remains the same. For initial inventories in Region 4, we inspect how the minimizers of function $EG(y;x)$ change. Obviously, the minimizer $\check{y}_1(y_2;x)$ is not influenced by the capacity increase in product 1, since $\mathbb{E}_{K_2}\Delta_{y_1}G_t(\check{y}_1(y_2;x), y_2 \wedge (x_2+K_2)) = 0$. Below, we argue that the curve $\check{y}_2(y_1;x)$ shifts downwards, as shown in Figure 2.15.

Figure 2.15: Sensitivity of the Production Targets to K_1

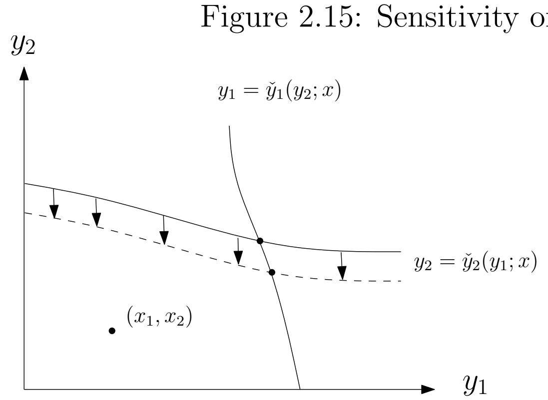

The function $G_t(y_1 \wedge (x_1+K_1), y_2)$ is supermodular in (y_2, K_1). Hence, the value of $\mathbb{E}_{K_1}\Delta_{y_2}G_t(y_1 \wedge (x_1+K_1), y_2)$ is increased due to Lemma 2.1. This implies that the minimizer $\check{y}_2(y_1;x)$ decreases, that is, the curve $\check{y}_2(y_1;x)$ shifts downwards. Recall that the curves $\check{y}_1(y_2;x)$ and $\check{y}_2(y_1;x)$ have negative slopes slopes, greater than or equal to -1. Hence, the point of intersection of these curves shifts such that product 1 target is increased and product target 2 is decreased. □

Proof of Proposition 2.6. We show that $\Delta_{\tilde{y}_i\kappa}J_t(\tilde{y},\epsilon,\kappa) \leq 0$ and $(\Delta_{\tilde{y}_j\kappa}\Delta_{\tilde{y}_1\tilde{y}_2} - \Delta_{\tilde{y}_j\tilde{y}_j}\Delta_{\tilde{y}_i\kappa})J_t(\tilde{y},\epsilon,\kappa) \geq 0$ for $i = 1, 2$ and $j = 3 - i$. We consider only $i = 1$ and the results are valid for $i = 2$. We follow the same sequence as in the proof of Proposition 2.1.

Region A: $\underline{\min(\tilde{y}_1, \epsilon_1) + \min(\tilde{y}_2, \epsilon_2) \le \kappa}$. In this region, we have $\Delta_{\tilde{y}_i \kappa} J_t(\tilde{y}, \epsilon, \kappa) = 0$ for $i = 1, 2$. Hence, the desired properties hold by equality.

Region B: $\underline{s_i(\kappa) = \tilde{y}_i - \zeta_i(\tilde{y}_1 + \tilde{y}_2 - \kappa)}$.

$$
\begin{aligned}
\Delta_{\tilde{y}_1 \kappa} J_t(\tilde{y}, \epsilon, \kappa) &= -\beta \Delta_{x_1 x_2} V_{t+1}(\zeta_1, \zeta_2) \zeta_1' - \beta \Delta_{x_2 x_2} V_{t+1}(\zeta_1, \zeta_2) \zeta_2' \\
\Delta_{\tilde{y}_2 \kappa} J_t(\tilde{y}, \epsilon, \kappa) &= -\beta \Delta_{x_1 x_2} V_{t+1}(\zeta_1, \zeta_2) \zeta_1' - \beta \Delta_{x_2 x_2} V_{t+1}(\zeta_1, \zeta_2) \zeta_2' \\
\Delta_{\tilde{y}_2 \tilde{y}_2} J_t(\tilde{y}, \epsilon, \kappa) &= \beta \Delta_{x_1 x_2} V_{t+1}(\zeta_1, \zeta_2) \zeta_1' + \beta \Delta_{x_2 x_2} V_{t+1}(\zeta_1, \zeta_2) \zeta_2' \\
\Delta_{\tilde{y}_1 \tilde{y}_2} J_t(\tilde{y}, \epsilon, \kappa) &= \beta \Delta_{x_1 x_2} V_{t+1}(\zeta_1, \zeta_2) \zeta_1' + \beta \Delta_{x_2 x_2} V_{t+1}(\zeta_1, \zeta_2) \zeta_2'
\end{aligned}
$$

Because of the inductional assumptions, we have $\Delta_{\tilde{y}_1 \kappa} J_t(\tilde{y}, \epsilon, \kappa) \le 0$, and due to above expressions, we have $(\Delta_{\tilde{y}_2 \kappa} \Delta_{\tilde{y}_1 \tilde{y}_2} - \Delta_{\tilde{y}_2 \tilde{y}_2} \Delta_{\tilde{y}_1 \kappa}) J_t(\tilde{y}, \epsilon, \kappa) = 0$.

Region C: $\underline{s_1(\kappa) = \bar{s}_1(\kappa) \text{ and } s_2(\kappa) = \underline{s}_2(\kappa)}$. We consider two subregions.

Region C1: First, let $\epsilon_1 \le \tilde{y}_1$.

$$
\begin{aligned}
\Delta_{\tilde{y}_1 \kappa} J_t(\tilde{y}, \epsilon, \kappa) &= -\beta \Delta_{x_1 x_2} V_{t+1}(\tilde{y}_1 - \epsilon_1, \tilde{y}_2 + \epsilon_1 - \kappa) \\
\Delta_{\tilde{y}_2 \kappa} J_t(\tilde{y}, \epsilon, \kappa) &= -\beta \Delta_{x_2 x_2} V_{t+1}(\tilde{y}_1 - \epsilon_1, \tilde{y}_2 + \epsilon_1 - \kappa) \\
\Delta_{\tilde{y}_2 \tilde{y}_2} J_t(\tilde{y}, \epsilon, \kappa) &= \beta \Delta_{x_2 x_2} V_{t+1}(\tilde{y}_1 - \epsilon_1, \tilde{y}_2 + \epsilon_1 - \kappa) \\
\Delta_{\tilde{y}_1 \tilde{y}_2} J_t(\tilde{y}, \epsilon, \kappa) &= \beta \Delta_{x_1 x_2} V_{t+1}(\tilde{y}_1 - \epsilon_1, \tilde{y}_2 + \epsilon_1 - \kappa)
\end{aligned}
$$

Because of the inductional assumptions, we have $\Delta_{\tilde{y}_1 \kappa} J_t(\tilde{y}, \epsilon, \kappa) \le 0$, and from the above expressions, we have $(\Delta_{\tilde{y}_2 \kappa} \Delta_{\tilde{y}_1 \tilde{y}_2} - \Delta_{\tilde{y}_2 \tilde{y}_2} \Delta_{\tilde{y}_1 \kappa}) J_t(\tilde{y}, \epsilon, \kappa) = 0$.

Region C2: Next, let $\epsilon_1 > \tilde{y}_1$.

$$\begin{aligned}
\Delta_{\tilde{y}_1\kappa} J_t(\tilde{y},\epsilon,\kappa) &= -\beta\Delta_{x_2x_2}V_{t+1}(0,\tilde{y}_1+\tilde{y}_2-\kappa)\\
\Delta_{\tilde{y}_2\kappa} J_t(\tilde{y},\epsilon,\kappa) &= -\beta\Delta_{x_2x_2}V_{t+1}(0,\tilde{y}_1+\tilde{y}_2-\kappa)\\
\Delta_{\tilde{y}_2\tilde{y}_2} J_t(\tilde{y},\epsilon,\kappa) &= \beta\Delta_{x_2x_2}V_{t+1}(0,\tilde{y}_1+\tilde{y}_2-\kappa)\\
\Delta_{\tilde{y}_1\tilde{y}_2} J_t(\tilde{y},\epsilon,\kappa) &= \beta\Delta_{x_2x_2}V_{t+1}(0,\tilde{y}_1+\tilde{y}_2-\kappa)
\end{aligned}$$

Because of the inductional assumptions, we have $\Delta_{\tilde{y}_1\kappa} J_t(\tilde{y},\epsilon,\kappa) \le 0$, and due to above expressions, we have $(\Delta_{\tilde{y}_2\kappa}\Delta_{\tilde{y}_1\tilde{y}_2} - \Delta_{\tilde{y}_2\tilde{y}_2}\Delta_{\tilde{y}_1\kappa})J_t(\tilde{y},\epsilon,\kappa) = 0$.

Region D: $s_1(\kappa) = \underline{s}_1(\kappa)$ and $s_2(\kappa) = \bar{s}_2(\kappa)$. Same as Region C.

Now we are ready to evaluate the effect of increase in the shared capacity. First, assume that the shared capacity is uncertain. From Theorem 2.1, $C_t(\tilde{y},\epsilon) = \mathbb{E}_\kappa J_t(\tilde{y},\epsilon,\kappa)$. Using Lemma 2.1 (stated in Appendix 2.11), the derivatives of $C_t(\tilde{y},\epsilon)$ with respect to $\tilde{y}_1$ and $\tilde{y}_2$ are decreased when the shared capacity is stochastically increased. Since $G_t(\tilde{y}) = \mathbb{E}_\epsilon C_t(\tilde{y},\epsilon)$, the derivatives of $G_t(\tilde{y})$ with respect to $\tilde{y}_1$ and $\tilde{y}_2$ decrease, as well. As a result, the minimizers of $G_t(\tilde{y})$ are increased, effectively increasing the sum of the inventory targets by Proposition 2.10. When the shared capacity is deterministic, $G_t(\tilde{y}) = \mathbb{E}_\epsilon J_t(\tilde{y},\epsilon,\kappa)$, where κ is now a fixed parameter. Since $\Delta_{\tilde{y}_i\kappa} J_t(\tilde{y},\epsilon,\kappa) \le 0$ and $(\Delta_{\tilde{y}_j\kappa}\Delta_{\tilde{y}_1\tilde{y}_2} - \Delta_{\tilde{y}_j\tilde{y}_j}\Delta_{\tilde{y}_i\kappa})J_t(\tilde{y},\epsilon,\kappa) \ge 0$ for $i = 1,2$ and $j = 3-i$, the targets are increased for both products due to Proposition 2.10. □

Proof of Proposition 2.7. We only prove for product 1. Following the same procedure in Proposition 2.6, we can verify that $\Delta_{\tilde{y}_1\epsilon_1} J_t(\tilde{y},\epsilon,\kappa) \le 0$, $\Delta_{\tilde{y}_1\epsilon_1} J_t(\tilde{y},\epsilon,\kappa) \le \Delta_{\tilde{y}_2\epsilon_1} J_t(\tilde{y},\epsilon,\kappa)$, $(\Delta_{\tilde{y}_2\epsilon_1}\Delta_{\tilde{y}_1\tilde{y}_2} - \Delta_{\tilde{y}_2\tilde{y}_2}\Delta_{\tilde{y}_1\epsilon_1})J_t(\tilde{y},\epsilon,\kappa) \ge 0$ and $(\Delta_{\tilde{y}_1\epsilon_1}\Delta_{\tilde{y}_1\tilde{y}_2} - \Delta_{\tilde{y}_1\tilde{y}_1}\Delta_{\tilde{y}_2\epsilon_1})$ $J_t(\tilde{y},\epsilon,\kappa) \le 0$. Recall that $G_t(\tilde{y}) = \mathbb{E}_{\epsilon,\kappa} J_t(\tilde{y},\epsilon,\kappa)$. First observe that the minimizers of $G_t(.)$ behave as follows. As ϵ_1 is stochastically increased, $\bar{y}_1(\tilde{y}_2)$ increases. As a consequence, in Region 2, where only product 1 is produced, product 1 target increases, while the product 2 target remains the same. In Region 3, product 1 target remains the same. To analyze Region 4, we need to consider global minimizers (y_1^0, y_2^0). Due to Proposition 2.10, part *(iii)*, y_1^0 increases, while y_2^0 decreases. □

Proof of Proposition 2.8. We only consider sensitivity to p_1. Sensitivity to p_2, h_1 and h_2 are similar. Following the same procedure in Proposition 2.6, we can verify that $\Delta_{\tilde{y}_1 p_1} J_t(\tilde{y}, \epsilon, \kappa) \leq 0$ and $\Delta_{\tilde{y}_2 p_1} J_t(\tilde{y}, \epsilon, \kappa) \geq 0$. Since $G_t(\tilde{y}) = \mathbb{E}_{\epsilon,\kappa} J_t(\tilde{y}, \epsilon, \kappa)$, and due to Proposition 2.10, we conclude that product 1 target increases, while the product 2 target decreases. □

2.10. Appendix: Model Extensions

In this section, we provide model extensions to address infinite-horizon and Markov-modulated settings.

Infinite Time Horizon: Assuming that all parameters are stationary, and with a discount factor of $0 \leq \beta < 1$, the optimality equations are stated below.

$$\textbf{Phase One :} \quad V^*(x) = \min_{y \geq x} \mathbb{E}_{K,\epsilon} C^*(y \wedge (x+K), \epsilon), \tag{2.18}$$

$$\textbf{Phase Two :} \quad C^*(y,\epsilon) = \min_{s_i(.)} \Big\{ \mathbb{E}_\kappa \{ h[y - s(\kappa)] + p[\epsilon - s(\kappa)] + \beta V^*(y - s(\kappa)) \} \tag{2.19}$$

$$s.t \quad \underline{s}_i(\kappa) \leq s_i(\kappa) \leq \bar{s}_i(\kappa), \tag{2.20}$$

$$s_1'(\kappa) + s_2'(\kappa) \leq 1, \tag{2.21}$$

$$s_i'(\kappa) \geq 0 \ ; \quad \forall \kappa \geq 0, \ i = 1,2 \Big\} \tag{2.22}$$

The following proposition establishes the infinite-horizon solution and shows that the optimal policy structure is the same as its finite-horizon counterpart.

Proposition 2.9. *There exists a function* $V^*(x) = \lim_{t \to \infty} V_t(x)$*, which solves (2.18)-(2.22) and satisfies the second-order properties, and* $p_i + h_i + \beta \Delta_{x_i} V^*(x) \geq 0$ *for* $i = 1, 2$.

Proof of Proposition 2.9. Consider the following reformulation based on

Theorem 2.1 and Proposition 2.1. Notice that periods are now indexed backwards in time, and $V_0(x) = x$.

$$\textbf{Stage One :} \qquad V_t(x) = \min_{y \geq x} \mathbb{E}_{K,\kappa,\epsilon} J_t(y \wedge (x+K), \epsilon, \kappa), \tag{2.23}$$

$$\textbf{Stage Two :} \quad J_t(\tilde{y}, \epsilon, \kappa) = \min_{(z_1,z_2)} \Big\{ hz + p[\epsilon - \tilde{y} + z] + \beta V_{t-1}(z) \tag{2.24}$$

$$s.t \quad 0 \leq z_i \leq \tilde{y}_i \;\; i = 1,2 \tag{2.25}$$

$$\tilde{y}_i - z_i \leq \epsilon_i \;\; i = 1,2 \tag{2.26}$$

$$z_1 + z_2 \geq \tilde{y}_1 + \tilde{y}_2 - \kappa \Big\} \tag{2.27}$$

We focus our discussion on the infinite-horizon version of (2.23)–(2.27). To establish the infinite-horizon solution, we verify the conditions (a)–(d) of Theorem 8-14 in Heyman and Sobel (1984). Conditions (b) and (d) are straightforward. Condition (b) holds, since single-period costs are non-negative for any initial state and any feasible decision. Condition (d) holds, since continuity is a consequence of convexity. To show that condition (a) holds, it suffices to show that $V_t(x)$ is bounded above by a function $B(x)$ for all t, as $V_t(x) \leq V_{t+1}(x)$ due to $V_0(x) = 0$. This is straightforward by considering a suboptimal policy, which sets $y = x$ and $s_i(\kappa) = 0$ at the current period t and onwards, (i.e., no production and no calibration of final products). Define $B_t(x)$ as the cost of such policy. Define $B(x) = \lim_{t\to\infty} B_t(x) = \frac{1}{1-\beta}\Big(h_1x_1 + h_2x_2 + p_1\mathbb{E}[\epsilon_1] + p_2\mathbb{E}[\epsilon_2]\Big)$. Obviously, $V_t(x) \leq B_t(x) \leq B(x)$, hence condition (a) holds.

To show condition (c), we show that the set of feasible decisions is compact. This is obvious for stage two decisions, as the linear constraints define a closed and bounded region. For stage one, there is no loss of optimality by replacing $y \geq x$ with $x \leq y \leq \max(U, x)$ with $U_i < \infty$ sufficiently large (see Example 8–33 in Heyman and Sobel, 1984). Therefore, compactness of the feasible set in stage one can be claimed without loss of optimality, hence condition (c) holds.

To show that $V^*(x)$ satisfies the second-order properties, we follow a similar procedure as in Hu et al. (2008) and restate the second-order conditions as follows: for any $\epsilon > 0$, $\Delta_{x_1} V^*(x_1+\epsilon, x_2) \geq \Delta_{x_1} V^*(x_1, x_2+\epsilon)$ and $\Delta_{x_2} V^*(x_1, x_2+\epsilon) \geq \Delta_{x_2} V^*(x_1+\epsilon, x_2)$. Since $V_t(x)$ satisfies the second-order conditions, it suffices to show that $\Delta_{x_i} V_t(x)$ converges to $\Delta_{x_i} V^*(x)$. From Theorem 8-14 of Heyman and Sobel (1984), $V_t(x)$ converges uniformly to $V^*(x)$, hence $\Delta_{x_i} V^*(x) = \lim_{t\to\infty} \Delta_{x_i} V_t(x)$. □

Proposition 2.9 implies that it is possible to solve a finite-horizon problem for a sufficiently large horizon length and obtain the numerical solution to an infinite-horizon problem in a practical manner. Figure 2.16 illustrates convergence of the optimal production targets with initial inventory (0,0), as a test case.

Figure 2.16: Production targets increase in the length of horizon

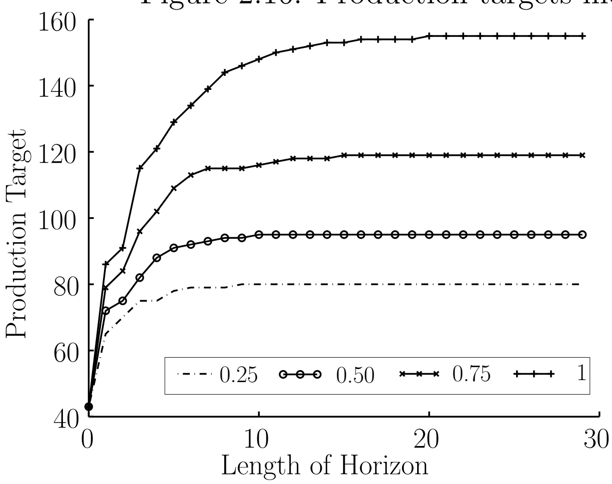

Note: Identical products configuration with $p = 0.99$ and $h = 0.01$, $\mathbb{E}[\epsilon_i] = 24$, $cv(\epsilon_i) = 0.8$, $\mathbb{E}[\kappa] = 60$, $cv(\kappa) = 0.50$, $\mathbb{E}[K_i] = 24$, legend shows different values for $cv(K_i)$.

Markov-Modulated Process: Our model assumed that demands and capacities are independent across periods. This assumption can be relaxed employing a Markov chain driving the demand and capacity distributions. This enables the original model to capture the real-life situations involving for example, seasonal demand, or the possibility that capacity loss in a given period is an indicator of potential capacity loss in the next few periods to come.

Suppose that there are N states of the world with a probability transition matrix $P = [p_{\omega,\omega'}]$. When the state of the world is ω, dedicated capacity distributions are

$K_\omega = (K_{1\omega}, K_{2\omega})$, shared capacity distribution is κ_ω and the demand distributions are given by $\epsilon_\omega = (\epsilon_{1\omega}, \epsilon_{2\omega})$. With this, the model is reformulated below.

$$\textbf{Phase One :} \quad V_t(x;\omega) = \min_{y \geq x} \mathbb{E}_{K_\omega, \epsilon_\omega} C_t(y \wedge (x + K_\omega), \epsilon_\omega; \omega),$$

$$\begin{aligned}\textbf{Phase Two :} \quad C_t(y,\epsilon;\omega) &= \min_{s_i(.)} \Big\{ \mathbb{E}_{\kappa_\omega, \omega'} \{ h[y - s(\kappa_\omega)] + p[\epsilon - s(\kappa_\omega)] \\ &\qquad + \beta V_{t+1}(y - s(\kappa_\omega); \omega') | \omega \} \\ &\qquad s.t \quad \underline{s}_i(\kappa) \leq s_i(\kappa) \leq \bar{s}_i(\kappa), \\ &\qquad\qquad s_1'(\kappa) + s_2'(\kappa) \leq 1, \\ &\qquad\qquad s_i'(\kappa) \geq 0 \ ; \qquad \forall \kappa \geq 0, \ i = 1, 2 \Big\}\end{aligned}$$

The revised formulation has the same optimal policy structure, however, the optimal policy itself depends on the state of the world ω. For example, if one production process i has lost full capacity for a number of periods, then product $3 - i$ follows a modified base-stock policy, which is a function of the existing stock of product i.

2.11. Appendix: Technical Results Needed for Sensitivity Analysis

Our sensitivity results depend on the following technical results below. The first result is a special case of Theorem 3.10.1 of Topkis (1978).

Lemma 2.1 (Topkis)**.** *For a submodular function $\phi(x, \epsilon)$ (i.e., $\Delta_{x\epsilon}\phi(x, \epsilon) \leq 0$), let $x^* = \arg\min_x \mathbb{E}[\phi(x, \epsilon)]$. If ϵ is increased stochastically, then x^* increases.*

Lemma 2.1 establishes the monotonicity of the minimizer of $\mathbb{E}[\phi(x, \epsilon)]$. Analogous conclusions can be drawn when $\phi(.)$ is supermodular. Next, we consider optimizing a function in two variables.

Proposition 2.10. *Let function $\phi(y; \gamma)$ satisfy the second-order properties in $y = (y_1, y_2)$ (Definition 2.2) and let $(y_1^*, y_2^*) = \arg\min_{y_1, y_2} \mathbb{E}_\gamma[\phi(y; \gamma)]$. If γ increases stochastically, then:*

i) y_i^* *increases if* $\Delta_{y_i\gamma}\phi(y;\gamma) \leq 0$ *and* $\Delta_{y_i\gamma}\phi(y;\gamma) \leq \Delta_{y_{3-i}\gamma}\phi(y;\gamma)$.

ii) y_i^* *increases and* y_{3-i}^* *decreases if* $\Delta_{y_i\gamma}\phi(y;\gamma) \leq 0 \leq \Delta_{y_{3-i}\gamma}\phi(y;\gamma)$.

iii) $y_1^* + y_2^*$ *increases if* $\Delta_{y_i\gamma}\phi(y;\gamma) \leq 0$ *for each* $i = 1, 2$. *Both* y_1^* *and* y_2^* *increase if, in addition,* $(\Delta_{y_{3-i}\gamma}\Delta_{y_1y_2} - \Delta_{y_{3-i}y_{3-i}}\Delta_{y_i\gamma})\phi(y;\gamma) \geq 0$ *for each* $i = 1, 2$.

Proof of Proposition 2.10. Before we prove the main result of this proposition, consider a function $\psi(y;\lambda)$ that satisfies the second-order properties in (y_1, y_2). Let $\hat{y}(\lambda) = \arg\min_y \psi(y;\lambda)$. This is the deterministic counter-part of our problem. $\hat{y}(\lambda)$ satisfies the following set of equations.

$$\Delta_{y_1}\psi(\hat{y}(\lambda);\lambda) = 0 \tag{2.28}$$

$$\Delta_{y_2}\psi(\hat{y}(\lambda);\lambda) = 0 \tag{2.29}$$

Taking the derivatives of the equations (2.28) and (2.29) with respect to λ and rearranging the terms, we obtain the following expressions.

$$\hat{y}_i'(\lambda) = \left(\frac{\Delta_{y_{3-i}\lambda}\Delta_{y_1y_2} - \Delta_{y_{3-i}y_{3-i}}\Delta_{y_i\lambda}}{\Delta_{y_1y_1}\Delta_{y_2y_2} - \Delta_{y_1y_2}^2}\right)\psi(\hat{y}(\lambda);\lambda) \tag{2.30}$$

$$\hat{y}_1'(\lambda) + \hat{y}_2'(\lambda) = -\left(\frac{\Delta_{y_1\lambda}(\Delta_{y_2y_2} - \Delta_{y_1y_2}) + \Delta_{y_2\lambda}(\Delta_{y_1y_1} - \Delta_{y_1y_2})}{\Delta_{y_1y_1}\Delta_{y_2y_2} - \Delta_{y_1y_2}^2}\right)\psi(\hat{y}(\lambda);\lambda) \tag{2.31}$$

Using expressions (2.30) and (2.31), as well as the second-order properties, we can readily verify properties 1–3 for $\psi(y;\lambda)$ below, which are the deterministic counterparts of properties (i)–(iii).

1. If $\Delta_{y_i\lambda}\psi(y;\lambda) \leq 0$ and $\Delta_{y_i\lambda}\psi(y;\lambda) \leq \Delta_{y_{3-i}\lambda}\psi(y;\lambda)$, then $\hat{y}_i'(\lambda) \geq 0$.

2. If $\Delta_{y_i\lambda}\psi(y;\lambda) \leq 0 \leq \Delta_{y_{3-i}\lambda}\psi(y;\lambda)$, then $\hat{y}_i'(\lambda) \geq 0$ and $\hat{y}_{3-i}'(\lambda) \leq 0$.

3. If $\Delta_{y_i\lambda}\psi(y;\lambda) \leq 0$ for each $i = 1, 2$, then $\hat{y}_1'(\lambda) + \hat{y}_2'(\lambda) \geq 0$. If, in addition,

$(\Delta_{y_{3-i}\lambda}\Delta_{y_1 y_2} - \Delta_{y_{3-i}y_{3-i}}\Delta_{y_i\lambda})\psi(y;\lambda) \geq 0$ for each $i = 1, 2$, then $\hat{y}_i'(\lambda) \geq 0$ for both $i = 1, 2$.

Consider now the original optimization problem $(y_1^*, y_2^*) = \arg\min_{y_1,y_2} \mathbb{E}_\gamma[\phi(y;\gamma)]$. Consider two random variables with $\gamma_1 \leq_{s.t} \gamma_2$. Denote by $F_1(.)$ and $F_2(.)$ the cumulative probability distribution functions of γ_1 and γ_2 respectively. Due to stochastic dominance, $F_1^{-1}(u) \leq F_2^{-1}(u)$ for all $0 \leq u \leq 1$.

