Title
Last Time Buy Problems with Sequential Capacity Reservations

Author
Bertelli, Erik P

Publication Date
2022

Peer reviewed|Thesis/dissertation

Last Time Buy Problems with Sequential Capacity Reservations

by

Erik P Bertelli

A dissertation submitted in partial satisfaction of the

requirements for the degree of

Doctor of Philosophy

in

Engineering – Industrial Engineering and Operations Research

in the

Graduate Division

of the

University of California, Berkeley

Committee in charge:

Professor Candace Arai Yano, Chair
Professor Philip Kaminsky
Associate Professor Jose Guajardo

Fall 2022

Last Time Buy Problems with Sequential Capacity Reservations

Abstract

Last Time Buy Problems with Sequential Capacity Reservations

by

Erik P Bertelli

Doctor of Philosophy in Engineering – Industrial Engineering and Operations Research

University of California, Berkeley

Professor Candace Arai Yano, Chair

Many modern consumer electronics firms design their own products but outsource the production to contract manufacturers. Some of these products are also multi-generational, with short product life cycles and updated versions released on a regular schedule. Therefore, firms must eventually make an end-of-production decision for each product generation. We consider a new version of the last time buy problem—traditionally, a procurement quantity problem for the last possible purchase of a product—facing such a consumer electronics firm. In our problem setting, the contract manufacturer requires the firm to make sequential capacity reservations to retain the option to procure new units, a contract feature that commonly arises when the contract manufacturer has a high opportunity cost of capacity. The presence of the sequential capacity reservation requirement also creates the need to decide the timing of the last time buy, prior to which orders can be placed in each period. We formulate the problem as a dynamic program and identify properties of the optimal policy that are different and more complex than under the usual simpler assumption that capacity reservations (or analogously, production setups) do not need to be for sequential periods.

The dynamic program to find the optimal strategy for any problem instance is computationally intensive. From a numerical study, we observed that most optimal strategies have up to two order-up-to levels, a low-to-moderate value—appropriate for satisfying this period's demand—to be used with low starting inventory levels, and a higher order-up-to level—appropriate for serving as the last time buy—to be used with starting inventory levels above a threshold. We developed a heuristic solution procedure based on these insights that performs near optimally (with an average optimality gap of 0.03% for our set of problem instances) and takes very little computing time.

We also consider a variant of our original problem in which the firm must commit to the timing of the last time buy at the beginning of the problem horizon. From a numerical study, we found that the value of the option to dynamically extend the capacity reservation, as in our original problem variant, is quite small due to the effects of the sequential capacity reservation requirement. We also consider an extension with an option to buy back units from the firm's customers after the last time buy has been made, which is especially applicable when all remaining demand is for warranty replacement units and the buy-back units can be refurbished to satisfy warranty claims. We consider both the case in which the customer response to the buy-back offer is deterministic and the case in which it is stochastic. From a numerical study, we find that the introduction of the buy-back offer has the potential to greatly reduce the firm's expected costs, and this is true whether the customer response to the buy-back offer is deterministic or stochastic. Furthermore, our results suggest that, in a sequential capacity reservation setting, the buy-back option is more valuable than the flexibility afforded by a period-to-period capacity reservation contract.

To Mom and Dad

Contents

List of Figures

List of Tables

Acknowledgments

I am immensely grateful for the help and support of my colleagues, friends, and family throughout my studies. This dissertation would not have been possible without all of you.

First, I would like to thank my advisor, Candi Yano. Throughout my time at Berkeley, you have always supported me as a researcher and student. I especially appreciate the freedom you gave me to pursue my own research topic and your patience as I tried many different projects. Most of all, I am grateful for your dedication to your students; you have been more generous with your time than any other professor I have known.

I would also like to thank some other members of the faculty at Berkeley. Professor Jose Guajardo, thank you for your thoughtful comments on my research; this dissertation was greatly improved by your reminder to focus on "why." Professor Phil Kaminsky, thank you for your suggestion of the stochastic buy-back model and for including me in your weekly research meetings; I learned so much from you and from your other students (and the free pizza was great too). Finally, Professor Dorit Hochbaum, thank you for giving me the opportunity to join your research group as an undergrad—even though I was a sophomore who had never even seen a network graph.

The day-to-day operations of the IEOR department would not be possible without the work of its dedicated staff. I would especially like to thank Rebecca Pauling, Heather Iwata, Keith McAleer, Diana Salazar, Anayancy Paz, Ginnie Sadil, Erica Diffenderfer, and Goldie Negelev. Thank you for helping me navigate the many forms and deadlines required to complete this degree. Thank you all for the work you do to support the graduate students by bringing us together for great panels, lunches, and end-of-semester parties.

Berkeley IEOR has been a great place to call home for the last six years, and no small part of that has been the amazing graduate students in the department. I have learned more from all of you than from any class I have ever taken. I am especially grateful to Kevin Li for his advice and friendship over the years. I would like to thank Jared Bauman, Brent Eldridge, Angel Yang, Wen Gao, and Guang Yang for helping me survive the first-year curriculum; I would not have passed the prelim exam without our study group. Cheng Lyu, thank you for agreeing to work with me when I was a mere undergrad researcher and for showing me the basics of network graphs. Quico Spaen, thank you for your commitment to organizing the IEOR social club to bring us all together. Mark Velednistksy, thank you for your larger-than-life personality and the many fun days spent playing board games. I would also like to thank everyone who made the IEOR graduate student office a fun and friendly place to spend the day. Haoyang Cao, thank you for being a friendly desk neighbor who never tired of my (frequent) requests for help when I was stuck. Yonatan Mintz, Matt Olfat, Dean Grosbard, Salar Fattahi, and Alfonso Lobos, thank you for the many animated discussions over the years when I needed a break from thinking about math. Junyu Cao, I am grateful for our online chess games and chats about movies during the pandemic. Renyuan Xu, Han Feng, Mo Zhou, Nan Yang, Jiaying Shi, Shiman Ding, Rebecca Sarto, Auyon Siddiq, Carlos Deck, Hansheng Jiang, Tugce Gurek, Heejung Kim, Arman Jabbari, Stewart Liu, Sheng Liu, Chen Chen, Julie Mulvaney-Kemp, Pedro Hespanhol, Andres Gomez, and Jiung Lee, thank you all for making Berkeley IEOR a wonderful place to work.

I must also thank my many friends who tolerated my complaints and never pointed out when my overly-optimistic estimated graduation dates would come and go. Kyle, Tyler, and Leah, thank you for providing much-needed stress relief in the form of our frequent gaming sessions. Hamilton and Sally, I am grateful for our many trips to the movies and backyard barbeques. Michael, Helen, Frank, Phoebe, and Elizabeth, thank you for our biweekly Gloomhaven nights; it was my honor to serve as your tank. Lila, Sherman, Shouvik, Josh, and many, many others, thank you all for your friendship, you have all enriched my life more than I can ever repay. Finally, I would like to thank the members of my Christ Church community group over the years: Tyler, Allie, Ethan, Kate, David, Kathryn, Meghan, Caleb, Michelle, Jia Ahn, Cheng-Yin, Gordon, Sarah, Meredith, Bre, Libby, Caroline, Jenna, Megan, Andrew, Cody, and Stevie. Thank you all for supporting me through friendship and prayer.

I also want to thank the amazing teachers who laid the foundation for my education. Mrs. Burton, thank you for instilling in me a lifelong love of learning; the independence you gave me in K-8 gave me the confidence to tackle any problem. Mr. Sanchez, thank you for developing my love of math; the times you had us teach a topic to the class forced me to really understand the material more than any problem set could have. Mrs. Sanchez, thank you for introducing me to the world of statistics; whenever I get stuck, I remember your advice to “keep moving forward!” Mr. Harding, thank you for your impartial introduction to government and politics; I have kept a love for national politics to this day. Ms. Gonzales, thank you for encouraging my critical thinking and tolerating my antics; I still look back fondly on the many animated discussions held in your classroom. Finally, Pastor Spence, thank you for treating me like an equal at a time when most other adults just saw me as a kid.

Above all, I would like to thank my family for supporting me. Kaari and Vincent, thank you for the thoughtful discussions and for bringing the wonderful Walter into our lives. Grandma and Grandpa, thank you for the financial support and for being the best landlords I could ever ask for; your generosity made it possible for me to live in one of the most expensive places in the world and complete this dissertation. Dad, thank you for encouraging me to follow my interests and for showing me what it means to be a man; every year I grow more impressed with everything you have accomplished. Mom, thank you for always listening; whenever I was frustrated or just wanted to talk, you always knew just what to say.

Chapter 1

Introduction

“Smart” devices, products that include some form of computer to perform various tasks, have become a ubiquitous part of the modern economy. These products range from large, complicated machines such as automobiles to small and previously simple products such as light bulbs. The research in this dissertation is particularly relevant to consumer electronic products that are intermediate in their complexity and cost. Examples include smartphones, wearable fitness trackers, and small Internet-of-Things connected devices such as smart switches or smart plugs. A key feature of these modern consumer electronic devices is that they are multi-generational, with short product life cycles. Consumers expect that updated versions, with new and improved features, will be released on a regular schedule. For these product lines, new generations are introduced, and older versions are eventually phased out. The timing of each introduction and phase-out is often dictated by market forces, such as the need to match a competitor’s new product launch or to have the next-generation item in stock for the holiday season. Further, eventual obsolescence usually occurs due to changes in the supporting hardware and/or software technology. We are interested in the inventory management problem facing a consumer electronics firm that offers such multi-generational products and has control over the timing of the end of production for models that are being phased out of retail sales.

Any consideration of this inventory management problem facing a consumer electronics firm must account for the method by which it produces the product under consideration. A common feature of the consumer electronics industry is the widespread use of contract manufacturers. In 1981, in the very early days of the consumer electronics industry, IBM introduced the IBM Personal Computer and chose to outsource the processor to Intel and the operating system to Microsoft rather than develop them in-house. As pointed out by Lüthje [33] and Sturgeon [47], IBM also chose to outsource the assembly of the motherboards to a firm called SCI, which at the time was a supplier in the aerospace industry. While Intel and Microsoft became household names, SCI is less well known, but went on to become a large player in the growing contract manufacturing industry. According to company financial statements, Sanmina (the successor to SCI following a series of mergers) had revenues of $6.7 billion in 2021 [43]. However, this pales in comparison to the revenues of the company that dominates the electronics contract manufacturing industry, Foxconn (also known as Hon Hai Precision Industry), which had 2021 revenues of $215 billion [17], more than either Microsoft ($168 billion [35]) or Intel ($79 billion [13]). It is clear that contract manufacturing is now an established part of the consumer electronics industry.

A contract manufacturer often devotes a set of resources, such as assembly lines, machinery, and personnel, exclusively to a customer's product. At the same time, the contract manufacturing industry is dominated by a small number of large corporations, and therefore a typical contract manufacturer has multiple customers simultaneously, even firms that are direct competitors in the same product segment. The contract manufacturer therefore has a large opportunity cost of setting aside resources for a given product. As any product approaches the end of its selling horizon, the number of units being produced begins to decline, and, consequently, the contract manufacturer would prefer to shift its limited resources to other more lucrative, higher-volume products. Therefore, the contract manufacturer will seek a contract mechanism that requires its customers to bear some or all of these opportunity costs. One simple mechanism is a fixed cost per period that does not depend on the quantity produced, which is a cost structure that also arises in many production/inventory problems, alongside variable costs that are a function of the quantity produced. In the case of the contract manufacturer, the opportunity cost would be a part of the fixed cost. We consider a setting in which a contract manufacturer imposes a per-period fixed cost on its customer for every period that the customer retains access to the production line. We characterize this fixed cost as a *capacity reservation fee*.

It is within this context that we introduce a new variant of the last time buy problem in this dissertation. A detailed literature review of the last time buy problem is presented in Chapter 2, but the essence of the traditional problem is that a firm must determine how many units to purchase to satisfy all remaining demand for a given product in view of the economic tradeoffs between the cost of shortages and the cost of excess inventory that is no longer useful. In our variant of the problem, the firm must not only determine the *size* of the last time buy, but also its *timing*. The firm is motivated to end the production of a product to avoid continuing to pay the per-period capacity reservation fee which is charged by the contract manufacturer. In contrast to the more standard setup cost often included in other inventory models (i.e., a fixed charge that must be paid only if production is positive), this capacity reservation fee is assessed if the firm wishes to reserve the right to place future orders, regardless of whether an order is placed in a given period. Further, if the firm does not pay the capacity reservation fee, the contract manufacturer will terminate production of the current active product and switch its use to another product, possibly one sold by another firm. One important consequence of these two conditions is that all periods in which the production line is available for use must be sequential.

The sequential requirement for the capacity reservation is a key feature that differentiates our models from others in the literature. Within this framework, we consider three major variants. In the first variant, which we refer to as the pre-commit variant, the duration of the capacity reservation must be decided by the firm at the beginning of an appropriate planning horizon. This variant applies in settings in which the contract manufacturer has a good deal of market power relative to the firm and can therefore demand that the firm commit to the timing of the last time buy well in advance. The pre-commit variant also serves as a baseline against which to compare the other variants of our problem.

The second variant, which we refer to as the dynamic variant, is more flexible: the firm can extend the capacity reservation on a period-by-period basis. This situation is more likely to arise if the firm has a large amount of market power relative to that of the contract manufacturer. For example, a large multinational may order such large volumes across a wide range of product segments that it is able to demand flexible contracting terms. We model this variant of the problem as a dynamic program, with periodic options to extend the capacity reservation. However, the contract manufacturer still requires the firm to pay sequential per-period capacity reservation fees to compensate it for the large opportunity cost mentioned earlier. The difference between the expected costs of the dynamic and pre-commit variants also provides an estimate of the value of the contract flexibility.

Finally, the third variant, which we refer to as the dynamic variant with a one-time buy-back option, is identical to the previous dynamic variant with one additional feature: the firm may make a one-time buy-back offer to existing customers after ending the capacity reservation with the intent of refurbishing the buy-back units. This variant is applicable when the late-stage demand is for units to satisfy warranty claims. In the case of consumer electronics, the combination of short product life cycles and long warranty periods—often required by law, such as in the European Union [12]—creates a situation in which warranty replacement demand extends far past the time of the last retail sale. However, warranties on consumer electronics normally allow the replacement of a defective or failed unit with an item of similar quality. This opens up the possibility of using refurbished units to satisfy warranty claims. The opportunity to make a one-time buy-back offer to customers allows the firm to acquire a potentially substantial number of refurbishable units after the contracted production of the product has ended. This option may allow the firm to end the capacity reservation earlier or reduce the size of the last time buy, or a combination of the two.

The remainder of this dissertation is organized as follows. In Chapter 2, we provide a review of pertinent literature. In Chapter 3, we present a formal problem statement detailing the problem setting as well as the assumptions underlying our models. In Chapter 4, we formulate the pre-commit variant of the problem. In Chapter 5, we formulate the dynamic variant of the problem and derive some properties of the optimal policy structure. In Chapter 6, we present a heuristic method for solving the dynamic variant of the problem. In Chapter 7, we present a numerical study that compares the performance of the heuristic with that of optimal strategies for both the pre-commit variant and the dynamic variant. In Chapter 8, we consider the dynamic variant with buy-back and explore how the addition of the one-time buy-back option changes the optimal solutions in a numerical study. Finally, in Chapter 9, we conclude with a summary of our contributions and a discussion of possible future work.

Chapter 2

Literature Review

Our work builds primarily upon the literature on last time buy problems. We first present a comprehensive literature review of last time buy problems categorized according to the nature of the extension of the basic problem. We then present a brief overview of results on lost sales inventory problems as well as inventory problems with capacity reservation requirements that are directly relevant to our research.

2.1 Previous Work on the Last Time Buy Problem

The problem of deciding how much of a product to order or produce so as to satisfy all remaining demand until the end of the planning horizon is commonly called the last time buy problem, but it is also known as the lifetime buy, the end-of-life buy, the final order problem, or the all-time requirement problem. This problem is most commonly motivated by settings in which a supplier is discontinuing an essential component and a manufacturer that utilizes the component must decide on a last-time-buy order quantity so that the manufacturer may satisfy all anticipated future demand, either with a high probability or in a manner that limits expected shortages. For instance, a manufacturer may be informed that a supplier is discontinuing a specific component, and therefore must decide how many of the discontinued components to order to enable the manufacturer to continue production until an alternative component can be identified and incorporated into the product. A similar situation arises when the component is a repair part for a machine and the firm operating the machine must decide how many spare parts to purchase in order to keep the machine operating for a desired length of time. A third example of a scenario in which this problem arises is a manufacturer that must decide how many units to produce and hold in inventory in order to satisfy all remaining demand that may arise from future warranty claims for a discontinued product.

Moore [36] provides the earliest study of the last time buy problem. He is primarily interested in forecasting the *all-time requirement*, which is the demand for a part from the present for all time periods into the future. He uses this forecast as the deterministic demand input for an n-stage dynamic inventory model that accounts for ordering and inventory holding costs to satisfy this requirement. He compares the cost of the optimal solution to his dynamic model to that incurred by using the economic order quantity (EOQ) and shows that, although the dynamic model results in a larger number of orders being placed, it eliminates the obsolescence cost due to leftover units at the end of the time horizon and lowers the inventory holding costs. Although the main emphasis in Moore's paper is on the forecasting method, the problem described in this paper serves as a foundation for this entire area of research.

Stochastic demand is the first generalization of this problem to be considered in the literature. In cases with stochastic demand, researchers either assume that the service level is specified as a constraint or the service level is optimized by minimizing the total expected cost, including shortage costs, inventory holding costs, and costs of production. Fortuin [16] was the first to consider the predefined service level version of the problem in the setting of a manufacturer seeking to manage spare parts inventory. He assumes that demand comes from a non-stationary and uncorrelated Gaussian process with an exponentially decreasing mean. This allows him to represent the total demand over the time horizon as a single Gaussian random variable and find the optimal last-time-buy quantity at the start of the time horizon for a given service level. Fortuin [15] then extends this model by allowing the desired service level to vary over time. He observes that, in order to maintain a constant service level over the entire time horizon, the actual service levels in earlier periods turn out to be considerably higher than the corresponding required levels. It is only in the final period that the actual service level approaches the desired service level. By allowing the desired service level to decrease over time, the manufacturer can realize considerable savings relative to maintaining a constant service level.

One of the first cost minimization approaches with stochastic demand was proposed by Teunter and Haneveld [50]. They consider this problem in the context of a single machine with multiple essential components, each requiring a separate last-time-buy decision to be made at the beginning of the time horizon. The owner of the machine desires for it to operate until a given time, but the machine is considered inoperable at the earliest time at which a spare part is no longer available to repair a broken component. If the machine becomes inoperable prior to this desired horizon, the owner of the machine incurs a fixed penalty cost as well as a variable cost that is proportional to the length of time that the machine is inoperable. The demand for each of the components is generated by a Poisson process that is not necessarily independent of the demand for the other components. Teunter and Haneveld present a model to minimize the expected cost, which consists of the purchasing and inventory holding costs of the essential components as well as the penalty costs due to the machine becoming inoperable prior to the desired end-of-service period. They show that the multi-component problem can be approximately decomposed into many single-component problems if the out-of-service penalties are very large relative to the ordering and inventory holding costs of the components. They then derive optimality conditions for the last-time-buy quantities.