For a constant parameter $0 \leq \lambda \leq 1$, define a random variable $\gamma(\lambda)$ such that its inverse cumulative distribution function is given by $F^{-1}(u;\lambda) = (1-\lambda)F_1^{-1}(u) + \lambda F_2^{-1}(u)$ for all $0 \leq u \leq 1$. As a result, the inverse functions satisfy $F_1^{-1}(u) \leq F^{-1}(u;\lambda) \leq F_2^{-1}(u)$ for all $0 \leq u, \lambda \leq 1$. In other words, $\gamma_1 \leq_{s.t} \gamma(\lambda) \leq_{s.t} \gamma_2$. Also note that $\gamma(\lambda_1) \leq_{s.t} \gamma(\lambda_2)$ for $0 \leq \lambda_1 \leq \lambda_2 \leq 1$.

To show that the monotonicity properties hold when γ is stochastically increased from $\gamma = \gamma_1$ to $\gamma = \gamma_2$, it suffices to show that the monotonicity results hold for $\gamma = \gamma(\lambda)$ when λ is increased marginally. For that purpose, define $\psi(y;\lambda) = \mathbb{E}_{\gamma(\lambda)}[\phi(y;\gamma(\lambda))]$. First, observe that $\psi(y;\lambda)$ satisfies the second-order properties. Therefore, we can use the results obtained above.

We prove only the second part of property (iii), which is the least straightforward case. All other properties can be established similarly. We show that $(\Delta_{y_2\lambda}\Delta_{y_1y_2} - \Delta_{y_2y_2}\Delta_{y_1\lambda})\psi(y;\lambda) \geq 0$, hence $\hat{y}_1'(\lambda) \geq 0$. First, we derive $\Delta_\lambda\psi(y;\lambda)$ as follows. (Below, we change the variable of integration. We let $x = F^{-1}(u;\lambda)$.)

$$\begin{aligned}
\psi(y;\lambda) &= \mathbb{E}_{\gamma(\lambda)}[\phi(y;\gamma(\lambda))] = \int \phi(y;x)dF(x;\lambda) \\
&= \int_0^1 \phi(y;F^{-1}(u;\lambda))du \\
&= \int_0^1 \phi(y;(1-\lambda)F_1^{-1}(u) + \lambda F_2^{-1}(u))du
\end{aligned}$$

$$\Delta_\lambda\psi(y;\lambda) = \int_0^1 \Delta_\gamma\phi(y;F^{-1}(u;\lambda))[F_2^{-1}(u;\lambda) - F_1^{-1}(u;\lambda)]du$$

The above expression is useful in deriving $\Delta_{y_i\lambda}\psi(y;\lambda)$ for each i. As a result, we can immediately derive $(\Delta_{y_2\lambda}\Delta_{y_1y_2} - \Delta_{y_2y_2}\Delta_{y_1\lambda})\psi(y;\lambda)$ as follows.

$$\begin{aligned}(\Delta_{y_2\lambda}\Delta_{y_1y_2} - \Delta_{y_2y_2}\Delta_{y_1\lambda})\psi(y;\lambda) &= \int_0^1 \Big(\Delta_{y_2\gamma}\phi(y;F^{-1}(u;\lambda))\Delta_{y_1y_2}\mathbb{E}_{\gamma(\lambda)}[\phi(y;\gamma(\lambda))] \\ &\quad - \Delta_{y_1\gamma}\phi(y;F^{-1}(u;\lambda))\Delta_{y_2y_2}\mathbb{E}_{\gamma(\lambda)}[\phi(y;\gamma(\lambda))]\Big) \ldots \\ &\quad [F_2^{-1}(u;\lambda) - F_1^{-1}(u;\lambda)]du\end{aligned}$$

To show that $(\Delta_{y_2\lambda}\Delta_{y_1y_2} - \Delta_{y_2y_2}\Delta_{y_1\lambda})\psi(y;\lambda) \geq 0$, it suffices to show that the statement $\Delta_{y_2\gamma}\phi(y;\gamma)\Delta_{y_1y_2}\mathbb{E}_{\gamma(\lambda)}[\phi(y;\gamma(\lambda))] - \Delta_{y_1\gamma}\phi(y;\gamma)\Delta_{y_2y_2}\mathbb{E}_{\gamma(\lambda)}[\phi(y;\gamma(\lambda))] \geq 0$ holds. But, this immediately follows, since $(\Delta_{y_2\gamma}\Delta_{y_1y_2} - \Delta_{y_2y_2}\Delta_{y_1\gamma})\phi(y;\gamma) \geq 0$ and since $F_2^{-1}(u;\lambda) \geq F_1^{-1}(u;\lambda)$. $\square$

Chapter 3

Strategic Behavior of Suppliers in the Face of Production Disruptions

3.1. Introduction

As companies become more integrated and their supply chains more complex, they are exposed to many risks often driven by the geographical locations in which they operate. Including natural disasters, economic and political crises, labor strikes, currency devaluation, and pandemics, the total number and cost of both natural and man-made disasters has increased dramatically over the last decade (Tang 2006). Well-publicized examples of natural disasters include Hurricane Mitch in 1998, when flooding destroyed banana plantations in Honduras, Guatemala, and Nicaragua, which affected 10% of the world's banana crop. During the spring of 2001, mad cow disease broke out in England, resulting in mass destructions of cattle and shortages of European hides for leather goods' manufacturers. In March 2000, a fire in the Philips Semiconductor plant in Albuquerque, New Mexico disabled production, which led to a supply shortage for both Ericsson and Nokia. Today, many firms recognize these and other supply risks are a major threat to their businesses. Firms including Cisco, Walmart, Dell, and Apple, have publicly reported that their profitability can be influenced by supply disruptions due to supplier non-performance or natural disasters.

Potential for significant losses has led many practitioners and researchers to question how to build more resilient supply chains (Martha and Subbakrishna 2002, Sheffi 2005, Hendricks and Singhal 2005, Tang 2006, Tomlin 2006, Babich et al. 2007).

Given that some disruptions cannot be avoided, what distinguishes firms is the way they handle the disruptions. For example, Nokia was able to source from alternative suppliers during their March 2000 disruption, whereas Ericsson "did not have a Plan B" and lost at least $400 million in potential revenue (Latour 2001). And, as Chiquita leveraged alternative sources of bananas to maintain deliveries, Dole suffered revenue declines and struggled to find alternatives sources of supply. During the outbreak of mad cow disease in Europe, Natale, Gucci, and Wilson Leather were locked into supply contracts. Naturalizer, Danier, and Justin Boot relied on inventories. Etienne Aigner shifted purchases to other regions, but faced stiff cost increases (Martha and Subbakrishna 2002). Tang (2006) has proposed several mitigation strategies based on his observations of successful business practices, including production postponement, maintaining strategic inventories, building a flexible supply base, and ensuring flexible transportation. Among these strategies, the present paper focuses on strategic inventories and a flexible supply base (i.e., sourcing strategy).

Our work is motivated by firms that contract with various suppliers and assign different roles to them, that is, primary versus backup supplier. For example, Kolbus, a German manufacturer of bookbinding machinery, sources the majority of its parts from offshore suppliers. In addition, for 70% of its purchased products, the firm maintains backup suppliers from a network of local suppliers to hedge against both demand-side and supply-side uncertainties. Kolbus's COO stated that, "If we actually need the local suppliers as an extended workbench on a full scale, we compensate them for their deliveries generously. And they are prepared, still knowing that they cannot reckon on us for steadily incoming orders" (Sting and Huchzermeier 2010). Prior literature that has investigated the use of backup suppliers has indeed considered the profitability of using more expensive suppliers as backup in the presence of supply disruptions (e.g., Tomlin 2006, Sting and Huchzermeier 2010). It is usually assumed in those studies, however, that the manufacturer is the sole decision maker and that supplier behavior is exogenous. In such an environment, conventional wisdom suggests–and many models confirm–that a manufacturer cannot be worse off by having backup suppliers. Our study, however, treats suppliers as active decision makers aware of their strengths and weaknesses and evaluates the costs and benefits associated with flexible sourcing, resulting from strategic supplier behaving as price-setters.

We consider a manufacturer that can source from either a perfectly reliable supplier, an unreliable supplier, or both.[1] The option of sourcing from more than one supplier is labeled as a flexible sourcing strategy. The unreliable supplier faces occasional production disruptions, which temporarily stop the flow of materials to the manufacturer. If the manufacturer sources primarily from the unreliable supplier, he may choose to buffer the disruptions with inventory, use the reliable supplier as a backup source, or use both strategies. The reliable supplier may play one of two roles: serve as the primary supplier or provide backup capacity to the manufacturer when the unreliable supplier experiences a disruption (possibly in addition to some guaranteed primary capacity, as we explain later). These are significantly different roles, which may justify different contracts and drive the supplier to make different upfront capacity decisions.

The primary objective of the present study is to analyze the suppliers' strategic behavior when they compete for a manufacturer's business. We focus on the suppliers' pricing strategy and the resulting sourcing strategy of the manufacturer in the equilibrium. Our study also evaluates the effects of the capacity of the reliable supplier and examines the case when demand is non-stationary. We seek to answer the following questions: How do the suppliers set price under various reliability profiles? What sourcing strategy does the manufacturer adopt in the equilibrium, assuming the equilibrium exists? What pricing schemes will prevail in the equilibrium? Under

[1] In our extensions, we also discuss the case when both suppliers are unreliable.

what circumstances do the manufacturer and the suppliers benefit from a flexible sourcing strategy?

To answer these questions, we propose a model in which the suppliers offer terms of delivery (wholesale prices) to the manufacturer. The manufacturer then chooses the suppliers and the roles they will play, as well as decides the inventory policy that will be used. We start by characterizing the inventory policy of the manufacturer and show that it can be described by the number of periods to *"cover,"* and it is independent of values of demand, even when demand is non-stationary. To evaluate strategic supplier behavior, we consider two pricing games in which suppliers compete. In the single-wholesale-price game, each supplier announces a single (wholesale) price. In the contingent-pricing game, the reliable supplier offers wholesale prices contingent on whether she serves as the primary or the backup supplier. The unreliable supplier offers two wholesale prices: one for on-time deliveries and one for late deliveries. The difference between prices can equivalently be interpreted as a penalty for late deliveries in the form of a supplier rebate or a charge-back, which are commonly observed in practice. We find that the single-wholesale price game leads to a conflict of incentives in terms of the roles suppliers want to play and the amount of business they get. This is formally confirmed as non-existence of pure-strategy Nash equilibria in most practical situations. In reality, the reliable supplier may wish to quote different wholesale prices depending on the role she plays. Similarly, the unreliable supplier may offer incentives to win the manufacturer's business. The contingent-pricing game, therefore, corresponds to a more intuitive relationship and has a unique pure-strategy Nash equilibrium. We derive conditions for the manufacturer's various sourcing strategies and corresponding inventory policies, and we describe the resulting equilibrium sourcing outcomes. Except for cases with a significant cost advantages for one of the suppliers, the manufacturer uses the (less expensive) unreliable supplier as well as the (more expensive) reliable supplier. The economic benefits, however, are less obvious. The conventional wisdom is that the manufacturer should never be worse off by having backup suppliers. With endogenously determined wholesale prices, however, the manufacturer does not necessarily benefit from the existence of a backup supplier and, in fact, is typically worse off. Consequently, an up-front commitment to sole-sourcing and using simple wholesale price contracts may be beneficial, as opposed to opening up the opportunity for multi-sourcing and more flexible contracts. Interestingly, suppliers may benefit from flexible sourcing even though the manufacturer does not; indeed, the reliable supplier always benefits from offering backup capacity, whereas the unreliable supplier might benefit in some situations from a reliable supplier's backup capacity despite reduced business volume. From a system perspective, a flexible sourcing strategy may degrade the supply chain's performance.

We characterize the sourcing policy by deriving the conditions under which the manufacturer (1) sole sources from the reliable supplier; (2) the conditions under which he sole sources from the unreliable supplier; and (3) the conditions under which he sources primarily from the unreliable supplier, using the reliable supplier as backup. We show that even when the manufacturer uses both suppliers, it is optimal for him to maintain a safety stock of inventory, because the reliable supplier charges a high price for providing backup capacity. For a special case with stationary demand, we study the choice of capacity level by the reliable supplier. The problem can be decoupled and interpreted based on two independent roles that the reliable supplier can play: regular supplier and backup. Although these roles could be played simultaneously, we show that the reliable supplier will never want to play both roles, that is, to supply only a portion of total demand as a primary supplier.

Our extensions include the case of two unreliable suppliers. We identify conditions when one of the suppliers can be treated as perfectly reliable. We also investigate how the predictability of recovery times influences competitive outcomes and equilibrium profits. Although reduced variability is usually considered favorable in the operations literature, we find that the unreliable supplier may achieve higher profits with unpredictable (more variable) disruptions, resulting from dampening the competition due to the availability of backup supplier.

The remainder of the paper is organized as follows. Section 3.2 reviews the literature. The model and assumptions are introduced in Section 3.3. The manufacturer's optimal sourcing and inventory policies for given wholesale prices are derived in Section 3.4. Section 3.5 analyzes the wholesale-price game and the contingent-pricing game and presents a collaborative framework. Section 3.6 describes the extensions we derive from the primary scenario. Finally, Section 3.7 contains the paper's summary and conclusions.

3.2. Literature Review

The earliest papers dealing with disruptions concentrated on supply disruptions within a single facility or from a given supplier. Meyer et al. (1979) modeled a production facility with stochastic failures and repairs. Hopp and Spearman (1991) and Berg et al. (1994) considered similar settings with machine breakdowns and internal disruptions within a facility. Bielecki and Kumar (1988) derived the conditions under which zero-inventory policies were optimal for a manufacturing facility subject to random failures. Parlar and Perry (1995), Parlar (1997), Gupta (1996), and Arreola-Risa and DeCroix (1998) analyzed disruptions at upstream suppliers. In principle, the analysis is the same as analyzing the disruptions within a manufacturing facility, provided the downstream buyer has perfect access to the upstream supplier's operational status in real-time. Song and Zipkin (1996) analyzed a single-supplier model with a more general supply process (a generalized version of Kaplan 1970) and concluded that under no-order-crossing assumption, the optimal ordering policy is independent of the state of the outstanding orders. Due to generality of the framework, any single-supplier problem in a multi-period setting is likely to fit into Song and Zipkin's framework.

An increasing number of papers have considered scenarios that are more complex than a single-supplier single-buyer relationship. The benefits of multiple-supplier sourcing have been studied in the context of risks such as price reductions resulting from competition among suppliers (Elmaghraby 2000) and variable supplier lead times (Minner 2003). The first papers to consider multiple suppliers to mitigate disruption risks are Parlar and Perry (1996) and Gurler and Parlar (1997). Both papers consider identical-cost infinite-capacity suppliers that were subject to exponentially distributed failure and repair times with fixed ordering costs. Paper by Tomlin (2006) is the most relevant to our work. He assumes that demand is constant and, in a periodic-review setting, two suppliers can serve the manufacturer; one is reliable and the other is unreliable. The unreliable supplier faces no capacity constraints between disruptions but has zero capacity during a disruption. The reliable supplier has a strict capacity constraint and a positive lead-time needed to start production. The major difference between Tomlin's work and ours is that in his model, the wholesale prices are exogenously set, whereas we assume that the suppliers are price-setters and compete for the manufacturer's business. By modeling suppliers as price-setters, we can analyze how suppliers' behavior influences the structure of the supply contract and the manufacturer's sourcing strategy. In addition to the strategies analyzed in Tomlin (2006), Tomlin (2009a) also considers demand switching as a potential lever to mitigate supply disruptions.

Another stream of literature related to supply chain disruptions is one that considers supplier yield uncertainty. Yano and Lee (1995) provides a comprehensive review of the yield uncertainty literature. Among the papers in this group, those that considered multiple suppliers are Gerchak and Parlar (1990), Yano (1991), Anupindi and Akella (1993), Agrawal and Nahmias (1997), Federgruen and Yang (2009), Gurnani et al. (2000), and Babich et al. (2007). A review of these papers can be found in Tomlin (2006) and Babich et al. (2007). Among these papers, Babich et al. (2007) is most relevant to our work. It models a single-period procurement problem with multiple uncertain suppliers. Ordering decisions are made before the supply uncertainty is resolved. Supply disruptions are correlated across the suppliers – each of suppliers can satisfy the order fully or produce nothing, according to a Bernoulli yield distribution. The suppliers are price-setters, as they are in our model. Babich et al. (2007) derives the equilibrium wholesale prices for the two-supplier case with deterministic and stochastic demand. Rather than single wholesale-price contracts, we consider a contingent-price framework and allow for different roles that the suppliers can play (which are not considered in the Babich et al. model). Furthermore, our model is a multi-period one with inventories carried across periods. This enables us to evaluate a more realistic and intuitive pricing strategy and to study the efficacy of broader set of sourcing strategies when inventory is a potential mitigation strategy for disruptions. Tang and Kouvelis (2011) also deals with disruptions that occur in the form of stochastically proportional yield. These authors consider two competing (identical) manufacturers that have an option to source from two competing (identical) suppliers in a single-period setting. The paper focuses on how the yield correlation across suppliers influences the desirability of dual-sourcing for the manufacturers. In contrast, we consider a multi-period model and a supply process that models random failures and recoveries. Inventory is a potential mitigation strategy under this setting. We consider supplier asymmetries with respect to production costs and reliability profiles and allow suppliers to play different roles by offering contracts contingent on that role. This approach leads to different insights. While the aforementioned papers assume that the supplier yield distribution is known to the buyer, Tomlin (2009b) studies the effect of Bayesian learning on optimal sourcing and inventory decisions.

Babich (2006) considers a slightly different setting, one in which a manufacturer has an option to source from two unreliable suppliers, with unequal production lead-times, that actively set their wholesale prices in a single-period setting. The supplier with a shorter lead-time (the faster supplier) can commence production later than the one with a longer lead time (the slower supplier). This gives the manufacturer an option to order from the slower supplier first and wait until the faster supplier must start production. Significantly, the paper finds that the manufacturer could be worse off with flexible sourcing when the suppliers set prices strategically, and that the suppliers could be better off. We do not allow the reliable supplier (faster supplier in the Babich setting) to postpone her pricing decision and announce a wholesale price, if the unreliable supplier (slower supplier in the Babich setting) faces a disruption. Interestingly, even when the suppliers must make price commitments upfront, we find that the manufacturer could still be worse off with flexibility, whereas the suppliers could be better off. This is not expected *a priori*.

Some of the papers that concentrate on bargaining and principle-agent models with asymmetric supply reliability information are also related to our work. Gurnani and Shi (2006) considers a bargaining setting involving a supplier and a buyer with asymmetric information on supplier reliability. Focusing on asymmetric supply reliability information, Yang et al. (2009) considers a single-period model with one manufacturer and one supplier. The supplier could be one of two types, either high-reliability or low-reliability. The supplier can choose to pay a penalty for not being able to deliver or use a backup option. The manufacturer (the principal) designs contracts to maximize expected profits, whereas the supplier (the agent) truthfully reveals her private information by choosing the contract that maximizes her payoff. The manufacturer, having the full power to design the contract, extracts all the rents in the absence of information asymmetry. In our paper, suppliers offer wholesale prices (rather than the manufacturer offering a menu of contracts), which reflects a different type of power for the suppliers. Furthermore, by focusing on a one-period model, inventory cannot be used in Yang et al. model to mitigate risk.

Wan and Beil (2009) analyzes a supply base diversification problem to mitigate cost shocks to procurement, where the buyer, due to lack of bargaining power, asks for bids from suppliers located in various geographical regions. The primary focus of this study was to evaluate diversification strategies that minimize the cost of products plus (random) transportation costs. In contrast, our focus is on the structure of pricing strategies. Also, we consider complete information games only and allow for inventory to be used as a mitigating factor.

3.3. Model and Assumptions

We consider a finite-horizon periodic-review inventory model, in which a downstream firm (manufacturer) faces a deterministic, not necessarily stationary, demand throughout the planning horizon, denoted by d_t for $1 \leq t \leq T$.[2] The manufacturer has an option to source from an unreliable supplier U, who is subject to disruptions; from a reliable supplier R; or from both. In each period, U is in one of two states, ON or OFF. During an ON state, U can satisfy the order instantaneously, whereas he cannot produce anything when in an OFF state. The players are risk-neutral, and their objective is to maximize their expected profits. The contract is signed at the beginning of the horizon (Phase I below), and the terms of the contract are not subject to change during the life of the contract. The sequence of events is as follows.

Phase I: The suppliers announce their pricing schemes before the planning horizon begins. The manufacturer determines which supplier will be his primary source, and whether any supplier will be chosen as a backup supplier.

Phase II: At the beginning of any given period in the planning horizon, the unreliable supplier's state is revealed. Thereafter, the manufacturer decides how much to order from each supplier. The products are received, if any. Finally, the demand is satisfied subject to available inventory.

Supplier U's state is governed by a discrete-time Markov chain X_t. We define θ_f

[2] The deterministic demand assumption is appropriate when long-term supply fluctuations are considered more important compared to short-term demand fluctuations. In addition, our model is appropriate for situations in which the demand in the short term is quite predictable. Other disruption papers that make deterministic demand assumptions include Tomlin (2006), Parlar (1997), Parlar and Perry (1995), Parlar and Perry (1996), and partly Babich et al. (2007).

and θ_r as the probability of failure and recovery respectively, that is, $\text{Prob}(X_{t+1} = \text{OFF}|X_t = \text{ON}) = \theta_f$ and $\text{Prob}(X_{t+1} = \text{ON}|X_t = \text{OFF}) = \theta_r$. The Markov chain is in steady state in the beginning of the planning horizon (i.e., $t = 0$). U does not face any capacity constraints and after recovering from an OFF-state, U "catches up" with unsatisfied demand in one period. If R serves as the primary supplier, orders she receives are perfectly predictable, and R has exact capacity dedicated to the orders. Additionally, R promises up to β units of capacity per period for backup (emergency) orders, which represents her backup capacity that she reserves for all other uses. Furthermore, suppliers do not hold inventories and supply the products with zero lead times. Production costs are linear and denoted by c_u and c_r for the unreliable and reliable suppliers, respectively. Similarly, suppliers' wholesale prices are w_u and w_r. The manufacturer sells at price p. Unsatisfied demand is backlogged. The manufacturer incurs a goodwill penalty π_b for backlogs and holding cost h for inventories.

It is useful here to introduce some simplified notation related to the state of the unreliable supplier. An alternative definition is a Markov chain X'_t with states ON and OFF(k) for $k \geq 1$, where $OFF(k)$ means that X_t has been in the OFF state for exactly k periods. The steady-state probabilities for X'_t are $\pi_{ON} = 1 - \pi_{OFF} = \frac{\theta_r}{\theta_f + \theta_r}$ and $\pi_{OFF(k)} = \theta_r(1-\theta_r)^{k-1}\pi_{OFF}$. Define $F(t) = \sum_{k=0}^{t} \pi_{OFF(t)} = 1 - \pi_{OFF}(1-\theta_r)^t$ as the cumulative distribution function of the long-run average length of a disruption and $\bar{F}(t) = 1 - F(t)$. We note that $F(0) = \pi_{ON}$ describes the proportion of time with no supplier disruption. Finally, we define the inverse function $F^{-1}(x) := \min\{k \mid F(k) \geq x\}$.

3.4. Phase II: Optimal Sourcing Strategy and Inventory Policies under Exogenous Prices

This section contributes to the existing literature by generalizing the optimal policy for non-stationary and deterministic demand. Most important, it provides building blocks for analyzing the pricing games considered in the next section. Whereas the primary model assumes demand backlogging, at the end of the section we provide generalizations for various customer responses toward stock-outs.

We start with analyzing Phase II where the manufacturer, with knowledge of the suppliers' prices (revealed in Phase I), decides on an optimal sourcing strategy (which depends on the state of the unreliable supplier and inventory level). We first allow the manufacturer to source only from U. We then consider the possibility of sourcing from R in addition to U.

3.4.1 Sole-Sourcing from the Unreliable Supplier

We first consider a special case with demand only in the last period. Solving the last-period policy provides us with insights to derive the solution to the entire problem, including when demand is non-stationary.

Simplified Problem: *The manufacturer faces demand only in the last period, $d_t = 0$ for $t < T$. In the case of disruption, customers continue to wait until the supplier becomes available.*

We solve the simplified problem using a dynamic programming formulation with the objective to minimize the expected holding and penalty costs.[3] Denote by $V_t(x, s)$ the manufacturer's expected cost from period t onward, given the starting inventory level x and U's state $s \in \{\text{ON,OFF}\}$. Clearly, $x \leq d_T$ and $V_T(x, \text{ON}) = 0$. (If the supplier is ON, no cost is incurred.) In addition, $V_T(x, \text{OFF}) = \frac{\pi_b}{\theta_r}(d_T - x)$, where $\frac{\pi_b}{\theta_r}$ is the expected cost of backlogging unsatisfied demand until the disruption ends. For $t < T$, the optimality equations are:

$$V_t(x, \text{ON}) = \max_{x \leq y \leq d_T} G_t(y) \tag{3.1}$$

$$G_t(y) = hy + (1 - \theta_f)V_{t+1}(y, \text{ON}) + \theta_f V_{t+1}(y, \text{OFF}) \tag{3.2}$$

$$V_t(x, \text{OFF}) = \theta_r V_{t+1}(x, \text{ON}) + (1 - \theta_r)V_{t+1}(x, \text{OFF}) \tag{3.3}$$

[3]The revenue, p, and the procurement cost, w_u, can be excluded from the formulation.

The intermediate cost-to-go function $G_t(y)$ involves only the holding cost in the current period and future expected costs. The optimal policy is described in Proposition 3.1.

Proposition 3.1. (i) *A base-stock policy is optimal for the simplified problem. Let* $\kappa = F^{-1}\left(\frac{\pi_b}{\pi_b+h}\right)$. *The base stock at period* t *is* $s_t^* = 0$ *for* $t < T - \kappa$ *and* $s_t^* = d_T$ *for* $t \geq T - \kappa$.

(ii) *At* $t = 0$, *the probability that the manufacturer faces a supply disruption is* $\bar{F}(\kappa)$ *and the total expected holding and penalty cost at* $t = 0$ *is* $L(\kappa)d_T$, *where* $L(k) := \left(\frac{\pi_b}{\theta_r}\right)\bar{F}(k) + H(k)$ *and* $H(k) := hk + \left(\frac{h}{\theta_r}\right)\bar{F}(k) - \frac{h}{\theta_r}\pi_{OFF}$. *The coverage* κ *is the minimizer of function* $L(k)$.

To meet the last period's demand, the manufacturer should place an order κ periods in advance. The optimal κ does not depend, however, on the volume of demand, d_T. Thus, κ can be interpreted as a time buffer that balances the cost of disruption with the cost of holding inventory.

The solution to *the simplified problem* readily extends to the general case, in which the manufacturer must meet non-stationary demand across T periods. By isolating each period and assuming FIFO inventory use, we solve multiple simplified problems. Because the economic parameters (holding cost and penalty) remain constant, and due to feasibility of such a policy, the optimal policy turns out to be one in which a given period's requirement is procured exactly κ periods in advance, provided the supplier is in an operating state. That is, the manufacturer carries inventories that will be used in the current period and during the next κ periods. In the case of a disruption, the manufacturer catches up with the desired trajectory as quickly as possible.

Theorem 3.1. (i) *There exists a positive integer* κ, *which we refer to as* coverage, *such that in period* t, *it is optimal to raise the inventory up to* $s_t^* = d_t + d_{t+1} + \cdots + d_{t+\kappa}$.

The optimal coverage κ is independent of the volume of demand in each period. The optimal coverage is $\kappa = F^{-1}\left(\frac{\pi_b}{\pi_b+h}\right)$. (ii) *At $t = 0$, the manufacturer's expected profit per customer is $p - w_u - L(\kappa)$.*

We note that a zero-inventory policy is optimal (i.e., $\kappa = 0$) if $\frac{\pi_b}{\theta_r} \leq \frac{h}{\theta_f}$, that is, when the expected penalty cost per customer is less than the expected holding cost per unit held. Given that the inventory policy does not depend on the demand information beyond κ periods ahead, effectively we do not need the deterministic demand assumption for the entire horizon providing the demand can be reliably predicted for a sufficient number (at least κ) of periods.

While we stated above the results for the case when all unsatisfied customers are backlogged, the analysis extends to various types of customer responses toward stock-outs. In reality, some customers may switch to a competitor or some could only tolerate waiting for a limited number of periods. To account for such behavior, we define the following customer responses toward stock-outs and present modifications corresponding to an optimal policy for the manufacturer. We assume that the manufacturer incurs π_l per lost customer.

Partial Lost Sales (PLS(α)): A fraction α of the customers, with $0 \leq \alpha \leq 1$, are willing to wait until the product becomes available, whereas the remaining $1 - \alpha$ are lost immediately.

Gradually Lost Sales (GLS(α)): Customers who are backlogged leave the system gradually. During each period, a fraction $1 - \alpha$ of all of the backlogged customers are lost.

Theorem 3.2. *The optimal coverage is $\kappa = F^{-1}\left(\frac{\sigma}{\sigma + \frac{h}{\theta_r}}\right)$, where*

i) Under PLS(α), $\sigma = \alpha\left(\frac{\pi_b}{\theta_r}\right) + (1-\alpha)(p + \pi_l - w_u)$.

ii) Under GLS(α), $\sigma = \left(\frac{\alpha\theta_r}{1-\alpha+\alpha\theta_r}\right)\left(\frac{\pi_b}{\theta_r}\right) + \left(\frac{1-\alpha}{1-\alpha+\alpha\theta_r}\right)(p + \pi_l - w_u)$.

The "optimal number of periods to stock," κ, satisfies a newsvendor trade-off, where σ is the cost of underage, and $\frac{h}{\theta_r}$ is the cost of overage (i.e., the expected holding cost during a disruption).

The key insight developed in this subsection is that the production function is a shifted version of the demand function by a constant time κ.

By decreasing the size of a period, the periodic-review model readily extends to a continuous-time model with exponentially distributed supplier up-time and supplier down-time. The demand can be modeled as discrete realizations in time, through demand rate function $d(t)$, or through a combination of both.

3.4.2 Sourcing Strategies with a Reliable Alternative Source

We assume now that the manufacturer has an alternative supplier, R, which is completely reliable. R can be used as the primary (and sole) supplier, as the backup supplier to U, or R can serve both functions. In the last case, a portion of the total demand is satisfied by R, whereas the remainder is satisfied by U. For this remainder, R may also serve as a back-up to U in the event of a disruption.

If R serves as the primary supplier, the orders she receives are perfectly predictable. Thus, we assume that R has an exact capacity dedicated to orders placed on a regular basis. Additionally, R promises up to β units of capacity per period for emergency orders. To derive the optimal policy, we assume stationary demand, $d_t = 1$ for all t and $0 \leq \beta \leq 1$. When the demand is non-stationary, the conclusions of Proposition 3.2 are valid for the two extremes of the backup capacity: $\beta = 0$ and $\beta \geq \sup_t d_t$.

Proposition 3.2. *Let* $\kappa_0 = F^{-1}\left(\frac{\pi_b}{\pi_b+h}\right)$ *and* $\kappa_1 = F^{-1}\left(\frac{w_r-w_u}{w_r-w_u+\frac{h}{\theta_r}}\right)$. *The manufacturer's optimal sourcing policy is as follows.*

(i) If $w_r < w_u$*, it is optimal to source solely from* R*.*

(ii) *If* $w_u \le w_r \le w_u + L(\kappa_0)$*, it is optimal to split the contract and source* $1-\beta$ *units of the demand from* R *and* β *units from* U *on a regular basis, while covering* κ_1 *periods (resulting in a safety stock of* $\kappa_1\beta$*). If a disruption lasts longer than* κ_1 *periods, the manufacturer sources additional* β *units per period from* R*.*

(iii) *If* $w_u + L(\kappa_0) \le w_r \le w_u + \frac{\pi_b}{\theta_r}$*, it is optimal to split the contract as follows.* $1-\beta$ *units of the demand is sourced solely from* U*. The remainder is sourced from* U *when* U *is operational and sourced from* R *when* U *faces a disruption. The manufacturer covers* κ_0 *periods for the portion sourced solely from* U *(resulting in a safety stock of* $\kappa_0(1-\beta)$*) and covers* κ_1 *periods for the remaining portion (resulting in a safety stock of* $\kappa_1\beta$*). As a result, the manufacturer keeps* $(1-\beta)\kappa_0 + \beta\kappa_1$ *units in safety stock and sources from* R *if a disruption at* U *lasts longer than* κ_1 *periods.*

(iv) *If* $w_r > w_u + \frac{\pi_b}{\theta_r}$*, it is optimal to source solely from* U *and cover* κ_0 *periods.*

Note that cases *(ii)* and *(iii)* have an intuitive interpretation. Consider two independent and identical manufacturers, except that Manufacturer 1 faces a constant demand β and uses a backup source, whereas Manufacturer 2 faces a constant demand $1-\beta$ and does not keep a backup source. Then, Manufacturer 1 should cover κ_1 periods (order up to $\kappa_1\beta$) and Manufacturer 2 should cover κ_0 periods (order up to $\kappa_0(1-\beta)$). Proposition 3.2 proves that such decoupling, with artificial roles assigned to two manufacturers, is indeed optimal. Consequently, the profits of the suppliers and the manufacturer are linear in β; therefore, the reliable supplier will offer either no backup capacity or as much backup capacity as possible, if profitable.

This section has treated the suppliers' wholesale prices as exogenously given. In practice, suppliers actively compete on price. We are now ready to address the main question of the paper and explore the impact of reliability on the pricing strategies of the suppliers. We seek to answer the following questions. How do the suppliers price under various reliability profiles? What sourcing strategy does the manufacturer adopt in the equilibrium, assuming it exists? What pricing schemes will prevail in the equilibrium? Under what circumstances do the manufacturer and the suppliers benefit from a flexible sourcing strategy?

3.5. Price Competition of the Suppliers

The key questions we address in this section are how the suppliers set prices and what the manufacturer's resulting sourcing strategy is. For general, non-stationary demand, we analyze two extreme cases. In the first case, R does not offer backup capacity, $\beta = 0$, and in the second case, R offers infinite backup capacity, $\beta = \infty$. The intermediate case with finite backup capacity and stationary demand is a straightforward extension of these two extremes due to Proposition 3.2. We adopt the continuous-time view for inventory replenishment.

Two pricing games are considered: the single-wholesale-price game and the contingent-pricing game. In the single-wholesale-price game, the reliable and unreliable suppliers simultaneously announce single wholesale prices. We find that a pure-strategy Nash equilibrium does not exist, except for very special cases. This is because the price that the reliable supplier wishes to charge as the primary supplier is not the same as the price she would charge as the backup supplier. The contingent-pricing game addresses this shortcoming (and reflects realities that suppliers face). In this scheme, the reliable supplier offers two wholesale prices, one as a primary supplier and one as a backup supplier. In response, the unreliable supplier offers two wholesale prices, one for on-time deliveries and one for delayed deliveries. In contrast, we find that the contingent-pricing game has a pure-strategy Nash equilibrium, and derive conditions for different sourcing strategies of the manufacturer. We also evaluate the efficacy of the contingent-pricing game using a centrally-managed supply chain as a benchmark.

3.5.1 Single-Wholesale-Price Game

Here we analyze two separate cases. In the first case, R does not offer any backup capacity, and in the second case, R offers backup capacity.

• Assume that R cannot offer any backup capacity, $\beta = 0$.

Proposition 3.3. *Assume that R does not offer backup capacity, $\beta = 0$. Let $\kappa_0 = F^{-1}\left(\frac{\pi_b}{\pi_b+h}\right)$. If $c_r \le c_u + L(\kappa_0)$, then, R is awarded the entire contract with $(w_r^*, w_u^*) = (c_u + L(\kappa_0), c_u)$. Otherwise, U is awarded the entire contract with $(w_r^*, w_u^*) = (c_r, c_r - L(\kappa_0))$.*

The effective cost of sourcing from U turns out to be $w_u + L(\kappa_0)$. Hence, in the equilibrium, R can charge a premium over the cost of U equal to $L(\kappa_0)$. This is the expected holding and backlogging costs that the manufacturer would incur when he sets the inventory coverage to κ_0.

• Assume that R can serve as a backup supplier with infinite capacity, $\beta = \infty$. If $w_r \le w_u$, the manufacturer clearly uses R as the primary supplier. If $w_r > w_u$, then the manufacturer uses the less expensive supplier, U, as his primary supplier and may use R as his backup supplier. Thus, R either sets $w_r = w_u$ and serves as the primary supplier or sets $w_r > w_u$ and serves as the backup supplier. Denote by $w_{rb}^*(w_u)$ the optimal wholesale price that R would charge conditional on serving as the backup supplier, that is, $w_r > w_u$.

Proposition 3.4. *If $w_u < c_r - \frac{\pi_b}{\theta_r}$, then $w_{rb}^*(w_u) = c_r$ and the manufacturer does not source from R. If $w_u \ge c_r - \frac{\pi_b}{\theta_r}$ and $w_r > w_u$, R serves as the backup supplier and*

i) If $w_u < c_r + \frac{h}{\theta_r}$, $w_{rb}^*(w_u) = w_u + \frac{\pi_b}{\theta_r}$ and the manufacturer's inventory coverage is κ_0.

ii) If $w_u \ge c_r + \frac{h}{\theta_r}$, $w_{rb}^*(w_u) = w_u + \min\left\{\frac{\pi_b}{\theta_r}, \frac{h}{\theta_f}\right\}$ and the manufacturer holds no inventory.

The presence of backup supplier may either drive the manufacturer's inventory down to zero or leave it unchanged, compared to the case where R is not available as a backup supplier. When inventory is held in equilibrium, R charges a premium on U's wholesale price, equal to the expected penalty costs for backlogging the demand. Otherwise, when no inventory is held, R's premium is equal to the expected cost of holding one unit of inventory until a disruption occurs or the expected penalty costs for backlogging, whichever is lower. The smaller premium is beneficial for R, because it corresponds to a no-inventory policy and, consequently, increases the portion of R's business. The following proposition characterizes the equilibrium prices in the single-wholesale-price game.

Proposition 3.5. *Let* $\delta = \frac{\theta_f}{\theta_r} \min \left\{ \frac{\pi_b}{\theta_r}, \frac{h}{\theta_f} \right\}$. *Assume that* R *offers backup capacity,* $\beta = \infty$. *If* $c_u \geq c_r + \delta$, R *is awarded the entire contract with* $(w_r^*, w_u^*) = (c_u, c_u)$. *If* $c_u < c_r + \delta$, *no pure-strategy Nash equilibrium exists for the single-wholesale-price game.*

The single-wholesale-price game has a unique pure-strategy Nash equilibrium only when R has significant cost advantage, $c_r \leq c_u - \delta$. This condition may not hold under settings in which the reliability is associated with higher production costs. Comparing Propositions 3.3 and 3.5 to the case where $c_r \leq c_u - \delta$, R wins the contract in both cases and charges $w_r^* = c_u$ if $\beta = \infty$ and $w_r^* = c_u + L(\kappa_0)$ if $\beta = 0$. Interestingly, R charges a higher price when she cannot offer backup capacity. We briefly discuss the reasons for this seemingly counter-intuitive result. If $\beta = \infty$, the possibility of offering backup capacity prevents R from charging a premium for its reliability. Unless R sets a lower price than U, the manufacturer will chose U as the primary supplier and R as the backup supplier. Without back-up capacity, R can charge a higher price than U. Therefore, if R has plenty of unutilized capacity, she may find herself desperate to use it and unable to charge a premium. If, however, R can credibly dedicate her capacity to other customers in the event she does not get the contract, then she can take advantage of being reliable and charge a premium.

In the case $c_r > c_u - \delta$, we observe from the proof of Proposition 3.5 that non-existence of pure-strategy Nash equilibria is driven by the fact that R wishes to charge different prices for acting as a primary supplier versus acting as a backup supplier. If the price decreases significantly due to competition, R finds it optimal to increase her price and become the backup supplier. However, in response, U finds it optimal to increase prices, and hence, firms do not reach equilibrium. Similarly, they would not reach equilibrium even if they were allowed to revise prices during the horizon. To account for the possibility that R is able to price based on her role, we introduce the contingent-pricing game.

3.5.2 Contingent-Pricing Game

In practice, one would expect the price to differ based on the supplier's role, the expected volume, and the expected predictability of using supplier capacity. In this section, we consider the contingent-pricing game, in which R is able to announce wholesale prices contingent upon her role. In the beginning of the horizon, R quotes two contingent prices, (w_r, w_{rb}). The wholesale price, w_r, is for regular orders and w_{rb} is for the emergency (backup) orders. We assume that R is able to verify whether an order is placed on regular basis or on an emergency basis. Common knowledge of demands and supplier state makes this a natural assumption in our model. U is also allowed to offer two wholesale prices, (w_u, w_{ud}). The wholesale price, w_u, is for units delivered on time, whereas w_{ud} is the wholesale price for delayed orders. The difference $\phi = w_u - w_{ud}$ can be interpreted as a penalty (or rebate) that U pays for delayed orders. We represent U's decision as (w_u, w_{ud}) and (w_u, ϕ) interchangeably. The backup capacity, β, continues to play a significant role.

- For $\beta = 0$, the contingent-pricing game reduces to the single-wholesale-price game. R does not set a price for backup availability (or $w_{rb} = \infty$). In response, U can achieve any outcome with a single wholesale price and Proposition 3.3 applies.

• Assume that R can serve as a backup supplier with infinite capacity, $\beta = \infty$. To derive the equilibrium outcomes, we first derive R's best response to U given (w_u, ϕ).

Lemma 3.1. *Let* $\kappa_{cp}(\phi) = F^{-1}\left(\frac{\pi_b - \theta_r\phi}{\pi_b + h - \theta_r\phi}\right)$. *Two thresholds exist,* $w_u^l(\phi) \leq w_u^h(\phi)$ *that characterize R's best response.*

If $w_u \leq w_u^l(\phi)$, *R sets* $w_r = w_{rb} = c_r$ *and does not serve the manufacturer.*

If $w_u^l(\phi) \leq w_u \leq w_u^h(\phi)$, *$R$ serves as the backup supplier with* $w_{rb} = w_u - \phi + \frac{\pi_b}{\theta_r}$ *and* $w_r = c_r + (w_{rb} - c_r)\,\bar{F}(\kappa_{cp}(\phi))$.

If $w_u \geq w_u^h(\phi)$, *R serves as the sole supplier with* $w_r = w_u + L(\kappa_{cp}(\phi)) - \phi\bar{F}(\kappa_{cp}(\phi))$ *and* $w_{rb} = c_r + \frac{w_r - c_r}{\bar{F}(\kappa_{cp}(\phi))}$.

Figure 3.1: Sourcing outcome when R best responds to U

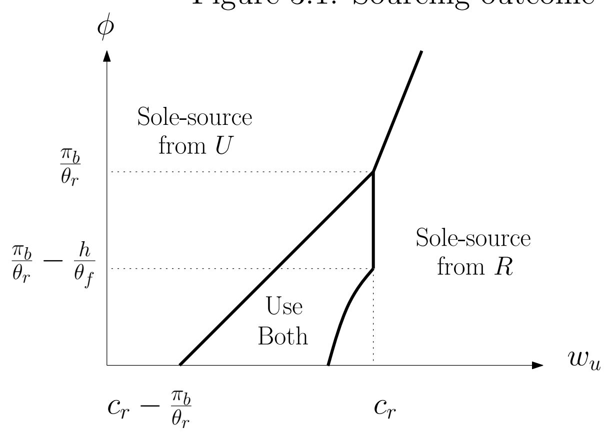

Figure 3.1 shows the resulting sourcing outcomes when U's wholesale prices are exogeneously determined and R is best-responding to U's wholesale prices. Given how R best responds to a contract offered by U, we characterize the equilibrium outcomes below.

Proposition 3.6. *Let* $\kappa_0 = \kappa_{cp}(0)$ *and* $\Delta = \frac{H(\kappa_0)}{F(\kappa_0)}$, *satisfying* $0 \leq \Delta \leq L(\kappa_0)$. *There exists a unique equilibrium of the contingent-pricing game. The equilibrium outcomes are:*

(i) *If* $c_r - c_u \leq \Delta$, *the manufacturer sole-sources from R:* $w_r^* = c_u + L(\kappa_0)$, $w_u^* = w_{ud}^* = c_u$.

(ii) *If* $\Delta < c_r - c_u < \frac{\pi_b}{\theta_r}$, *the manufacturer primarily sources from* U *and uses* R *as the backup supplier:* $w_u^* = c_u + \phi^*$, $w_{ud}^* = c_u$, $w_{rb}^* = c_u + \frac{\pi_b}{\theta_r}$. *The equilibrium penalty,* ϕ^*, *is weakly increasing in* $c_r - c_u$*; hence, equilibrium inventory,* $\kappa_{cp}(\phi^*)$, *is weakly decreasing in* $c_r - c_u$.

(iii) *If* $c_r - c_u \geq \frac{\pi_b}{\theta_r}$, *the manufacturer sole-sources from* U*:* $w_u^* = c_r - \frac{\pi_b}{\theta_r} + \phi^*$, $w_{ud}^* = c_r - \frac{\pi_b}{\theta_r}$, $w_r^* = w_{rb}^* = c_r$. *The equilibrium penalty,* ϕ^*, *and inventory,* $\kappa_{cp}(\phi^*)$, *do not depend on* $c_r - c_u$.

It is easier to interpret Proposition 3.6 when $\frac{\pi_b}{\theta_r} \leq \frac{h}{\theta_f}$, that is, when the inventory holding cost is too high to allow the manufacturer to carry inventory. Because a zero-inventory policy is optimal in equilibrium, the manufacturer's orders can be classified into two categories: regular orders that arise with a frequency π_{ON} and emergency (or delayed) orders that arise with a frequency π_{OFF}. With this separation of orders, suppliers effectively engage in two pricing games with the flexibility of quoting different wholesale prices for each game. In equilibrium, the manufacturer's cost for each type of order (including a backlog penalty, if any) is given by his second best option: for regular orders, the manufacturer incurs $\max\{c_r, c_u\}$ and for emergency (or delayed orders), he incurs $\max\left\{c_r, c_u + \frac{\pi_b}{\theta_r}\right\}$. (Sourcing from U results in penalty costs for delayed orders.)

The equilibrium outcomes are clearly influenced by the cost advantage of the unreliable supplier, $c_r - c_u$. Except for the extreme cases, the manufacturer uses the less expensive unreliable supplier and the more expensive reliable supplier. The equilibrium outcomes are also influenced by other factors such as the manufacturer's holding and penalty costs, as well as the frequency and length of disruptions. Figure 3.2 illustrates how the equilibrium sourcing outcome depends on (a) the average uptime and downtime and (b) the manufacturer's inventory holding and penalty costs.

Figure 3.2(a) has an intuitive interpretation. For disruptions with sufficiently short average length, in equilibrium, U is able to secure the contract as the sole supplier.

Figure 3.2: Equilibrium Sourcing Outcome

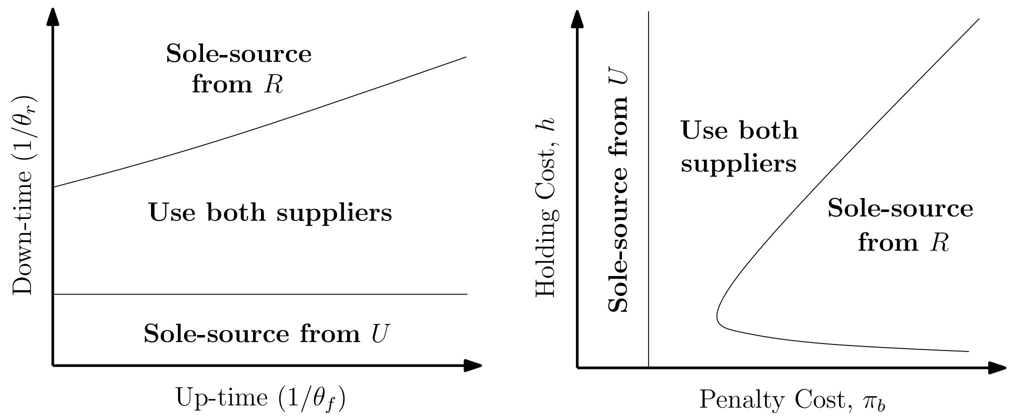

(a) Parameters π_b, h and $c_r > c_u$ are fixed.

(b) Parameters θ_f,θ_r and $c_r > c_u$ are fixed.

Otherwise, depending on the average length of disruptions, the manufacturer either sole-sources from R or keeps R as the backup source. The latter case implies lost orders for U whenever he faces disruptions that are longer than the manufacturer's inventory coverage.

Figure 3.2(b), on the other hand, is less intuitive and reflects the effect of strategic supplier behavior. When the penalty cost is negligible, U is able to secure the contract as the sole supplier. When the penalty cost is moderately high, the manufacturer uses both suppliers. These results are intuitive. However, when the penalty cost increases, the manufacturer's sourcing strategy is influenced as well by holding cost, but in a non-monotonic way. If the holding cost is negligible, the manufacturer can mitigate disruptions by holding a significant amount of inventories. Consequently, R cannot compete aggressively to serve as the primary supplier. When the holding cost is very high, the manufacturer is unable to hold inventory economically and has an incentive to maintain a backup supplier. In this case, R takes advantage of being reliable and becomes the backup supplier, not because she cannot compete aggressively to be the sole supplier, but because it is more profitable to be the backup. Finally, if the holding cost is in the middle range, it is more profitable for R to compete more aggressively and secure the entire contract as the sole supplier. The dynamics we observe for high penalty costs are driven purely by strategic supplier behavior in the presence of backup capacity.

Exploring the equilibrium outcomes of the contingent-pricing game with and without backup capacity enables us to address the focal questions of our paper. Does the manufacturer benefit from flexible sourcing by allowing the reliable supplier to offer backup capacity at a higher wholesale price? Which supplier(s) benefit from a flexible sourcing arrangement? In the absence of strategic behavior (that is, with fixed wholesale prices), the manufacturer will always benefit from flexible sourcing, whereas the supplier U could potentially be hurt, because his orders are lost with flexible sourcing. Supplier R may or may not benefit, depending on the original arrangement without flexible sourcing. Below, we show that this intuition does not necessarily carry over when strategic supplier behavior is taken into account. To address how flexible sourcing influences profits fully, we need the following lemma, in which we prove monotonicity of profits in suppliers' production costs. We use subscript m to denote the manufacturer. Subscripts u and r are already defined and refer to the suppliers.

Lemma 3.2. *Denote by $\Pi_i^*(c_r, c_u)$ player i's equilibrium profit in the contingent-pricing game with $\beta = \infty$. Let $\Pi_{sc}^*(c_r, c_u)$ denote the total supply chain profit. Then, $0 \leq \frac{\partial \Pi_i^*(c_r,c_u)}{\partial c_j} \leq 1$ and $-1 \leq \frac{\partial \Pi_i^*(c_r,c_u)}{\partial c_i} \leq 0$ for $i, j \in \{u, r\}$ with $i \neq j$. Also, $-1 \leq \frac{\partial \Pi_i^*(c_r,c_u)}{\partial c_j} \leq 0$ for $i \in \{m, sc\}$ and $j \in \{u, r\}$.*

The results of Lemma 3.2 are intuitive. When the supplier's cost is increased, that supplier's equilibrium profit decreases; in addition, the manufacturer and the supply chain performance suffer. On the other hand, the profit of the competing supplier improves. The next proposition explores how backup capacity influences profits, where we compare the case of backup availability ($\beta = \infty$) and no backup availability ($\beta = 0$).

Proposition 3.7. • *Let* $\frac{\pi_b}{\theta_r} \le \frac{h}{\theta_f}$*, leading to zero inventory in equilibrium. If* $0 \le c_r - c_u \le \frac{\pi_b}{\theta_r}$*, with backup capacity, the profits of both suppliers and the total supply chain profit are higher, whereas the manufacturer's profit is lower. Otherwise, profits are not influenced by backup capacity.*

• *Let* $\frac{\pi_b}{\theta_r} > \frac{h}{\theta_f}$*. If* $c_r - c_u \le \Delta$*, profits are not influenced by backup capacity. Let* $c_r - c_u > \Delta$*. R's profit is higher with backup capacity when* $c_r - c_u \le \frac{\pi_b}{\theta_r}$*, and her profit does not depend on backup capacity when* $c_r - c_u > \frac{\pi_b}{\theta_r}$*. With backup capacity, there exist* $\bar{c}_u$*,* $\bar{c}_m$ *and* $\bar{c}_{sc}$*, such that the profits of* U*, the manufacturer, and the supply chain are higher if* $c_r - c_u \le \bar{c}_u$*,* $c_r - c_u \ge \bar{c}_m$ *and* $c_r - c_u \le \bar{c}_{sc}$*, respectively, and lower otherwise.* $L(\kappa_0) \le \bar{c}_u, \bar{c}_m, \bar{c}_{sc} \le \frac{\pi_b}{\theta_r}$ *and* $\bar{c}_u \le \bar{c}_m$*.*

Proposition 3.7 shows that R always (weakly) benefits from offering backup capacity in the contingent-pricing game. This is in contrast to the single-wholesale-price game, in which R could be worse off with backup capacity. Thus, the contingent-pricing game enables R to make better pricing decisions and does not penalize R for maintaining backup capacity.