Researchers have considered many extensions to the last time buy problem with stochastic demand. Bradley and Guerrero [5, 6] and Hur et al. [24] consider the case of a last time buy problem in which the decision must be made jointly for two or more different components, rather than just for a single component. Pourakbar and Dekker [39] analyze a last time buy problem in which customers are segmented into classes with different associated profits and costs and develop strategies to ration inventory among these classes. Leifker et al. [31] consider a situation in which, at the conclusion of the current service contract, there is a state-dependent probability that the customer may request an extension (with prearranged terms) to which the firm may or may not agree. They seek to optimize the size of the last time buy in the presence of this mutual option. However, unlike the three papers mentioned above, the vast majority of extensions incorporate alternative sources for satisfying demand other than inventory from the last time buy itself. These alternatives may include a different supplier that charges a higher marginal cost or, in the spare parts setting, allowing the firm to repair rather than replace defective items. However, a few extensions also involve considering the timing of the last time buy itself. In particular, the firm may extend the production run for some period of time, allowing the firm to save on inventory holding costs.

The extensions to the basic last time buy problem that involve alternative methods for satisfying demand can be organized into five categories: (1) allowing the firm to contract with additional production sources to supplement the initial inventory purchased in the last time buy, (2) allowing the firm to harvest additional spare parts from returned items, (3) allowing the firm to repair some defective items rather than replace the entire item from new finished goods inventory, (4) allowing the firm to make a product buy-back or trade-in offer to source additional items that can be used to satisfy future demand, and (5) allowing the firm to decide the timing of the last time buy itself. Some models involve only one type of extension, while others combine two different extensions in a more complicated decision structure.

Our research offers two major contributions. First, we are the first to introduce a production/inventory model in which the production line must remain available for sequential periods. That is, the firm cannot place an order in a period by paying a setup cost for that period, but instead can only retain access to the production line by paying per-period capacity reservation payments in sequential periods. This unique cost structure is what incentivizes the firm to eventually end the capacity reservation. We make use of the pre-commit and dynamic models to explore how the cost structure affects the firm's optimal strategy under different assumptions about the structure of the capacity reservation contract. In the pre-commit model, we require the firm to commit in advance to a specific number of periods of capacity reservation, while in the dynamic model we allow the firm to choose to extend the capacity reservation on a period-to-period basis.

Second, to the best of our knowledge, our work is the first to consider the combination of features (4) and (5), which we explore in our third model. Of the four non-timing extensions to the last time buy problem, the combination of a buy-back offer with a decision about the timing of the last time buy is the most relevant to the consumer electronics industry, where it is often impractical to resume the capacity reservation with the initial contract manufacturer or to switch to another contract manufacturer after the capacity reservation with the first contract manufacturer has ended. In either case, the contract manufacturer would incur expensive fixed costs and might require considerable time to set up a production line to produce the product in question. Therefore, either option would not be sensible given the short product life cycles. Consumer electronic devices are often too small and complex to harvest spare parts from or repair in a cost-effective manner; however, functioning units can be refurbished to an acceptable level of quality for use as warranty replacement units. These functioning units can be sourced from either product returns or buy-back/trade-in offers. In the consumer electronics market, customers are rarely contractually required to return their used items at pre-determined times (unlike in the market for capital equipment, for instance). However, voluntary buy-back or trade-in offers are common. Finally, the buy-back is especially complementary in the context of the sequential capacity reservation whose cost structure incentivizes the firm to make the last time buy earlier than it would in a model with a traditional setup cost, while the buy-back allows the firm to preserve the option to acquire units later in the horizon, if necessary.

Table 2.1 presents a list of papers on the last time buy problem in chronological order and classified by the same five categories of extensions described above. We first review models in which the timing of the last time buy has been predetermined, focusing on the alternative sources of inventory as categorized in the first four columns of the table. We then turn to models in which the timing of the last time buy can be jointly optimized by the firm in conjunction with the size of the last time buy; pertinent papers have check marks in the final column of the table.

Last Time Buy with Pre-Determined Timing

We now consider the literature on last time buy or similar problems in which the timing of the last time buy has been predetermined. We organize this subsection by type of extension listed in the first four columns of Table 2.1.

Additional Source of Supply

One of the most straightforward extensions to the last time buy problem is to incorporate the opportunity to source more units from an alternate supplier. The alternate supplier typically would be an aftermarket supplier whose product is often differentiated by a higher price (either per unit or due to additional fixed costs) than that of the original supplier. The alternate source of supply may also be a different component that is still in active production, but this may entail a costly redesign of the entire product. Finally, there may also be an option to change the service strategy entirely, and instead satisfy demand with a different product or through cash payments in the case of warranty claims. A common feature of all these alternatives is the ability to satisfy demand after the last time buy through a more costly option, which we have grouped together under the title “additional source of supply.”

The first consideration of the additional source of supply option was studied by Guerts and Moonen [22], who allow for a single order with zero setup cost at the start of the time horizon as well as follow-up orders in subsequent periods that incur a positive setup cost but

Table 2.1: Selected Papers On Last Time Buy Problems

Paper	Additional Source of Supply	Product Returns	Repair	Product Trade-Ins	Timing of LTB
Moore (1971) [36]					✓
Fortuin (1980) [16]					
Fortuin (1981) [15]					
Geurts and Moonen (1992) [22]	✓				
Teunter and Haneveld (1998) [50]					
Teunter and Fortuin (1999) [48]		✓			
Teunter and Haneveld (2002) [49]	✓				
Cattani and Souza (2003) [8]					✓
Kleber and Inderfurth (2007) [28]		✓			
Inderfurth and Mukherjee (2008) [26]	✓	✓			
van Kooten and Tan (2009) [51]			✓		
Kleber et al. (2012) [29]			✓	✓	
Leifker et al. (2012) [30]	✓				
Pourakbar et al. (2012) [40]	✓		✓		
Inderfurth and Kleber (2013) [25]	✓	✓			
van der Heijden and Iskandar (2013) [23]			✓		
Pourakbar et al. (2014) [41]		✓			
Shen and Willems (2014) [45]	✓				
Behfard et al. (2015, 2018) [1] [2]			✓		
Cole et al. (2015, 2016) [10, 11]				✓	
Frenk et al. (2019) [18, 19, 20]	✓				
Shi and Liu (2020) [46]	✓				
Ozyoruk et al. (2022) [37]	✓				✓
Our contribution				✓	✓

the same per-unit procurement cost. Thus, the additional source of supply in this model, i.e., the follow-up orders, only differ from the original order in that they incur positive setup costs. They assume a fixed time horizon with stationary Poisson demand. They present a dynamic programming formulation of the problem and find the (near-) optimal initial order and subsequent reorder policies for a range of parameter values. They also consider the sensitivity of these solutions to uncertainty in the parameter values to determine which parameters have a large impact on the optimal solution. The dynamic version of our problem (see Chapter 5) is somewhat similar to that addressed by Guerts and Moonen but our problem differs in three distinct ways: (1) our problem has a positive capacity reservation cost in the first period; (2) we require continuity of the capacity reservation; and (3) our model admits arbitrary demand distributions. The first and third differences do not add substantial complications, but, as we will see later, the continuity requirement for the capacity reservation leads to qualitatively different characteristics of optimal and good heuristic policies.

Shen and Willems [45] consider the problem facing a manufacturer who has been informed that a component of its product is being phased out. The manufacturer has the option of using an alternate supplier with larger fixed and variable costs of ordering. The authors also introduce the option to eventually transition to an entirely new substitute component via a product redesign after some positive lead time; obtaining the new component incurs both fixed and variable costs. Therefore, they instead characterize the last time buy as a "bridge buy" that enables the manufacturer to continue selling the current product until the system transitions to the newly redesigned product. They show that the expected profit function is non-concave but are able to determine the optimal "bridge buy" quantity numerically. Shi and Liu [46] also consider the option to redesign the product. They consider a firm producing a product that relies on a component part being discontinued at a deterministic date in the future. They formulate a two-stage dynamic program that has the form of an optimal stopping problem. In each period, the firm must decide whether to initiate a product redesign that makes use of an alternative component. Depending on that decision, the firm must also replenish or dispose of the spare part inventory for the legacy component and (after the redesign has been initiated) the alternative component. They find that a threshold policy is optimal, in which the design refresh is initiated once inventory drops below a period-specific threshold.

Some versions of the last time buy problem with additional sources of supply are continuous-time models. Teunter and Haneveld [49] consider a continuous-time model in which continued production is allowed after the last time buy but at a higher per-unit production cost with no setup costs. They assume demand is characterized by a Poisson process with a constant intensity throughout the problem horizon. The continuous-time framework allows them to find the optimal last-time-buy order-up-to level as well as the optimal order-up-levels for the higher-cost production option applicable to specific time intervals after the last time buy. They note that the optimal order-up-to levels after the last time buy are decreasing as the end of the horizon approaches. Leifker et al. [30] also consider a continuous-time model with an indefinite time horizon in which the firm generates revenue by selling spare parts to its customers. They consider two variants of the problem. In the first, the firm has the option to procure additional units as needed at a higher per-unit cost after the last time buy, while in the second the firm instead pays a lost sales penalty directly to the customer for any unmet demand. They show that, in the first variant with the option for additional production, the overall profit function is concave with respect to the size of the initial last time buy. For the second variant with a lost sales penalty, they are unable to confirm concavity but develop an upper bound that allows them to narrow down the number of potential solutions and thereby find the optimal last-time-buy quantity numerically. There are other papers (e.g., [51]) with continuous-time models, but they include other strategies that are more complicated; we will discuss them later in this literature review.

Another form of alternative sourcing involves switching to a different service strategy. In a series of papers, Frenk et al. consider this option in a continuous time framework with demand modeled as a non-homogeneous Poisson process. The manufacturer first places the last time buy at the beginning of the time horizon. Then, at a time of the manufacturer's choosing, it may discard all inventory and instead satisfy demand via an alternative strategy (such as replacement with a newer generation product) at a higher per-unit cost. In [20] they consider the static variant of the problem to find the optimal last-time-buy quantity and switching time. In [18] and [19] they consider the dynamic variant as an optimal stopping time problem and solve for the optimal policy via dynamic programming. A number of other papers ([25, 26, 37, 40]) allow for additional production in conjunction with other strategies that are more complicated, and therefore we will discuss them later in this literature review.

Product Returns

Another source of inventory is product returns from customers. Teunter and Fortuin [48] introduced this feature in the context of a spare parts system by incorporating an outside process that contributes a random number of units in each period. They assume that all returned units can be instantly remanufactured at zero cost and returned to the supply of spare parts. They also allow any number of available units to be discarded in any period to save on inventory holding costs. They find a near-optimal last-time-buy quantity as well as an optimal dispose-down-to level in each period after the last time buy. Kleber and Inderfurth [28] build upon [48] by including a per-unit remanufacturing cost. In their model, at the beginning of the problem horizon, a single last-time-buy order must be placed. In every subsequent period, the manufacturer may remanufacture any number of available returned units. The authors assume inventory holding costs are incurred on spare parts but not on returned items that have not yet been remanufactured. They also assume complete backordering with backordering costs incurred in intermediate periods, while shortage costs are incurred on any unfulfilled demand at the end of the horizon. They use a newsvendor-like heuristic to determine a near-optimal last-time-buy quantity as well as near-optimal remanufacture-up-to levels in subsequent periods.

Inderfurth and Mukherjee [26] extend the Kleber and Inderfurth [28] model by adding the option of an outside manufacturer with a higher per-unit production cost. They first propose a decision tree model with a limited time horizon and a small number of scenarios. However, they point out that detailed production and remanufacturing decisions necessitate a decision tree with an "almost exponentially increasing" size. Therefore, they also propose a stochastic dynamic programming formulation of the problem to determine the number of additional units to procure as well as the number of units to remanufacture. Although the dynamic programming formulation is intractable for realistic problem sizes, they suggest a simple time-dependent order-up-to and remanufacture-up-to policy based on a newsvendor-type analysis as a heuristic. They then compare the cost of their heuristic policy to that of the optimal solution for an example problem and find that the heuristic policy is only about 5% higher in cost than the more complex optimal policy. Inderfurth and Kleber [25] build upon this work to explore the performance of the two-parameter order-up-to heuristic developed in [26]. They carry out a numerical study and find that it yields near-optimal results in almost all instances. In Pourakbar et al. [41], the authors consider a slightly different setting in which a capital goods manufacturer has contracted with its customers for pre-planned phaseouts of the product. As customers purchase newer equipment, they return the older generation to the manufacturer. The manufacturer can harvest parts from these returns to be used for repairs of in-service equipment. The authors initially model the return timing and yields as deterministic but also extend the model to allow for stochastic return times and quantities. Importantly, the authors also note that these returns shrink the maximum warranty demand by removing products from circulation.

Repair

In the context of a machine that may require spare parts to continue operating, a natural extension of the last time buy problem is allowing the manufacturer to repair some parts rather than replace them. Van Kooten and Tan [51] consider a situation with the possibility of repairing failed components and returning them to the supply of spare parts in order to use those repaired units to satisfy future demand for components. They model these repairs as taking an exponential time as well as having a random yield. Any successful repairs result in a good-as-new spare part and any units whose repairs fail are discarded. This allows the authors to model the system using a continuous-time transient Markov chain for a given last-time-buy quantity. As each warranty claim arrives, one spare part is removed from inventory to satisfy the warranty demand and a repairable part is gained. After the repair lead time, a unit is moved from the repairable supply to the spare parts supply (if the repair succeeds) or it is lost permanently (if the repair fails). There are two absorbing states in which no more warranty claims can be satisfied, both of which the firm views as unacceptable in terms of customer service. In the first absorbing state, there are no spare parts or repairable units remaining, and therefore no further demand can be satisfied. In the second absorbing state, the repair backlog is above a threshold at which customer service is regarded as unacceptable. They find the distribution of time until the system reaches an absorbing state, which allows them to find a last-time-buy quantity that guarantees a specified service level during the time horizon. They also show that, rather than analyzing the Markov chain directly, they can use a binomial approximation of the time to reach the absorbing state that can be computed more efficiently. Behfard et al. [1, 2] consider a similar problem for a product with a single component with a repair option, with a known constant yield and positive lead time for repair. They assume a time-dependent base stock policy for repair, under which defective components are only repaired when the inventory of replacement parts falls below the base stock level in that period. Using a numerical search, they find a near-optimal last-time-buy quantity and base stock policy for the repair option.

In contrast to the aforementioned models that allow for repairs of components, other models allow for repairs of the entire defective units. For instance, van der Heijden and Iskandar [23] consider the case of a product with an increasing failure rate. The manufacturer must place the last time buy at the beginning of the decision horizon and pay a per-unit purchase cost. It may also choose to repair any failed units at a per-unit cost which is lower than the production cost. However, because the repair is assumed to be minimal, the unit is returned to operational status with the same distribution of remaining life as before the repair. Therefore, the repaired product's failure rate is unchanged, and there is a critical age threshold above which it is no longer cost-effective to repair the product. They present a model to jointly optimize the final order quantity and the product-age replacement threshold and find near-optimal values using a numerical search.

Finally, Pourakbar et al. [40] consider the last time buy with a repair option in the context of consumer electronics, which is similar to our setting. Due to price erosion, they anticipate that, at some point before the end of the time horizon, the unit price of the product will drop below the repair cost. Therefore, they consider an alternative policy of offering customers a newer version of the product or a discount on a next-generation product. They seek to jointly optimize the size of the last time buy for spare parts and a switching time, at which all remaining spare part inventory is scrapped and after which demand will be satisfied by the alternative policy. At any time before this switching time, the firm may attempt to repair defective devices. If the repair fails, the defective device is replaced via the alternative policy, and an additional penalty is paid. The authors consider both a static policy, optimizing the order quantity and switching time at the beginning of the problem horizon, and a dynamic policy, in which the switching time can be chosen based on new information as time progresses.

Product Trade-Ins

Multi-generational products enable product buy-backs or trade-ins (for the next generation product) to serve as another method of satisfying warranty demand after the last-time-buy decision has been made. While other authors have allowed for returned products that are then remanufactured, they all treat the product returns as a random process over which the manufacturer has no control. In reality, a manufacturer may offer a cash-back incentive or product discount to existing customers to entice them to trade in their older products in order to increase the supply of returns available for remanufacturing. The first paper to model this possibility is by Kleber et al. [29]. In their deterministic model, a firm sells spare parts to repair the units held by its existing customer base. In each period, after demand for spare parts occurs, the manufacturer may choose to buy back failed products in their entirety rather than repair them by replacing a broken part, incurring a per unit buy-back cost and forgoing the revenue from selling a spare part. However, the failed product can then be remanufactured at a per-unit cost in order to yield an additional spare part that can be used in the future. They present a mixed integer linear program (MILP) formulation and find the optimal final order quantity, buy-back quantity, and remanufacturing quantity in a multi-period system with deterministic product failure rates and deterministic remanufacturing yields. They also note that the buy-backs reduce the total possible demand by removing products from circulation.

Cole et al. [10, 11] also consider a trade-in offer, with the distinction that the offer is made to customers before their items have failed. They assume that customer valuations of the items they already own are uniformly distributed and ordered by the age of the product. They further assume that any customer will accept any trade-in offer that is higher than their personal valuation. Therefore, given the uniformly distributed valuations of the customers, a larger trade-in offer would correspond to a linear increase in trade-in volume; customers with older items take advantage of lower trade-in offers and progressively more customers with newer items take advantage of the trade-in offer as its value increases. They analyze a last time buy coupled with two different trade-in policies. The first is a "full trade-in" policy, in which a one-time offer is made to all customers, while the second is a "matching trade-in" policy, in which only a portion of customers are offered a trade-in in each period in order to match the rate of supply with the rate of demand. Due to the assumption that customer valuation is a function of the age of the product, the "matching trade-in" policy sends customers trade-in offers in order of their warranty expiration date. Using a deterministic analysis, they evaluate the cost of these policies and find that a trade-in may be advantageous for both consumers and the manufacturer. The consumers benefit from the opportunity for some of them to receive a cash payment greater than their personal valuation of the product, and the manufacturer benefits from the reduction in potential warranty demand that may need to be satisfied. The authors then briefly comment on the effect of uncertainty in demand or trade-in yield. They conclude that a one-time trade-in offer made to a sizable portion of the warranty population, with the potential for follow-up offers later in the time horizon, is likely to be preferable to exclusively pursuing either a "full trade-in" policy or a "matching trade-in" policy. Making a large initial offer that is accepted by a large portion of the population is preferable because it causes a significant drop in the warranty population, and therefore reduces future warranty claims.

Last Time Buy with Choice of Timing

We now turn to the literature on models in which the timing of the last time buy is a decision variable. Only a few authors consider this option. The earliest paper in this area is by Moore [36] who considers a case with deterministic demand. Accounting for the fixed and variable costs of production as well as inventory holding costs, Moore presents an optimization model designed to determine the timing of the last time buy along with production quantities up to and including the last time buy. Cattani and Souza [8] consider the possibility of postponing the last time buy rather than placing it at time zero in a case with stochastic demand in a make-to-order system. They assume the manufacturer pays a variable cost of production but no fixed costs of ordering. On the other hand, the supplier faces a maintenance cost that is non-decreasing in the duration of the postponement of the last time buy, which reflects the costs required to keep the production line available. Under various demand assumptions, they conclude that any delay of the last time buy benefits the manufacturer at the expense of the supplier due to both the maintenance costs incurred by the supplier. They also observe that an earlier last time buy necessitates a larger total order quantity because it needs to account for greater total uncertainty of demand. They then perform a game theoretic analysis to show that a benefit-sharing mechanism offered by the manufacturer is necessary to induce the supplier to agree to delay the last time buy and thereby achieve a mutually beneficial outcome.