Surprisingly, U and the manufacturer may or may not benefit from backup capacity. This is a major departure from our intuition for the case of non-strategic suppliers. Even more surprising, there is no case when all supply chain players improve their profits with backup capacity simultaneously; indeed, in some of those situations, the supply chain performance degrades.

The findings of Proposition 3.7 are illustrated in Figure 3.3, when equilibrium policy results in zero inventory, and Figure 3.4, with positive inventory. When inventory holding cost is high (Figure 3.3), both suppliers benefit from the existence of backup capacity. When both suppliers are used, in the special case with zero inventory, the equilibrium wholesale prices are driven by the following dynamics.

• The unreliable supplier might offer a higher price, w_u, for regular deliveries and a lower, discounted price, w_{ud}, for delayed deliveries. With a discount equal to

Figure 3.3: Sourcing Outcomes and Benefits of Backup Capacity with $\frac{\pi_b}{\theta_r} \leq \frac{h}{\theta_f}$

Sourcing outcome	Regions with respect to $c_r - c_u$		
with $\beta = 0$	R	U	
with $\beta = \infty$	R	Use both suppliers R and U	U
	Benefits of Backup Capacity		
Reliable Supplier	**0**	+ + + + + + + + +	**0**
Unreliable Supplier	**0**	+ + + + + + + + +	**0**
Manufacturer	**0**	− − − − − − − − −	**0**
Supply Chain	**0**	+ + + + + + + + +	**0**

Note. "+": benefits, "−": does not benefit, "**0**": is indifferent

the expected penalty, w_{ud} is easily below U's cost and U loses money on delayed orders. Because the backup business is separate from on-time deliveries, U is better off charging at cost, $w_{ud} = c_u$, and letting R serve as the backup supplier.

• The reliable supplier is interested in the backup business and can charge a premium corresponding to the expected backlogging penalty, $w_{rb} = c_u + \frac{\pi_b}{\theta_r}$. Alternatively, R can serve as the primary (and only) supplier and charge a price that makes the manufacturer indifferent to the cost of buying from an unreliable supplier (at w_u) plus the cost of backup delivery from R.

• This is where w_u becomes critical. As w_u decreases, R's benefit from serving as the primary supplier diminishes and falls below her benefit from serving as a backup supplier. Competitive forces drive w_u down until suppliers reach a point ($w_u = c_r$) where R benefits more from serving as the backup supplier. When this occurs, both suppliers make a larger profit. See Appendix 3.8 for a numerical example.

This is a striking result, which shows that the manufacturer is worse off by seeking flexible sourcing, whereas the suppliers are better off. The flexibility of playing two separate games removes the conflict of incentives that arise in the single-wholesale-price game in terms of the roles that suppliers play and the amount of business they get. It also allows the suppliers to quote prices based on their strengths and weaknesses: when U's strength is low cost and his weakness is long disruptions, then he competes to be the primary supplier and simply forgoes orders during disruptions.

Figure 3.4: Sourcing Outcomes and Benefits of Backup Capacity with $\frac{\pi_b}{\theta_r} > \frac{h}{\theta_f}$

Sourcing outcome	Regions with respect to $c_r - c_u$		
with $\beta = 0$	R	U	
with $\beta = \infty$	R	Use both suppliers R and U	U

Benefits of Backup Capacity

Reliable Supplier	**0**	+ + + + + + + + + +	**0**
Unreliable Supplier	**0**	+ + + + + + + + +	− − −
Manufacturer	**0**	− − − − − − − − −	+ +
Supply Chain	**0**	+ + + + + + + + +	− − −

Note. "+" : benefits, "−" : does not benefit, "**0**" : is indifferent

The situation may be slightly different when inventory is inexpensive. As shown in Figure 3.4, in some cases, the manufacturer can benefit from flexible sourcing. When the manufacturer benefits from a flexible sourcing strategy, however, supplier U is worse off. Moreover, the entire supply chain's performance is worse compared to the case when flexible sourcing is not allowed ($\beta = 0$).

Although supply chain performance often improves, typically, a coordinated outcome is not achieved.[4] Figure 3.5 illustrates the total supply chain profit and inventory level under three scenarios (assuming $c_u = 0$): (1) coordinated outcome, (2) contingent-pricing game with $\beta = 0$ and (3) contingent-pricing game with $\beta = \infty$. The inventory decision is distorted with the contingent-pricing game with $\beta = \infty$.

To see why distortion occurs, note that in equilibrium (Proposition 3.6) U charges his cost for delayed orders, $w^*_{ud} = c_u$. In response, the equilibrium wholesale price of the (reliable) backup supplier is $w^*_{rb} = c_u + \frac{\pi_b}{\theta_r}$, which does not depend on c_r. However, U's equilibrium wholesale price, w^*_u, is increasing in c_r. Thus, the manufacturer faces an increasing wholesale price (as a function of c_r) for on-time deliveries and a

[4]To obtain a coordinated outcome, it suffices to set $w_r = w_{rb} = c_r$ and $w_u = w_{ud} = c_u$ and use our earlier results to derive the optimal policy.

constant wholesale price for emergency (backup) orders. The manufacturer also has less incentive to carry inventory, while he has more incentive to rely on the backup supplier. Consequently, the distortion of inventory takes place with adverse effects on the supply chain.

Figure 3.5: Supply Chain Performance

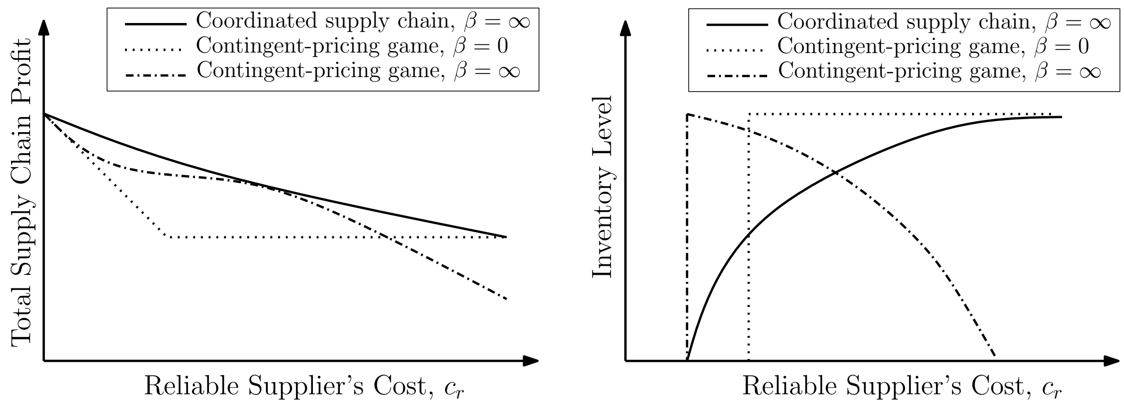

(a) Total supply chain profit in equilibrium

(b) Inventory Level

In summary, whereas the manufacturer is most often worse off by seeking flexible sourcing, situations exist in which it is not the case: when the difference between supplier costs is significant and inventory is inexpensive.

Incorporating strategic supplier behavior shows that opening up opportunities for more flexibility typically hurts the manufacturer. Consequently, the manufacturer may not adopt a flexible sourcing strategy even when it is beneficial to the supply chain. Instead, a manufacturer may ask the suppliers to quote single wholesale prices with an upfront commitment to a sole sourcing strategy.

3.6. Extensions

In this section, we discuss potential directions for extending our model, including the case when both suppliers are unreliable, the possibility of non-memoryless recovery times, and positive lead-times.

- Both Suppliers are Unreliable

We index suppliers by $i = 1, 2$, with the production costs of the suppliers $c_1 \leq c_2$. We assume that disruptions across suppliers are independent and suppliers have infinite capacity. Let θ_{fi} and θ_{ri} be the probabilities that supplier i faces a disruption and recovers from a disruption in the next period, respectively, and $F(t) := 1 - \left(\frac{\theta_{fi}}{\theta_{fi}+\theta_{ri}}\right)(1-\theta_{ri})^t$.

Proposition 3.8. *If* $c_1 \leq c_2 \leq c_1 + \frac{\pi_b}{\theta_{r1}}$, *state-dependent coverage policy is optimal, and it is optimal to use supplier 2 as backup. There are three optimal coverages (one for each state):* κ_{ON-ON}, κ_{ON-OFF}, *and* κ_{OFF-ON}. κ_{ON-ON} *and* κ_{ON-OFF} *are increasing in* $c_2 - c_1$ *and* κ_{OFF-ON} *is decreasing in* $c_2 - c_1$.

Figure 3.6 illustrates how the optimal policy depends on the cost difference $c_2 - c_1$. As supplier 2 becomes more costly, the manufacturer holds more safety stock. When supplier 1 faces disruption, however, and only supplier 2 is available, the manufacturer delays the purchase of inventory and, consequently, maintains lower inventory coverage. Figures 3.6 (a) and (b) together illustrate that the optimal coverage in one state is not necessarily greater or smaller than the optimal coverage in another state.

It is interesting to consider how strategic supplier pricing affects the manufacturer's equilibrium sourcing strategy. The suppliers' strategy space expands, because both suppliers can play the role of the primary supplier or the backup supplier. The optimal inventory policy (state-dependent coverage) can only be computed through backward induction, rather than closed-form expressions; furthermore, supplier payoffs in the pricing games cannot be expressed in closed form (which was possible in the original model). Although this limits the analysis we can perform, below we derive the conditions under which one of the suppliers can be treated as *perfectly reliable* for the model we analyze in the present paper.

We consider a special case, in which supplier 2's state in every period is given by independent Bernoulli trials, where the supplier is OFF with probability ν and ON with probability $1 - \nu$, independent of the supplier's state in the previous periods.

Figure 3.6: Illustration of State-Dependent Optimal Coverage

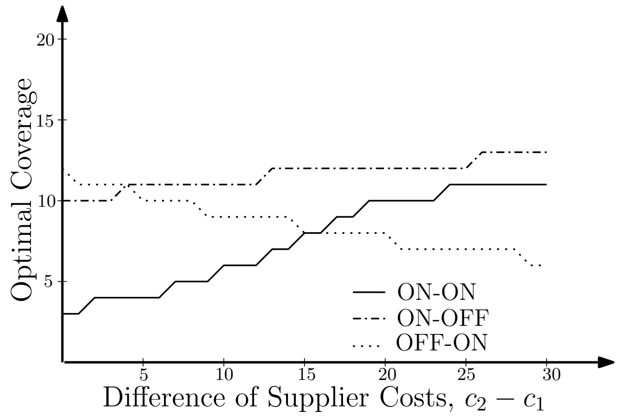

(a) Supplier 2 is less reliable, $(\theta_{f1}, \theta_{f2}) = (0.05, 0.10)$, $(\theta_{r1}, \theta_{r2}) = (0.15, 0.10)$.

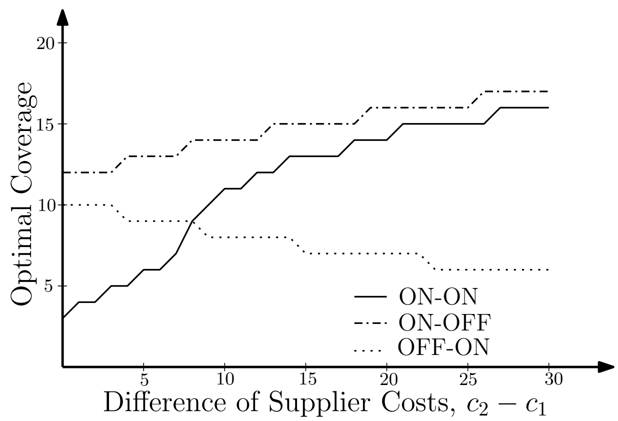

(b) Supplier 2 is more reliable, $(\theta_{f1}, \theta_{f2}) = (0.10, 0.05)$, $(\theta_{r1}, \theta_{r2}) = (0.10, 0.15)$.

Note. For all numerical examples, we assume $\pi_b = 99$ and $h = 1$.

Therefore, $\theta_{f_2} = 1 - \theta_{r2} = \nu$.

Proposition 3.9. *If* $c_1 \leq c_2$ *and* $\frac{\pi_b}{\pi_b + h} \leq 1 - (1 - \theta_{r1})\nu$, *then the manufacturer maintains a certain inventory coverage when supplier 1 is available (independent of supplier 2's state) and sources from supplier 2 only when supplier 1 is not available and the inventory level drops to zero. The optimal coverage satisfies* $\kappa = \min \left\{ k \mid F_1(k) \geq \frac{\sigma\nu + (c_2 - c_1)(1-\nu)}{\sigma\nu + (c_2 - c_1)(1-\nu) + \frac{h}{\theta_{r1}}} \right\}$, *where* $\sigma = \frac{\pi_b}{1-(1-\theta_{r1})\nu} + \frac{(1-\nu)(1-\theta_{r1})(c_2-c_1)}{1-(1-\theta_{r1})\nu}$.

Proposition 3.9 provides a sufficient condition when the optimal policy is characterized by a single number, κ, as in the case with a *perfectly reliable supplier*, where the manufacturer sources from the reliable supplier only when inventory drops to zero. This case applies when supplier 2 faces infrequent and short disruptions. In this case, the supplier can be treated as *perfectly reliable* and the conclusions of the original model apply. The logic of coverage applies except the cost of underage now needs to include the possibility that supplier 2 could also face a disruption.

Although the derivation of equilibrium strategies of the suppliers is generally not straightforward with two unreliable suppliers, the equilibrium outcomes can still be derived if inventory holding cost is prohibitively high. Because either supplier can serve as the primary source, with the other serving as backup, we allow each supplier to quote three wholesale prices: wholesale price for regular deliveries, w_i; wholesale price for backup availability, w_i^b; and wholesale price for delayed deliveries, provided that an order is not placed with the other supplier who is operational, w_i^d. Here, we refer to supplier i's opponent as supplier $-i$.

Proposition 3.10. (i) *Assume that suppliers do not offer backup availability. Let* $i = \arg\min_{i=1,2}\left\{c_i + \frac{\pi_b}{\theta_{ri}}\pi^i_{OFF}\right\}$. *Then, supplier i serves as the sole supplier and charges a wholesale price,* $w_i = c_{-i} + \frac{\pi_b}{\theta_{r,-i}}\pi^{-i}_{OFF} - \frac{\pi_b}{\theta_{ri}}\pi^i_{OFF}$.

(ii) *Assume that the suppliers offer backup.*

Let $\rho_i = \pi^i_{OFF} - \pi^1_{OFF}\pi^2_{OFF}\left(\frac{\theta_{ri}}{\theta_{r1}+\theta_{r2}-\theta_{r1}\theta_{r2}}\right)$. *If* $c_2 > c_1 + \frac{\pi_b}{\theta_{r1}}$, *supplier 1 serves as the sole supplier in equilibrium with* $w_1 = c_2 + \left(\frac{\pi_b}{\theta_{r2}(1-\rho_1)}\right)\pi^2_{OFF}$ *and* $w^1_d = c_2 - \frac{\pi_b}{\theta_{r1}}$. *Otherwise, if* $c_1 \leq c_2 \leq c_1 + \frac{\pi_b}{\theta_{r1}}$, *then supplier 1 serves as the primary supplier and supplier 2 serves as the backup supplier, with* $w_1 = c_2 + \frac{\pi_b}{\theta_{r2}}\left(\frac{\rho_2}{1-\rho_1}\right)$ *and* $w^2_b = c_1 + \frac{\pi_b}{\theta_{r1}}$.

(iii) *With backup capacity, the profits of both suppliers are higher, whereas the manufacturer's profit is lower.*

With two unreliable suppliers, flexible sourcing still leads to larger profits for the suppliers and lower profit for the manufacturer. Extensive numerical study shows that similar conclusions would apply when both suppliers are unreliable and inventory is allowed. To illustrate this, consider a numerical example with two identical suppliers, having equal production costs ($c_1 = c_2 = \$10$) and equal probabilities of failure and recovery ($\theta_{f1} = \theta_{f2} = 0.01$ and $\theta_{r1} = \theta_{r2} = 0.10$). For the manufacturer, let $p = \$100$, $h = \$0.20$, $\pi_b = \$5$ and $d_t = 1$ for all t. If only sole sourcing is allowed, the suppliers compete against one another very aggressively, and because they are identical, their wholesale prices are equal to their costs, \$10. Thus, in equilibrium, the suppliers would make zero profit. In this case, the manufacturer's inventory is 9, and his average profit per period is \$86.55. If flexible sourcing is allowed, then the suppliers compete and undercut their wholesale prices until they reach a point where serving as a primary supplier and serving as a backup supplier are equally profitable. At this point, the competition stops, and the suppliers charge \$11.79 for regular orders, \$60 for emergency orders, and \$10 for delayed orders. In equilibrium, one supplier is used as the primary supplier, and the other supplier is used as backup. The manufacturer still carries 9 units of inventory when the primary supplier is operational. Although the suppliers make equal profits, on average \$1.86, the manufacturer's profit is \$84.69. Thus, the manufacturer's profit is lower with flexible sourcing, whereas the suppliers' profits are larger.

- Non-memoryless Recovery Times

Our model assumes exponential up-times and down-times, which is a common assumption in most studies dealing with supply disruptions. This assumption not only provides us with a simpler analytical framework, but it is also appropriate if the disruptions are very unpredictable. To understand the effect of less variability in the recovery times, we examine the boundary case with deterministic recovery times.

As before, assume that the demand is constant and occurs continuously at a rate of 1. The time until a disruption (i.e., uptime) follows an exponential distribution with mean $U = \frac{1}{\theta_f}$, and a disruption (i.e., downtime) lasts for D time units. Focusing first on single wholesale prices for each supplier, we derive the manufacturer's optimal inventory policy when $w_r > w_u$.

The condition for the optimality of zero inventory policy is the same as for exponential disruptions. In deterministic disruptions, however, the manufacturer may stop sourcing from R toward the end of a disruption, unlike the case with exponential disruptions. Also, if R is able to provide infinite backup capacity, the maximum premium that R can charge to guarantee manufacturer sourcing during a disruption is lower compared to the memoryless case. Denote by t_e as the time left until the end of the current disruption.

Proposition 3.11. *If the manufacturer sole sources from U, the optimal coverage is $\kappa_d = \max\left(0, \frac{\pi_b D - hU}{\pi_b + h}\right)$.*

Assume that R serves as the backup supplier with infinite capacity.

(i) *Let $\pi_b D \leq hU$. Then, a zero inventory policy is optimal, and it is optimal to source from the backup supplier if $\{w_r < w_u + \pi_b D$ and $t_e \leq \frac{w_r - w_u}{\pi_b}\}$.*

(ii) *Let $\pi_b D > hU$. If $w_r > w_u + \frac{h(U+D)\pi_b}{\pi_b + h}$, the manufacturer sole sources from U and holds $\kappa_d = \frac{\pi_b D - hU}{\pi_b + h}$ units of inventory. Otherwise, the manufacturer holds $\frac{w_r - w_u - hU}{h}$ units of inventory and sources from R when $t_e \leq \frac{w_r - w_u}{\pi_b}$.*

Based on numerical experiments, Figures 3.3 and 3.4 remain structurally unchanged when disruptions are deterministic. The thresholds and the profits, however, are obviously influenced. We therefore compare the effect of variability on both suppliers' and manufacture's profits.

Reducing variability is usually considered favorable. In the context of disruptions, we expect that more variable disruptions hurt the manufacturer and the supplier U in a competitive setting. This setting benefits R, however, because R holds a competitive advantage due to being reliable. It is natural, therefore, to expect that R prefers exponential disruptions, while U and the manufacturer prefer deterministic disruptions. Following, we show that these hold when R does not offer backup capacity. Through numerical examples, however, we demonstrate that this logic does not hold when backup capacity is available.

Proposition 3.12. *Assume that R does not offer backup capacity. For the same average length of disruptions, R is weakly better off with memoryless disruptions, whereas U and the manufacturer are weakly worse off.*

When backup capacity is available, the benefits are distributed differently. Numerical evidence suggests that the intuition is correct for the manufacturer, who is always worse off and for R, who is always better off with memoryless disruptions. Contrary to the intuition, U may be better or worse off. Figure 3.7 shows that the difference of supplier costs, $c_r - c_u$, and the manufacturer's holding cost influence the preference between memoryless and deterministic disruptions for U.

Figure 3.7: U's preferred disruption type

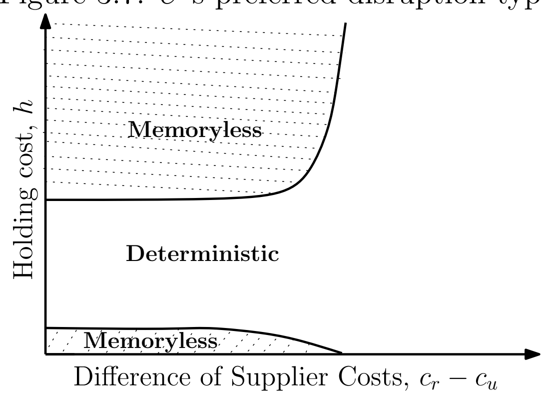

When U has a significant cost advantage ($c_r - c_u$ is very large), the manufacturer sole sources from U under both disruption types. In this case, U fulfills all of the orders, but at a higher wholesale price during deterministic disruptions; hence he is better off with deterministic disruptions. For smaller values of cost difference, U serves as the primary supplier, along with R as backup. In this case, U's portion of business is determined not only by the length of the disruptions and time between disruptions, but also by the inventory that the manufacturer holds. In addition to inventory, a secondary driver is the equilibrium wholesale price that U is able to charge for on-time deliveries during both of the disruption types.

Consider first the case when inventory holding cost is so high that the manufacturer carries little or no inventory for both disruption types. If the difference of supplier costs is small, R can profitably serve as both the primary and backup supplier. With deterministic disruptions, however, the benefits of serving as the backup supplier decreases compared to memoryless disruptions, and R is better off serving as the primary supplier. U, in order to become the primary supplier, must lower his price until R becomes indifferent between the two roles. As a result, contrary to conventional wisdom, U's profit is lower with deterministic disruptions (upper-left region in Figure 3.7). In other words, more variable disruptions may result in larger profits for U because it dampens the competition: if U had an option to reduce the variability of disruptions, he would be better off not doing so. The situation differs when the when inventory holding cost is so large that R could still be competitive as a primary supplier with memoryless disruptions, but is not competitive as a primary supplier with deterministic disruptions. With deterministic disruptions, U can take advantage of the situation and increase the price until R is indifferent between the two roles. Thus, U's profit increases with deterministic disruptions. With high holding costs, the effect of inventory changes is very small, and they do not influence the aforementioned logic.

However, with low holding costs, the manufacturer extensively uses inventory to buffer disruptions. With deterministic disruptions, inventory is lower compared to memoryless case. Thus, the profitability of serving as a backup supplier compared to serving as primary supplier is higher for R. Because U must lower his price until R becomes indifferent between the two roles, R's indifference occurs at a higher price with deterministic disruptions. Thus, deterministic disruptions lead to a higher wholesale price for U in equilibrium. Inventories are lower with deterministic disruptions, however, leading to a loss of business for U. Typically, the loss of business is very small compared to the increase in wholesale price, except when the holding cost is also very small.

• Lead-times

Our model applies to situations in which production lead-times are negligible, which is often the case; indeed, production lead times may be on the scale of days, whereas transportation lead-times are on the scale of weeks. Although our analysis assumes zero lead-times for transportation, the model can readily accommodate positive transportation lead-times when the reliable supplier's lead-time does not exceed that of the unreliable supplier. The manufacturer's ordering policy simply needs to be shifted to account for the transportation lead time.

3.7. Summary and Conclusions

In this chapter, we have considered a manufacturer's choice to source from either a perfectly reliable supplier, an unreliable supplier, or both where suppliers are active decision makers. We also evaluated the costs and benefits associated with flexible sourcing considering the suppliers' strategic price-setting behavior. We first showed that the manufacturer's optimal inventory policy can be described based on how many periods' worth of stock to keep, regardless of the magnitude of the demand in each period. By linking the optimal inventory policy to the classical newsvendor problem, we identified the trade offs that lead to the optimal number periods to cover for various types of stock-outs.

To evaluate the effect of strategic supplier behavior on the benefits of flexible sourcing, we considered two pricing games. We found that the single-wholesale price game led to a conflict of incentives in terms of the roles the suppliers want to play and the amount of the business they are awarded. We formally confirmed the non-existence of pure-strategy Nash equilibria in most practical situations. In contrast, second game, the contingent-pricing game, reflected a more intuitive relationship. Here, we showed that a unique pure-strategy Nash equilibrium always exits. Except for cases that result in significant enough cost advantages for one of the suppliers, we found that the manufacturer uses the less expensive, unreliable supplier and the more expensive, reliable supplier.

One of our findings is that with endogenously determined wholesale prices, the manufacturer does not necessarily benefit from having a backup supplier and, in fact, is typically worse off. Thus, an upfront commitment to sole sourcing and using simple wholesale pricing contracts may actually be beneficial, as opposed to creating an opportunity for one supplier to serve as a backup through more flexible contracts. Interestingly, suppliers may benefit from flexible sourcing even though the manufacturer does not. The reliable supplier always benefits from maintaining backup capacity, whereas the unreliable supplier might benefit, in some situations, from a reliable supplier's backup capacity despite reduced business volume. From a system perspective, a flexible sourcing strategy may degrade the supply chain performance.

Finally, we extended our results in two dimensions: (1) the possibility of having two unreliable suppliers and (2) the possibility of predictable recovery times. For the case of two unreliable suppliers, we derived the conditions under which one supplier can be treated as perfectly reliable. Although reducing variability is usually considered favorable in the operations literature, we find that the unreliable supplier may achieve higher profits with unpredictable (more variable) disruptions, which results from dampening the competition due to the availability of a backup supplier.

3.8. Appendix: A Numerical Example for the Contingent-Pricing Game

Consider demand of 1 unit per day, an average up-time $1/\theta_f$ of 120 days and average down-time $1/\theta_r$ of 30 days, with a probability of being operational, $\pi_{ON} = 80\%$. Assume that the selling price is $p = \$10$, U's cost is $c_u = \$0$, R's cost is $c_r = \$4$, and backlogging cost is $\pi_b = \$1$ per customer per day. During a disruption, the manufacturer expects to incur $(30)(\$1) = \$30$ per customer and, on average, manufacturer incurs $(0)(\pi_{ON}) + (30)\pi_{OFF} = \$6$ per customer for backlogs. For example, if the manufacturer sole-sources from U and the wholesale price is $w_u = \$1$, the manufacturer's average (per-period) profit is $10 - 1 - 6 = \$3$. If suppliers compete and R does not offer backup capacity, R wins and charges $w_r = \$6$, while U charges $w_u = \$0$ (and does not win).

Consider the case, when R offers backup capacity. Suppose that U charges $w_u = \$5$. To discount for delays, he would need to charge $\$5 - \$30$ for late orders. He prefers to set $w_{ud} = \$0$. In response R can charge $w_{rb} = \$0 + \$30$. To set w_r, R would be evaluating manufacturer's long-run average cost: $w_u \pi_{ON} + w_{rb} \pi_{OFF} =$ (5)(0.80) + (30)(0.20) =10 and charge $w_r = \$10$. The profit as a backup supplier is $(0.20)(\$30 - \$4) = \$5.2$, while profit when being primary supplier is $(1.0)(\$10 - \$4) = \$6$. In this case, R would prefer to be primary supplier. If, however, U chooses price $w_u \leq \$4$, R prefers to be backup supplier. At $w_u = \$4$, profit of unreliable supplier is $(0.8)(4) = \$3.2$ (versus \$0 when R does not offer backup capacity) and profit of reliable one is \$5.2 (versus 6-4=\$2 when R does not offer backup capacity). On the other hand, manufacturer's profit is 10-(4)(0.80)-(30)(0.20)=\$0.80 (versus 10-6=\$4 when R does not offer backup capacity).

3.9. Appendix: Mathematical Proofs

Proof of Proposition 3.1. *(i)* Proof is by induction. We normalize demand d_T to 1. It can be shown inductively that $G_t(y)$ is linear in y and $V_t(x, s)$ is linear in x for each s, where $x, y \in [0, 1]$. As a result, if it is optimal to order in period t, it is optimal to order all requirements. Since the optimal decision to order does not depend on x, we assume that $x = 0$ in the beginning of the horizon. As a result, we have $V_t(0, ON) = \min\{G_t(0), G_t(1)\}$. Clearly, $G_t(1) = h(T - t)$ and we derive $G_t(0)$ as follows. In the last period, $s_T^* = 1$. Assume that the $s_{t+1}^* = \cdots = s_T^* = 1$. Denote by $\gamma = 1 - \theta_r$. Clearly, $V_{t+1}(0, \mathrm{ON}) = h(T - t - 1)$.

$$\begin{aligned} V_{t+1}(0, \mathrm{OFF}) &= \sum_{k=1}^{T-t-1} (1-\theta_r)^{k-1}\theta_r[h(T-t-1-k)] + \gamma^{T-t-1}\left(\frac{\pi_b}{\theta_r}\right) \\ &= h(T-t-1) - \frac{h}{\theta_r} + \gamma^{T-t-1}\left(\frac{\pi_b + h}{\theta_r}\right) \end{aligned}$$

Thus, $V_{t+1}(0, \text{OFF})$ captures the fact that we are going to place an order as soon as the supplier becomes available. Since $G_t(0) = (1-\theta_f)V_{t+1}(0, \text{ON}) + \theta_f V_{t+1}(0, \text{OFF})$, we immediately have the following.

$$G_t(1) - G_t(0) \;=\; h - \theta_f \left[\gamma^{T-t-1}\left(\frac{\pi_b + h}{\theta_r}\right) - \frac{h}{\theta_r}\right]$$

We have $s_t^* = 1$ if $G_t(1) < G_t(0)$, or equivalently, $F(T-t-1) > \frac{\pi_b}{\pi_b+h}$. Let k be the period satisfying $F(T-k-1) \leq \frac{\pi_b}{\pi_b+h} < F(T-k-2)$. We will show that $s_t^* = 0$ for all $t \leq k$. The $n-$step transition matrix is derived in Lewis (2005), see below, where 0 and 1 denote OFF and ON. It can be shown that $p_{10}^{(n)} \leq n\theta_f$ (omitted).

$$\begin{pmatrix} p_{11}^{(n)} & p_{10}^{(n)} \\ p_{01}^{(n)} & p_{00}^{(n)} \end{pmatrix} \;=\; \frac{1}{\theta_f + \theta_r}\left\{\begin{pmatrix} \theta_r & \theta_f \\ \theta_r & \theta_f \end{pmatrix} + [1-\theta_f-\theta_r]^n \begin{pmatrix} \theta_f & -\theta_f \\ -\theta_r & \theta_r \end{pmatrix}\right\}$$

Assume that $s_{t+1} = \cdots = s_k^* = 0$ and $s_{k+1}^* = \cdots = s_T^* = 1$. Consider period t. Since it is not optimal to order until period $k+1$, the following holds.

$$\begin{aligned} G_t(y) &= h(t-k)y + p_{11}^{(t-k)} V_{k+1}(y, \text{ON}) + p_{10}^{(t-k)} V_{k+1}(y, \text{OFF}) \\ G_t(1) - G_t(0) &= h(t-k) - p_{10}^{(t-k)}\left[\gamma^{T-k-1}\left(\frac{\pi_b + h}{\theta_r}\right) - \frac{h}{\theta_r}\right] \;\geq\; 0 \end{aligned}$$

The last inequality is due to $p_{10}^{(t-k)} \leq (t-k)\theta_f$ and $F(T-k-1) \leq \frac{\pi_b}{\pi_b+h}$, therefore, $s_t^* = 0$. As a conclusion, it is optimal to order only during the last κ periods, where $F(\kappa) \geq \frac{\pi_b}{\pi_b+h} > F(\kappa-1)$.

(ii) The manufacturer faces a disruption if $X'_T = OFF(k)$, where $k \geq \kappa$. Since X'_t is in steady state at $t = 0$, the probability of that happening is $\bar{F}(\kappa)$ and the total expected holding and penalty cost is $L(\kappa) = \pi_{ON} V_{T-\kappa}(0, \text{ON}) + \pi_{OFF} V_{T-\kappa}(0, \text{OFF})$. That is, the manufacturer will not order until period $T - \kappa$. Using the expressions for the value function in part *(i)*, we obtain the desired statement for $L(\kappa)$. By

investigating the difference $L(k+1) - L(k) = h - (\pi_b + h)(1 - F(k))$, we can readily show that $L(k)$ is discrete-convex in k and minimized at $k = \kappa$. $\square$

Proof of Theorem 3.1. *(i)* Since the demand is deterministic, we view the sourcing problem as when to procure inventory for each future period. The optimal procurement problem for a particular future period can be viewed as a simplified problem described earlier. As a consequence of Proposition 3.1, it is optimal to procure inventory for the current period and next κ periods that follow. Since so defined myopic policy is feasible (inventory is never above the target), by the logic of Veinott (1966) it is optimal. *(ii)* This follows from the fact that the periods can be decoupled and that X_t' is in steady state at $t = 0$. $\square$

Proof of Theorem 3.2. The analysis in Proposition 3.1 can be repeated by letting $V_T(x, \text{OFF}) = \sigma(d_T - x)$, where σ is the expected cost for not satisfying demand. Under PLS(α), we clearly have $\sigma = \alpha\left(\frac{\pi_b}{\theta_r}\right) + (1-\alpha)(p+\pi_l - w_u)$, accounting for the lost profit and goodwill penalty. For GLS(α), $\sigma = E_\tau\left[\sum_{t=1}^{\tau} \alpha^t \pi_b\right] + (1 - E_\tau[\alpha^\tau])(p+\pi_l-w_u)$. Since τ is a geometric random variable, we have $E[\alpha^\tau] = \frac{\alpha\theta_r}{1-\alpha+\alpha\theta_r}$, leading to $\sigma = \left(\frac{\alpha\theta_r}{1-\alpha+\alpha\theta_r}\right)\left(\frac{\pi_b}{\theta_r}\right) + \left(\frac{1-\alpha}{1-\alpha+\alpha\theta_r}\right)(p + \pi_l - w_u)$. $\square$

Proof of Proposition 3.2. Let $\beta = 0$. Then, only sole-sourcing strategies are available. Manufacturer's profit is $p - w_r$ and $p - w_u - L(\kappa_0)$ when he sole-sources from R and U respectively. The optimal strategy is to solely source from U and cover κ_0 periods if $w_r > w_u + L(\kappa_0)$. Let $\beta = 1$. If $w_r < w_u$, it is optimal to solely source from R. If $w_r - w_u \geq \frac{\pi_b}{\theta_r}$, it is not optimal to source from R, at all. Otherwise, R is used as backup supplier and the inventory level is κ_1, using Theorem 3.2 with $\sigma = w_r - w_u$.

Let $0 < \beta < 1$. Assume that R is available only as a backup source. During a disruption, the manufacturer may start sourcing from R before running out of inventory. In that case, the manufacturer uses R's inventory first to satisfy demand, therefore, R's inventory is not carried from one period to another, which is simply relabeling how the inventory on-hand is used. Without loss of generality, we consider an infinite-horizon setting. Let x_k and z_k denote the fractions of demand satisfied from the safety stock and from R respectively in the k^{th} period of a disruption. Thus, $1 - x_k - z_k$ units of demand is backlogged. Manufacturer's safety stock is $\sum_{k=1}^{\infty} x_k$, where x_k's are zero for sufficiently large k.

Let $x = (x_1, x_2, \cdots)$ and $x = (z_1, z_2, \cdots)$. At the end of the k^{th} period of a disruption, the manufacturer will be left with $I_k(x) := \sum_{j=k+1}^{\infty} x_j$ units of inventory, while he will accumulate a backlog of $B_k(x,z) := \sum_{j=1}^{k}(1 - x_j - z_j)$ customers. The long-run average cost of the manufacturer is $C(x,z) = \pi_{ON} h I_0(x) + \sum_{k=1}^{\infty} \pi_{OFF(k)} \left[h I_k(x) + \pi_b B_k(x,z) + (w_r - w_u) z_k \right]$. By rearranging the terms, and letting $a_k = -h + \left(\frac{\pi_b + h}{\theta_r}\right) \pi_{OFF(k)}$ and $b_k = \left(w_u + \frac{\pi_b}{\theta_r} - w_r\right) \pi_{OFF(k)}$, the manufacturer's problem is stated as follows.

$$\begin{aligned} \min \quad & C(x,z) = \frac{\pi_b}{\theta_r} \pi_{OFF} - \sum_{k=1}^{\infty} [a_k x_k + b_k z_k] \\ s.t \quad & 0 \le x_k \le 1,\ 0 \le z_k \le \beta,\ x_k + z_k \le 1,\ k = 1, 2, \cdots \end{aligned}$$

Since $C(x,z)$ is linear and separable in x_k and z_k, the problem reduces to solving an LP for each k, where the objective is to maximize $a_k x_k + b_k z_k$ subject to the same set of constraints. It can be readily shown that $a_k \ge b_k$ only for $k \le \kappa_1$, $a_k \ge 0$ only for $k \le \kappa_0$, and $b_k \ge 0$ for all k, provided that $w_r \le w_u + \frac{\pi_b}{\theta_r}$. As a result, the solution to the LP for any k is described as follows. (Denote by x^* and z^* the optimal decisions).

i) Let $w_r \le w_u + \frac{\pi_b}{\theta_r}$. If $\kappa_0 = 0$, then $a_k < 0 \le b_k$ for all k, and hence, $x_k^* = 0$ and $z_k^* = \beta$. Let $\kappa_0 > 0$. For $k \le \kappa_1$, $a_k \ge b_k \ge 0$, hence, $x_k^* = 1$ and $z_k^* = 0$. For $\kappa_1 < k \le \kappa_0$, $0 \le a_k \le b_k$, hence $x_k^* = 1 - \beta$ and $z_k^* = \beta$. Finally, for $k > \kappa_0$, $a_k < 0 \le b_k$, hence $x_k^* = 0$ and $z_k^* = \beta$. As a result, the manufacturer carries $I(x^*) = (1-\beta)\kappa_0 + \beta\kappa_1$ units in the safety stock and sources from R if a disruption takes longer than κ_1 periods. *ii)* Let $w_r > w_u + \frac{\pi_b}{\theta_r}$. As a result, $z_k^* = 0$ for all k since $b_k < 0$. In addition, $x_k^* = 1$ for $k \leq \kappa_0$ and $x_k^* = 0$ otherwise as in case *(i)*. As a result, the manufacturer carries $I(x^*) = \kappa_0$ units in the safety stock and does not source from R.

This suggests that we do not lose optimality by splitting demand into two portions with sizes β and $1-\beta$ respectively, and providing only the first portion access to the backup supplier. This separation allows us to compare sourcing alternatives in a convenient way. For each portion of demand, the optimal policy is as presented for $\beta = 0$ and $\beta = 1$ above. $\square$

Proof of Proposition 3.4. Denote by $\Pi_{rb}(w_r, w_u)$ R's profit. The subscript b stands for backup. The maximal price that R can charge is $w_u + \frac{\pi_b}{\theta_r}$, hence, if $w_u < c_r - \frac{\pi_b}{\theta_r}$, $w_{rb}^*(w_u) = c_r$, and R does not serve as backup supplier. Let $w_u \geq c_r - \frac{\pi_b}{\theta_r}$. We consider two sub-cases. If $\frac{\pi_b}{\theta_r} \leq \frac{h}{\theta_f}$, a zero-inventory policy is optimal for all $w_u \leq w_r \leq w_u + \frac{\pi_b}{\theta_r}$, and $\Pi_{rb}(w_r, w_u) = \pi_{OFF}(w_r - c_r)$, implying that $w_{rb}^*(w_u) = w_u + \frac{\pi_b}{\theta_r}$. Let $\frac{\pi_b}{\theta_r} > \frac{h}{\theta_f}$. Since a zero-inventory policy is optimal when $w_r \leq w_u + \frac{h}{\theta_f}$, $w_{rb}^*(w_u) \geq w_u + \frac{h}{\theta_f}$. Hence, we consider the range $w_u + \frac{h}{\theta_f} \leq w_r \leq w_u + \frac{\pi_b}{\theta_r}$. For such w_r, manufacturer's inventory satisfies $F(\kappa) = \frac{w_r - w_u}{w_r - w_u + \frac{h}{\theta_r}}$.We equivalently view R's price setting problem as one where she sets inventory κ, implying a wholesale price $w_r = w_u + \frac{h}{\theta_r}\left(\frac{F(\kappa)}{1-F(\kappa)}\right)$. Using this relationship, we state R's profit as a function of inventory.

$$\Pi_{rb}(w_r, w_u) \;=\; [1 - F(\kappa)](w_r - c_r) \;=\; w_u - c_r + \left(c_r + \frac{h}{\theta_r} - w_u\right) F(\kappa)$$

If $w_u < c_r + \frac{h}{\theta_r}$, R's profit is decreasing in κ, hence the price should be set such that $\kappa = 0$, implying, $w_r = w_u + \frac{h}{\theta_f}$. If $w_u \geq c_r + \frac{h}{\theta_r}$, R profit is increasing in κ, hence R should set the highest possible wholesale price, $w_r = w_u + \frac{\pi_b}{\theta_r}$. Consequently, inventory satisfies $F(\kappa) = \frac{\pi_b}{\pi_b + h}$. $\square$

Proof of Proposition 3.5. We first derive R's best response for a given w_u.

Denote by $\Pi_r^*(w_u)) = w_u - c_r$ R's profit as the primary supplier with $w_r = w_u$. Let $\Pi_{rb}^*(w_u) = \max_{w_r > w_u} \Pi_{rb}(w_r, w_u)$. We compare R's profits under each role. If $\frac{\pi_b}{\theta_r} \leq \frac{h}{\theta_f}$, $w_{rb}^*(w_u) = w_u + \frac{\pi_b}{\theta_r}$ by Proposition 3.4 and manufacturer holds no inventory. As a result, $\Pi_{rb}^*(w_u) = \left(w_u + \frac{\pi_b}{\theta_r} - c_r\right)\pi_{OFF}$ and $\Pi_r^*(w_u) \geq \Pi_{rb}^*(w_u)$ if and only if $w_u \geq c_r + \frac{\pi_b \theta_f}{\theta_r^2} = c_r + \frac{\theta_f}{\theta_r} \min\left\{\frac{\pi_b}{\theta_r}, \frac{h}{\theta_f}\right\}$. Let $\frac{\pi_b}{\theta_r} > \frac{h}{\theta_f}$. We consider two sub-cases. If $w_u \leq c_r + \frac{h}{\theta_r}$, $w_{rb}^*(w_u) = w_u + \frac{\pi_b}{\theta_r}$ and the inventory satisfies $F(\kappa_0) = \frac{\pi_b}{\pi_b + h}$. Therefore, $\Pi_{rb}^*(w_u) = [1 - F(\kappa_0)](w_r - c_r) = \left(\frac{h}{\pi_b + h}\right)\left(w_u + \frac{\pi_b}{\theta_r} - c_r\right)$, and

$$\begin{aligned}
\Pi_{rb}^*(w_u) - \Pi_r^*(w_u) &= \frac{\pi_b}{\theta_r}\left(\frac{h}{\pi_b + h}\right) - \left(\frac{\pi_b}{\pi_b + h}\right)(w_u - c_r) \\
&\geq \frac{\pi_b}{\theta_r}\left(\frac{h}{\pi_b + h}\right) - \left(\frac{\pi_b}{\pi_b + h}\right)\frac{h}{\theta_r} = 0
\end{aligned}$$

Thus, serving as the backup supplier is more profitable. If $w_u \geq c_r + \frac{h}{\theta_r}$, $w_{rb}^*(w_u) = w_u + \frac{h}{\theta_f}$ and manufacturer holds zero inventory. Therefore, $\Pi_{rb}^*(w_u) = \left(w_u + \frac{h}{\theta_f} - c_r\right)\pi_{OFF}$, and, $\Pi_r^*(w_u) - \Pi_{rb}^*(w_u) = \left(w_u - c_r - \frac{h}{\theta_r}\right)\pi_{ON} \geq 0$. Thus, serving as the primary supplier is more profitable if $w_u \geq c_r + \frac{h}{\theta_r} = c_r + \frac{\theta_f}{\theta_r}\min\left\{\frac{\pi_b}{\theta_r}, \frac{h}{\theta_f}\right\}$.With this, we proceed to equilibrium analysis. Denote by $\delta = \frac{\theta_f}{\theta_r}\min\left\{\frac{\pi_b}{\theta_r}, \frac{h}{\theta_f}\right\}$.

i) Let $c_u \geq c_r + \delta$,. Since $w_u \geq c_u \geq c_r + \delta$, R prefers to serve as the primary supplier, hence $w_r \leq w_u$. Clearly, $(w_r^*, w_u^*) = (c_u, c_u)$ is a Nash-equilibrium. To show uniqueness, assume that there exists another pure-strategy Nash equilibrium, (w_r, w_u). Let $w_r < w_u$. In that case, R can profitably increase w_r to $w_r + \epsilon$ for small enough $\epsilon > 0$. So, $w_r = w_u$. If $c_u < w_r = w_u$, then U can profitably decrease w_u to $w_u - \epsilon$ for $\epsilon > 0$ small enough. Hence, $w_r = w_u = c_u$ and it is the unique equilibrium.

ii) Let $c_u < c_r + \delta$. To reach a contradiction, assume that there exists a pure-strategy Nash equilibrium . If $c_u \leq w_r < w_u$, then, U can profitably decrease w_u to w_r, thus, $w_r \geq w_u$. Also, $w_u \leq c_r + \delta$, otherwise, R could profitably undercut U's price and become the primary supplier. When $w_u \leq c_r + \delta$, R serves as the backup supplier and sets $w_r = w_u + \frac{\pi_b}{\theta_r}$. In this case, U can profitably increase price, for example, set wholesale price equal to w_r. Therefore, there exists no pure-strategy Nash equilibrium when $c_u < c_r + \delta$. $\square$

Proof of Lemma 3.1. Due to rebate ϕ, manufacturer's effective backlogging cost is $\pi_b - \theta_r\phi$, hence results for optimal policy in the single-wholesale price game can be used with minor modifications. Let $\kappa := \kappa_{cp}(\phi)$ for conciseness in the rest of the proof. Clearly, if the manufacturer sole-sources from U, his inventory level is κ and his expected cost is $w_u F(\kappa) + w_{ud}\overline{F}(\kappa) + L(\kappa)$. R's objective is to maximize profits, while setting a price that is acceptable to the manufacturer. Thus, if R serves as the primary supplier, she sets $w_r = w_u F(\kappa) + w_{ud}\overline{F}(\kappa) + L(\kappa) = w_u + L(\kappa) - \phi\overline{F}(\kappa)$. On the other hand, if she serves as the backup supplier, due to Proposition 3.4, she sets $w^*_{rb} = w_{ud} + \frac{\pi_b}{\theta_r}$, if $w_{ud} < c_r + \frac{h}{\theta_r}$, and $w^*_{rb} = w_{ud} + \min\left\{\frac{\pi_b}{\theta_r}, \frac{h}{\theta_f}\right\}$ otherwise. When R can profitably play both roles, she sets the optimal price for the more profitable role, while setting sufficiently high price for the other role, resulting in the same profit. Clearly, R cannot make profit serving as the backup supplier if $w_u - \phi + \frac{\pi_b}{\theta_r} \leq c_r$. Similarly, R cannot make profit by serving as the primary supplier if $w_u + L(\kappa) - \phi\overline{F}(\kappa) \leq c_r$. If both inequalities hold, R cannot compete at all, thus, she sets $w_r = w_{rb} = c_r$.

To derive R's best response, we compare the profitability of each strategy: serve as primary supplier vs. serve as backup supplier. Denote by $\Pi^*_r(w_u, \phi) = w_u + L(\kappa) - \phi\overline{F}(\kappa) - c_r$ and $\Pi^*_{rb}(w_u, \phi) = \bar{F}(\kappa)(w_{rb} - c_r)$ R's profit as primary supplier and backup supplier respectively, when she best responds. First, let $\phi \geq \frac{\pi_b}{\theta_r}$, in which case, $\kappa = 0$, $w_{rb} = w_{ud} + \frac{\pi_b}{\theta_r}$, $\Pi^*_r(w_u, \phi) = w_u + \left(\frac{\pi_b}{\theta_r} - \phi\right)\pi_{OFF} - c_r$ and $\Pi^*_{rb}(w_u, \phi) = \pi_{OFF}\left(w_u + \frac{\pi_b}{\theta_r} - \phi - c_r\right)$. Serving as primary supplier is not profitable if $w_u + \left(\frac{\pi_b}{\theta_r} - \phi\right)\pi_{OFF} < c_r$. Also, serving as a backup supplier is not profitable if $w_u + \frac{\pi_b}{\theta_r} - \phi < c_r$. Assume that both roles are profitable. Since, $w_u + \left(\frac{\pi_b}{\theta_r} - \phi\right)\pi_{OFF} \geq c_r$ and $\phi \geq \frac{\pi_b}{\theta_r}$, we must have $w_u \geq c_r$. Since $\Pi^*_r(w_u, \phi) - \Pi^*_{rb}(w_u, \phi) = (w_u - c_r)\pi_{ON} \geq 0$, serving as the primary supplier is more profitable. Therefore, $w^l_u(\phi) = w^h_u(\phi) = c_r + \left(\phi - \frac{\pi_b}{\theta_r}\right)\pi_{OFF}$.

Assume next that $0 \leq \phi \leq \frac{\pi_b}{\theta_r}$. First, let $w_{ud} = w_u - \phi > c_r + \frac{h}{\theta_r}$. In this case, $\kappa = 0$ and $w_{rb} = w_u + \min\left\{\frac{\pi_b}{\theta_r}, \frac{h}{\theta_f}\right\}$. If $\frac{\pi_b}{\theta_r} \leq \frac{h}{\theta_f}$, and a zero-inventory policy is optimal, $\Pi_r^*(w_u,\phi) - \Pi_{rb}^*(w_u,\phi) = \pi_{ON}(w_u - c_r) \geq 0$. Let $\frac{\pi_b}{\theta_r} > \frac{h}{\theta_f}$. In this case, $w_{rb} = w_u - \phi + \frac{h}{\theta_f}$. For this case, it can readily be shown that $\Pi_r^*(w_u,\phi) - \Pi_{rb}^*(w_u,\phi)$ is increasing in w_u and $\Pi_r^*(c_r + \frac{h}{\theta_r} + \phi, \phi) - \Pi_{rb}^*(c_r + \frac{h}{\theta_r} + \phi, \phi) \geq 0$, implying that $\Pi_r^*(w_u,\phi) \geq \Pi_{rb}^*(w_u,\phi)$ for $w_u - \phi \geq c_r + \frac{h}{\theta_r}$, that is, serving as the primary supplier is more profitable for R. Lastly, assume $w_u - \phi < c_r + \frac{h}{\theta_r}$,thus, $w_{rb} = w_u - \phi + \frac{\pi_b}{\theta_r}$. Then, $\Pi_r^*(w_u,\phi) = w_u + L(\kappa) - \phi\overline{F}(\kappa) - c_r$ and $\Pi_{rb}^*(w_u,\phi) = \bar{F}(\kappa)\left(w_u + \frac{\pi_b}{\theta_r} - \phi - c_r\right)$. Observing that $\Pi_r^*(w_u,\phi) - \Pi_{rb}^*(w_u,\phi)$ is increasing in w_u, there is a unique solution to the equation $\Pi_r^*(w_u,\phi) = \Pi_{rb}^*(w_u,\phi)$ for a fixed ϕ, denoted by $w_u^h(\phi)$. Rearranging the terms, we obtain $w_u^h(\phi) = c_r - \frac{H(\kappa)}{F(\kappa)}$. Thus, serving as the primary supplier is optimal only if $w_u \geq w_u^h(\phi)$. For all cases with $0 \leq \phi \leq \frac{\pi_b}{\theta_r}$, serving as backup supplier is not profitable for R when $w_u - \phi + \frac{\pi_b}{\theta_r} \leq c_r$. Defining $w_u^l(\phi) = c_r - \frac{\pi_b}{\theta_r} + \phi$, serving as backup supplier is optimal only when $w_u^l(\phi) \leq w_u \leq w_u^h(\phi)$.