A recent article examines a problem similar to ours, but with some differences in the assumptions about the decision structure. Ozyoruk et al. [37] consider the problem of dynamically optimizing decisions pertinent to the end-of-life phase of a product. In each period, the firm can choose one of three options: (1) place an order with a setup cost; (2) satisfy demand from inventory ("continuation"); or (3) dispose of all inventory, end the production option, and permanently transition to an outside supplier with a higher per-unit cost ("stopping"). The disposal option simplifies the problem by eliminating the need to account for the carry-over inventory. They formulate the problem as an optimal stopping time problem with additional decisions. They use Martingale theory to aid in the calculation of the value function and obtain a number of structural results. They show that, for the state space defined by the period and the starting inventory level, the optimal policy regions are disjoint but often intertwined and lack monotone boundaries. However, they are able to identify the ordering, continuation, and stopping regions. This result enables them to recursively calculate the distribution of the stopping time (when the system will transition to the outside supplier) for any given period and inventory level. They then report numerical results for a range of parameter settings. It is important to note that, in their model, the incentive for the firm to end production by transitioning to the outside supplier comes from the opportunity to avoid inventory holding costs through disposal. This is in contrast to our model, where the firm is motivated to end production in order to avoid the per-period capacity reservation costs that are required to keep the production line available.

2.2 Lost-Sales Inventory Problems

In periodic review inventory models, shortages are mainly handled in one of two ways: backordering or lost sales. In a model with backordering, any unmet demand is only temporarily left unsatisfied before eventually being satisfied with the next available unit of inventory. In some models, this process incurs a per-unit and/or per-unit per-period backordering cost, but the salient point is that the demand is eventually satisfied (except perhaps at the end of the horizon in a finite horizon problem). This allows the use of the inventory position (on-hand inventory plus on-order inventory minus backorders) to fully define the inventory state for a given period. However, in a lost-sales model, any demand that cannot be immediately satisfied from inventory is considered to be lost and may also incur a per-unit shortage cost. As pointed out in Bijvank and Vis [3], this means that the inventory position is no longer sufficient to fully define the inventory state, which instead must separately track the on-hand inventory as well as each order that has yet to arrive, resulting in a state space that is exponential in the lead time. For this reason, it is well known that lost-sales inventory models are more difficult to analyze, even though in real-world consumer environments the lost-sales assumption is often more appropriate than a backordering assumption [3].

As briefly mentioned in the introduction, our models have a type of fixed ordering costs in the form of the sequential capacity reservation fee as well as zero lead time and lost sales assumptions. In Bijvank and Vis's [3] review of lost-sales inventory models, they point out that, in lost-sales models with normal fixed order costs (i.e., without the requirement for continuity) and zero lead time, multiple authors have found a (R, s, S) policy to be optimal [9, 52, 53, 56]. In such a policy, R is the number of periods between reviews, s is the reorder point, and S is the order-up-to level. The reorder point s is required to avoid paying unnecessary fixed costs. However, in our model, the periods in which production may occur must be sequential, which is why we characterize them as capacity reservation costs rather than traditional fixed ordering costs. Therefore, the firm cannot avoid the capacity reservation cost if it plans to ever place an order in the future. To the best of our knowledge, no lost-sales inventory model makes use of our sequential capacity reservation cost structure.

Our main reason for assuming lost sales is computational convenience. After the last time buy, unmet demand is lost because no other sources of inventory are available (except in our problem variant with a buy-back option). By assuming lost sales rather than backordering in periods prior to the last time buy, we avoid the need to consider negative levels of inventory in the dynamic programming state space. To ensure that lost sales would be limited in practical applications, in our numerical studies, we utilize shortage costs that would lead to the moderate-to-high service levels typically observed in pertinent contexts.

2.3 Capacity Reservation Problems

There is a very substantial literature on capacity reservation problems in which the buyer reserves capacity in the target period typically at a constant cost per unit of capacity and then has the option to utilize the capacity at an additional constant cost per unit. Such a capacity reservation is often considered alongside one or both of the following options: (i) the opportunity to make an early purchase in the current period at a discounted cost per unit and (ii) the opportunity to make a "spot" purchase in the target period (when the demand occurs) either at a constant cost per unit that is known in advance or at an unknown spot price. Cachon [7] provides a comprehensive overview of contracting in supply chain problems, including capacity reservation contracts. We refer the reader to Wu et al. [55] and Martínez-Costa et al. [34] for surveys of this literature and to references in Qi et al. [42] and Li et al. [32] for examples of more recent research. The literature that we have mentioned thus far pertains to settings in which there is a single uncertain demand to be satisfied, usually representing the demand for the entire sales season of a seasonal product.

There is a small literature on multi-period capacity reservation problems. Costa and Silver [14] consider a multi-period inventory problem in which both the demand and supplier capacity in each period are stochastic. Prior to the first period, the firm must determine how much capacity to reserve in each period until the end of the planning horizon. The authors develop an exact solution based on a combination of branch-and-bound methods and dynamic programming. However, this problem differs from ours in that there is no requirement to reserve capacity in sequential periods. Furthermore, the authors require the entire capacity reservation to be made at the beginning of the planning horizon. Building on this work, Boulaksil et al. [4] consider a multi-period inventory problem in which an original equipment manufacturer (OEM) has outsourced production to a contract manufacturer (CM). The CM serves multiple OEMs and requires each OEM to place capacity reservations before ordering.

The CM then either accepts or partially rejects these reservations by some decision rule that is unknown to the OEMs. The authors present a stochastic dynamic programming model as well as some characteristics of the optimal policy. In general, they show that CM is able to achieve good results with relatively small reservation costs. However, this setting differs from ours in that the CM's production line is not used exclusively to satisfy the demand of one OEM, but is rather shared across many OEMs. This pooling allows the CM to be more flexible with the capacity reservation arrangement than our setting allows. A number of other papers (e.g., [26, 38, 44]) consider multi-period capacity reservation problems with varying assumptions; however, this literature is primarily concerned with the size of the capacity reservation, not the duration.

The nature of the capacity reservation in our problem context differs in several ways. First, the buyer reserves a pre-agreed increment of capacity (e.g., the capacity of a production line) that is large enough to accommodate any reasonable order size near the end of the product's life cycle. Therefore, there is no capacity quantity decision. The cost of the reservation accounts for the supplier's opportunity cost of capacity on the pertinent production facility. All variable procurement costs are included in the per-unit cost charged to the buyer. Second, the capacity reservation periods must be sequential because it is too difficult for the supplier to reconfigure the production line (if needed) and to train employees to produce another product, possibly just for a short period of time. This also helps to explain why our capacity reservation cost reflects the opportunity cost of capacity in any intervening periods in which the buyer chooses not to place an order.

Karmarkar et al. [27] consider an extension of the single-item dynamic lot-sizing problem in which the firm incurs a setup cost to switch a production machine from "off" to "on" as well as a fixed "reservation" cost in any subsequent period in which the machine remains on. They mention that this arrangement is more appropriate for products in which the production run lengths are commonly longer than the period lengths. They point out that, if the setup cost is zero, their problem becomes the classic Wagner-Whitin dynamic lot-sizing problem [54]. Karmarkar et al. present a dynamic programming algorithm for the uncapacitated case. However, in their problem the firm can restart production at any time after the machine has been turned off, possibly after producing some units, by paying the setup cost again, while in our problem any shutdown is permanent. The nature of the reservation in the Karmarkar et al. paper is the same as in our model but we do not have setup costs in our model. They highlight the key difference between *setup* and *reservation* costs, the latter of which represents the cost of "tying up the machine so that other products cannot be made on it." [27] These same underlying economic and operational phenomena motivate our use of capacity reservation fees, although the requirement for sequential reservations and the presence of stochastic demand make our problem more challenging.

Chapter 3

Problem Statement

We consider a problem facing a firm that sells a multi-generational product near the end of its product life cycle. At some predetermined time, the firm will discontinue the older generation of the product and exclusively sell the newer generation of the product. This leaves the firm with the question of *when* to stop ordering the older generation item and *how much* inventory to have available at that time. This problem is a variant of the last time buy problem in which the firm has the choice of when to place the last time buy as well as how large the purchase should be.

We assume a fixed time horizon with a predetermined time after which the firm intends to stop selling the older generation of the product. This is consistent with our primary setting of interest, the high-tech consumer electronics industry, in which the discontinuation time of each generation is decided by the firm, usually many months prior to the end of the product's selling horizon. We also assume that, given the history of demand, the firm can construct a statistical distribution of demand in each time period until the last possible period in which demand can occur, which we represent as the end of the time horizon. Although in reality, the firm would continue to update the demand forecasts over time as new data arrives, for simplicity we assume that the demand distributions are fixed at the beginning of the time horizon and are not updated based on new information. The demand may represent some combination of retail demand and demand for warranty replacement units, but we do not differentiate between these in our base model.

We specifically consider situations in which the product is assembled or produced by a contract manufacturer that has set up a production line specifically for the product under consideration. Time and effort are required for the contract manufacturer to change over the production line to produce a different product, so, under normal conditions, the contract manufacturer does not do so until production of the current product is terminated by the firm. The contract manufacturer charges the firm a capacity reservation fee per period while the line is set up to produce the product under consideration. This means that the firm must continue to pay the per-period capacity reservation fee in order to keep the production line at the contract manufacturer available even if the firm does not intend to order any units for the current period. We assume this cost is linear in the number of periods and there is no discount or surcharge for making a reservation for a longer period of time. We make this assumption because this cost primarily represents the contract manufacturer's opportunity cost: keeping the production line available to the firm in question prevents the contract manufacturer from allocating that equipment or personnel to other products. It is worth emphasizing that the firm may choose to pay the capacity reservation fee entirely for the option of being able to order in a future period. This differentiates the capacity reservation fee in our model from the setup cost for ordering included in more traditional inventory models.

The decision horizon may begin at any point during the product life cycle, but it would typically begin when demand for the product has declined sufficiently that the firm may prefer to use the contract manufacturer's resources to produce another product, e.g., the next generation product. Our problem is a planning problem, which is primarily concerned with the capacity reservation decision. Therefore, we utilize a discrete-time model in which each period is sized to the length of a typical capacity reservation, e.g., a month or quarter.

We assume there are no restrictions on the order quantity during the time interval under consideration, as the product is near the end of its life cycle when demand is declining and therefore we assume that anticipated future demand in any period will be below production capacity. We also make the simplifying assumption that, in any period in which capacity has been reserved, there is no minimum order quantity. We assume the firm uses a state-dependent order-up-to policy, so in each period it must choose an order-up-level that may depend on the starting inventory level. Due to the complexity of the problem, which we discuss in more detail later, it is not possible to show that a simple state-dependent policy is guaranteed to be optimal. However, we believe that a state-dependent order-up-to policy is sufficiently general to allow us to find near-optimal solutions, as it is feasible for the firm to select a unique order-up-to level for every possible combination of starting inventory and period, which is a more complex policy than any firm could realistically implement in practice. We restrict consideration to this class of policies and seek optimal solutions within this class.

We assume that there is no lead time, so any units that are ordered arrive in the same period and are available to satisfy demand during that period. The periods in our model correspond to capacity reservations for durations of months or quarters. In our model, the contract manufacturer is primarily performing final assembly, which can be completed in a short period of time. Shipments from the contract manufacturer would typically occur weekly. Further, in the consumer electronics setting, the combination of high value and small size and weight of the items would lead the firm to prefer rapid transportation options, such as air freight, along with frequent, e.g., weekly, shipments. If capacity reservations are for monthly or quarterly periods and shipments occur weekly, then the lead time would be only a fraction of a reservation period. Therefore, the zero lead time assumption is a reasonable approximation while at the same time greatly simplifying our model by eliminating differences between the on-hand inventory and the inventory position. This allows us to focus on our primary interest, the timing of the last-time-buy decision.

In all three variants of our problem, at some point, the firm will have exhausted all possible sources of inventory, either by ending the capacity reservation or, in the case of our third problem variant, by making the one-time buy-back offer. As a result, any unmet demand after this point must be lost as there are no more opportunities to order units, and therefore we cannot use a model with full backordering. We make a simplifying assumption that any unmet demand in *any period* is lost permanently (with no backordering allowed), even in periods prior to the last time buy. As mentioned earlier in Section 2.2, there is an important technical reason for this assumption, as it allows us to avoid a potentially much larger state space that would arise if inventory levels were allowed to be negative.

As part of our lost-sales model, we assume the firm incurs a shortage cost for each unit of unmet demand. In the consumer electronics setting, the lost-sales assumption may be justified because the product in question is in a competitive category in which brand switching is very common. If this is the case, the shortage cost may represent the lost profit associated with losing a customer to a competitor. There are also other contexts, such as demand for spare parts for units under warranty, in which unmet demand represents a breach of contract, and therefore should incur an immediate penalty. In this case, the shortage cost may represent the cost of a cash payment to the customer in compensation. Alternatively, this shortage cost may represent the lost profit arising from replacing the customer's broken product with a new next-generation product that the customer may have otherwise purchased at a later time. All of the above interpretations support relatively high shortage costs, leading to high service levels and limited lost sales prior to the last time buy, further supporting our assumption of a lost-sales model throughout the problem horizon.

The firm's objective is to minimize the expected cost of satisfying demand over the remaining time horizon. Because the production process is well established by the point in time when the firm needs to make a last-time-buy decision, all costs should be well known to the firm. We assume that all costs are positive values. We also assume that all costs are constant over the time horizon of concern. This assumption is justified due to the relatively short time horizon of the problem. For that same reason, we do not discount costs but it is straightforward to do so if desired. In addition to the per-period capacity reservation fee, the firm must pay the per-unit variable procurement cost for any units ordered, a per-unit inventory holding cost for any units in inventory at the end of any period, and a per-unit shortage cost for any unmet demand in any period. We assume that no disposal of inventory is allowed in any period, with the sole exception of the end of the time horizon, when excess units are discarded with no salvage value and at zero cost. While in reality, older generations of products may have a small value in secondary markets, in general, the high-tech nature of these products means they become obsolete quickly, and therefore for simplicity, we assume their salvage value is zero. It is straightforward to generalize our model to include salvage values and disposal costs.

We study three different variants of the last-time-buy problem described above, all of which include the requirement that the firm must continue to pay periodic capacity reservation fees in order to keep the production line at a contract manufacturer available. All assumptions above apply to each variant unless otherwise specified. In Section 3.1 we consider the simplest version of the problem, in which the firm must pre-commit to the duration of the capacity reservation at the start of the time horizon. In Section 3.2 we consider a dynamic variant of the problem, in which the firm can decide in each period whether to extend the capacity reservation. Finally, in Section 3.3 we consider an extension of the dynamic variant of the problem, in which the firm has the option to make a one-time product buy-back offer in any period after the capacity reservation has been ended. In the other two variants of the problem, we do not explicitly differentiate between retail demand and demand for warranty replacement units. However, in this variant, we assume that all late-horizon demand is for warranty replacement units, and therefore the buy-back units can be refurbished to serve as an alternate source of units to satisfy warranty claims. We will discuss the details and assumptions of this buy-back offer in the relevant section.

3.1 Pre-Commit Variant

We now describe our first problem variant in which the firm must pre-commit at the beginning of the decision horizon to the remaining duration of the capacity reservation. However, in each intervening period, the firm may order any quantity of the product; only the timing of the last possible order is pre-committed. The sequence of events is as follows:

1. At the start of the decision horizon, select the final period of the capacity reservation and pay the per-period capacity reservation fees through the selected period.

2. In each period up to and including the last period of the capacity reservation:

    a) Order and pay for units.

    b) Receive units.

    c) Fulfill demand to the extent inventory is available.

    d) Incur any inventory holding or shortage costs.

    e) Proceed to the next period.

3. In each period after the conclusion of the capacity reservation (note that there are no decisions in this situation):

    a) Fulfill demand to the extent inventory is available.

    b) Incur any inventory holding or shortage costs.

    c) Proceed to the next period.

This pre-commit variant will primarily serve as a baseline for comparison with the dynamic variants of the model, which are detailed below. We note, however, that capacity reservation contracts do exist in which the term of the reservation must be decided in advance, possibly with some flexibility for modification at a cost.

3.2 Dynamic Variant

We now discuss the dynamic variant of this problem. Instead of pre-committing to reserve capacity for a number of periods, the firm contracts with the manufacturer on a period-to-period basis, paying the capacity reservation fee to keep the product line available. Once the firm chooses not to pay the capacity reservation fee, the firm will no longer be able to order units from the contract manufacturer. Decisions concerning both the capacity reservation and, if applicable, the order quantity for the current period are made at the beginning of each period. In practical application, the capacity reservation decision may actually be made in the latter part of the preceding period. Our assumption regarding the timing of the decision is an approximation, but it is not far from reality in instances where the contract manufacturer's contribution to the production is primarily manual assembly. Assuming the firm chooses to continue the capacity reservation, the sequence of events is as follows:

1. Pay the capacity reservation fee.
2. Order and pay for units.
3. Receive units.
4. Fulfill demand to the extent inventory is available.
5. Incur any inventory holding or shortage costs.
6. Proceed to the next period.

In the first period in which the firm chooses not to pay the capacity reservation fee, and in all subsequent periods, the firm faces the following sequence of events (note that there are no decisions in this situation):

1. Fulfill demand to the extent inventory is available.
2. Incur any inventory holding or shortage costs.
3. Proceed to the next period.

The optimal solution for the pre-commit variant described in Section 3.1 will yield an upper bound on the optimal expected cost of the dynamic problem described here, because the cost elements are the same but the decision space is more constrained. This upper bound is interesting because the difference between it and the optimal expected cost of the dynamic problem represents the value of the flexibility afforded by the option to decide dynamically whether to continue the capacity reservation or not. Some firms facing competition in securing contract manufacturing capacity face the pre-commit version of the problem, while others with more power relative to the contract manufacturer may be able to take advantage of the flexibility afforded by the dynamic variant.

3.3 Dynamic Variant with One-Time Buy-Back Option

We now discuss a variant of the dynamic version of this problem that includes the option for a one-time product buy-back offer. This variant is especially relevant in the context of warranty demand for a discontinued product. We assume that, at this stage in the product life cycle, there is no future retail demand and the only remaining demand is from warranty claims. The goal of the buy-back is to source additional units which may be refurbished and used to satisfy future warranty demand.

We assume that, at some point before the end of the time horizon, the firm may make a buy-back offer to existing customers who own the product in question. We assume that the firm can choose the buy-back price and can estimate the number of customers who will take advantage of the offer at any given buy-back price. For simplicity, we initially assume that the total number of units acquired from customers is a deterministic function of the buy-back price. However, later in this section, we introduce an extension in which we relax this assumption by allowing the total number of units to be a stochastic function of the buy-back price.

We assume that the buy-back offer can occur only after the firm has ended the capacity reservation at the contract manufacturer. This assumption is not too restrictive because the firm must offer a buy-back price that is high enough to induce some customers to part with their devices, and the customer's expectation for the value of the item is tied to the retail price they paid, which is often double the production cost or more. Therefore, the unit cost of the buy-back will be higher than the per-unit procurement cost of ordering them from the contract manufacturer, so the firm would not utilize the buy-back offer if the option to order from the contract manufacturer were still available. Requiring that the buy-back offer occurs after the last time buy allows us to simplify the decision structure without greatly compromising the model's realism.

We assume that the buy-back offer can only be made once. Our rationale for this assumption can be explained as follows. Suppose for a moment that multiple buy-back offers were possible. If the customer's personal valuation of the product is assumed to be constant over the time horizon (which would be reasonable in the short term), then any follow-up buy-back offer after a previous offer would only yield incremental units if a higher per-unit buy-back price were offered. This means that the buy-back yield of each subsequent offer would depend not just on the current value of the offer but also on the values of all previous buy-back offers made. This would cause the decision structure to be very complicated, and for that reason, we restrict the buy-back offer to be a one-time option only. In practice, the combination of the administrative inconvenience of offering a buy-back and the relatively short time between the depletion of newly-produced parts available for warranty (i.e., the point at which a buy-back offer first becomes useful) and the end of the time horizon means that the firm will realistically have only one opportunity to organize a buy-back with its customers, so this assumption is not very restrictive.

Obviously, new units may be used to satisfy warranty claims. However, most warranties only require the replacements to be of a similar quality to the used, in-service units. Therefore, we assume that the firm is able to refurbish the buy-back units to a quality level that is adequate to satisfy warranty requirements, which may be lower than that of new units. For example, their durability may be lower. We assume the inventory holding costs of new and refurbished units are not sufficiently dissimilar to distinguish them, so as an approximation, we treat them as equal.

We assume that the firm has a good estimate of the cost and yield of this refurbishing process due to the mature stage of the product within its life cycle as well as the firm's experience with previous warranty claims. We also assume that a deterministic fraction of units can be refurbished. We combine the cost to acquire the buy-back units and the cost to refurbish them into a single "buy-back cost," and we express the total cost of a buy-back offer as a function of the net number of buy-back units added to inventory after refurbishment.

After considering this deterministic model, we also consider an extension in which the total number of units acquired from customers is a stochastic function of the buy-back price. In this extension, the firm chooses a per-unit buy-back price to offer to its customers, but the number of customers who respond to the offer is random, with the expected number of responding customers increasing with the buy-back price. However, we retain the assumption that a deterministic fraction of buy-back units can be refurbished at a constant per-unit cost and therefore continue to represent the cost to acquire and refurbish each unit as a single "buy-back cost." This per-unit "buy-back cost" is incurred for a random number of units, causing the overall buy-back cost to be a stochastic function of the buy-back price selected by the firm. We discuss the implementation of the stochastic buy-back model in detail in Section 8.5.

We now discuss the sequence of events facing the firm. In periods in which the firm chooses to continue the capacity reservation with the contract manufacturer, the sequence of events is identical to that described in Section 3.2:

1. Pay the capacity reservation fee.
2. Order and pay for units.
3. Receive units.
4. Fulfill demand to the extent inventory is available.
5. Incur any inventory holding or shortage costs.
6. Proceed to the next period.

However, after the capacity reservation has been terminated, the firm now has the option to make a one-time buy-back offer and therefore faces the following sequence of events in each period:

1. If the one-time buy-back option has not yet been exercised and the firm chooses to exercise it in the current period:

   a) Choose a buy-back quantity, collect those items, and then refurbish and return those items to inventory.

   b) Pay the corresponding buy-back and refurbishing costs.

2. Fulfill demand to the extent inventory is available.

3. Incur any inventory holding or shortage costs.

4. Proceed to the next period.

After the one-time option has been exercised, it can no longer be exercised in any subsequent period. Therefore, in each period after the buy-back has been offered, the firm has no decisions remaining and faces the following sequence of events:

1. Fulfill demand to the extent inventory is available.

2. Incur any inventory holding or shortage costs.

3. Proceed to the next period.

The solution to the dynamic variant described in Section 3.2 provides an upper bound on the optimal expected cost of this variant. The difference between the respective expected costs represents the value of the flexibility offered by the buy-back option.

In the following chapters, we present formulations of the three problem variants described above and characterize properties of the corresponding optimal solutions. We also present a heuristic for the dynamic variant of the problem and contrast its performance with that of the optimal solution found via dynamic programming.

In the next three chapters, we present formulations of our various problem variants. Due to the presence of binary decision variables for extending the capacity reservation at a fixed cost per period, binary decision variables for selecting the buy-back option, interrelations among certain binary variables, as well as inventory-related decisions and costs under stochastic demand and buy-back yields (as applicable), our formulations are nonlinear mixed-integer dynamic programs. As such, some of our formulations include constructs and constraints that are more complicated than in many production/inventory models.

Chapter 4

Pre-Commit Model

We consider capacity reservation and procurement decisions facing a firm that sells a product over a fixed time horizon of N periods. The periods are numbered in chronological order, with period 1 being the first period of the problem horizon and period N being the final period in which demand occurs. The production line for the product is assumed to be available for reservation at a contract manufacturer at the start of the problem horizon and the firm has a known number of units, I_1, in stock, where I_n denotes the number of units in inventory at the beginning of period n.

In the pre-commit variant of the problem, at the beginning of the first period, the firm must decide how many consecutive periods it would like to reserve capacity at the contract manufacturer. Although there are many ways to represent this decision, for mathematical convenience, we represent this decision with the vector $\mathbf{y}$, which consists of binary variables $y_n : n = 1, ..., N$, which are set equal to 1 if production capacity is reserved in period n, and 0 otherwise. We define a dummy variable $y_0 = 1$ that indicates capacity was reserved immediately prior to the beginning of the problem horizon, which ensures that we have the option of reserving capacity in the first period and possibly continuing. For each period the firm reserves capacity it must pay a capacity reservation fee c_f, so a choice of n periods would incur a cost of nc_f.

Let $S_n(I_n)$ represent the order-up-to level in period n if there are I_n units on hand at the beginning of the period. For compactness, we will often represent the order-up-to variable simply as S_n. This value is a constrained order-up-to level, with the restriction that $S_n \geq I_n$, because disposal is not allowed by assumption, and, in practice, firms rarely dispose of inventory when they anticipate being able to use it within a short period of time. The firm must pay a per unit procurement cost c_p. Therefore, the total ordering cost in period n is $c_p(S_n - I_n)$. Unmet demand is permanently lost and incurs a per-unit shortage cost c_s. Any excess units in inventory at the end of any period incur a per-unit inventory holding cost c_h, except at the end of the horizon, when any excess units are disposed of with zero salvage value and no disposal cost. It is straightforward to generalize our model and methodology to incorporate positive salvage values or disposal costs. In each period n, demand is represented by the random variable D_n with a corresponding probability density function f_n.

We seek to optimize the firm's decisions over the time horizon. We do this by constructing a dynamic program which expresses the cost in each period as a function of the current period costs plus the expected cost-to-go assuming optimal decisions are made in all subsequent periods. These costs depend on the two-dimensional state of the system at the time, defined by the starting inventory level and the binary variable indicating whether or not the production line was reserved for the current period.

Throughout the dissertation, the term "policy" is used when describing the general form of the solution (such as an order-up-to-S inventory policy), whereas "strategy" is used when referring to the full, state-contingent dynamic programming solution for a specific problem instance. In the pre-commit variant of the problem, the optimal strategy consists of: (1) the optimal duration of the sequential capacity reservation, and (2) the optimal state-dependent order-up-to level for each period for which capacity was reserved.

We next present the dynamic programming formulation. In our formulations, we separate the decision variables from the state variables using a pipe ($|$) to make clear the distinction.

4.1 Model Formulation

We use $C_n(S_n|I_n, \mathbf{y})$ to denote the value function corresponding to producing up to S_n units in period n when the firm has pre-committed to the capacity reservation defined by the vector of decision variables $\mathbf{y}$ and starts period n with I_n units on hand. Throughout, we use an asterisk (*) to signify the optimal expected cost-to-go function. This represents the cost associated with the firm making the optimal decision in the current period and all subsequent periods for the given state variables. Therefore, we use $C_n^*(I_n, \mathbf{y})$ to denote the optimal expected cost-to-go associated with the state variables $(I_n, \mathbf{y})$.

We can now recursively define the value and cost-to-go functions for all periods. The terminal value function is:

$$C_{N+1}(\cdot) = 0 \tag{4.1}$$

because all inventory is disposed of at the end of the horizon with no salvage value and no disposal cost. It is clear that we also have:

$$C_{N+1}^*(\cdot) = 0 \tag{4.2}$$

because there are no decisions to make at the end of the horizon. However, because the terminal value function does not impact the firm's decisions, we exclude it from the formulation in the remainder of this chapter.

For period N, we have the value function:

$$C_N(S_N|I_N, \mathbf{y}) = c_p(S_N - I_N) + c_s \int_{S_N}^{\infty} (x - S_N) f_N(x) dx \tag{4.3}$$

with the obvious constraint $S_N \geq I_N$. In the final period, the only costs that apply are the procurement costs, if any, and the expected shortage costs. We also have the optimal expected cost-to-go:

$$C_N^*(I_N, \mathbf{y}) = \min_{I_N \leq S_N \leq I_N + My_N} C_N(S_N | I_N, \mathbf{y}) \tag{4.4}$$

where M is a very large constant. The upper bound in the constraint on S_N is a *big-M constraint* which specifies that an order can be placed only if allowed by the pre-committed capacity reservation; this constraint links the binary capacity reservation decision with the order quantity decision. The lower bound in the constraint on S_N ensures that the selected order-up-to level is greater than or equal to the starting inventory level. Therefore, if the firm has not pre-committed up to and including period N, $y_N = 0$ and consequently, $S_N = I_N$.

For any period $n \in \{2, 3, ...N-1\}$, we have the value function:

$$\begin{aligned} C_n(S_n | I_n, \mathbf{y}) =& c_p(S_n - I_n) + c_h \int_0^{S_n} (S_n - x) f_n(x) dx + c_s \int_{S_n}^{\infty} (x - S_n) f_n(x) dx \\ &+ \int_0^{S_n} C_{n+1}^*(S_n - x, \mathbf{y}) f_n(x) dx + \int_{S_n}^{\infty} C_{n+1}^*(0, \mathbf{y}) f_n(x) dx \end{aligned} \tag{4.5}$$

with the constraint $S_n \geq I_n$. In (4.5), the first term is the variable procurement cost, the second and third terms are the expected inventory holding and shortage costs, respectively, in period n, and the final two terms are collectively the optimal expected cost-to-go in period $n+1$. Due to our lost sales assumption, the inventory at the end of period n is zero if demand exceeds S_n, and therefore we separate the expectation by whether the demand results in positive or zero inventory at the end of period n after demand has occurred, which is also the inventory level at the beginning of period $n+1$. We also have the optimal expected cost-to-go:

$$C_n^*(I_n, \mathbf{y}) = \min_{I_n \leq S_n \leq I_n + My_n} C_n(S_n | I_n, \mathbf{y}) \tag{4.6}$$

where the constraint ensures that an order can be placed only if allowed by the pre-committed capacity reservation and that the selected order-up-to level must be greater than or equal to the starting inventory level.

Finally, for period 1, we have the value function:

$$\begin{aligned} C_1(S_1, \mathbf{y} | I_1) =& c_f \sum_{n=1}^{N} y_n + c_p(S_1 - I_1) \\ &+ c_h \int_0^{S_1} (S_1 - x) f_1(x) dx + c_s \int_{S_1}^{\infty} (x - S_1) f_1(x) dx \\ &+ \int_0^{S_1} C_2^*(S_1 - x, \mathbf{y}) f_1(x) dx + \int_{S_1}^{\infty} C_2^*(0, \mathbf{y}) f_1(x) dx \end{aligned} \tag{4.7}$$

with the constraint $S_1 \geq I_1$. The first term represents the cost of pre-committing to the capacity reservation for the selected number of periods, the second term is the variable procurement cost of ordering up to S_1 units, the third and fourth terms are the expected inventory holding and shortage costs, respectively, in period 1, and the final two terms are the optimal expected cost-to-go in period 2 accounting for cases in which the final inventory in period 1 is positive or zero after demand has occurred, respectively. We also have the optimal expected cost-to-go:

$$C_1^*(I_1) = \min_{y_n \in \{0,1\}, y_n \leq y_{n-1} \,\forall\, n \in \{1,2...N\}} \left\{ \min_{I_1 \leq S_1 \leq I_1 + My_1} C_1(S_1, \mathbf{y}|I_1) \right\} \tag{4.8}$$

In the inner minimization, the constraint once again ensures an order can be placed only if the capacity reservation has been extended for the current period. For the outer minimization, the first constraint limits $\mathbf{y}$ to binary choices while the second ensures that all periods for which capacity is reserved are sequential.

4.2 Properties of the Optimal Solution

Throughout the remainder of this section, when we refer to the optimal order-up-to level, we mean the optimal order-up-to level satisfying the constraint $S_n(I_n) \geq I_n$. We will first show that at least one such optimal constrained order-up-to level exists.

Proposition 4.1. *An optimal constrained order-up-to $S_n^*(I_n)$ exists that minimizes the value function $C_n(S_n|I_n, \boldsymbol{y})$ for arbitrary n and I_n.*

Proof. See Appendix A. □

With the existence of an optimal order-up-to level established, we will now establish a few results that are helpful in calculating the optimal strategy. We will first consider the cost-to-go function for all periods after the capacity reservation has been ended.

Proposition 4.2. *Suppose period m is the final period of the pre-committed capacity reservation. Then, for all $n \in \{m+1, m+2, ...N\}$ the optimal expected cost-to-go $C_n^*(I_n, \mathbf{y})$ is convex in the starting inventory level I_n.*

Proof. With the capacity reservation ended, we have $y_n = 0$ for $n > m$, and the constraint in equation (4.6) is binding, so no further units can be purchased. Therefore the only remaining costs are the linear inventory holding and shortage costs for all remaining periods. It is well known that this results in an expected cost function which is convex in I_n. □

With the convexity of all periods after the capacity reservation established, we now consider the value function in the final period of the capacity reservation.

Theorem 4.3. *Suppose period n is the final period of the pre-committed capacity reservation. The value function $C_n(S_n|I_n, \boldsymbol{y})$ for a given starting inventory level I_n is convex in the order-up-to level S_n for $S_n \geq I_n$.*

Proof. Consider the value function $C_n(S_n|I_n, \mathbf{y})$ defined in (4.5). The first term is linear in S_n. The expected inventory holding cost and shortage cost in terms two and three, respectively, are convex in S_n. Finally, the integrand in each of the final two terms is convex in the first argument by Proposition 4.2, and, because convexity is preserved in expectation, the final two terms are also convex in S_n. Therefore the overall value function is convex in the order-up-to level S_n. □

Theorem 4.3 allows us to find the optimal order-up-to level through a procedure such as bisection search for a given starting inventory level I_n when period n is the final period of the capacity reservation. In Section 6.2 we make use of this property in our construction of the heuristic for the dynamic variant. However, Theorem 4.3 only allows us to find the optimal state-dependent order-up-to level in the final period of the capacity reservation itself. For the pre-commit model overall, we find the optimal strategy numerically, as detailed in Chapter 7.

Chapter 5

Dynamic Model

We now consider the dynamic version of the problem, in which the firm contracts with the contract manufacturer on a period-to-period basis. Much of the setting and notation are the same as in the previous model, with the main difference being that instead of a pre-committing to a production availability schedule $\mathbf{y}$ ahead of time, the firm can decide in each period if it would like to keep the production line available. The sequential capacity reservation requirement still holds: provided that the production line was reserved in the previous period, the firm may pay the capacity reservation fee c_f and keep the option to order in the current period. Once again we define the binary decision variable y_n to be equal to one if the production line is reserved in period n, and zero otherwise. We also introduce a new binary state variable α_n, which is equal to one if the capacity reservation was extended in period $n-1$, and zero otherwise. We assume that $\alpha_1 = 1$ so that it is feasible to reserve capacity in period 1 and possibly beyond that. All other variables and parameters are unchanged from the pre-commit version of the problem.

In the remainder of this chapter, we present our formulation and initial analysis.

5.1 Model Formulation

After the capacity reservation has been terminated (i.e., $\alpha_n = 0$) there are no more decisions to make and therefore the value function in any period n only depends on the remaining state variable I_n. Let $C_n(I_n, \alpha_n = 0)$ be the value function in period n starting with I_n units of inventory and with the capacity reservation ended before the beginning of the period.

As in Chapter 4, the terminal value function is:

$$C_{N+1}(\cdot) = 0 \tag{5.1}$$

because all inventory is disposed of at the end of the horizon with no salvage value and no disposal cost. It is clear that we also have:

$$C^*_{N+1}(\cdot) = 0 \tag{5.2}$$

because there are no decisions to make at the end of the horizon. However, because the terminal value function does not affect the firm's decisions, we exclude it from the formulation in the remainder of this chapter.

We begin with period N and then define $C_n(I_n, \alpha_n = 0)$ recursively using the fact that $\alpha_{n+1} = \alpha_{n+2} = \ldots = \alpha_{N-1} = 0$. The value function in period N if the capacity reservation has been ended and the initial inventory is I_N units consists only of the expected shortage costs:

$$C_N(I_N, \alpha_N = 0) = c_s \int_{I_N}^{\infty} (x - I_N) f_N(x) dx \tag{5.3}$$

Similarly, we can define the value function after the capacity reservation has been ended for an arbitrary period $n < N$ with I_n units of starting inventory as follows:

$$\begin{aligned} C_n(I_n, \alpha_n = 0) = {} & c_h \int_0^{I_n} (I_n - x) f_n(x) dx + c_s \int_{I_n}^{\infty} (x - I_n) f_n(x) dx \\ & + \int_0^{I_n} C^*_{n+1}(I_n - x, \alpha_{n+1} = 0) f_n(x) dx + \int_{I_n}^{\infty} C^*_{n+1}(0, \alpha_{n+1} = 0) f_n(x) dx \end{aligned} \tag{5.4}$$

The first two terms are, respectively, the expected inventory holding cost and the expected shortage cost in period n. The third and fourth terms together are the expected cost-to-go in period $n+1$. We also note that, because there are no decisions to make in this situation, the optimal expected cost-to-go is equivalent to the value function, and therefore in our notation, for an arbitrary period n:

$$C^*_n(I_n, \alpha_n = 0) = C_n(I_n, \alpha_n = 0) \tag{5.5}$$

Although there are no decisions to make, we are still interested in properties of the cost-to-go function that will help to characterize the optimal strategy. We next show that $C^*_n(I_n, \alpha_n = 0)$ is convex.

Proposition 5.1. *The optimal expected cost-to-go after the capacity reservation has been ended, for any arbitrary period n, $C^*_n(I_n, \alpha_n = 0)$ is convex in the starting inventory level I_n.*

In the final period N, this result follows directly from the linear inventory holding and shortage costs. We next show by induction that convexity also holds for any arbitrary period n.

Proof. Consider the period-N value function as defined in (5.3). In this case, the shortage cost function is linear, and therefore the value function is convex in the starting inventory I_N. As the optimal expected cost-to-go is equivalent to the value function in (5.5), the optimal expected cost-to-go function is also convex in the starting inventory I_N. We next show by induction that the convexity property carries over to any arbitrary period n.

Assume for the purpose of induction that the expected cost-to-go function in period $n+1$ is convex in the starting inventory level I_{n+1}. Consider the value function for a period

$n < N$ as defined in (5.4). Once again, the inventory holding and shortage costs are linear and therefore the expected inventory holding and shortage costs are convex in the starting inventory I_n. The third and fourth terms together are the optimal expected cost-to-go in period $n+1$, which is convex in the initial inventory by assumption, and convexity is preserved under expectation. Therefore, by induction, $C_n(I_n, \alpha_n = 0)$ is convex in the initial inventory I_n for any period n. □

The form of the value function when the firm has the option to extend the capacity reservation depends on the period. If it is the final period N, the value function includes only expected shortage costs in addition to the fixed and variable procurement costs. Therefore, the value function if the firm extends the capacity reservation and orders up to S_N units in period N is:

$$C_N(S_N, y_N | I_N, \alpha_n = 1) = c_f y_N + c_p(S_N - I_N) + c_s \int_{S_n}^{\infty} (x - S_n) f_n(x) dx \tag{5.6}$$

with the constraint $I_N \leq S_N \leq I_N + M y_N$. On the other hand, for any other period $n < N$, the value function associated with extending the capacity reservation and producing up to S_n units in period n is:

$$\begin{aligned} C_n(S_n, y_n | I_n, \alpha_n = 1) &= c_f y_n + c_p(S_n - I_n) \\ &+ c_h \int_0^{S_n} (S_n - x) f_n(x) dx + c_s \int_{S_n}^{\infty} (x - S_n) f_n(x) dx \\ &+ \int_0^{S_n} C^*_{n+1}(S_n - x, \alpha_{n+1} = y_n) f_n(x) dx + \int_{S_n}^{\infty} C^*_{n+1}(0, \alpha_{n+1} = y_n) f_n(x) dx \end{aligned} \tag{5.7}$$

with the constraint $I_n \leq S_n \leq I_n + M y_n$. The first two terms in the above expression represent the cost of ordering up to S_n units, consisting of the capacity reservation fee plus the variable procurement cost. The third and fourth terms are the expected inventory holding and shortage costs in the current period. Finally, the fifth and sixth terms together represent the optimal expected cost-to-go from period $n+1$ and onward, given the resulting end-of-period inventory level and the fact that the capacity reservation was extended in period n. The fifth term accounts for cases in which demand is less than or equal to the order-up-to level, while the sixth term accounts for cases in which demand exceeds the order-up-to level.