Finally, we show the monotonicity of $w_u^h(\phi)$. For $\phi \geq \frac{\pi_b}{\theta_r}$, $w_u^h(\phi) = c_r + \left(\phi - \frac{\pi_b}{\theta_r}\right)\pi_{OFF}$, is increasing in ϕ, with a slope equal to $\pi_{OFF} \leq 1$. For $\frac{\pi_b}{\theta_r} - \frac{h}{\theta_f} \leq \phi \leq \frac{\pi_b}{\theta_r}$, $w_u^h(\phi) = c_r$ is a constant. Let $0 \leq \phi \leq \frac{\pi_b}{\theta_r} - \frac{h}{\theta_f}$. For analytical convenience, we assume continuous inventory replenishment with exponential up-times and down-times, thus, $F(\kappa) = 1 - \pi_{OFF}e^{-\theta_r\kappa}$. As a result, $F'(\kappa) = \theta_r\bar{F}(\kappa)$ and $H'(\kappa) = hF(\kappa)$, where $H(\kappa) = h\kappa + \frac{h}{\theta_r}\bar{F}(\kappa) - \frac{h}{\theta_r}\pi_{OFF}$ is the expected holding cost. Defining $W(\kappa) = c_r - \frac{H(\kappa)}{F(\kappa)}$, we have $w_u^h(\phi) = W(\kappa_{cp}(\phi))$. Denote by $\mathcal{D}_x$ the derivative operator with respect to a variable x. We need $\mathcal{D}_\phi w_u^h(\phi) = \mathcal{D}_\kappa W(\kappa_{cp}(\phi))\mathcal{D}_\phi\kappa_{cp}(\phi)$. With straightforward calculation, we have $\mathcal{D}_\kappa W(\kappa) = \frac{H(\kappa)\theta_r\bar{F}(\kappa)}{F^2(\kappa)} - h$ and $\mathcal{D}_\phi\kappa_{cp}(\phi) = -\frac{1}{\pi_b + h - \theta_r\phi} \leq 0$. To show that $\mathcal{D}_\phi w_u^h(\phi) \geq 0$, it suffices to show that $\mathcal{D}_\kappa W(\kappa) \leq 0$, or alternatively, $H(\kappa)\theta_r\bar{F}(\kappa) \leq hF^2(\kappa)$. Define $R(\kappa) := [hF^2(\kappa) - H(\kappa)\theta_r\bar{F}(\kappa)]/h = 1 - \pi_{OFF}e^{-\theta_r\kappa}(\kappa\theta_r + 1 + \pi_{ON})$. Since $R'(\kappa) = \theta_r\pi_{OFF}e^{-\theta_r\kappa}(\kappa\theta_r + \pi_{ON}) \geq 0$, and $R(0) = \pi_{ON}^2 \geq 0$, we have that $R(\kappa) \geq 0$ for all $\kappa \geq 0$, implying that $\mathcal{D}_\phi w_u^h(\phi) \geq 0$. Next, we show that

$\mathcal{D}_\phi w_u^h(\phi) \leq \pi_{OFF}$. It can be shown that the threshold $w_u^h(\phi)$ is convex in ϕ following routine arguments as above. Given that $w_u^h(\phi)$ is convex and increasing in ϕ, $\mathcal{D}_\phi w_u^h(\phi) \leq \mathcal{D}_\phi w_u^h\left(\frac{\pi_b}{\theta_r} - \frac{h}{\theta_f}\right) = \pi_{OFF}$ for all $\phi \in \left[0, \frac{\pi_b}{\theta_r} - \frac{h}{\theta_f}\right]$. $\square$

Proof of Proposition 3.6. To show that $0 \leq \Delta \leq L(\kappa_0)$, assume first that $\frac{\pi_b}{\theta_r} \leq \frac{h}{\theta_f}$, implying that $\kappa_0 = 0$, $\Delta = 0$ and $L(\kappa_0) = \frac{\pi_b}{\theta_r}\pi_{OFF}$. Hence, $0 \leq \Delta \leq L(\kappa_0)$ holds. Let $\frac{\pi_b}{\theta_r} > \frac{h}{\theta_f}$, hence, $\kappa_0 > 0$ and $F(\kappa_0) = \frac{\pi_b}{\pi_b+h}$. Clearly, $\Delta \geq 0$. Note that $w_u^h(0) = c_r - \Delta$. Due to Lemma 3.1, we have $c_r - \Delta = w_u^h(0) \geq w_u^l(0) = c_r - \frac{\pi_b}{\theta_r}$, implying that $\Delta \leq \frac{\pi_b}{\theta_r}$. With straightforward calculation, $L(\kappa_0) - \Delta = \bar{F}(\kappa_0)\left(\frac{\pi_b}{\theta_r} - \Delta\right) \geq 0$, hence $\Delta \leq L(\kappa_0)$.

Without loss of generality, we assume that $0 \leq \phi = w_u - w_{ud} \leq \frac{\pi_b}{\theta_r}$ in the equilibrium, since zero-inventory policy is optimal for $\phi \geq \frac{\pi_b}{\theta_r}$, and an alternative equilibrium with $\phi \leq \frac{\pi_b}{\theta_r}$ can be constructed, where all the players make the same profit. Let $c_u > c_r$. For any $w_u \geq c_u$ and $0 \leq \phi \leq \frac{\pi_b}{\theta_r}$, R's best response is to set prices such that she serves as the sole supplier, due to Lemma 3.1. Therefore, U cannot compete with R when $c_u > c_r$, hence, R serves as the sole-supplier. Let $c_r \leq c_u$. To be competitive, U cannot set a price larger than c_r, that is, $w_u \leq c_r$, also implying that $w_{ud} \leq c_r$. Whenever R serves as the backup supplier, R's optimal backup price is $w_{rb} = w_{ud} + \frac{\pi_b}{\theta_r}$, due to Lemma 3.1. The equilibrium possesses the following properties (Properties I and II).

Property I: If the manufacturer primarily sources from U and uses R as the backup supplier in the equilibrium, then $w_{ud}^* = c_u$.

Proof of Property I: Let $w_{ud} > c_u$ in the equilibrium. R's best response is $w_{rb} = w_{ud} + \frac{\pi_b}{\theta_r}$, resulting in an inventory of $\kappa = \kappa_{cp}(\phi)$. In such a case, U has an incentive to increase the penalty ϕ slightly to $\phi + \epsilon$, where $\epsilon > 0$ sufficiently (infinitesimally) small. In this case, the manufacturer sole-sources from U and still holds the same amount of inventory, κ. U's profit increases by $(w_{ud} - c_u)F(\kappa)$. Therefore, we cannot have $w_{ud} > c_u$, if both suppliers are used in the equilibrium. Thus, $w_{ud} = c_u$.

Property II: If the manufacturer sole-sources from U, then $w_{ud}^* = c_r - \frac{\pi_b}{\theta_r}$.

Proof of Property II: If U is the sole-supplier in the equilibrium, then, we must have $w_u - \phi \leq c_r - \frac{\pi_b}{\theta_r}$, due to Lemma 3.1. Let $w_u - \phi < c_r - \frac{\pi_b}{\theta_r}$. In this case, U can improve profit by keeping ϕ fixed, but slightly increasing w_u. This causes no change in the manufacturer's inventory, and it results in improvement in U's profit. Hence, $w_u - \phi = w_{ud} = c_r - \frac{\pi_b}{\theta_r}$ in the equilibrium.

(i) Let $c_r \leq c_u + \Delta$. To reach a contradiction, assume that U is the sole-supplier, thus U sets $w_{ud} = c_r - \frac{\pi_b}{\theta_r} < c_u$ due to Property II, hence R's best response is to set $w_r = w_{rb} = c_r$. When that is the case, U has incentive to slightly increase w_{ud} (keeping w_u unchanged) so that manufacturer uses R as the backup source. By doing so, manufacturer's inventory remains the same, and U does not have to incur loss during a disruption, significantly improving profits, reaching a contradiction. Similarly, assume that U is the primary supplier in equilibrium. In this case, $w_{ud} = c_u$ due to Property I. Since $w_u^h(\phi)$ is increasing in ϕ with a slope less than 1 (due to Lemma 3.1), for any (w_u, ϕ) with $w_u - \phi = c_u$, we have that $w_u > w_u^h(\phi)$, where R's best response is to serve as the sole supplier, reaching a contradiction. Since U is unable to serve the manufacturer, U sets $w_u = w_{ud} = c_u$, while R sets $w_r^* = c_u + L(\kappa_0)$ and $w_{rb}^* = \infty$ (or sufficiently large). Clearly, none of the suppliers have incentive to unilaterally deviate from the equilibrium.

(ii) Let $c_u + \Delta < c_r - c_u < \frac{\pi_b}{\theta_r}$. U does not serve as the sole supplier in equilibrium, by the exact same reasoning as in part *(i)*. Therefore, U serves as the primary supplier along with R as backup. Due to Property I, we have $w_{ud}^* = c_u$ and $w_{rb}^* = c_u + \frac{\pi_b}{\theta_r}$. To derive w_u^*, define the set $\mathcal{F}_u = \{(w_u, \phi) \mid w_u - \phi = c_u,\ w_u^l(\phi) \leq w_u \leq w_u^h(\phi)\}$. Due to Lemma 3.1 and Property I, $(w_u^*, \phi^*) \in \mathcal{F}_u$. Since $w_{rb}^* = c_u + \frac{\pi_b}{\theta_r}$ does not depend on w_u^*, (w_u^*, ϕ^*) is one that maximizes U's profit on the set $\mathcal{F}_u$. For any $(w_u, \phi) \in \mathcal{F}_u$ and $w_{rb} = c_u + \frac{\pi_b}{\theta_r}$, manufacturer's inventory is $\kappa_{cp}(\phi)$ due to Lemma 3.1, therefore, U's profit is $(w_u - c_u)F(\kappa_{cp}(\phi))$, and $(w_u^*, \phi^*) = \arg\max_{(w_u,\phi)\in\mathcal{F}_u}(w_u -$

$c_u)F(\kappa_{cp}(\phi))$. Let $\bar{\phi} = \sup\{\phi \mid (w_u, \phi) \in \mathcal{F}_u\}$. Since $w_u - \phi = c_u$, equivalently, $\phi^* = \arg\max_{0 \le \phi \le \bar{\phi}} \phi F(\kappa_{cp}(\phi))$, leading to $w_u^* = c_u + \phi^*$ and $\kappa^* = \kappa_{cp}(\phi^*)$. Next, we describe how (w_u^*, ϕ^*) is actually derived. Define $\Pi_u(\phi) := \phi F(\kappa_{cp}(\phi))$ as U's profit.

- If $\frac{\pi_b}{\theta_r} \le \frac{h}{\theta_f}$, then a zero-inventory policy is always optimal, thus, $F(\kappa_{cp}(\phi)) = F(0) = \pi_{ON}$. Therefore, the function $\Pi_u(\phi) = \phi\pi_{ON}$ is linear and increasing in ϕ. In addition, $w_u^h(\phi) = c_r$ for all $0 \le \phi \le \frac{\pi_b}{\theta_r}$. As a result, $w_u^* = c_r$ and $\phi^* = \bar{\phi} = c_r - c_u$. See Figure 3.8(a) for an illustration.

Figure 3.8: U's equilibrium pricing

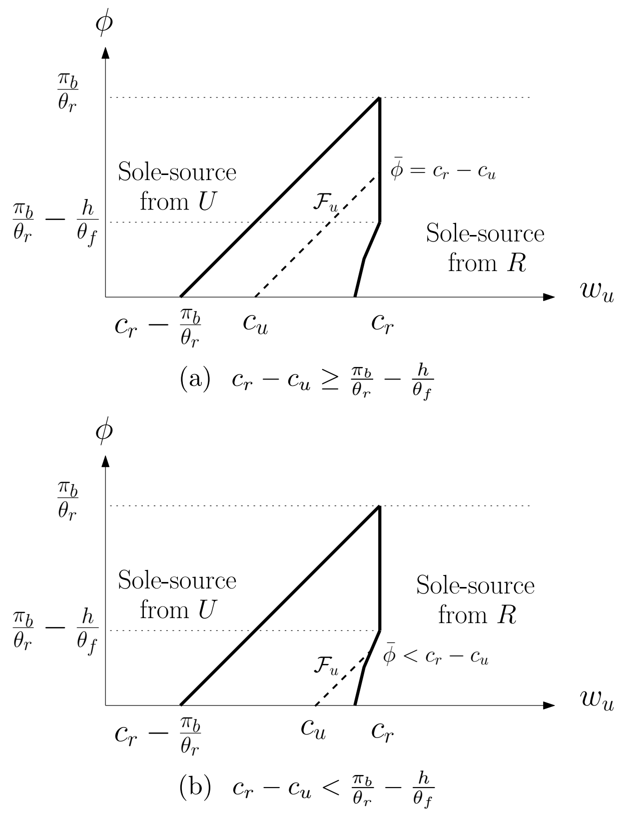

- If $\frac{\pi_b}{\theta_r} > \frac{h}{\theta_f}$, then a zero-inventory policy is optimal only when $\phi \ge \frac{\pi_b}{\theta_r} - \frac{h}{\theta_r}$, leading to $w_u^h(\phi) = c_r$. If $c_r - c_u \ge \frac{\pi_b}{\theta_r} - \frac{h}{\theta_r}$, then $\bar{\phi} = c_r - c_u$. Otherwise, $\bar{\phi} < c_r - c_u$, as illustrated in Figure 3.8. When $\phi \ge \frac{\pi_b}{\theta_r} - \frac{h}{\theta_r}$, $\tilde{\Pi}_u(\phi) = \phi\pi_{ON}$. On the other hand, when $\phi < \frac{\pi_b}{\theta_r} - \frac{h}{\theta_r}$, we have $F(\kappa) = \frac{\pi_b - \theta_r\phi}{\pi_b + h - \theta_r\phi}$ and $\tilde{\Pi}_u(\phi) = \frac{(\pi_b - \theta_r\phi)\phi}{\pi_b + h - \theta_r\phi}$. $\tilde{\Pi}_u(\phi)$ is concave in ϕ and $\tilde{\Pi}'_u(0) = \frac{\pi_b}{\pi_b + h} > 0$. U's optimal wholesale price is computed as follows. Let $\phi_s = \arg\max\left\{\tilde{\Pi}_u(\phi) \mid 0 \le \phi \le \frac{\pi_b}{\theta_r} - \frac{h}{\theta_r}\right\}$. If $\bar{\phi} < c_r - c_u$, U's profit is fully characterized by the function $\tilde{\Pi}_u(\phi)$, therefore, $\phi^* = \min(\phi_s, \bar{\phi})$ due to concavity. If, however, $\bar{\phi} = c_r - c_u$, then, we inspect U's profit for $\phi = \phi_s$ and $\phi = c_r - c_u$, that is, we compare $\tilde{\Pi}_u(\phi_s)$ and $(c_r - c_u)\pi_{ON}$, and choose the one that maximizes U's profit. Note that if $\phi_s = \frac{\pi_b}{\theta_r} - \frac{h}{\theta_f}$ (a corner solution), this means that $\tilde{\Pi}_u(w_u, \phi)$ is increasing over $\left[0, \frac{\pi_b}{\theta_r} - \frac{h}{\theta_f}\right]$. Since U's profit is linear and increasing over $\left[\frac{\pi_b}{\theta_r} - \frac{h}{\theta_f}, \bar{\phi}\right]$, U's profit is maximized at $\phi = \bar{\phi}$. Whenever U's optimal penalty is $\phi^* = \phi_s$, we must have $\phi_s < \frac{\pi_b}{\theta_r} - \frac{h}{\theta_f}$ as an interior solution. Also note that, whenever $\phi^* = \bar{\phi}$, $w_u^* = w_u^h(\phi^*)$, implying that R is indifferent between serving as the backup supplier versus the sole supplier. In addition, the manufacturer is indifferent between sole-sourcing from R versus using both suppliers.

To establish the equilibrium, we need to show that no supplier has incentive to unilaterally deviate from the prices we have derived. Due to Lemma 3.1, we have $w_r^* = c_r + \left(c_u + \frac{\pi_b}{\theta_r} - c_r\right)\bar{F}(\kappa^*)$. Clearly, R does not have incentive to unilaterally deviate, as R's strategy is derived by choosing the best strategy given U's strategy. Assume that U hopes to unilaterally deviate to a better strategy, (w_u, ϕ). We consider two cases. In the first case, $w_u - \phi \ge c_u$, in which case, the manufacturer still uses R as the backup supplier. In the second case, $w_u - \phi < c_u$, hence, U serves as the sole supplier. Consider first $w_u - \phi \ge c_u$. Let ψ be such that $w_u = c_u + \psi$. Manufacturer's cost of underage is $w_{rb} - w_u = \frac{\pi_b}{\theta_r} - \psi$. Therefore, manufacturer's optimal inventory is $\kappa_{cp}(\psi)$. U's profit is, then, $(w_u - c_u)F(\kappa) = \psi F(\kappa_{cp}(\psi))$. U's problem is stated as follows.

$$
\begin{aligned}
\max \quad & \psi F(\kappa_{cp}(\psi)) && (3.4)\\
s.t \quad & w_u - \phi \ge c_u \\
& w_u - c_u = \psi \\
& w_u F(\kappa_{cp}(\psi)) + \left(c_u + \frac{\pi_b}{\theta_r}\right)\bar{F}(\kappa_{cp}(\psi)) \le w_r^*
\end{aligned}
$$

The last inequality ensures that the manufacturer does not switch to R after U revises his prices. Note that ϕ does not influence U's profit, as ϕ is never paid to the manufacturer in case of a disruption. U can influence profit only by increasing or decreasing w_u. **Case 1:** Let $\phi^* = \bar{\phi}$ in the current solution, in which case, the manufacturer is indifferent between sole-sourcing from R and sourcing primarily from U along with R as backup. If U increases w_u, then, the manufacturer sole-sources from R. If U decreases w_u, U's profit is decreased. Therefore, U does not have incentive to deviate. **Case 2:** Let $\phi^* = \phi_s < \frac{\pi_b}{\theta_r} - \frac{h}{\theta_f}$ in the current solution. Since ϕ_s is also a maximizer of the objective function in (3.4), U does not have incentive to deviate.

Next, consider the case of $w_u - \phi < c_u$, in which case U serves as the sole supplier. The manufacturer's optimal inventory is $\kappa_{cp}(\phi)$. U's problem is stated as follows.

$$\begin{aligned} \max \quad & w_u - c_u - \phi\bar{F}(\kappa_{cp}(\phi)) \\ s.t. \quad & w_u - \phi \leq c_u \\ & w_u - \phi\bar{F}(\kappa_{cp}(\phi)) \leq w_r^* \end{aligned}$$

First constraint is binding in the optimal solution, $w_u - \phi = c_u$. Thus, U's problem can be restated as, max $\phi F(\kappa_{cp}(\phi))$ s.t. $c_u + \phi F(\kappa_{cp}(\phi)) \leq w_r^*$. Since $w_u - \phi = c_u$, U's profit as the sole supplier is the same as his profit when he is only the primary source. Thus, U does not have incentive to deviate.

Finally, we show that the equilibrium inventory is weakly decreasing in the cost differential, $c_r - c_u$. Clearly, $\bar{\phi}$ is weakly increasing in $c_r - c_u$, since $w_u^h(\phi)$ is an increasing function of ϕ. Therefore, the optimal penalty, $\phi^* = \arg\max_{0 \leq \phi \leq \bar{\phi}} \phi F(\kappa_{cp}(\phi))$ is weakly increasing in $c_r - c_u$ (increasing the upper bound of optimization cannot make the optimal solution smaller). Since $\kappa_{cp}(\phi)$ is a decreasing function of ϕ, $\kappa_{cp}(\phi^*)$ is weakly decreasing in ϕ^*, that is, a larger value of ϕ^* leads to lower inventory.

(iii) Let $c_r - c_u \geq \frac{\pi_b}{\theta_r}$. Towards a contradiction, assume that U serves as the primary supplier along with R as backup, therefore, $w_{ud} = c_u$ due to Property I, and $w_{rb} = \max\left(c_r, w_{ud} + \frac{\pi_b}{\theta_r}\right) = c_r$. In this case, manufacturer actually sole-sources from U, reaching a contradiction. Furthermore, U can improve profits by setting $w^*_{ud} = c_r - \frac{\pi_b}{\theta_r}$. Thus, U is the sole supplier, and $w^*_r = w^*_{rb} = c_r$. Since $w_u - \phi = c_r - \frac{\pi_b}{\theta_r}$, U's profit, $w_u - \bar{F}(\kappa_{cp}(\phi))\phi - c_u$ can be expressed as a function of ϕ, $\Pi_u(\phi) = c_r - \frac{\pi_b}{\theta_r} - c_u + F(\kappa_{cp}(\phi))\phi$. Clearly, the choice of optimal ϕ does not depend on c_r and c_u. To obtain the optimal ϕ, consider two cases.

- If $\frac{\pi_b}{\theta_r} \leq \frac{h}{\theta_f}$, then a zero-inventory policy is always optimal, thus, $F(\kappa_{cp}(\phi)) = F(0) = \pi_{ON}$. Therefore, the function $\Pi_u(\phi)$ is linear and increasing in ϕ. The optimal penalty is $\phi^* = \frac{\pi_b}{\theta_r}$, hence, $w^*_u = c_r$.

- If $\frac{\pi_b}{\theta_r} > \frac{h}{\theta_f}$, then a zero-inventory policy is only when $\phi \geq \frac{\pi_b}{\theta_r} - \frac{h}{\theta_f}$. Hence, $\phi = \frac{\pi_b}{\theta_r}$ dominates all ϕ such that $\frac{\pi_b}{\theta_r} - \frac{h}{\theta_f} \leq \phi < \frac{\pi_b}{\theta_r}$. Let $\phi < \frac{\pi_b}{\theta_r} - \frac{h}{\theta_f}$, in which case, $F(\kappa) = \frac{\pi_b - \theta_r\phi}{\pi_b + h - \theta_r\phi}$ and $\tilde{\Pi}_u(\phi) = c_r - \frac{\pi_b}{\theta_r} - c_u + \frac{(\pi_b - \theta_r\phi)\phi}{\pi_b + h - \theta_r - \phi}$. It can be readily shown that $\tilde{\Pi}_u(\phi)$ is concave for $\phi \in \left[0, \frac{\pi_b}{\theta_r} - \frac{h}{\theta_f}\right]$ by inspecting the second derivatives. Therefore, the unique maximizer of $\tilde{\Pi}_u(\phi)$ can be obtained over the interval $\left[0, \frac{\pi_b}{\theta_r} - \frac{h}{\theta_f}\right]$ using the first order condition. However, $\Pi_u(\phi)$ is not concave over the entire domain of ϕ, one needs to compare the two solutions, one in the interval $\left[0, \frac{\pi_b}{\theta_r} - \frac{h}{\theta_f}\right]$, and the other being $\phi = \frac{\pi_b}{\theta_r}$, using the exact same procedure as in *(ii)*. Clearly, R does not have incentive to deviate. U does not have incentive to deviate, either, since U's wholesale prices maximize his profit given that R does not deviate from $(w^*_r, w^*_{rb}) = (c_r, c_r)$.

□

Proof of Lemma 3.2. • For $c_r - c_u \leq \Delta$, R is the sole supplier and $\Pi^*_r(c_r, c_u) = c_u + L(\kappa_0) - c_r$, $\Pi^*_u(c_r, c_u) = 0$, $\Pi^*_m(c_r, c_u) = p - c_u - L(\kappa_0)$ and $\Pi^*_{sc}(c_r, c_u) = p - c_r$. Therefore, desired properties hold.

• For $c_r - c_u \geq \frac{\pi_b}{\theta_r}$, U is the sole supplier. From Proposition 3.6, equilibrium penalty, ϕ^*, does not depend on c_r and c_u. The profits are, $\Pi^*_r(c_r, c_u) = 0$, $\Pi^*_u(c_r, c_u) =$

$c_r - \frac{\pi_b}{\theta_r} - c_u + \phi^* F(\kappa_{cp}(\phi^*))$, $\Pi^*_m(c_r, c_u) = p - c_r + \frac{\pi_b}{\theta_r} - \phi^* F(\kappa_{cp}(\phi^*)) - L(\kappa_{cp}(\phi^*))$, and $\Pi^*_{sc}(c_r, c_u) = p - c_u - L(\kappa_{cp}(\phi^*))$. Therefore, desired properties hold.

• Let $\Delta \leq c_r - c_u \leq \frac{\pi_b}{\theta_r}$. Recall that $\phi_s = \arg\max \left\{ \tilde{\Pi}_u(\phi) \mid 0 \leq \phi \leq \frac{\pi_b}{\theta_r} - \frac{h}{\theta_r} \right\}$, $\bar{\phi} = \sup\{\phi \mid w_u - \phi = c_u,\ w^l_u(\phi) \leq w_u \leq w^h_u(\phi)\}$ and $\phi^* = \min(\phi_s, \bar{\phi})$. By its definition, we have $w^h_u(\bar{\phi}) - \bar{\phi} = c_u$. Since $w^h_u(\phi) = c_r - \frac{H(\kappa_{cp}(\phi))}{F(\kappa_{cp}(\phi))}$, $\bar{\phi}$ is the unique solution to the equation $c_r - c_u = \bar{\phi} + \frac{H(\kappa_{cp}(\bar{\phi}))}{F(\kappa_{cp}(\bar{\phi}))}$. Due to Lemma 3.1, $0 \leq \frac{dw^h_u(\phi)}{d\phi} \leq 1$, hence, $0 \leq \frac{\partial \bar{\phi}}{\partial c_r} \leq 1$ and $-1 \leq \frac{\partial \bar{\phi}}{\partial c_u} \leq 0$.

For $c_r - c_u = \Delta = \frac{H(\kappa_0)}{F(\kappa_0)}$, we readily have $\bar{\phi} = 0$, therefore, $\phi^* = 0$. Let $\bar{c}$ be the value of $c_r - c_u$ such that $\bar{\phi} = \phi_s$. As a consequence, $\phi^* = \bar{\phi}$ if $\Delta \leq c_r - c_u \leq \bar{c}$.

▶ First, consider the case of $\Delta \leq c_r - c_u \leq \bar{c}$. In this subregion, U's profit is $\Pi^*_u(c_r, c_u) = \tilde{\Pi}_u(\bar{\phi})$, R's profit is $\Pi^*_r(c_r, c_u) = \left(c_u + \frac{\pi_b}{\theta_r} - c_r\right) \bar{F}(\kappa_{cp}(\bar{\phi}))$, and finally, manufacturer's profit is $\Pi^*_m(c_r, c_u) = p - (c_u + \bar{\phi})F(\kappa_{cp}(\bar{\phi})) - \left(c_u + \frac{\pi_b}{\theta_r}\right) \bar{F}(\kappa_{cp}(\bar{\phi})) - H(\kappa_{cp}(\bar{\phi})) = p - c_u - \bar{\phi}F(\kappa_{cp}(\bar{\phi})) - L(\kappa_{cp}(\bar{\phi}))$.