If the firm chooses to not extend the capacity reservation in period n, that corresponds to setting $y_n = 0$ and $S_n = I_n$. In that case, (5.7) becomes equivalent to (5.4) and the value function is the same as if the period had started with the capacity reservation not available. We can now consider the optimal expected cost-to-go.

Let $C^*_n(I_n, \alpha_n = 1)$ be the optimal expected cost-to-go from period n onward when the capacity reservation can be extended:

$$C^*_n(I_n, \alpha_n = 1) = \min_{y_n \in \{0,1\}, I_n \leq S_n \leq I_n + M y_n} C_n(S_n, y_n | I_n, \alpha_n = 1) \tag{5.8}$$

The constraints in the minimization account for the two requirements that (1) the order quantity can be positive only if the capacity reservation is extended in the same period and (2) the chosen order-up-to level must be at least as large as the starting inventory (i.e., no disposal). The sequential capacity requirement is enforced through the setting of $\alpha_{n+1} = y_n$ in the value function defined in (5.7). Once again, notice that a solution of $y_n = 0$ forces $S_n = I_n$ and $\alpha_{n+1} = 0$, in which case (5.8) would be equivalent to (5.4).

We now consider the optimal order-up-to level assuming that the firm decides to extend the capacity reservation. Let $S_n^*(I_n)$ be an optimal state-dependent order-up-to level that minimizes $C(S_n, y_n = 1|I_n, \alpha_n = 1)$ subject to the constraint $S_n \geq I_n$. First, we show that such an order-up-to level exists:

Proposition 5.2. *For an arbitrary period n and starting inventory level I_n, there exists at least one optimal constrained order-up-to level $S_n^*(I_n)$ that minimizes the value function $C_n(S_n, y_n = 1|I_n, \alpha_n = 1)$.*

Proof. See Appendix A. □

Because at least one optimal constrained order-up-to level exists, the firm can choose between (1) extending the reservation and ordering up to $S_n^*(I_n)$, and (2) not extending the reservation. So we have an equivalent expression to (5.8):

$$C_n^*(I_n, \alpha_n = 1) = \min \left\{ C_n(S_n^*(I_n), y_n = 1|I_n, \alpha_n = 1)\ ,\ C_n(I_n, \alpha_n = 0) \right\} \tag{5.9}$$

where the optimal costs of both alternatives, as shown in the braces of (5.9), depend on the starting inventory level I_n. Our analysis will focus on the relationship between these two functions for a given period n as a function of the starting inventory level I_n.

We will now explore some properties of the optimal solution.

5.2 Properties of the Optimal Solution

Throughout the remainder of this section, when we refer to an optimal order-up-to level, we mean an optimal *constrained* order-up-to level satisfying the constraint $S_n(I_n) \geq I_n$. Due to the structure of the DP value function, we cannot rule out the possibility of multiple optimal order-up-to levels for a given starting inventory. Consider a two-period problem with a deterministic demand of 100 units per period, a per-unit per-period inventory holding cost of 1, a per-period capacity reservation cost of 100, a per-unit shortage cost of 100, and a per-unit procurement cost of 1. In this example, the shortage cost is so large relative to the procurement cost that it is optimal to satisfy all demand. If we assume there is no inventory available at the beginning of the horizon, the two reasonable solutions are to either: (1) order 100 units in each period or (2) order 200 units in the first period and hold the remaining 100 units in inventory until the second period. Because the capacity reservation fee is equal to the cost to hold 100 units in inventory for one period, these two solutions have the same cost. As a result, the optimal order-up-to level for any given starting inventory level is not necessarily unique.

Our numerical results show that, in practice, there are often two types of optimal order-up-to levels that occur across starting inventory levels: either a "small" order-up-to level that would satisfy the current period's demand with a high probability or a "large" order-up-to level that is sufficient to serve as the last time buy. In the problem instance described in the previous paragraph, 100 would be the "small" order-up-to level and 200 would be the "large" order-up-to level. Furthermore, in cases where both types of order-up-to levels are optimal in a given period but for different starting inventory levels, the "small" order-up-to level is generally optimal for smaller starting inventory levels while the "large" order-up-to level is generally optimal for larger starting inventory levels. We now use an example problem from our numerical study (see Chapter 7) containing both types of optimal order-up-to levels to further illustrate some properties of the optimal solution.

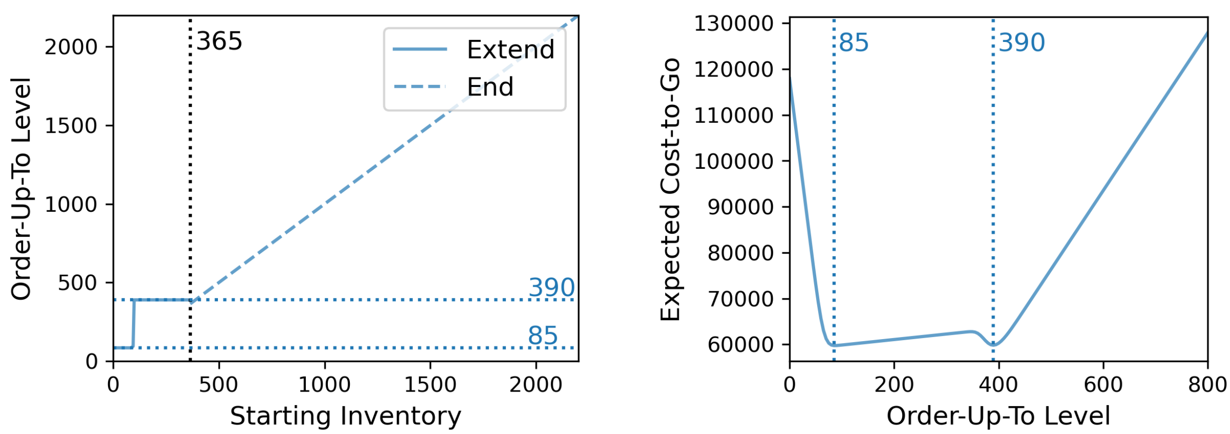

(a) Optimal Order-Up-To Levels in Period 6

(b) Value Function in Period 6 as a Function of S_n for $I_6 = 0$

Figure 5.1: Example of an Optimal Order-Up-To Strategy

Figure 5.1 was generated from a discretized problem instance in our numerical study. Figure 5.1a plots the optimal order-up-to level as a function of the starting inventory level and clearly shows that there are two prominent optimal order-up-to levels that result in positive order quantities. For starting inventory levels between 0 and 85, it is optimal to order up to 85 units. However, for starting inventory levels between 100 and 365, it is optimal to order up to 390 units. (Note that for starting inventory levels between 85 and 100, it is optimal to extend the capacity reservation but order nothing, a strategy we will discuss in greater detail in Chapter 7.) Both of the optimal order-up-to levels that result in positive order quantities (85 and 390) are visible as local minima of the value function (equation (5.7)) plotted in Figure 5.1b with respect to the order-up-to level, S_n. This implies that there may exist some starting inventory level for which the value function associated with these two order-up-to levels are equal, especially in a case with continuous demand quantities and starting inventory levels. This complicates attempts to characterize the optimal order-up-to levels. However, we can make the following claim:

Theorem 5.3. *If the capacity reservation is extended in an arbitrary period n, the smallest optimal constrained order-up-to level $S_n^*(I_n)$ is monotonically non-decreasing in the starting inventory level I_n.*

Without loss of generality—because the DP value function incorporates all of the consequences of the current decision on expected future costs—we can focus on the smallest optimal order-up-to level for each I_n. We will establish our claim by considering an arbitrary period n and starting inventory level I_n along with the corresponding smallest optimal order-up-to level $S_n^*(I_n)$ and showing that for any larger starting inventory level, any smaller order-up-to level is either (1) not feasible, or (2) more costly. As part of the proof, we will also prove the following corollary:

Corollary 5.3.1. *If the capacity reservation is extended in any arbitrary period n, for any starting inventory level I_n with its corresponding smallest optimal constrained order-up-to level $S_n^*(I_n)$, then $S_n^*(\tilde{I}_n) = S_n^*(I_n)$ for all starting inventory levels $\tilde{I}_n$ such that $I_n \leq \tilde{I}_n \leq S_n^*(I_n)$.*

This corollary is especially useful because it means that, for any arbitrary starting inventory level that we use as a point of reference and its corresponding optimal constrained order-up-to level, the same order-up-to level is also optimal for all starting inventory levels between the referenced starting inventory level and its corresponding order-up-to level. Note that it may not be optimal to extend the capacity reservation for some (or all) of those intermediate starting inventory levels, but, if the firm chooses to extend the capacity reservation, this corollary implies that the optimal order-up-to policy across all starting inventory levels will be a set of order-up-to levels, each of which is optimal over a compact interval of starting inventory levels. This finding informs our heuristic presented in Chapter 6, in which we attempt to find near-optimal candidate order-up-to levels prior to solving the dynamic program.

We now present the proofs of Theorem 5.3 and Corollary 5.3.1.

Proof. Without loss of generality, assume the starting inventory level in period n is $\hat{I}_n$. Let $S_n^*(\hat{I}_n)$ be the order-up-to level S_n that minimizes $C_n(S_n, y_n = 1|\hat{I}_n, \alpha_n = 1)$ subject to $S_n \geq I_n$. As we have mentioned, it is possible there are multiple optima; therefore, let $\hat{S}_n$ be the *smallest* value of $S_n^*(\hat{I}_n)$.

Consider any starting inventory level $\tilde{I}_n$ such that $\tilde{I}_n \geq \hat{I}_n$ with an $\tilde{S}_n$ defined analogously to $\hat{S}_n$. We will now prove that $\tilde{S}_n \geq \hat{S}_n$ for any $\tilde{I}_n$ such that $\tilde{I}_n \geq \hat{I}_n$.

Recall that the order-up-to levels are *constrained*, and therefore we know that $\hat{S}_n \geq \hat{I}_n$ and $\tilde{S}_n \geq \tilde{I}_n$. This fact, combined with our requirement that $\tilde{I}_n \geq \hat{I}_n$, means it is sufficient to consider two mutually exclusive and collectively exhaustive cases:

Case 1: $\tilde{I}_n > \hat{S}_n$. The first case is straightforward. In this case, the fact that the feasible range of values for S_n is restricted to $S_n \geq I_n$ is enough to guarantee that $\tilde{S}_n \geq \tilde{I}_n > \hat{S}_n$ and therefore $\tilde{S}_n > \hat{S}_n$.

Case 2: $\hat{S}_n \geq \tilde{I}_n$. The second case is more complicated. In this case, we will not only prove that the optimal order-up-to levels are monotonically non-decreasing, but we will also prove that they are identical. As a result, the proof for *Case 2* also establishes Corollary 5.3.1. Notice that in the definition of $C_n(S_n, y_n = 1|I_n, \alpha_n = 1)$ in (5.7) only the second term depends on I_n. Therefore, for a given (I_n, S_n), it is easy to see that the following is true for any I such that $S_n \geq I \geq I_n$:

$$C_n(S_n, y_n = 1|I_n, \alpha_n = 1) = C_n(S_n, y_n = 1|I, \alpha_n = 1) + c_p(I - I_n) \tag{5.10}$$

By the optimality of $\hat{S}_n$ for $\hat{I}_n$, we know that the value function at $\hat{S}_n$ must be less than or equal to the value function at any other feasible value of S_n, and therefore for any $S_n \geq \hat{I}_n$ we have:

$$C_n(\hat{S}_n, y_n = 1|\hat{I}_n, \alpha_n = 1) \leq C_n(S_n, y_n = 1|\hat{I}_n, \alpha_n = 1) \tag{5.11}$$

Now consider the restricted range $S_n \geq \tilde{I}_n$. Because $\hat{S}_n \geq \tilde{I}_n$ and $\tilde{I}_n \geq \hat{I}_n$, we know the above inequality must also hold for all $S_n \geq \tilde{I}_n$ (i.e., the restriction of the range does not make $\hat{S}_n$ infeasible). Using (5.10) with $\tilde{I}_n$ as our value of I, we expand the above inequality for any $S_n \geq \tilde{I}_n$ into:

$$C_n(\hat{S}_n, y_n = 1|\tilde{I}_n, \alpha_n = 1) + c_p(\tilde{I}_n - \hat{I}_n) \leq C_n(S_n, y_n = 1|\tilde{I}_n, \alpha_n = 1) + c_p(\tilde{I}_n - \hat{I}_n) \tag{5.12}$$

which simplifies to:

$$C_n(\hat{S}_n, y_n = 1|\tilde{I}_n, \alpha_n = 1) \leq C_n(S_n, y_n = 1|\tilde{I}_n, \alpha_n = 1) \tag{5.13}$$

Because (5.13) applies to any $S_n \geq \tilde{I}_n$, this means that $\hat{S}_n$ is the value of S_n that minimizes $C_n(S_n, y_n = 1|\tilde{I}_n, \alpha_n = 1)$ subject to $S_n \geq \tilde{I}_n$. Therefore, $\hat{S}_n$ is also the optimal order-up-to level for the starting inventory level $\tilde{I}_n$, and, as a result, we have $\tilde{S}_n = \hat{S}_n$, as desired.

Together, these two cases prove that $\tilde{S}_n \geq \hat{S}_n$ for any $\tilde{I}_n \geq \hat{I}_n$, and therefore $S^*_n(I_n)$ is monotonically non-decreasing in I_n. Note that this result required that we break potential ties in case of multiple optimal order-up-to levels by selecting the *smallest* optimal order-up-to level, as mentioned in Theorem 5.3.

We have also proved the result of Corollary 5.3.1 in this proof. Note that Corollary 5.3.1 applies to all starting inventory levels I such that $I_n \leq \tilde{I}_n \leq S^*_n(I_n)$, which is equivalent to *Case 2* in this proof. □

Corollary 5.3.1 enables us to better understand the optimal order-up-to levels plotted in Figure 5.1. For the starting inventory interval from 0 to 85, the optimal order-up-to level is 85 units, just as we would expect. Similarly, the order-up-to level of 390 is optimal over a starting inventory interval of 100 to 390. For the starting inventory levels between 85 and 100, it is optimal to not order any units. This is a special case in which $S_n^*(I_n) = I_n$ and, therefore, also satisfies Corollary 5.3.1. The key result is that any optimal order-up-to level that results in a positive order quantity (i.e., any $S_n^*(I_n) > I_n$) also defines the upper threshold of a starting inventory interval for which it is an optimal order-up-to level (i.e., $S_n^*(I_n)$ is an optimal S_n for all starting inventory levels from I_n to $S_n^*(I_n)$). We will further characterize these intervals and order-up-to levels in our numerical study presented in Chapter 7.

Chapter 6

Heuristic Method for the Dynamic Variant

While examining solutions for the problem in Chapter 5 for moderately-sized problem instances with discretized demand and inventory levels, we noticed that the structure of the optimal solutions often follows a pattern. Specifically, given that the firm has chosen to extend the capacity reservation, the optimal order-up-to level is often one of the following values that depend upon the firm's expected future behavior: (i) if the firm expects to extend the capacity reservation for at least one more period but is low on inventory, it will choose a "small" order-up-to level; (ii) if it expects to extend the capacity reservation for at least one more period but has sufficient inventory for the current period, it will place an order for zero units but extend the capacity reservation, thereby preserving its options in the future; and (iii) if it does not expect to extend the capacity reservation in the future, it will order-up-to a "large" level that it expects to suffice for the remainder of the horizon.

Using these observations, we develop a heuristic policy. Rather than consider all possible order-up-to levels, the firm considers only the three strategies described above along with the option of not extending the reservation. More specifically, the firm can order nothing and not extend the capacity reservation, it can order nothing but extend the capacity reservation, it can order up to a *continuing-buy* order-up-to level, which we denote as S_n^c, or it can order up to a *last-time-buy* order-up-to level, which we denote as S_n^ℓ. All other aspects of the problem are unchanged; we are only restricting the set of possible order-up-to levels. This heuristic can be applied to the dynamic variant of the problem defined in Chapter 5 or (with modification) to the variant with a buy-back defined in Chapter 8. However, for the remainder of this chapter, we will focus on the dynamic variant as defined in Chapter 5.

In our heuristic policy, the value function in any period n when the capacity reservation can be extended (originally defined in equation (5.8)) becomes:

$$C_n^*(I_n, \alpha_n = 1) = \min_{y_n \in \{0,1\}, I_n \le S_n \le I_n + My_n, S_n \in \{I_n, S_n^c, S_n^\ell\}} C_n(S_n, y_n | I_n, \alpha_n = 1) \quad (6.1)$$

where the constraint has been modified to restrict the S_n variable to one of the three heuristic options. In a numerical analysis described in Chapter 7, we show that the heuristic achieves good results when the values of S_n^c and S_n^ℓ are chosen appropriately. A key simplification in the heuristic is that both of these values depend only on the period and not on the starting inventory level, i.e., they are state-independent. Therefore, the computational effort required to find these candidate order-up-to levels is much less than it would be to find the state-dependent optimal solution by solving the full dynamic program.

6.1 Calculating the *Continuing-Buy* Order-Up-To Level

The *continuing-buy* order-up-to level represents the target amount of inventory when the firm plans to extend the capacity reservation for at least one more period after the current period. As such, the inventory level can be determined by a newsvendor-type analysis that trades off the cost of holding additional inventory against the cost of shortages. In doing so, we make a few simplifying assumptions.

The newsvendor shortage cost in our model is simply $c_s - c_p$, that is, the shortage penalty from a lost sale minus the avoided variable procurement cost. This accurately reflects the true cost incurred by the firm due to the lost sales assumption.

The newsvendor overage cost is more complicated. In situations with stationary demand and an infinite horizon, after a possible transient period in which there is excess inventory, it will never be the case that leftover inventory at the end of one period exceeds the optimal starting inventory in the next period. In such a case, one does not have to account for any consequent inventory-related costs and therefore the overage cost can be set to the single period inventory holding cost c_h. However, in other instances, the starting inventory may exceed the order-up-to levels for one or more periods, in which case the overage cost for an incremental unit will include a per-period inventory holding cost of c_h in each period until it is sold, and if it is not sold until the end of the horizon, then also the loss on disposal, c_p. However, in practical settings, it is unlikely that an excess unit of inventory is held for many periods, partly due to the use of an order-up-to policy that accounts for on-hand inventory. Moreover, a *continuing buy* tends to be the best option early in the horizon when the odds of any given unit of inventory never being used are the lowest, thus leaving only the sum of per-period inventory holding costs, which are small (particularly in comparison to the shortage cost) in practice, as the main component of the overage cost. Although we recognize that it is an underestimate of an expected overage cost that actually varies from period to period, for simplicity, we will use the same overage cost of c_h for all periods.

Given this simplification, we can easily define the *continuing-buy* order-up-to level S_n^c in any period n as:

$$S_n^c = F_n^{-1}\left(\frac{c_s - c_p}{c_s - c_p + c_h}\right) \tag{6.2}$$

where $F_n^{-1}(\cdot)$ is the inverse c.d.f. of demand in period n.

6.2 Calculating the *Last-Time-Buy* Order-Up-To Level

In contrast to the heuristic *continuing-buy* order-up-to level discussed in the last section, our proposed heuristic *last-time-buy* order-up-to level is more complicated to calculate. We can simplify the problem by assuming that the firm will not extend the capacity reservation after ordering up to this large *last-time-buy* order-up-to level. This turns the choice of S_n^ℓ for a given n into a more standard last-time-buy problem. However, due to our lost sales assumption, we cannot directly use the results of any previous paper.

If the firm does not expect to extend the capacity reservation in period $n+1$, then it will choose to purchase up to the *last-time-buy* order-up-to level in period n, and the value function can be written as:

$$C_n(S_n|I_n, \alpha_n = 1) = c_f + c_p(S_n - I_n) + C_n^*(S_n, \alpha_n = 0) \tag{6.3}$$

where $C_n^*(S_n, \alpha_n = 0)$ is the optimal cost-to-go with S_n units of inventory in period n when the firm can no longer extend the capacity reservation. This expands to:

$$\begin{aligned} C_n(S_n|I_n, \alpha_n = 1) = c_f + c_p(S_n - I_n) \\ + c_h \sum_{i=n}^{N-1} \int_0^{S_n} (S_n - x) f_{ni}(x) dx + c_s \int_{S_n}^{\infty} (x - S_n) f_{nN}(x) dx \end{aligned} \tag{6.4}$$

where $f_{ni}(\cdot)$ is the p.d.f. of the *cumulative* demand from period n through period i with $n \leq i \leq N-1$ and $f_{nN}(\cdot)$ is the p.d.f. of the *cumulative* demand from period n through period N. Note that we have used uncensored cumulative demand distributions in the integral expressions in (6.4). Although this would not provide an accurate representation of costs under our assumption of lost sales if there were ongoing replenishments, the expression in (6.4) is for the case of a last time buy in period n with an order-up-to level S_n. In this setting, inventory declines as each demand is observed until all inventory is depleted, if that eventually occurs, say, in period n'. The expressions for inventory holding costs in period n' and all subsequent periods in (6.4) become equal to zero, accurately capturing inventory holding costs after any depletion. Similarly, shortages must be tabulated from period n' onward. The last term accurately accounts for all shortages because the total supply at the beginning of period n is S_n and the density of aggregate demand from period n onward is reflected in f_{nN}.

After differentiating (6.4) with respect to S_n we have:

$$\frac{\partial}{\partial S_n} C_n(S_n|I_n, \alpha_n = 1) = c_p + c_h \sum_{i=n}^{N-1} F_{ni}(S_n) - c_s(1 - F_{nN}(S_n)) \tag{6.5}$$

where $F_{ni}(\cdot)$ is the c.d.f. of the *cumulative* demand from period n through period i with $n \leq i \leq N-1$ and $F_{nN}(\cdot)$ is the c.d.f. of the *cumulative* demand from period n through period N. After differentiating again with respect to S_n we have:

$$\frac{\partial^2}{\partial S_n^2} C_n(S_n | I_n, \alpha_n = 1) = c_h \sum_{i=n}^{N-1} f_{ni}(S_n) + c_s f_{nN}(S_n) \tag{6.6}$$

Clearly, the above expression is non-negative, and therefore the value function is convex in S_n. Therefore, we can use the first-order necessary condition to find the optimal value of S_n. Setting expression (6.5) equal to zero and rearranging terms yields:

$$c_p + c_h \sum_{i=n}^{N-1} F_{ni}(S_n) + c_s F_{nN}(S_n) = c_s \tag{6.7}$$

Given the convexity of the value function, we can utilize a bisection search to find the optimal S_n in period n, which we use as our heuristic *last-time-buy* order-up-to level, S_n^ℓ. We should also note that the above expression does not depend on I_n in any way, so the same candidate solution can be used for any starting inventory value in period n.

6.3 Implementation of the Heuristic

Given the candidate order-up-to values S_n^c and S_n^ℓ, we can now use the heuristic to find an approximate solution to the problem. We do so by recursively solving the simplified dynamic program using the heuristic value function defined for all periods $1 \leq n \leq N$ and starting inventory levels I_n:

$$C_n^*(I_n, \alpha_n = 1) = \min_{y_n \in \{0,1\}, I_n \leq S_n \leq I_n + M y_n, S_n \in \{I_n, S_n^c, S_n^\ell\}} C_n(S_n, y_n | I_n, \alpha_n = 1) \tag{6.8}$$

where $C_n(S_n, y_n | I_n, \alpha_n = 1)$ is defined as in equation (5.7). In practice, it is necessary to discretize the demand and inventory levels to a consistent level of granularity in order to make the problem tractable (as is also the case with the full dynamic program). However, because this heuristic involves only two state-independent candidate order-up-to levels, the computing time to find the approximate solution is much less than the time required to solve the full dynamic program for all possible state-dependent order-up-to levels.

Chapter 7

Numerical Study

In the previous three chapters, we presented the pre-commit and dynamic variants of our problem, and a heuristic for solving the dynamic variant. Due to the non-convexity of the various DP value functions, the optimal strategy for each problem instance must be determined numerically. (We use the term *strategy* to mean the full specification of the state-dependent solution for the DP.) In this chapter, we present the results of a comprehensive numerical study, the purpose of which is twofold: (1) to explore characteristics of the optimal strategies, including the form of the optimal strategy as a function of the starting inventory level and time period, (2) to compare the performance of the heuristic for the dynamic variant against the optimal solutions for both the pre-commit and dynamic variants, and to understand how the performance gaps are affected by problem parameters.

7.1 Description of Problem Parameters

In our numerical study, we use a 12-period horizon, which is long enough to observe differences in the timing of the last time buy if demand and cost parameters are chosen suitably.

We use five different demand patterns, where their descriptors characterize the pattern of the mean demand over time. All have the same total demand during the 12-period horizon but they differ in the shapes of their trajectories. The *Flat* demand pattern has a constant mean and provides a baseline for comparison. We have two declining demand patterns, one with *Linear Decline* and another with *Exponential Decline*; these are examples of what might occur near the end of a product's warranty horizon following the end of retail sales. Our fourth demand pattern, *Single Peak* has one peak and reflects a situation in which new product adoption is gradual and then wanes, or in the case of warranty demand, it progressively increases as the installed base increases and then progressively decreases as units pass their warranty horizon and/or consumers stop using their devices because they have moved on to a newer product. Our fifth demand pattern, *Double Peak*, has two peaks. In the case of new product demand, the pattern would parallel that of a product with two seasonal peaks, such as one associated with the holiday buying season and another associated with the beginning of a sports season. In the case of warranty demand, one peak could be the result of a spike in demand at product launch combined with infant mortality as is common with electronic products, and the later peak could be due to built-in obsolescence. We report the vector of mean demands (by period) for these demand patterns in Table 7.1 and show them graphically in Figure 7.1. We report the standard deviation of demand (by period) in Table 7.2 for the five demand patterns, and throughout the numerical study, we assume that demands are Normally distributed and statistically independent across periods. We selected the standard deviations for each period to be the largest possible without creating negative demand values in the demand discretization process that we describe in Section 7.2. These large standard deviations allow us to examine challenging, high-variance scenarios.

Table 7.1: Demand Patterns - Expected Values by Period

	Period											
Demand Pattern	1	2	3	4	5	6	7	8	9	10	11	12
Flat	100	100	100	100	100	100	100	100	100	100	100	100
Linear Decline	155	145	135	125	115	105	95	85	75	65	55	45
Exponential Decline	275	205	155	115	85	65	50	50	50	50	50	50
Single Peak	50	60	80	100	125	185	185	125	100	80	60	50
Double Peak	50	75	100	300	75	50	75	200	100	75	50	50

Table 7.2: Demand Patterns - Standard Deviations by Period

	Period											
Demand Pattern	1	2	3	4	5	6	7	8	9	10	11	12
Flat	20	20	20	20	20	20	20	20	20	20	20	20
Linear Decline	30	30	30	30	20	20	20	20	10	10	10	10
Exponential Decline	60	50	30	20	20	10	10	10	10	10	10	10
Single Peak	10	10	20	20	30	40	40	30	20	20	10	10
Double Peak	10	10	20	70	10	10	10	50	20	10	10	10

We chose cost parameters that collectively enable us to generate a set of problem parameters with a range of anticipated last-time-buy periods using the *Flat* demand pattern. With some back-of-the-envelope calculations detailed in Appendix B, we were able to determine that with a mean demand of 100 per period, holding the unit procurement cost fixed at a (normalized) 100 and the per unit per period inventory holding cost fixed at 12, capacity reservation costs of 2000, 4000, 8000 and 16,000 would lead to the last time buy being placed in periods 11, 9, 6, and 1, respectively. Recall that each additional period of demand covered by the last time buy must be held for a progressively longer duration, and thus the resulting savings in the capacity reservation costs must be large enough to compensate for that inventory cost increase; otherwise the firm will extend the capacity reservation.

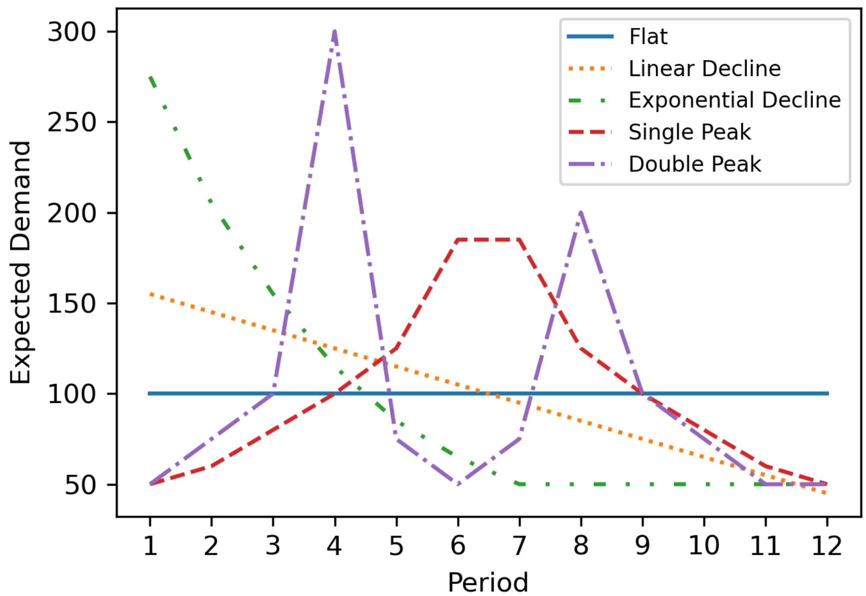

Figure 7.1: Expected Demand by Period for Each Demand Pattern

We chose shortage costs of 200, 400, 600, 800, and 1000, which correspond to service levels (newsvendor critical fractiles) of 50%, 75%, 83%, 88%, and 90%, respectively, in the last period of the horizon when leftover units would need to be disposed with zero salvage value. In earlier periods, leftover units would not need to be disposed of and would only incur the usual inventory holding cost, and for these periods, the service levels corresponding to the five shortage costs turn out to be 89.3%, 96.3%, 97.7%, 98.3%, and 98.7%, respectively. Both sets of service levels are within typical ranges utilized in practice.

In Table 7.3 we summarize the parameters for our numerical study. Additionally, we note that we vary the same cost parameters as those selected by Ozyoruk et al. [37] in their numerical study, although we also note that their model includes a few other costs (e.g., salvage costs) that are not considered in our model, although it would be easy to incorporate such costs. Between the five demand patterns, five shortage costs, and four capacity reservation fees, we have a total of one hundred unique parameter combinations.

7.2 Implementation Details

For computational tractability, we discretized the Normal distributions (for demands) by defining a range of possible demand values for each period in a manner inspired by Galloway [21]. In this method, the discretized random variable, X, has M possible demand values, $x_1, x_2, \ldots, x_M$ and the components of the probability density function $P(X = x_i), i =$

Table 7.3: Numerical Study - All Parameter Values

Parameters	Set of Values
N number of periods	12
Total expected demand	1200
c_p per-unit procurement cost	100
c_h per-unit per-period inventory holding cost	12
c_s per-unit shortage cost	$\{200, 400, 600, 800, 1000\}$
c_f per-period capacity reservation cost	$\{2000, 4000, 8000, 16000\}$
Demand Patterns	*5 options, see Tables 7.1 and 7.2*

$1, \ldots, M$ are defined by

$$P(X = x_i) = \frac{f(x_i)}{C} \tag{7.1}$$

where $f(\cdot)$ is the probability density function of a Normal distribution with mean μ and standard deviation σ, and

$$C = \sum_{i=1}^{M} f(x_i) \tag{7.2}$$

In our implementation, we set $x_i = \mu - 4\sigma + 0.5(i-1)\sigma$, $i = 1, \ldots, M$ with $M = 17$, i.e., values from four standard deviations below the mean to four standard deviations above the mean in increments of one-half of a standard deviation. Table 7.4 shows the discretization based on the standard Normal. Clearly, the mean is exactly zero and the standard deviation is also quite accurate, with a value of 0.9998. This method of discretization differs from the traditional one in which the Normal cumulative density is partitioned using equal-interval segments, but it approximates the variance of the distribution more accurately. For example, in the case of our chosen anchor points, the traditional discretization would lead to a standard deviation of 1.0103.

As a consequence of our discretization of the Normal distribution and our choices of the (time-dependent) means and variances of the distributions in our numerical study, all of the pertinent demand quantities are multiples of 5. Because we are interested in solving the problems starting with zero initial inventory, it is sufficient to consider order-up-to levels that are multiples of 5, and our solution algorithms are implemented accordingly.

We implemented our optimal and heuristic solution algorithms in Python™ and performed the computations on a desktop computer with an AMD Ryzen™ 7 3800X processor running at 3.9 GHz. CPU times are approximately 25 minutes to find the optimal dynamic solution and approximately 161 minutes to find the optimal pre-commit solution due to the need to search over the full range of possible commitment durations and find the conditionally optimal strategy for each. Our heuristic for the dynamic policy, on the other hand, required only about 6 seconds of computing time. For the full results of the numerical analysis for all 100 problem instances, see Tables B.1 through B.5 in Appendix B.

Table 7.4: Discrete Approximation of the Standard Normal

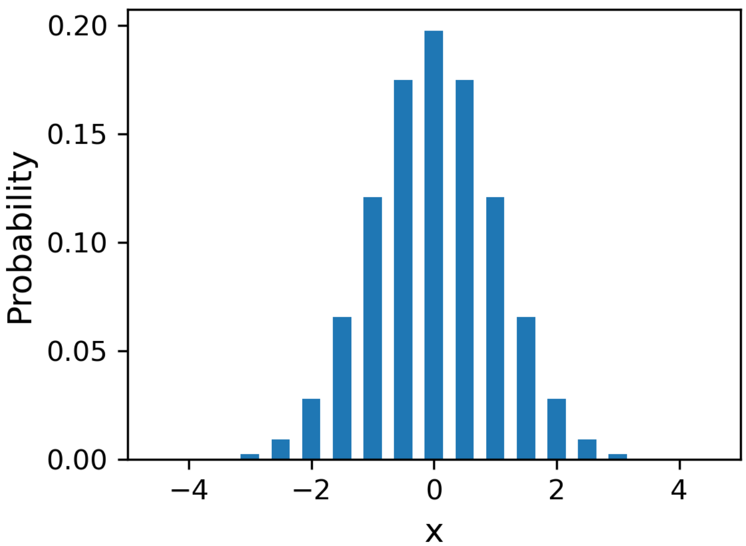

(a) Probability Mass Function

(b) PMF Values

Values	Probability
-4	0.0001
-3.5	0.0004
-3	0.0022
-2.5	0.0088
-2	0.0270
-1.5	0.0648
-1	0.1210
-0.5	0.1760
0	0.1995
0.5	0.1760
1	0.1210
1.5	0.0648
2	0.0270
2.5	0.0088
3	0.0022
3.5	0.0004
4	0.0001
μ	0
σ	0.9998

7.3 Characteristics of the Optimal Strategy

We are interested in the characteristics of the optimal strategy for the dynamic variant of our problem, as well as the performance of the heuristic relative to that of the dynamic variant, and the penalty a firm would incur due to the loss of flexibility from using the pre-commit policy. We focus on the characteristics of the optimal strategy for the dynamic variant for the remainder of this section and discuss the other performance differences in the subsequent section.

In Corollary 5.3.1, we showed that, if the firm chooses to extend the capacity reservation, the optimal order-up-to strategy across all starting inventory levels will be a set of order-up-to levels, each of which is optimal over a compact interval of starting inventory levels. In this numerical study, we show that, in practice, this property of the optimal strategy often leads to two easily identifiable order-up-to levels. Borrowing terminology from the heuristic presented in Chapter 6, we refer to these order-up-to levels as the *continuing-buy* and the *last-time-buy* order-up-to levels.

The *continuing-buy* order-up-to level is optimal when the firm expects to extend the capacity reservation for at least one more period. This situation normally occurs when the starting inventory level is below what is needed to service demand in the current period. On the other hand, the *last-time-buy* order-up-to level is optimal when the firm does not expect to extend the capacity reservation again. Note that, in the dynamic variant, it is possible that, after placing what was expected to be (with high probability) the *last time buy*, the firm may still choose to extend the capacity reservation in the next period. In our numerical study, this only occurs when demand is much larger than expected in the period in which the *last time buy* was just placed. Thus, it is more accurate to refer to the *last time buy* as the ***anticipated*** *last time buy*, but for brevity we will continue to refer to it as simply the *last time buy.* The same could be said for the *continuing buy*, although because the firm maintains the option to extend the capacity reservation in the next period (even if it may not necessarily choose to exercise it), we do not believe this nomenclature requires the same clarification.

Structure of Optimal Strategies within a Single Period

We begin this subsection by briefly describing the most common form of optimal strategy that we observed in the numerical results to provide a backdrop for discussing the more complex optimal strategies that we observed. We emphasize that due to our choices of the demand parameters, our discretization of the pertinent Normal densities, and our assumption of zero inventory at the beginning of the horizon, all possible optimal order-up-to levels are multiples of 5, and in our solution procedures, we either explicitly or implicitly consider all multiples of 5. As such, all of our calculations for the dynamic variant of the problem are exact, up to the level of precision of numerical representations in the computer. Thus, the strategies that we discuss here are optimal, not the result of numerical imprecision or rounding issues.

In the most commonly occurring form of optimal strategy, which we discuss in more detail later, the strategy can be characterized by three thresholds of starting inventory and two order-up-to levels, where these values are time-dependent. Letting I_a, I_b, and I_c represent the three thresholds of starting inventory and S_L and S_H represent the low and high order-up-to levels, respectively, and dropping the time-dependence for ease of exposition, the common form of the optimal strategy is: (a) for starting inventory below I_a, continue the capacity reservation and order up to S_L (which we call a *continuing buy* for short); (b) for starting inventory at or above I_a and below I_b, continue the capacity reservation but do not order (which we call *extend only* for short); (c) for starting inventory between at or above I_b and below I_c, continue the reservation and place an anticipated last time buy by ordering up to S_H (which we call *last time buy* for short); and (d) for starting inventory at or above I_c, do not continue the capacity reservation and do not order (which we call *end capacity reservation* for short).