<u>*Unreliable Supplier:*</u> Recall that $\tilde{\Pi}'_u(0) = \frac{\pi_b}{\pi_b + h} \leq 1$ from Proposition 3.6 and $\tilde{\Pi}'_u(\phi)$ is concave. Therefore, $\tilde{\Pi}'_u(\phi) \leq 1$ for all $\phi \geq 0$. Hence, $\frac{\partial \Pi^*_u(c_r, c_u)}{\partial c_r} = \tilde{\Pi}'_u(\bar{\phi}) \frac{\partial \bar{\phi}}{\partial c_r} \geq 0$ and ≤ 1, while $\frac{\partial \Pi^*_u(c_r, c_u)}{\partial c_u} = \tilde{\Pi}'_u(\bar{\phi}) \frac{\partial \bar{\phi}}{\partial c_u} \leq 0$ and ≥ -1.

<u>*Reliable Supplier:*</u> If $\kappa_{cp}(\bar{\phi}) = 0$, then $\bar{F}(\kappa_{cp}(\bar{\phi})) = \pi_{OFF}$ and $\Pi^*_r(c_r, c_u) = \left(c_u + \frac{\pi_b}{\theta_r} - c_r\right) \pi_{OFF}$. If, however, $\kappa_{cp}(\bar{\phi})$, then, $\bar{F}(\kappa_{cp}(\bar{\phi})) = \frac{h}{\pi_b + h - \theta_r \bar{\phi}}$. We readily have $0 \leq \frac{d\bar{F}(\kappa_{cp}(\bar{\phi}))}{d\bar{\phi}} \leq 1$. Since $0 \leq \frac{\partial \bar{\phi}}{\partial c_r} \leq 1$ and $-1 \leq \frac{\partial \bar{\phi}}{\partial c_u} \leq 0$, we have $0 \leq \frac{\partial \bar{F}(\kappa_{cp}(\bar{\phi}))}{\partial c_r} \leq 1$ and $-1 \leq \frac{\partial \bar{F}(\kappa_{cp}(\bar{\phi}))}{\partial c_u} \leq 0$. Hence desired monotonicity results for R hold.

<u>*Manufacturer and Supply Chain:*</u> Since $\phi^* = \bar{\phi}$, R is indifferent between serving as the sole supplier versus backup supplier in the equilibrium. Therefore, $\Pi^*_r(c_r, c_u) = w^*_u + L(\kappa_{cp}(\bar{\phi})) - \bar{\phi}\bar{F}(\kappa_{cp}(\bar{\phi})) - c_r$. Since $w^*_u = c_u + \bar{\phi}$, this reduces to $\Pi^*_r(c_r, c_u) = c_u - c_r + L(\kappa_{cp}(\bar{\phi})) + \bar{\phi}F(\kappa_{cp}(\bar{\phi}))$. As a consequence, manufacturer's profit is $\Pi^*_m(c_r, c_u) = p - c_r - \Pi^*_r(c_r, c_u)$, therefore, $\Pi^*_{sc}(c_r, c_u) = \Pi^*_r(c_r, c_u) + \Pi^*_u(c_r, c_u) + \Pi^*_m(c_r, c_u) = p - c_r + \Pi^*_u(c_r, c_u)$. We have already shown that $0 \leq \frac{\partial \Pi^*_r(c_r, c_u)}{\partial c_u} \leq 1$ and $-1 \leq \frac{\partial \Pi^*_r(c_r, c_u)}{\partial c_r} \leq 0$.

Thus, monotonicity properties for the manufacturer and the supply chain follow.

▶ Finally, consider the case $\bar{c} \leq c_r - c_u \leq \frac{\pi_b}{\theta_r}$. In this subregion, ϕ^* is either ϕ_s or $c_r - c_u$, whichever provides a higher profit for U. Thus, we are comparing $\tilde{\Pi}_u(\phi_s)$, which does not depend on c_r and c_u and $(c_r - c_u)\pi_{ON}$. If $\tilde{\Pi}_u(\phi_s) \geq \frac{\pi_b}{\theta_r}\pi_{ON}$, $\phi^* = \phi_s$ for all $c_r - c_u \geq \bar{c}$. Otherwise, there exists a threshold, $\tilde{c}$, where $\phi^* = \phi_s$ for $c_r - c_u \leq \tilde{c}$ and $\phi^* = c_r - c_u$ for $c_r - c_u \geq \tilde{c}$. Let $\bar{c} \leq c_r - c_u \leq \tilde{c}$, hence $\phi = \phi_s$. Then, U's profit is a constant in this range of costs. R's profit is $\Pi_r^*(c_r, c_u) = \left(c_u + \frac{\pi_b}{\theta_r} - c_r\right)\bar{F}(\kappa_{cp}(\phi_s))$, where $\bar{F}(\kappa_{cp}(\phi_s))$ is a constant less than 1. Finally, $\Pi_m^*(c_r, c_u) = p - c_u - \phi_s F(\kappa_{cp}(\phi_s)) - L(\kappa_{cp}(\phi_s))$. Monotonicity properties clearly hold for all players and the supply chain. Let $\tilde{c} \leq c_r - c_u \leq \frac{\pi_b}{\theta_r}$, in which case, $\phi = c_r - c_u$ and $\kappa_{cp}(\phi) = 0$. Then, U's profit is $\Pi_u^*(c_r, c_u) = (c_r - c_u)\pi_{ON}$, R's profit is $\Pi_r^*(c_r, c_u) = \left(c_u + \frac{\pi_b}{\theta_r} - c_r\right)\pi_{OFF}$, and finally, manufacturer's profit is $\Pi_m^*(c_r, c_u) = p - c_r\pi_{ON} - c_u\pi_{OFF} - \frac{\pi_b}{\theta_r}\pi_{OFF}$. Clearly, monotonicity properties for the manufacturer and the supply chain follow. □

Proof of Proposition 3.7. Let $\Pi_i^\beta(c_r, c_u)$ denote the equilibrium profit for $i \in \{u, r, m, sc\}$ when backup capacity is β. Consider first $\frac{\pi_b}{\theta_r} > \frac{h}{\theta_f}$. • Let $c_r - c_u \leq \Delta$. R serves as the sole supplier and charges $w_r^* = c_u + L(\kappa_0)$ regardless of the backup capacity. Thus, profits of all players are the same for $\beta = 0$ and $\beta = \infty$.

• Let $\Delta \leq c_r - c_u \leq L(\kappa_0)$. When $\beta = 0$, R serves as the sole supplier and charges $w_r^* = c_u + L(\kappa_0)$. When $\beta = \infty$, U serves as the primary supplier and R serves as the backup supplier. Denote by ϕ^* U's equilibrium penalty, thus, $w_u^* = c_u + \phi^*$ and $w_{rb}^* = c_u + \frac{\pi_b}{\theta_r}$. Denote by κ^* the equilibrium inventory with $\beta = \infty$.

<u>*Unreliable Supplier:*</u> When $\beta = 0$, U's profit is zero, hence, U's profit is improved with backup capacity.

<u>*Reliable Supplier:*</u> Since $\Pi_r^0(c_r, c_u) = c_u + L(\kappa_0) - c_r$, we have $\frac{\partial \Pi_r^0(c_r,c_u)}{\partial c_r} = -1$ and $\frac{\partial \Pi_r^0(c_r,c_u)}{\partial c_u} = 1$. From Lemma 3.2, we have $0 \leq \frac{\partial \Pi_r^\infty(c_r,c_u)}{\partial c_u} \leq 1$ and $-1 \leq \frac{\partial \Pi_r^\infty(c_r,c_u)}{\partial c_r} \leq 0$. Since $\Pi_r^\infty(c_r, c_u)$ is determined by $c_r - c_u$, we compare R's profit at the extreme values for the cost difference. For $c_r - c_u = \Delta$, $\Pi_r^\infty(c_r, c_u) = \Pi_r^0(c_r, c_u)$. For $c_r - c_u = L(\kappa_0)$,

we have $\Pi_r^\infty(c_r, c_u) \geq \Pi_r^0(c_r, c_u) = 0$. Since $\frac{\partial \Pi_r^\infty(c_r,c_u)}{\partial c_u} \leq 1 = \frac{\partial \Pi_r^0(c_r,c_u)}{\partial c_u}$ and $\frac{\partial \Pi_r^\infty(c_r,c_u)}{\partial c_r} \geq -1 = \frac{\partial \Pi_r^0(c_r,c_u)}{\partial c_r}$, we have $\Pi_r^\infty(c_r, c_u) \geq \Pi_r^0(c_r, c_u)$ for all $\Delta \leq c_r - c_u \leq L(\kappa_0)$.

Manufacturer: We compare the manufacturer's profit with no backup capacity and with infinite backup capacity: $\Pi_m^0(c_r, c_u) - \Pi_m^\infty(c_r, c_u) = \phi^* F(\kappa^*) + L(\kappa^*) - L(\kappa_0) \geq 0$. The inequality holds because $L(\kappa)$ is minimized when $\kappa = \kappa_0$. Thus, the manufacturer is worse-off with infinite backup capacity.

Supply Chain: We similarly compare the profits: $\Pi_{sc}^\infty(c_r, c_u) - \Pi_{sc}^0(c_r, c_u) = (c_r - c_u)F(\kappa^*) - H(\kappa^*)$. If $\phi^* = \bar{\phi}$, then, we have that $c_r - c_u = \bar{\phi} + \frac{H(\kappa^*)}{F(\kappa^*)}$, as shown in Lemma 3.2, therefore, $\Pi_{sc}^\infty(c_r, c_u) - \Pi_{sc}^0(c_r, c_u) = \bar{\phi} F(\kappa^*) \geq 0$, thus, the supply chain is better-off with backup capacity. If $\phi^* = \phi_s$, that is, $\bar{c} \leq c_r - c_u \leq \tilde{c}$, then, the supply chain is better-off with backup capacity, since $\Pi_{sc}^\infty(c_r, c_u) - \Pi_{sc}^0(c_r, c_u) \geq \bar{c} F(\kappa^*) - H(\kappa^*) = \phi_s F(\kappa^*) \geq 0$. Finally, if $\kappa^* = 0$, that is, $c_r - c_u \geq \tilde{c}$, then, $\Pi_{sc}^\infty(c_r, c_u) - \Pi_{sc}^0(c_r, c_u) = (c_r - c_u)\pi_{ON} \geq 0$.

- Let $c_r - c_u \geq \frac{\pi_b}{\theta_r}$. When $\beta = 0$, U serves as the sole supplier and charges $w_u^* = c_r - L(\kappa_0)$. When $\beta = \infty$, U still serves as the sole supplier.

Unreliable Supplier: Recall that U's equilibrium prices, (w_u^*, ϕ^*), are derived by solving the following.

$$\max \quad w_u - \phi \bar{F}(\kappa_{cp}(\phi)) - c_u \tag{3.5}$$

$$s.t. \quad w_u - \phi \bar{F}(\kappa_{cp}(\phi)) \leq c_r - L(\kappa_0) \tag{3.6}$$

$$w_u - \phi \leq c_r - \frac{\pi_b}{\theta_r} \tag{3.7}$$

From Proposition 3.6, $\Pi_u^\infty(c_r, c_u) = c_r - \frac{\pi_b}{\theta_r} - c_u + \phi^* F(\kappa^*)$, where $\kappa^* = \kappa_{cp}(\phi^*)$. Next, we obtain an upper bound on $\Pi_u^\infty(c_r, c_u)$ by dropping (3.7). The objective function value of the relaxed problem is $c_r - L(\kappa_0) - c_u$. Thus, $\Pi_u^\infty(c_r, c_u) = c_r - \frac{\pi_b}{\theta_r} - c_u + \phi^* F(\kappa^*) \leq c_r - L(\kappa_0) - c_u = \Pi_u^0(c_r, c_u)$, implying that U is worse-off with backup capacity. This also implies $\phi^* F(\kappa^*) \leq L(\kappa_0) + \frac{\pi_b}{\theta_r}$, which we use for manufacturer below.

<u>*Reliable Supplier:*</u> Backup capacity does not matter, as R makes zero profit in both cases.

<u>*Manufacturer:*</u> We compare manufacturer's profits: $\Pi_m^\infty(c_r, c_u) - \Pi_m^0(c_r, c_u) = \frac{\pi_b}{\theta_r} - \phi^* F(\kappa^*) - L(\kappa^*) \geq L(\kappa^*) - L(\kappa_0) \geq 0$. Thus, the manufacturer is better-off due to infinite backup capacity.

<u>*Supply Chain:*</u> We compare supply chain profits: $\Pi_{sc}^0(c_r, c_u) - \Pi_{sc}^\infty(c_r, c_u) = L(\kappa^*) - L(\kappa_0) \geq 0$.

• Let $L(\kappa_0) \leq c_r - c_u \leq \frac{\pi_b}{\theta_r}$. When $\beta = 0$, U serves as the sole supplier and charges $w_u^* = c_r - L(\kappa_0)$. When $\beta = \infty$, U serves as the primary supplier along with R as backup, thus, $w_u^* = c_u + \phi^*$ and $w_{rb}^* = c_u + \frac{\pi_b}{\theta_r}$.

<u>*Reliable Supplier:*</u> R makes no profit when $\beta = 0$. Thus, R's profit improves with backup capacity.

<u>*Unreliable Supplier:*</u> Since $\Pi_u^0(c_r, c_u) = c_r - L(\kappa_0) - c_u$, we have $\frac{\partial \Pi_u^0(c_r, c_u)}{\partial c_u} = -1$ and $\frac{\partial \Pi_u^0(c_r, c_u)}{\partial c_r} = 1$. From Lemma 3.2, we have $0 \leq \frac{\partial \Pi_u^\infty(c_r, c_u)}{\partial c_r} \leq 1$ and $-1 \leq \frac{\partial \Pi_u^\infty(c_r, c_u)}{\partial c_u} \leq 0$. Since $\Pi_u^\infty(c_r, c_u)$ is determined by $c_r - c_u$ only, we compare U's profit at the extreme values for the cost difference. For $c_r - c_u = L(\kappa_0)$, we have $\Pi_u^\infty(c_r, c_u) \geq \Pi_u^0(c_r, c_u) = 0$. For $c_r - c_u = \frac{\pi_b}{\theta_r}$, U's profit remains unchanged even if he serves as the primary supplier along with R as backup. As we shown above, U is worse-off, thus $\Pi_u^\infty(c_r, c_u) \leq \Pi_u^0(c_r, c_u)$ for $c_r - c_u = \frac{\pi_b}{\theta_r}$. Since $\frac{\partial \Pi_u^\infty(c_r, c_u)}{\partial c_r} \leq 1 = \frac{\partial \Pi_u^0(c_r, c_u)}{\partial c_r}$ and $\frac{\partial \Pi_u^\infty(c_r, c_u)}{\partial c_u} \geq -1 = \frac{\partial \Pi_u^0(c_r, c_u)}{\partial c_u}$, and $\Pi_u^\infty(c_r, c_u)$ and $\Pi_u^0(c_r, c_u)$ depend on c_r and c_u only through $c_r - c_u$, there exists a threshold $\bar{c}_u$ such that we have $\Pi_u^\infty(c_r, c_u) \geq \Pi_u^0(c_r, c_u)$ for all $L(\kappa_0) \leq c_r - c_u \leq \bar{c}_u$ and $\Pi_u^\infty(c_r, c_u) \leq \Pi_u^0(c_r, c_u)$ for all $\bar{c}_u \leq c_r - c_u \leq \frac{\pi_b}{\theta_r}$. Since $\Pi_u^0(c_r, c_u) - \Pi_u^\infty(c_r, c_u) = (c_r - c_u) - \phi^* F(\kappa^*) - L(\kappa_0)$, $\bar{c}_u = \phi^* F(\kappa^*) + L(\kappa_0)$, where (ϕ^*, κ^*) is the equilibrium when $c_r - c_u = \bar{c}_u$.

<u>*Manufacturer:*</u> Same logic follows for the manufacturer. Manufacturer is worse-off for both $c_r - c_u = L(\kappa_0)$ and better-off for $c_r - c_u = \frac{\pi_b}{\theta_r}$. Using a similar logic based

on slopes, we conclude that the there exists a threshold where the manufacturer is indifferent with or without backup capacity. Since $\Pi_m^0(c_r, c_u) - \Pi_m^\infty(c_r, c_u) = \phi^* F(\kappa^*) + L(\kappa^*) - (c_r - c_u)$, $\bar{c}_m = \phi^* F(\kappa^*) + L(\kappa^*)$, where (ϕ^*, κ^*) is the equilibrium when $c_r - c_u = \bar{c}_m$. Note that $\phi^* F(\kappa^*)$ is increasing in $c_r - c_u$ and $L(\kappa^*) \leq L(\kappa_0)$. Therefore $\bar{c}_u \leq \bar{c}_m$. When $\frac{\pi_b}{\theta_r} \leq \frac{h}{\theta_f}$, that is $\phi^* = \frac{\pi_b}{\theta_r}$ and $\kappa^* = 0$, $\bar{c}_m = \bar{c}_u = \frac{\pi_b}{\theta_r}$ readily follows.

Supply Chain: The supply chain is better-off for both $c_r - c_u = L(\kappa_0)$ and worse-off for $c_r - c_u = \frac{\pi_b}{\theta_r}$, as derived earlier. Since $\Pi_{sc}^0(c_r, c_u) - \Pi_{sc}^\infty(c_r, c_u) = H(\kappa^*) - L(\kappa_0) + (c_r - c_u)\bar{F}(\kappa^*)$, the threshold, $\bar{c}_{sc}$ satisfies $\bar{c}_{sc} = \frac{L(\kappa_0) - H(\kappa^*)}{\bar{F}(\kappa^*)}$, where, κ^* is the equilibrium inventory for $c_r - c_u = \bar{c}_{sc}$. When $\frac{\pi_b}{\theta_r} \leq \frac{h}{\theta_f}$, then, $\bar{c}_{sc} = \frac{\pi_b}{\theta_r}$.

Finally, the case of $\frac{\pi_b}{\theta_r} \leq \frac{h}{\theta_f}$ is a special case, where $\Delta = 0$ and $\bar{c}_m = \bar{c}_u = \bar{c}_{sc} = \frac{\pi_b}{\theta_r}$, and it can readily be shown that profits are not influenced with backup capacity when $c_r - c_u = \frac{\pi_b}{\theta_r}$. □

Proof of Proposition 3.8. We derive the optimal policy for the simplified problem with demand only in the last period, $d_T = 1$. Denote by s supplier states, where $s \in \{11, 10, 11, 01, 00\}$ corresponding to ON-ON, ON-OFF, OFF-ON and OFF-OFF respectively, and henceforth used in this order. Denote by V_t^s the expected procurement, holding and penalty costs in state s from period t onwards, when the initial inventory is 0. V_t is the column vector and $c = (c_1, c_1, c_2, \infty)'$ is the cost of procurement corresponding to each state. P denotes the probability transition matrix for the supplier states and $\mathbf{1}$ is a column vector of 1's. Denote by $\gamma_i = 1 - \theta_{ri}$. If $c_2 > c_1 + \frac{\pi_b}{\theta_{r1}}$, it is not optimal to source from supplier 2. Let $c_1 \leq c_2 \leq c_1 + \frac{\pi_b}{\theta_{r1}}$.

The main dynamic programming recursion is $V_t = \min\{c + h(T-t)\mathbf{1}, PV_{t+1}\}$. In the last period, $V_T = (c_1, c_1, c_2, c_1 + \sigma)'$, where σ is the cost of underage, including the expected backlogging cost and the extra cost of procurement if both suppliers are OFF in the last period. To derive σ, let τ_1 and τ_2 be geometric random variables, which stand for time to recover for each supplier. τ_i has a probability of success θ_{ri}. Also, $\tau = \min(\tau_1, \tau_2)$ is a geometric random variable with a probability of success

$1 - \gamma_1\gamma_2$.

$$\begin{aligned} V_T^{00} &= \pi_b \mathbb{E}[\tau] + c_1 \text{Prob}(\tau_1 \le \tau_2) + c_2 \text{Prob}(\tau_1 > \tau_2) \\ &= c_1 + \frac{\pi_b}{1-\gamma_1\gamma_2} + \frac{\theta_{r2}\gamma_1(c_2 - c_1)}{1-\gamma_1\gamma_2} \;:=\; c_1 + \sigma \end{aligned}$$

The optimal policy is established similar to the simplified problem with one supplier perfectly reliable with the above modification on the terminal values. To show the monotonicity of the optimal policy, we fix $c_1 = 0$ and inductively prove that $0 \le \frac{\partial V_t^s}{\partial c_2} \le 1$. Clearly, $0 \le \frac{\partial V_T^s}{\partial c_2} \le 1$. Assume that $0 \le \frac{\partial V_{t+1}^s}{\partial c_2} \le 1$. Denote by P_s the row corresponding to state s. To indicate the dependence on c_2, we write $V_t^s = V_t^s(c_2)$. First, consider $s = 11$, $V_t^{11}(c_2) = \min\{h(T-t), P_{11}V_{t+1}\}$. Clearly, $0 \le \frac{\partial P_{11}V_{t+1}}{\partial c_2} \le 1$. Since $V_t^{11}(c_2)$ is the minimum of a constant and an increasing function of c_2, it follows that $0 \le \frac{\partial V_t^{11}(c_2)}{\partial c_2} \le 1$. If it is optimal to order in period t for $c_2 = 0$, $h(T-t) \le P_{11}V_{t+1}$, then it is still optimal to order when c_2 is increased. If, however, it is not optimal to order for $c_2 = 0$, then increasing c_2 may result in optimality of ordering in period t. Therefore, the optimal inventory corresponding to state 11 is (weakly) increasing in c_2. Similar argument holds for the state 10.

Lastly, $V_t^{01}(c_2) = \min\{c_2 + h(T-t), P_{01}V_{t+1}\}$. Thus, $V_t^{01}(c_2)$ is obtained by taking the minimum of two functions of c_2, one having a slope of 1, and the other having a slope less than or equal to 1. Therefore, if it is not optimal to order for $c_2 = 0$, it is not optimal to order for any other value of c_2. If, however, it is optimal to order when $c_2 = 0$, then increasing c_2 may cause non-optimality of ordering in period t, since the latter function increases at a smaller rate. Hence, the optimal coverage corresponding to state 01 is (weakly) decreasing in c_2. Also, $0 \le \frac{\partial V_t^{01}(c_2)}{\partial c_2} \le 1$. □

Proof of Proposition 3.9. Since supplier 2's state is independent of the previous periods, the optimal coverage in states 11 and 10 are equal. To establish zero optimal coverage in state 01, consider period $T-1$. If supplier 1 is OFF and supplier 2 is ON in period $T-1$, it is not optimal to order if $c_2 + h \ge \theta_{r1}c_1 + (1-\theta_{r1})(1 -$

$\nu)c_2 + (1-\theta_{r1})\nu(c_1+\sigma)$. This relationship may or may not hold depending on how large c_2 is. To ensure that the optimal coverage is zero in state 01, we let $c_2 = c_1$. Therefore, the condition we obtain for $c_2 = c_1$, obtained as $\frac{\pi_b}{\pi_b+h} \leq 1-(1-\theta_{r1})\nu$, is sufficient for any $c_2 > c_1$ to ensure zero coverage.

To compute the optimal coverage, it suffices to compute the appropriate cost of underage. If supplier 1 is OFF in the last period, the cost of underage is $c_2 - c_1$ when supplier 2 is ON. Otherwise, when supplier 2 is also OFF, the cost of underage is σ, as derived in Proposition 3.8. Therefore, the overall cost of underage is conveniently obtained as $(c_2-c_1)(1-\nu)+\sigma\nu$. $\square$

Proof of Proposition 3.10. *(i)* The case of no backup capacity follows the same reasoning as Proposition 3.3.

(ii) Let ρ_i denote the probability of sourcing from supplier $-i$ if supplier i serves as the primary supplier. We derive ρ_1 as follows. When both suppliers are disrupted, probability that supplier 1 recovers first is $\text{Prob}(\tau_1 \leq \tau_2) = \frac{\theta_{r1}}{1-\gamma_1\gamma_2}$, which was derived in the proof of Proposition 3.8. In the long-run, supplier 2 is used as backup with the following probability: $\rho_1 = \pi^1_{OFF} - \pi^1_{OFF}\pi^2_{OFF}\text{Prob}(\tau_1 \leq \tau_2)$. Therefore, $\rho_i = \pi^i_{OFF} - \pi^1_{OFF}\pi^2_{OFF}\left(\frac{\theta_{ri}}{1-\gamma_1\gamma_2}\right)$ for each $i = 1,2$.

- If $c_2 \geq c_1 + \frac{\pi_b}{\theta_{r1}}$, supplier 2 is not competitive, and supplier 1 serves as the sole supplier. Supplier 2 prices down to cost, c_2. Supplier 1 can profitably serve as the sole supplier by setting $w^1_d = c_2 - \frac{\pi_b}{\theta_{r1}}$. Since, in the equilibrium, the manufacturer is indifferent between the two suppliers, we must have $w_1(1-\rho_1) + \left(w^1_d + \frac{\pi_b}{\theta_{r1}}\right)\rho_1 = c_2 + \frac{\pi_b}{\theta_{r2}}\pi^2_{OFF}$, leading to $w_1 = c_2 + \left(\frac{\pi_b}{\theta_{r2}(1-\rho_1)}\right)\pi^2_{OFF}$.
- Let $c_1 \leq c_2 \leq c_1 + \frac{\pi_b}{\theta_{r1}}$. In equilibrium, the backup supplier must be indifferent between serving as the sole supplier and serving as backup, and the sole supplier must weakly prefer serving as the sole supplier. Following the logic of Proposition 3.6, we have $w^i_d = c_i$, leading to $w^i_b = c_{-i} + \frac{\pi_b}{\theta_{r,-i}}$. Thus, no supplier can serve as the sole supplier. If supplier i serves as backup, his payoff is $\Pi^b_i = \left(c_{-i} + \frac{\pi_b}{\theta_{r,-i}} - c_i\right)\rho_{-i}$.