Of the 1200 (100 x 12) combinations of problem instances and periods within the horizon, only 34 exhibited optimal strategies that cannot be characterized by the optimal strategy described above. Below, we describe the two types of complex strategies that together account for all 34 observed instances.

The first complex strategy, observed for 17 problem-instance-period combinations, is only slightly more complicated than the most commonly occurring strategy detailed above. Rather than three thresholds of starting inventory and two order-up-to levels, this strategy has an additional threshold between I_b and I_c. This threshold does not lead to an additional order-up-to level, but it does introduce an *extend only* decision between the *last time buy* and *end capacity reservation* decisions. Letting I_d represent the new threshold, the optimal strategy is: (a) for starting inventory below I_a, a *continuing buy*, as before; (b) for starting inventory at or above I_a and below I_b, *extend only*, as before; (c) for starting inventory at or above I_b and below I_d, a *last time buy*; (d) for starting inventory at or above I_d and below I_c, *extend only*; and (e) for starting inventory at or above I_c, *end capacity reservation*, as before. This implies that, for some problem instances, the net benefit of extending the capacity reservation (and therefore maintaining the option to order in a future period) is positive even when the inventory is larger than the (anticipated) *last-time-buy* order-up-to level. In every problem instance with this complex strategy, the new threshold is very close to the *end-capacity-reservation* threshold (within 25 units, which is 25% of the average demand per period in our problem instances).

The other complex strategy, observed for the other 17 problem-instance-period combinations, deviates from the most commonly occurring strategy by having five or more thresholds of starting inventory and three or more order-up-to levels that result in positive order quantities. Without exception, the additional thresholds are clustered closely together and are near the *end-capacity-reservation* threshold. Furthermore, the optimal order-up-to levels associated with these thresholds are very close to the threshold inventory levels themselves, resulting in small order quantities. In Figure 7.2a we plot the optimal *constrained* order-up-to level as a function of the starting inventory for a problem instance with the *Single Peak* demand pattern, a capacity reservation cost of 4000, and a shortage cost of 400 that has such a complex optimal strategy. Also shown by a dotted vertical line is the threshold of starting inventory below which it is optimal to continue the capacity reservation. The dotted horizontal lines show the various order-up-to levels that result in positive order quantities. In Figure 7.2b we show an enlargement of the plot in the region of starting inventory levels near where the complexities arise.

For this problem instance, the optimal strategy detailed in Figure 7.2 has six starting inventory thresholds, leading to three intervals of starting inventory with associated order-up-to levels that result in positive order quantities and three intervals of starting inventory in which the *extend only* strategy is optimal. In Table 7.5 we provide the optimal strategy for all seven intervals of starting inventory demarcated by the six thresholds. Of particular note are the *extend only* strategies are optimal for two singleton starting inventory intervals (605 and 625). Both of these values correspond to one of the optimal order-up-to levels, so it would be equally correct to view each of these singleton points as the upper limit of a starting inventory interval in which it is optimal to order up to 605 (or 625, respectively).

Other problem instances have optimal strategies with even more thresholds, in some cases leading to up to five optimal order-up-to levels that result in positive order quantities. In all cases, the additional thresholds are close to each other and to the *end-capacity-reservation*

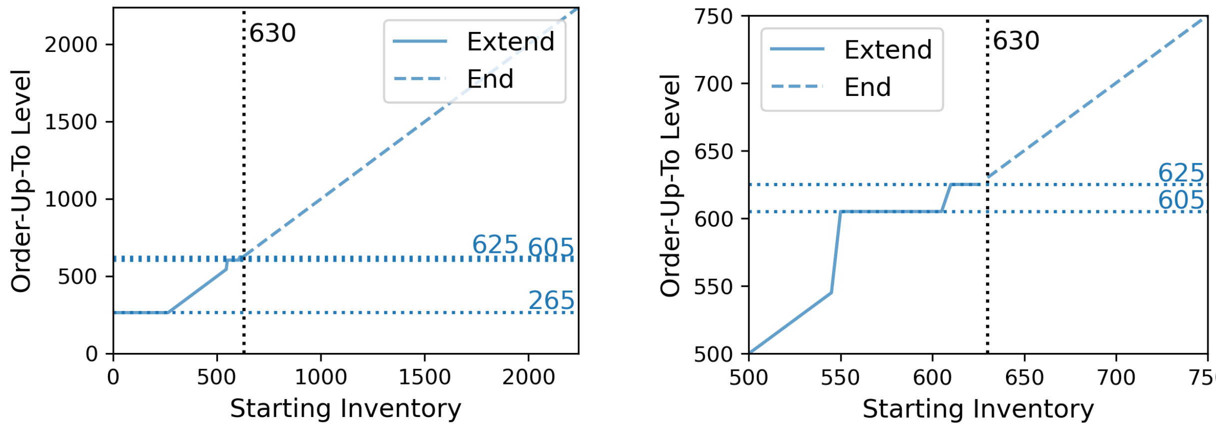

(a) Optimal Order-Up-To Levels in Period 7 (b) Optimal Order-Up-To Levels Detailed

Figure 7.2: An Example of a Complex Strategy - *Single Peak* Demand Pattern with $c_f = 4000$ and $c_s = 1000$

Table 7.5: All Starting Inventory Intervals of the Complex Optimal Strategy for the *Single Peak* Demand Pattern with $c_f = 4000$ and $c_s = 1000$ by Starting Inventory Interval

Starting Inventory	**Optimal Strategy**
$[0, 260]$	Order-Up-To 265
[265, 545]	*Extend Only*
[550, 600]	Order-Up-To 605
605	*Extend Only*
[610, 620]	Order-Up-To 625
625	*Extend Only*
$[630, \infty)$	*End Capacity Reservation*

threshold. Although such cases are rare, the existence of these optimal strategies greatly complicates any attempt to characterize the general form of the optimal policy.

In view of what we have observed in the examples discussed above, we cannot claim that the optimal policy is simple. In particular, it is possible that there are thresholds of starting inventory beyond the three that arise in the commonly-occurring pattern, and generally speaking, for each additional inventory threshold, there is either a different associated order-up-to level or an additional *extend only* interval. Fortunately, in every problem instance in our numerical study with a complex optimal policy, the additional inventory thresholds are close in value, as are the additional order-up-to levels, and the additional *extend only* intervals are small or even singletons. As such, if one is not concerned about implementing a truly optimal policy, it is likely to be sufficient to utilize sensible approximate values for inventory thresholds and order-up-to levels in the event that the starting inventory level happens to fall within an interval with a complex optimal policy. For example, in the problem instance corresponding to Figure 7.2 and Table 7.5, an order-up-to level of 625 could be used for the starting inventory range of $[550, 625]$, and additional costs would be incurred due to a slightly suboptimal order-up-to level only if the starting inventory happens to fall into the given interval. On average across the starting inventory range of $[550, 625]$, the expected cost of the simplified strategy is only 0.03% higher than the cost of the complex optimal strategy. As such, the operational benefits of the simplified strategy far outweigh the expected cost increase.

We now return to the most commonly-occurring form of the optimal strategy, as outlined earlier, and provide some intuition for its structure. In Figure 7.3a, we plot the optimal *constrained* order-up-to level as a function of starting inventory for a problem instance with *Flat* demand, a capacity reservation cost of 8000 and a shortage cost of 400. Also shown by a dotted vertical line is the threshold of starting inventory below which it is optimal to continue the capacity reservation. The dotted horizontal lines show the low and high order-up-to levels associated with the *continuing buy* and *last time buy*, respectively.

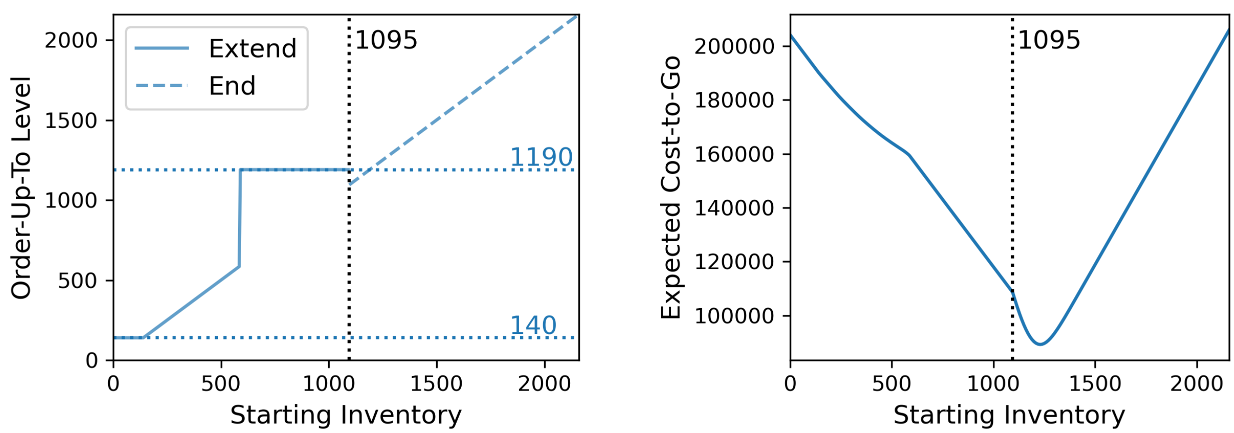

(a) Optimal Order-Up-To Levels in Period 1 (b) Expected Cost-to-Go in Period 1

Figure 7.3: An Example of a Typical Optimal Strategy - *Flat* Demand Pattern with $c_f = 8000$ and $c_s = 400$

The four intervals of starting inventory formed by the thresholds mentioned above are also evident in the expected cost-to-go function, which is shown in Figure 7.3b. In the interval of smallest starting inventory values in which the *continuing buy* is optimal, the function is linearly decreasing because each additional unit of starting inventory decreases the number of units that needs to be purchased. In the next larger interval of starting inventory, the *extend only* policy is optimal, implying a zero order. The expected cost-to-go function is convex decreasing because each additional unit of starting inventory decreases the number to be purchased in the future, but it also generates increasingly greater expected inventory holding costs because of the declining probability that it will be sold in the current period. In the next larger interval of starting inventory, the (anticipated) last time buy policy is optimal, i.e., the firm should order up to the second (higher) order-up-to level. Due to the order-up-to nature of the policy in this region, the expected cost to go is linearly increasing for the same reason as in the *continuing buy* interval. The expected cost-to-go in the *end capacity reservation* interval is convex as shown in Proposition 5.1. The non-convexity of the expected cost-to-go functions obviates the possibility of proving clean theoretical results. For some problem parameters, the expected cost-to-go functions are unimodal in the starting inventory, as in Figure 7.3b, and in most instances, we have found the functions to be sufficiently well behaved that it is not difficult to find optimal solutions numerically. However, for some problem parameters, the expected cost-to-go functions in some periods are not unimodal. As such, some care needs to be taken when solving the DP.

Above, we briefly explained why the expected cost-to-go function takes on a specific shape in each pertinent interval of starting inventory levels. Below, we provide further details on the economic tradeoffs that lead to the specifics of the optimal strategy for the numerical example under consideration.

Although the capacity reservation cost is moderately high in this problem instance, if the starting inventory is low, it is better to extend the reservation and place an order intended to satisfy the current period's demand than to order enough for the entire horizon (recall that Figure 7.3a shows the optimal strategy for period 1). Ordering enough to satisfy demand for a few periods is suboptimal because it would still be necessary to pay the capacity reservation costs if the firm wishes to order again.

For slightly higher values of starting inventory, the inventory is sufficient to provide the desired level of service for the current period but the inventory level is below or well below the desired amount to satisfy demand for the remainder of the horizon. In such a situation, an order in the current period would generate unnecessary inventory holding costs, but the firm almost surely needs the option to place another order and thus chooses to extend the capacity reservation but does not place an order.

For the next larger interval of starting inventory, the inventory is nearly as large as the last-time-buy inventory target would be if one were required to place a last-time-buy in that period. One might think it would be optimal not to place an order and perhaps only extend the capacity reservation. However, the purchase of a modest number of units is optimal if the incremental inventory holding costs are less than the expected savings from reducing the probability of needing to extend the capacity reservation for yet another period. In such situations, it is optimal to place an anticipated last time buy, up to a second (higher) order-up-to level. This higher order-up-to level arises due to the sequential capacity reservation requirement: if the firm were not required to pay the capacity reservation fee in the current period to maintain the option to order later, it would choose not to order now and then order in a later period, if needed. Facing the sequential capacity reservation requirement, the firm can be better off incurring additional inventory holding costs because it eliminates the need to extend the capacity reservation for potentially many periods. Note that we refer to the order as an *anticipated* last time buy because, in the dynamic model, the firm could extend the capacity reservation if an unexpectedly high demand warrants it.

In both this example and in solutions to all of our other problem instances, there is a single inventory threshold above which it is optimal to end the capacity reservation. Although we were not able to prove that this holds in general, we were unable to construct a counterexample. This is not surprising as the observation aligns with intuition: for a sufficiently high starting inventory level, it is not economically sensible to continue the capacity reservation. We note that because of the sequential capacity reservation requirement in our model, we would expect the inventory threshold for ending the capacity reservation to be higher than if the sequential requirement were relaxed.

Another consequence of the sequential capacity requirement is the noticeable discontinuity in the optimal order-up-to levels plotted in Figure7.3a. The optimal *last-time-buy* order-up-to level in the figure is 1190 units, but the end-capacity-reservation threshold is 1095 units, creating a discontinuity in the optimal order-up-to levels at a starting inventory level of 1095. The reason for this result is that for starting inventory levels between 1095 and 1190, it is not worthwhile to pay the capacity reservation fee of 8000 to order 95 or fewer units. Instead, it is less expensive not to order and to incur the expected shortage costs. For starting inventory levels below 1095, however, it is better to order up to 1190. This fundamental tradeoff also arises in standard inventory models with fixed costs of ordering, in which it is common for the optimal policy to include a threshold of on-hand inventory below which an order is placed and above which no order is placed.

The end-capacity-reservation threshold is difficult to characterize in general, as it depends on the cost parameters, period, and expected remaining demand. We offer a few observations on the effects of the cost parameters. We have observed in our numerical study that, holding all else constant, the threshold is smaller for larger capacity reservation fees, as we would expect. Extending the capacity reservation becomes a less attractive option as the fee increases, and, as a result, the magnitude of the discontinuity in the optimal order-up-to levels also becomes larger. On the other hand, as one increases the shortage cost with all other parameters held constant, the end-capacity-reservation threshold increases (i.e., diagrammatically, the dotted vertical line moves to the right), both in absolute terms and as a percentage of the *last-time-buy* order-up-to level. As a result, the magnitude of the discontinuity in the optimal order-up-to levels decreases as the shortage cost increases, as the higher shortage cost is a greater deterrent to starting a period with less than the unconstrained optimal inventory level. However, as the end-capacity-reservation threshold depends on the intersection point of the end-reservation and extend-reservation cost-to-go functions, it is difficult to characterize analytically, even if (as is true for our heuristic) we know the *last-time-buy* order-up-to level.

We now illustrate the observations described above with some results from our numerical study. In Figure 7.4 we plot the optimal order-up-to levels and end-capacity-reservation thresholds in the first period for four different problem instances. These four problem instances all have the *Flat* demand pattern, but the cost parameters vary. The capacity reservation fee is 4000 for the first row and 8000 for the second row, and the shortage cost is 400 for the first column and 600 for the second column. As expected, the end-capacity-reservation threshold decreases by row (i.e., as the capacity reservation fee increases) but increases by column (i.e., as the shortage cost increases). This is only one example, but it is representative of the results of our numerical study overall.

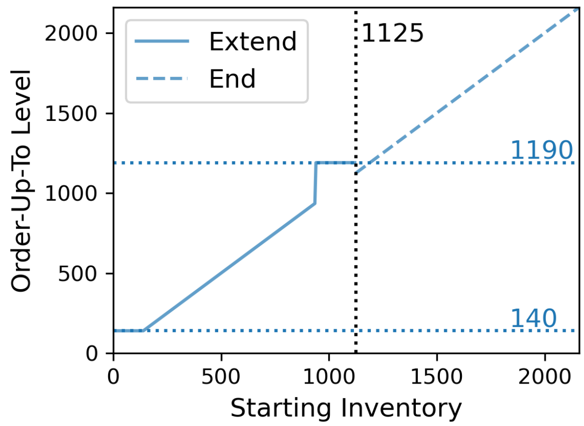

(a) $c_f = 4000$ and $c_s = 400$

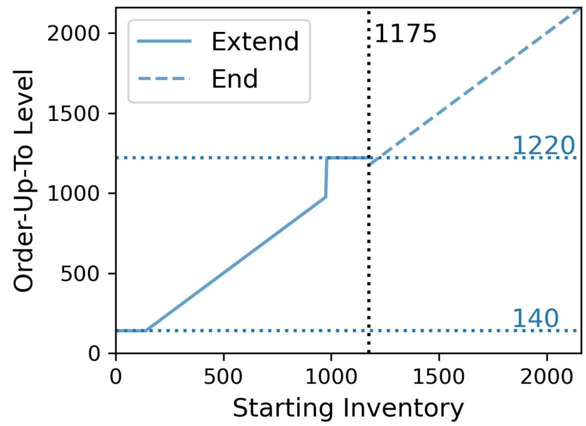

(b) $c_f = 4000$ and $c_s = 600$

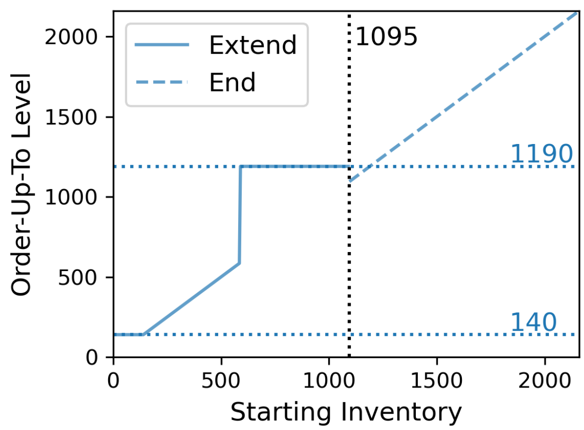

(c) $c_f = 8000$ and $c_s = 400$

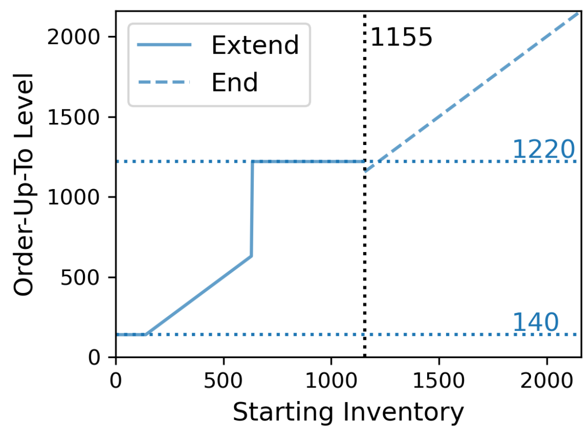

(d) $c_f = 8000$ and $c_s = 600$

Figure 7.4: Optimal Strategy in Period 1 for Various Cost Parameter Combinations - *Flat* Demand Pattern

Changes in the Optimal Strategy Across the Time Horizon

Thus far, we have only considered the optimal strategy in the first period as a function of the starting inventory level. We now briefly examine how the optimal strategy changes from period to period. Returning to the same problem instance used as the basis for Figure 7.3, we now plot the optimal order-up-to level curves for every period in Figure 7.5. However, it can be difficult to discern the various thresholds and optimal strategy types in the twelve-period chart. Therefore, we also provide a more compact graphical representation of the optimal strategy in Figure 7.6, but without the order-up-to levels. Figure 7.6 shows the optimal strategy type (*continuing buy*, *last time buy*, *extend only*, or *end capacity reservation*) for each period and starting inventory level combination. Not surprisingly, as time proceeds, the two choices that imply the need for future extensions of the capacity reservation, i.e., the *continuing buy* or *extend only*, become less viable and are eventually dominated by other choices.