Assume that supplier i charges w_i for regular deliveries. Manufacturer's average cost with supplier i as the primary supplier is $w_i(1-\rho_i)+\left(c_i+\frac{\pi_b}{\theta_{ri}}\right)\rho_i$. In equilibrium, manufacturer must be indifferent between the two suppliers, hence

$$w_1(1-\rho_1)+\left(c_1+\frac{\pi_b}{\theta_{r1}}\right)\rho_1 \;=\; w_2(1-\rho_2)+\left(c_2+\frac{\pi_b}{\theta_{r2}}\right)\rho_2$$

Supplier i's payoff if he serves as the primary supplier is $\Pi_i^p=(w_i-c_i)(1-\rho_i)\geq\Pi_i^b$. This leads to $w_i(1-\rho_i)\geq c_i(1-\rho_1-\rho_2)+\left(c_{-i}+\frac{\pi_b}{\theta_{r,-i}}\right)\rho_{-i}$ for each i. Define $\tilde{w}_i=w_i(1-\rho_i)$. Thus, the following must hold in equilibrium.

$$\begin{aligned}
\tilde{w}_1-\tilde{w}_2 &= \left(c_2+\frac{\pi_b}{\theta_{r2}}\right)\rho_2-\left(c_1+\frac{\pi_b}{\theta_{r1}}\right)\rho_1 := \alpha\\
\tilde{w}_1 &\geq c_1(1-\rho_1-\rho_2)+\left(c_2+\frac{\pi_b}{\theta_{r2}}\right)\rho_2 := \alpha_1\\
\tilde{w}_2 &\geq c_2(1-\rho_1-\rho_2)+\left(c_1+\frac{\pi_b}{\theta_{r1}}\right)\rho_1 := \alpha_2
\end{aligned}$$

We have $1-\rho_1-\rho_2=\pi_{ON}^1\pi_{ON}^2+\pi_{OFF}^1\pi_{OFF}^2\left(\frac{\theta_{r1}\theta_{r2}}{1-\gamma_1\gamma_2}\right)\geq 0$, and $\alpha_1-\alpha_2=\alpha-(c_2-c_1)(1-\rho_1-\rho_2)$. Since $c_1\leq c_2$, we have that $\alpha_1-\alpha_2\leq\alpha$. As suppliers compete by undercutting wholesale prices, we reach $\tilde{w}_1=\alpha_2+\alpha\geq\alpha_1$ and $\tilde{w}_2=\alpha_2$. At this point, supplier 2 is indifferent between serving as primary or backup, whereas supplier 1 strictly prefers serving as primary. Since supplier 1 can infinitesimally undercut, supplier 1 must serve as the primary supplier in equilibrium. Thus, the wholesale price for supplier 1 is $w_1^*=c_2+\frac{\pi_b}{\theta_{r2}}\left(\frac{\rho_2}{1-\rho_1}\right)$.

(iii) Assume that $0\leq c_2-c_1\leq\frac{\pi_b}{\theta_{r1}}\pi_{OFF}^1-\frac{\pi_b}{\theta_{r2}}\pi_{OFF}^2$. Then, supplier 2 is the sole supplier when suppliers do not offer backup capacity, and supplier 2 is the backup supplier when suppliers offer backup capacity. Since supplier 1 becomes the primary supplier with backup capacity in this region, supplier 1 clearly benefits. Denote by Π_i^β player i's payoff ($i\in\{1,2\}$) with backup capacity $\beta\in\{0,1\}$. Also, denote by C_m^β

the manufacturer's average cost.

$$
\begin{aligned}
\Pi_2^0 &= c_1 - c_2 + \frac{\pi_b}{\theta_{r1}}\pi_{OFF}^1 - \frac{\pi_b}{\theta_{r2}}\pi_{OFF}^2 \\
\Pi_2^1 &= \left(c_1 + \frac{\pi_b}{\theta_{r1}} - c_2\right)\rho_1 \\
\Pi_2^1 - \Pi_2^0 &= (c_2 - c_1)(1-\rho_1) - \frac{\pi_b}{\theta_{r1}}(\pi_{OFF}^1 - \rho_1) + \frac{\pi_b}{\theta_{r2}}\pi_{OFF}^2 \\
&\geq -\frac{\pi_b}{\theta_{r1}}(\pi_{OFF}^1 - \rho_1) + \frac{\pi_b}{\theta_{r2}}\pi_{OFF}^2 \;\geq\; 0
\end{aligned}
$$

The last inequality follows from the definition of ρ_1. Thus, supplier 2 benefits from backup capacity.

$$
\begin{aligned}
C_m^0 &= c_1 + \frac{\pi_b}{\theta_{r1}}\pi_{OFF}^1 \\
C_m^1 &= \left(c_2 + \frac{\pi_b}{\theta_{r2}}\left(\frac{\rho_2}{1-\rho_1}\right)\right)(1-\rho_1) + \left(c_1 + \frac{\pi_b}{\theta_{r1}}\right)\rho_1 + \pi_{OFF}^1\pi_{OFF}^2\left(\frac{\pi_b}{1-\gamma_1\gamma_2}\right) \\
C_m^1 - C_m^0 &= (c_2 - c_1)(1-\rho_1) + \frac{\pi_b}{\theta_{r2}} \;\geq\; 0
\end{aligned}
$$

Thus, the manufacturer is worse-off with backup capacity. Assume next that $\frac{\pi_b}{\theta_{r1}}\pi_{OFF}^1 - \frac{\pi_b}{\theta_{r2}}\pi_{OFF}^2 \leq c_2 - c_1 \leq \frac{\pi_b}{\theta_{r1}}$. Then, supplier 1 is the sole supplier when suppliers do not offer backup capacity, and supplier 1 is the primary supplier when suppliers offer backup capacity. Since supplier 2 becomes the backup supplier with backup capacity in this region, supplier 2 clearly benefits.

$$
\begin{aligned}
\Pi_1^0 &= c_2 - c_1 + \frac{\pi_b}{\theta_{r2}}\pi_{OFF}^2 - \frac{\pi_b}{\theta_{r1}}\pi_{OFF}^1 \\
\Pi_1^1 &= \left(c_2 + \frac{\pi_b}{\theta_{r2}}\left(\frac{\rho_2}{1-\rho_1}\right) - c_1\right)(1-\rho_1) \\
\Pi_1^1 - \Pi_1^0 &= -(c_2 - c_1)\rho_1 - \frac{\pi_b}{\theta_{r2}}(\pi_{OFF}^2 - \rho_2) + \frac{\pi_b}{\theta_{r1}}\pi_{OFF}^1 \\
&\geq -\frac{\pi_b}{\theta_{r2}}(\pi_{OFF}^2 - \rho_2) + \frac{\pi_b}{\theta_{r1}}(\pi_{OFF}^1 - \rho_1) \;=\; 0
\end{aligned}
$$

Thus, supplier 1 benefits from backup capacity. And similarly, manufacturer is worse-off with backup capacity, as shown below.

$$\begin{aligned}
C_m^0 &= c_2 + \frac{\pi_b}{\theta_{r2}}\pi_{OFF}^2 \\
C_m^1 &= \left(c_2 + \frac{\pi_b}{\theta_{r2}}\left(\frac{\rho_2}{1-\rho_1}\right)\right)(1-\rho_1) + \left(c_1 + \frac{\pi_b}{\theta_{r1}}\right)\rho_1 + \pi_{OFF}^1\pi_{OFF}^2\left(\frac{\pi_b}{1-\gamma_1\gamma_2}\right) \\
C_m^1 - C_m^0 &= -(c_2-c_1)\rho_1 - \frac{\pi_b}{\theta_{r2}}(\pi_{OFF}^2 - \rho_2) + \pi_{OFF}^1\pi_{OFF}^2\left(\frac{\pi_b}{1-\gamma_1\gamma_2}\right) + \frac{\pi_b}{\theta_{r1}}\rho_1 \\
&\geq -\frac{\pi_b}{\theta_{r2}}(\pi_{OFF}^2 - \rho_2) + \pi_{OFF}^1\pi_{OFF}^2\left(\frac{\pi_b}{1-\gamma_1\gamma_2}\right) \;=\; 0
\end{aligned}$$

Lastly, let $c_2 - c_1 \geq \frac{\pi_b}{\theta_{r1}}$. In both cases, supplier 1 is the sole supplier. Hence, supplier 2 is indifferent to backup capacity.

$$\begin{aligned}
\Pi_1^0 &= c_2 - c_1 + \frac{\pi_b}{\theta_{r2}}\pi_{OFF}^2 - \frac{\pi_b}{\theta_{r1}}\pi_{OFF}^1 \\
\Pi_1^1 &= \left(c_2 + \left(\frac{\pi_b}{\theta_{r2}(1-\rho_1)}\right)\pi_{OFF}^2 - c_1\right)(1-\rho_1) + \left(c_2 - \frac{\pi_b}{\theta_{r1}} - c_1\right)\rho_1 \\
\Pi_1^1 - \Pi_1^0 &= \frac{\pi_b}{\theta_{r1}}(\pi_{OFF}^1 - \rho_1) \geq 0
\end{aligned}$$

Thus, supplier 1 benefits from backup capacity. And similarly, manufacturer is worse-off with backup capacity, as shown below.

$$\begin{aligned}
C_m^0 &= c_2 + \frac{\pi_b}{\theta_{r2}}\pi_{OFF}^2 \\
C_m^1 &= \left(c_2 + \left(\frac{\pi_b}{\theta_{r2}(1-\rho_1)}\right)\pi_{OFF}^2\right)(1-\rho_1) + \left(c_2 - \frac{\pi_b}{\theta_{r1}}\right)\rho_1 + \frac{\pi_b}{\theta_{r1}}\pi_{OFF}^1 \\
C_m^1 - C_m^0 &= \frac{\pi_b}{\theta_{r1}}(\pi_{OFF}^1 - \rho_1) \;=\; 0
\end{aligned}$$

□

Proof of Proposition 3.11. Consider a renewal process, where renewals occur whenever U recovers from a disruption. A renewal cycle consists of an exponentially distributed up-time, and a deterministic down-time. For an inventory level $\kappa \leq D$, the expected cost during a cycle is $C(\kappa) = h\kappa U + \frac{h\kappa^2}{2} + \frac{\pi_b(D-\kappa)^2}{2}$. Note that inventory is depleted at a rate of 1 during a down-time and backlogs accumulate at a rate of 1 when all the inventory is depleted. The average cost per-unit time is given by $C(\kappa)/(U+D)$. By simple derivative arguments, $\kappa_d = \max\left(0, \frac{\pi_b D - hU}{\pi_b + h}\right)$ is the unique minimizer of $C(\kappa)$. We repeat a similar procedure when R offers backup capacity. Let $w_r = w_u + t\pi_b$ for some t. Clearly, the manufacturer will not source from R if $t > D$. Also, if $t < 0$, that is, $w_r < w_u$, the manufacturer will definitely source from R. Therefore, we assume $0 \leq t \leq D$. The manufacturer never sources from the backup source in the last t time units of the disruption. Hence, when $\kappa \geq D - t$, the backup supplier is never used. On the other hand, if $\kappa < D - t$, the manufacturer relies on inventory on-hand in the first κ time units of the disruptions, uses backup source in the next $D - \kappa - t$ time units of the disruption, and finally backlogs demand till the manufacturer recovers. Using this, we express $C(\kappa)$ (assuming $w_u = 0$ without loss of generality, thus $w_r = t\pi_b$) as follows.

$$C(\kappa) = \begin{cases} h\kappa U + \frac{h\kappa^2}{2} + \pi_b t(D-\kappa-t) + \frac{\pi_b t^2}{2} & \text{, if} \quad 0 \leq \kappa \leq D-t \\ \\ h\kappa U + \frac{h\kappa^2}{2} + \frac{\pi_b(D-\kappa)^2}{2} & \text{, if} \quad \kappa \geq D-t \end{cases}$$

$C(\kappa)$ is convex in κ and smooth at the break-point $\kappa = D-t$. If $C'(0) > 0$, or $hU > \pi_b t = w_r - w_u$, a-zero inventory policy is optimal. If $C'(D-t) = hU + h(D-t) - \pi_b t \leq 0$, then the optimal solution satisfies $\kappa \geq D - t$, in which case it is never optimal to source from the backup source. Therefore, if $w_r > w_u + \min\left(D, \frac{h(U+D)}{\pi_b+h}\right)\pi_b$, the backup source will not be used. As a result, the manufacturer carries κ_d units of inventory, as derived above. Let $t \leq \min\left(D, \frac{h(U+D)}{\pi_b+h}\right)$. Then, the optimal κ satisfies $C'(\kappa) = 0$ and $0 \leq \kappa \leq D - t$. Thus, $\kappa(w_r, w_u) = \frac{\pi_b t - hU}{h} = \frac{w_r - w_u - hU}{h}$. □

Proof of Proposition 3.12. Let d and e stand for deterministic and exponential disruptions respectively. Denote by $L_\theta(\kappa)$ the average holding and penalty costs that the manufacturer incurs per unit time when the disruption type is $\theta \in \{d, e\}$ and the manufacturer only sources from U. Define also κ_d and κ_e as the manufacturer's optimal inventory coverage when he sole-sources only from U.

When R does not offer backup capacity, the outcomes of the contingent-pricing game and the single-wholesale-price game are the same, and therefore, the structure of equilibrium outcomes remains the same: if $c_u + L_d(\kappa_d) \le c_r$, U wins with $w_u^* = c_r - L_d(\kappa_d)$, otherwise, R wins with $w_r^* = c_u + L_d(\kappa_d)$. It, then, suffices to show that $L_d(\kappa_d) \le L_e(\kappa_e)$. If $\pi_b D \le hU$, then, $\kappa_e = \kappa_d = 0$ and $L_e(\kappa_e) = \pi_b D \pi_{OFF}$, $L_d(\kappa_d) = \pi_b D \pi_{OFF}/2$. Let $\pi_b D > hU$. Recall that $L_e(\kappa) = h\kappa + (h + \pi_b)\left[1 - F(\kappa)\right]D - hD\pi_{OFF}$ and $L_d(\kappa) = \left(h\kappa U + \frac{h\kappa^2}{2} + \frac{\pi_b(D-\kappa)^2}{2}\right)/(D+U)$ due to Proposition 3.11. First, we prove that $\kappa_e \ge \kappa_d$.

To show $\kappa_e \ge \kappa_d$, it suffices to show $\frac{\pi_b}{\pi_b + h} = F(\kappa_e) \ge F(\kappa_d) = 1 - \pi_{OFF} e^{-\kappa_d/D}$. Without loss of generality, let $\pi_b + h = 1$ (or $\pi_b = 1 - h$), hence, $\kappa_d = D - h(U + D)$. Since $hU \le \pi_b D$ is assumed, we also have $h \le \pi_{OFF}$ for $\pi_b = 1 - h$. Thus, we need to show $\pi_{OFF} e^{1 - h/\pi_{OFF}} \ge h$. This, however, follows from the inequality that $e^{1-x} \ge x$ for all $0 \le x \le 1$, and it readily holds for $x = \pi_{OFF}/h$. With $\pi_b = 1 - h$ and appropriate simplifications, $L_d(\kappa_d)/h = D\pi_{ON} + \frac{\pi_{OFF} D}{2} - \frac{h(U+D)}{2}$ and $L_e(\kappa_e)/h = D\pi_{ON} + \kappa_e \;\ge\; D\pi_{ON} + \kappa_d$. To show $L_d(\kappa_d) \le L_e(\kappa_e)$, it suffices that $\kappa_d = D - h(U + D) \ge \frac{\pi_{OFF} D}{2} - \frac{h(U+D)}{2}$, or equivalently, $\pi_{OFF} - \frac{\pi_{OFF}^2}{2} \ge \frac{h}{2}$. Since $h \le \pi_{OFF}$, a sufficient condition is to show $\pi_{OFF} - \frac{\pi_{OFF}^2}{2} \ge \frac{\pi_{OFF}}{2}$, or $\pi_{OFF} \ge \pi_{OFF}^2$, which holds, since $0 \le \pi_{OFF} \le 1$. Therefore, $L_e(\kappa_e) \ge L_d(\kappa_d)$. $\square$

Chapter 4

Conclusions

This dissertation has examined the effects of internal and external disruptions on a firm's supply chain strategy and production/inventory policy. Chapter 2 explores the effects of internal disruptions that manifest themselves through capacity uncertainty, and investigate the effect of disruptions on the production and inventory policies of a firm producing multiple products in multiple stages. Chapter 3 explores the effects external disruptions on a firm's supply chain strategy and inventory policy, taking into account the active role suppliers play in defining the parameters of contracts. The focus is on the effects of strategic supplier behavior and the resulting sourcing practices.

In the context of internal disruptions (Chapter 2), we have considered optimal production and inventory control for a make-to-stock/calibrate-to-order system. The manufacturer has dedicated resources for each product in stage 1 and a common resource that all products share in stage 2. We fully characterized the optimal policy for the case of two products and proposed heuristic policies for the case of multiple products based on the optimal policy structure for two products. We numerically explored the effect of product asymmetries on the optimal policy and showed that depending on shared capacity level, three different modes of behavior are present. We also showed numerically that the performance of the heuristics is near-optimal when the number of products is sufficiently large and the shared capacity is large relative to the demand for individual products.

In Chapter 3, we explored suppliers' strategic behavior when external disruptions may disable their production. The main insight here was that with endogenously determined wholesale prices, the manufacturer does not necessarily benefit from flexible sourcing and, in fact, is typically worse off. Thus, an upfront commitment to sole sourcing and using simple wholesale price contracts may actually be beneficial, as opposed to opening the opportunity for one supplier to serve as a backup, through more flexible contracts. Interestingly, suppliers may benefit from flexible sourcing even though the manufacturer does not. The reliable supplier always benefits from maintaining backup capacity, whereas the unreliable supplier might benefit in some situations from the reliable supplier's backup capacity despite reduced business volume.

References

Agrawal, N., S. Nahmias. 1997. Rationalization of the supplier base in the presence of yield uncertainty. *Production and Oper. Management* **6**(3) 291–308.

Anupindi, R., R. Akella. 1993. Diversification under Supply Uncertainty. *Management Sci.* **39**(8) 944–963.

Arreola-Risa, A., G.A. DeCroix. 1998. Inventory management under random supply disruptions and partial backorders. *Naval Res. Logistics* **45**(7) 687–703.

Atali, A., O. Ozer. 2005. Multi-Item inventory systems with Markov-modulated demands and production quantity requirements. *Working Paper* .

Aviv, Y., A. Federgruen. 2001. Capacitated multi-item inventory systems with random and seasonally fluctuating demands: implications for postponement strategies. *Management Sci.* 512–531.

Babich, V. 2006. Vulnerable options in supply chains: Effects of supplier competition. *Naval Res. Logistics* **53**(7) 656–673.

Babich, V., A.N. Burnetas, P.H. Ritchken. 2007. Competition and diversification effects in supply chains with supplier default risk. *Manufacturing Service Oper. Management* **9**(2) 123–146.

Bashyam, S., M.C. Fu, B.K. Kaku. 1995. Application of perturbation analysis to multiproduct capacitated production-inventory control. *Proceedings of the American Control Conference*, vol. 3.

Berg, M., M.J.M. Posner, H. Zhao. 1994. Production-Inventory Systems with Unreliable Machines. *Oper. Res.* **42**(1) 111–118.

Bielecki, T., P.R. Kumar. 1988. Optimality of Zero-Inventory Policies for Unreliable Manufacturing Systems. *Oper. Res.* **36**(4) 532–541.

Bish, E.K., A. Muriel, S. Biller. 2005. Managing flexible capacity in a make-to-order environment. *Management Sci.* **51**(2) 167–180.

Cachon, G.P. 2003. Supply chain coordination with contracts. *Handbooks in operations research and management science* **11** 229–340.

Ciarallo, F.W., R. Akella, T.E. Morton. 1994. A periodic review, production planning model with uncertain capacity and uncertain demand-optimality of extended myopic policies. *Management Sci.* **40**(3) 320–332.

Clark, A.J., H. Scarf. 1960. Optimal policies for a multi-echelon inventory problem. *Management Sci.* **6**(4) 475–490.

Dayanik, S., J.S. Song, S.H. Xu. 2003. The Effectiveness of Several Performance Bounds for Capacitated Production, Partial-Order-Service, Assemble-to-Order Systems. *Manufacturing Service Oper. Management* **5**(3) 251.

De Vericourt, F., F. Karaesmen, Y. Dallery. 2000. Dynamic scheduling in a make-to-stock system: A partial characterization of optimal policies. *Oper. Res.* **48**(5) 811–819.

DeCroix, G.A., A. Arreola-Risa. 1998. Optimal production and inventory policy for multiple products under resource constraints. *Management Sci.* **44**(7) 950–961.

Duenyas, I., W.J. Hopp, Y. Bassok. 1997. Production quotas as bounds on interplant JIT contracts. *Management Sci.* **43**(10) 1372–1386.

Elmaghraby, W.J. 2000. Supply Contract Competition and Sourcing Policies. *Manufacturing Service Oper. Management* **2**(4) 350–371.

Evans, R.V. 1967. Inventory control of a multiproduct system with a limited production resource. *Naval Res. Logistics* **14**(2).

Federgruen, A., N. Yang. 2009. Optimal supply diversification under general supply risks. *Oper. Res.* **57**(6) 1451–1468.

Fu, K., V.N. Hsu, C.Y. Lee. 2006. Inventory and Production Decisions for an Assemble-to-Order System with Uncertain Demand and Limited Assembly Capacity. *Oper. Res.* **54**(6) 1137.

Gerchak, Y., M. Parlar. 1990. Yield randomness, cost trade-offs and diversification in the EOQ model. *Naval Res. Logistics* **37**(3) 341–354.

Gupta, D. 1996. The (Q, r) inventory system with an unreliable supplier. *Infor* **34**(2) 59–76.

Gurler, U., M. Parlar. 1997. An Inventory Problem with Two Randomly Available Suppliers. *Oper. Res.* **45**(6) 904–918.

Gurnani, H., R. Akella, J. Lehoczky. 2000. Supply management in assembly systems with random yield and random demand. *IIE Transactions* **32**(8) 701–714.

Gurnani, H., M. Shi. 2006. A bargaining model for a first-time interaction under asymmetric beliefs of supply reliability. *Management Sci.* **52**(6) 865.

Hendricks, K.B., V.R. Singhal. 2005. An empirical analysis of the effect of supply chain disruptions on long-run stock price performance and equity risk of the firm. *Production and Oper. Management* **14**(1) 35–52.

Heyman, DP, MJ Sobel. 1984. *Stochastic Models in Operations Research, Vol. II: Stochastic Optimization*. McGraw-Hill, New York.

Hopp, W.J., M.L. Spearman. 1991. Throughput of a constant work in process manufacturing line subject to failures. *International J. of Production Research* **29**(3) 635–655.

Hu, X., I. Duenyas, R. Kapuscinski. 2008. Optimal joint inventory and transshipment control under uncertain capacity. *Oper. Res.* **56**(4) 881–897.

Janakiraman, G., J.A. Muckstadt. 2009. A decomposition approach for a class of capacitated serial systems. *Oper. Res.* **57**(6) 1384–1393.

Kaplan, R.S. 1970. A Dynamic Inventory Model with Stochastic Lead Times. *Management Sci.* **16**(7) 491–507.

Latour, A. 2001. Trial by fire: A blaze in Albuquerque sets off major crisis for cell-phone giants. *Wall Street Journal* **1** 29.

Lewis, B.M. 2005. Inventory control with risk of major supply chain disruptions. *Georgia Institute of Technology* .

Martha, J., S. Subbakrishna. 2002. Targeting a just-in-case supply chain for the inevitable next disaster. *Supply Chain Management Review* **6**(5) 18–23.

Meyer, R.R., M.H. Rothkopf, S.A. Smith. 1979. Reliability and Inventory in a Production-Storage System. *Management Sci.* **25**(8) 799–807.

Minner, S. 2003. Multiple-supplier inventory models in supply chain management: A review. *International J. of Prod Economics* **81**(82) 265–279.

Muriel, A., A. Somasundaram, Y. Zhang. 2006. Impact of partial manufacturing flexibility on production variability. *Manufacturing Service Oper. Management* **8**(2) 192–205.

Nahmias, S., C.P. Schmidt. 1984. An efficient heuristic for the multi-item newsboy problem with a single constraint. *Naval Res. Logistics* **31**(3) 463–474.

Parker, R.P., R. Kapuscinski. 2004. Optimal policies for a capacitated two-echelon inventory system. *Oper. Res.* **52**(5) 739–755.

Parlar, M. 1997. Continuous-review inventory problem with random supply interruptions. *European J. of Operational Res.* **99**(2) 366–385.

Parlar, M., D. Perry. 1995. Optimal (Q, r, T) policies in deterministic and random yield models with uncertain future supply. *European J. of Operational Res.* **84** 431–443.

Parlar, M., D. Perry. 1996. Inventory models of future supply uncertainty with single and multiple suppliers. *Naval Res. Logistics* **43**(2) 191–210.

Plambeck, E.L., A.R. Ward. 2007. Note: A separation principle for a class of assemble-to-order systems with expediting. *Oper. Res.* **55**(3) 603–609.

Rockafellar, R.T. 1997. *Convex nalysis*. Princeton University Press.

Shaoxiang, C. 2004. The optimality of hedging point policies for stochastic two-product flexible manufacturing systems. *Oper. Res.* **52**(2) 312–322.

Sheffi, Y. 2005. Building a Resilient Supply Chain. *Harvard Business Review* **1**(8) 1–4.

Song, J.S., P.H. Zipkin. 1996. Inventory Control with Information about Supply Conditions. *Management Sci.* **42**(10) 1409–1419.

Song, J.S., P.H. Zipkin. 2003. Supply chain operations: Assemble-to-order systems. T. De Kok, S. Graves, eds., *Handbooks in Operations Research and Management Science*, vol. 30. North-Holland, Amsterdam, The Netherlands.

Sting, F.J., A. Huchzermeier. 2010. Ensuring Responsive Capacity: How to Contract with Backup Suppliers. *European J. of Operational Res.* **207** 725–735.

Tang, C.S. 2006. Robust strategies for mitigating supply chain disruptions. *International J. of Logistics* **9**(1) 33–45.

Tang, S.Y., P. Kouvelis. 2011. Supplier diversification strategies in the presence of yield uncertainty and buyer competition. *Manufacturing Service Oper. Management* **13**(4) 439–451.

Tomlin, B. 2006. On the Value of Mitigation and Contingency Strategies for Managing Supply Chain Disruption Risks. *Management Sci.* **52**(5) 639–657.

Tomlin, B. 2009a. Disruption-management strategies for short life-cycle products. *Naval Res. Logistics* **56**(4) 318–347.

Tomlin, B. 2009b. Impact of supply learning when suppliers are unreliable. *Manufacturing Service Oper. Management* **11**(2) 192–209.

Topkis, D.M. 1978. Minimizing a submodular function on a lattice. *Oper. Res.* 305–321.

Van Mieghem, J.A., N. Rudi. 2002. Newsvendor networks: Inventory management and capacity investment with discretionary activities. *Manufacturing Service Oper. Management* **4**(4) 313–335.

Veinott, A.F. 1966. The status of mathematical inventory theory. *Management Sci.* 745–777.

Wan, Z., D. R. Beil. 2009. Bargaining Power and Supply Base Diversification. *Working Paper* .

Yang, Z.B., G. Aydin, V. Babich, D.R. Beil. 2009. Supply Disruptions, Asymmetric Information, and a Backup Production Option. *Management Sci.* **55**(2) 192–2009.

Yano, C.A. 1991. Impact of Quality and Pricing on the Market Shares of Two Competing Suppliers in a Simple Procurement Model. *Technical Report 91-30, Department of Industrial and Operations Engineering, University of Michigan, Ann Arbor* .

Yano, C.A., H.L. Lee. 1995. Lot Sizing with Random Yields: A Review. *Oper. Res.* **43**(2) 311–334.