We have observed that some problem instances have optimal strategies whose changes across the time periods are not as simple as those we have seen thus far. Figure 7.7 illustrates the optimal strategy for the problem instance with the *Double Peak* demand pattern, a capacity reservation cost of 4000, and a shortage cost of 200. This example shows that the boundaries between strategy regions are not always monotonically decreasing (as they are in the case of the *Flat* demand pattern in Figure 7.6). This is a result of the trajectory of demand in the *Double Peak* demand pattern. The optimal *continuing-buy* order-up-to level increases over the first four periods as the expected demand also increases in those periods. In fact, the optimal strategies in periods one and two include only the *continuing buy*, *extend only*, and *end capacity reservation* strategies. Due to the highly variable demand in the *Double Peak* demand pattern, it is not in the firm's interest to place a *last time buy* early in the horizon. In period three, the *last time buy* appears in the optimal strategy for the first time, before eventually dominating all other strategies from period eight onward. As a result of the high variance demand pattern, the structure of the optimal policy is complex and changes significantly from one period to the next. However, in any one period, the optimal strategy still consists of the most common strategy described in the previous subsection: at most four intervals of starting inventory, with each corresponding to one of the *continuing buy*, *extend only*, *last time buy*, and *end capacity reservation* strategies.

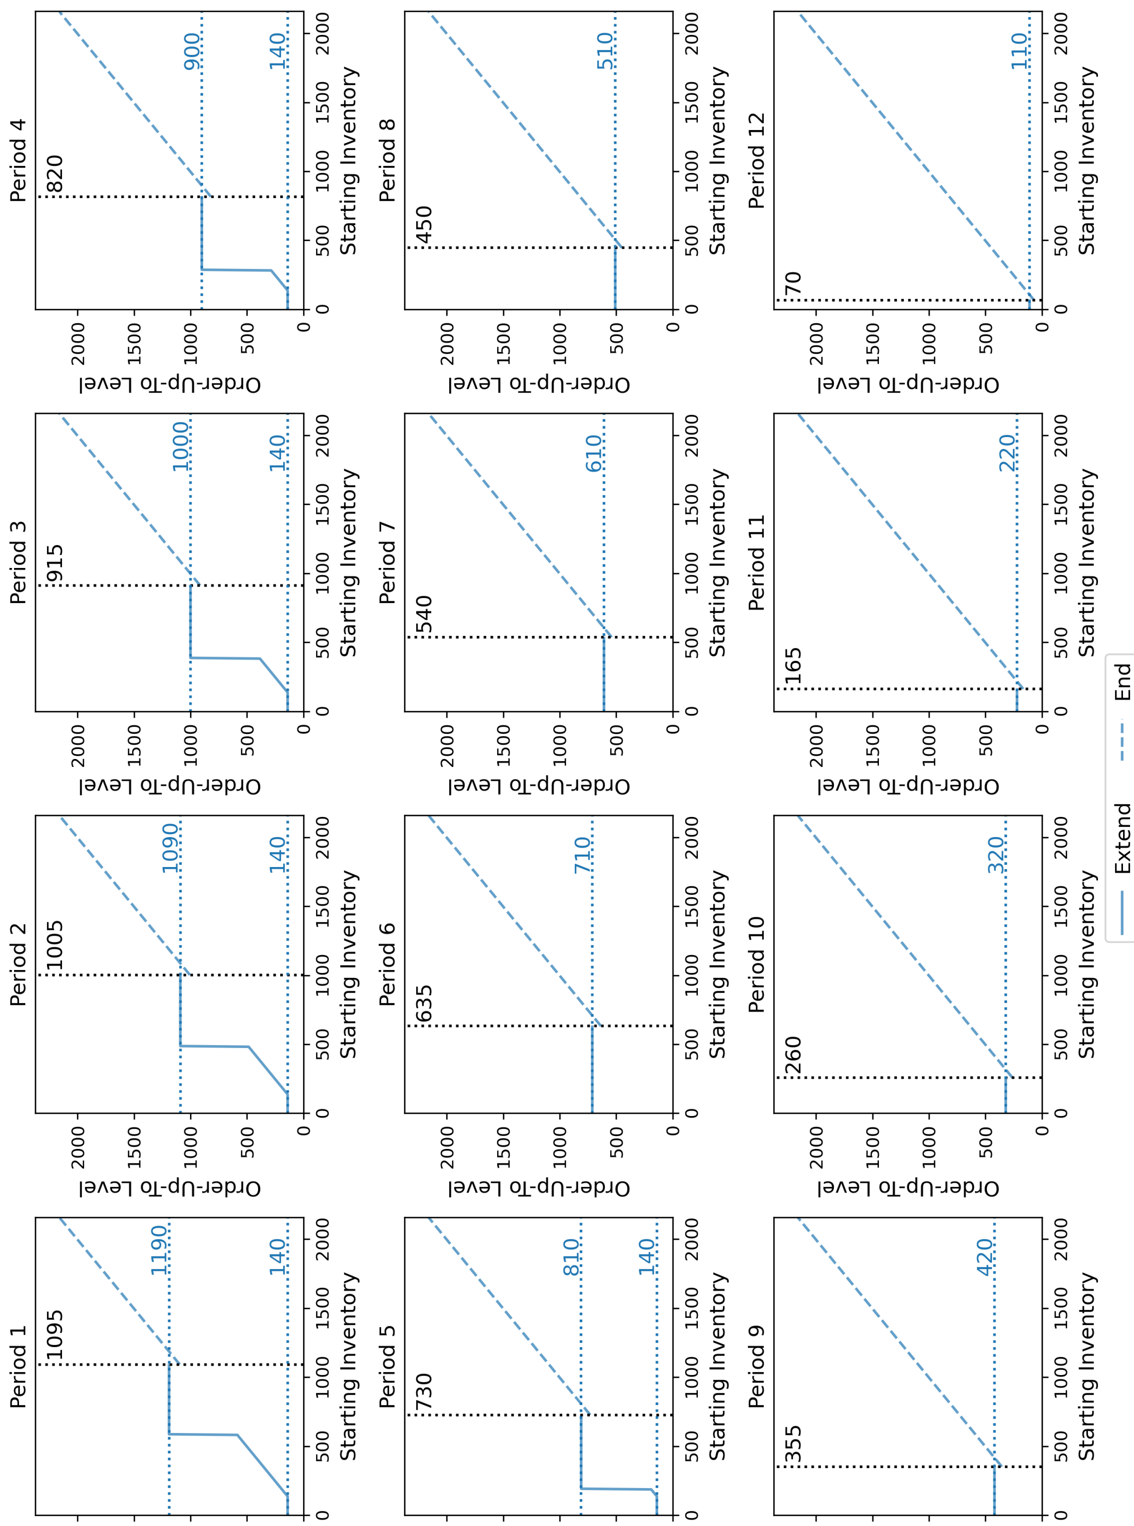

Figure 7.5: An Example of a Typical Optimal Strategy - All 12 Periods - *Flat* Demand Pattern with $c_f = 8000$ and $c_s = 400$

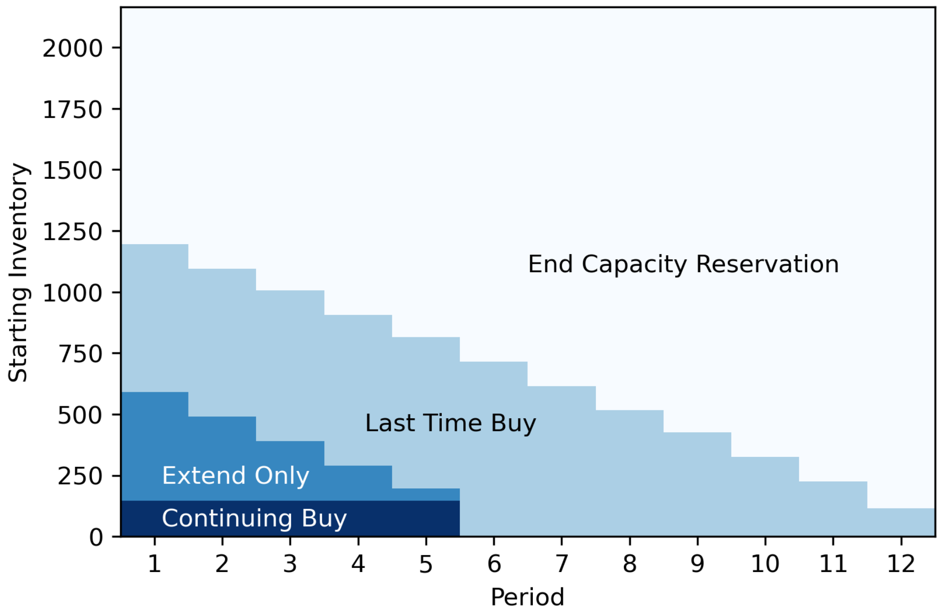

Figure 7.6: Optimal Strategy Boundaries - *Flat* Demand Pattern with $c_f = 8000$ and $c_s = 400$

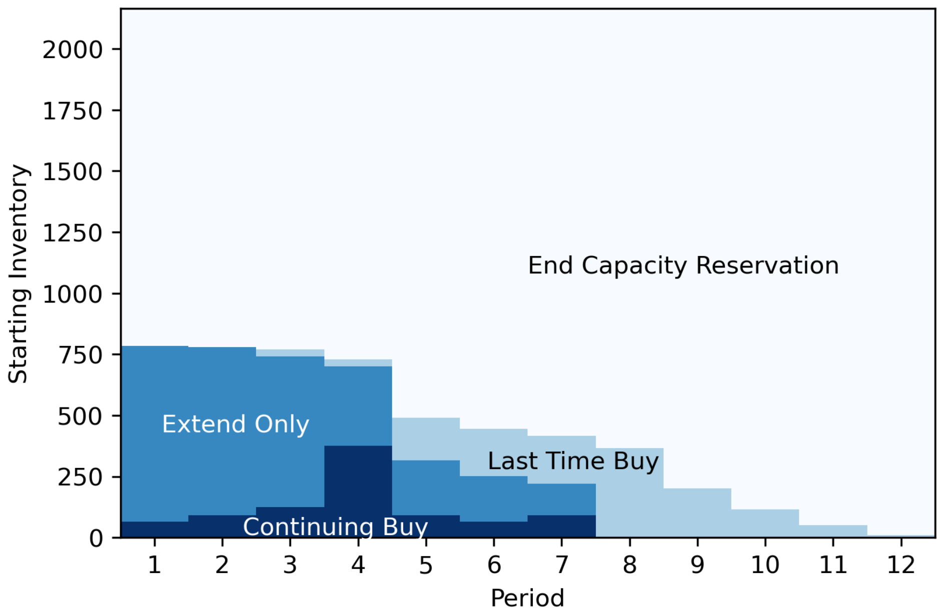

Figure 7.7: An Optimal Strategy with Non-Monotone Boundaries - *Double Peak* Demand Pattern with $c_f = 4000$ and $c_s = 200$

7.4 Performance Comparisons

We report results for our set of 100 problem instances assuming beginning-of-horizon inventory levels of zero. Both the heuristic and the optimal pre-commit solutions were very close to optimal. Table 7.6 presents summary results of optimality gaps, which we calculate in two ways: (i) the percentage gap considering all costs; and (ii) an estimated percentage gap considering controllable costs, where controllable costs do not include the variable costs to cover expected demand. The values in the table are the average, minimum, and maximum optimality gaps across the 100 problem instances.

Table 7.6: Numerical Study - Summary of Results

	Optimality Gap			**Adjusted Optimality Gap**		
Problem Variant	**Avg**	**Min**	**Max**	**Avg**	**Min**	**Max**
Dynamic Variant - Heuristic	0.03%	0.00%	0.27%	0.09%	0.00%	0.93%
Pre-commit Variant - Optimal	0.06%	0.00%	0.44%	0.19%	0.00%	1.01%

From these results, it appears that our heuristic methods for determining the *continuing-buy* and *last-time-buy* order-up-to levels and our approach for making the choice between them in each period, are more than adequate for finding excellent solutions, and require minimal computing times, however, there are a handful of problem instances with optimality gaps between heuristic and optimal dynamic strategies worth investigating. For one such instance, with the *Single Peak* demand pattern, a capacity reservation cost of 8000, and a shortage cost of 1000, the heuristic has an optimality gap of 0.16% and an adjusted optimality gap of 0.39%. In Figures 7.8 and 7.9, we plot for each period the order-up-to level curves for the optimal and heuristic strategies, respectively.

The two strategies have identical *continuing-buy* order-up-to levels and end-capacity-reservation thresholds in each period. However, a careful comparison reveals there are small differences in the *last-time-buy* order-up-to levels in periods five through ten. The optimal strategy is to place *continuing buys* in the first six periods followed by an (anticipated) *last time buy* up to 645 units in period seven, as detailed in Table 7.7. The optimal strategy in period eight is to extend the capacity reservation if the starting inventory level is below 415 units and to end the capacity reservation otherwise. This implies that, if demand in period seven exceeds 230 units (which, for our discretized demand distribution with a mean of 185 and standard deviation of 40, occurs with a 10% probability) the firm should extend the capacity reservation once again in period eight. This is an excellent example of a problem instance in which the *last-time-buy* order-up-to level is actually an ***anticipated*** *last-time-buy* order-up-to level; there is a small but meaningful probability that the firm would extend the capacity reservation in the dynamic problem.

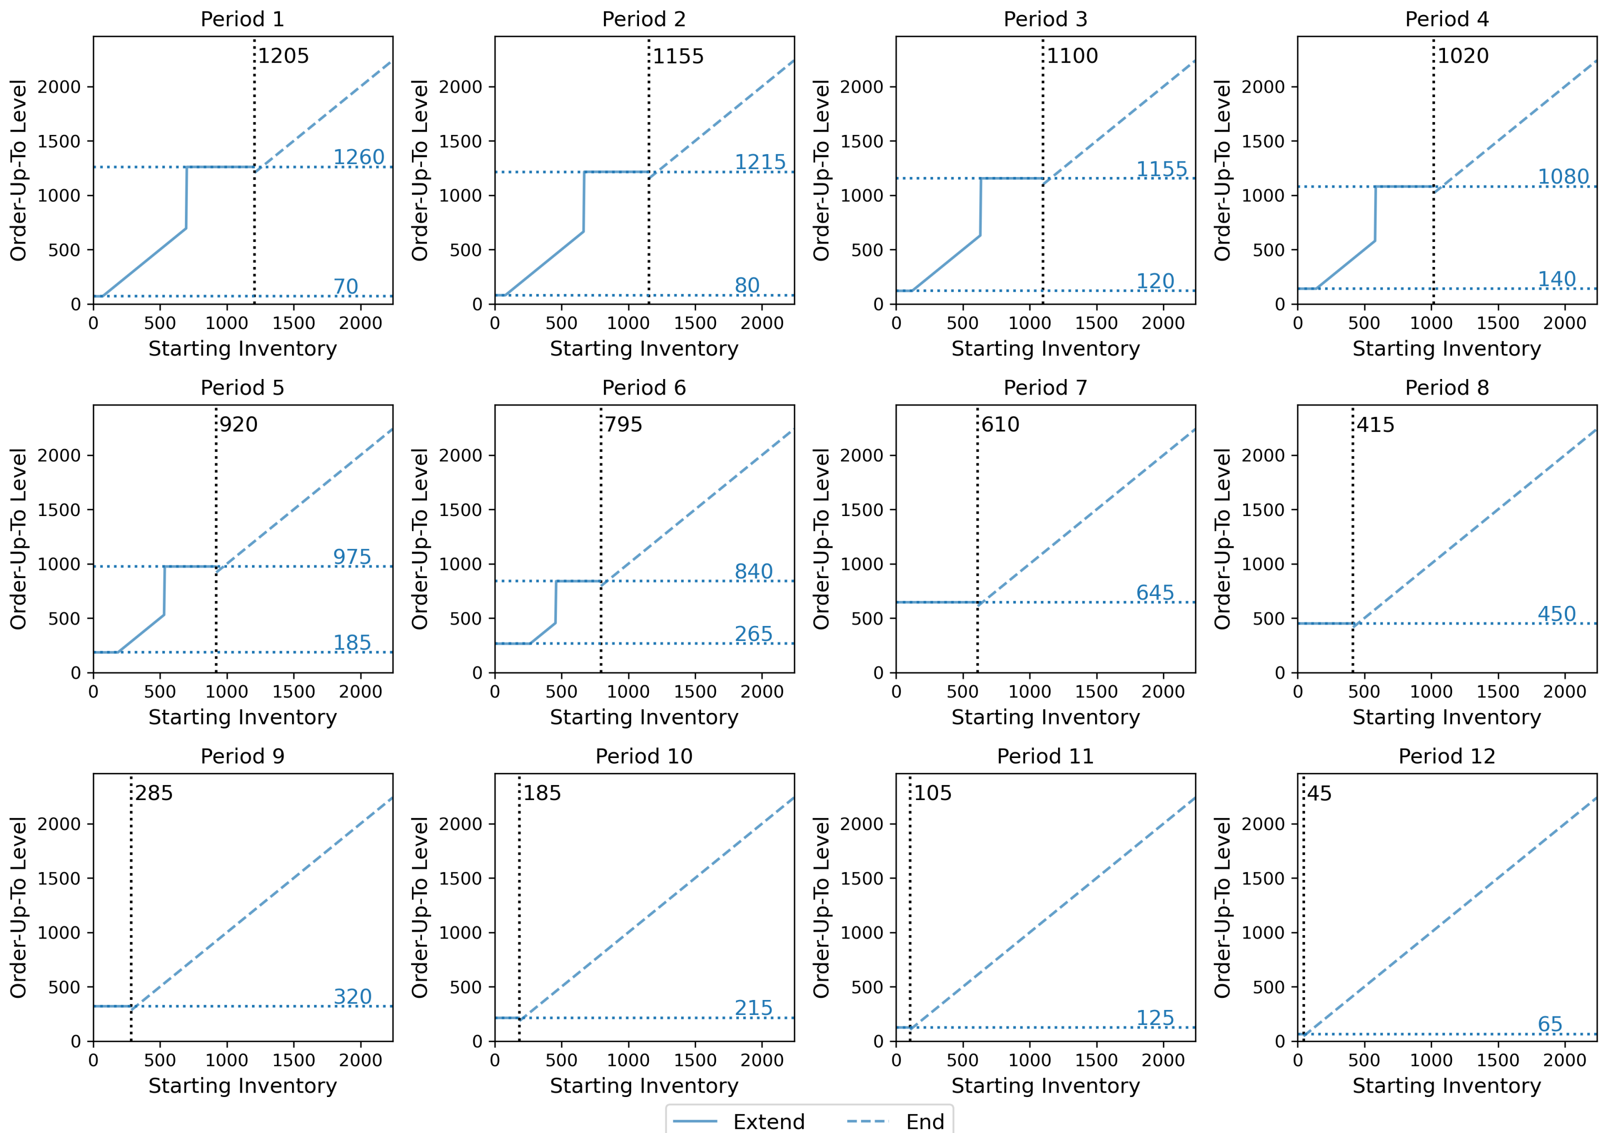

Figure 7.8: Comparison of Optimal and Heuristic Strategies - Optimal Strategy - *Single Peak* Demand Pattern with $c_f = 8000$ and $c_s = 1000$

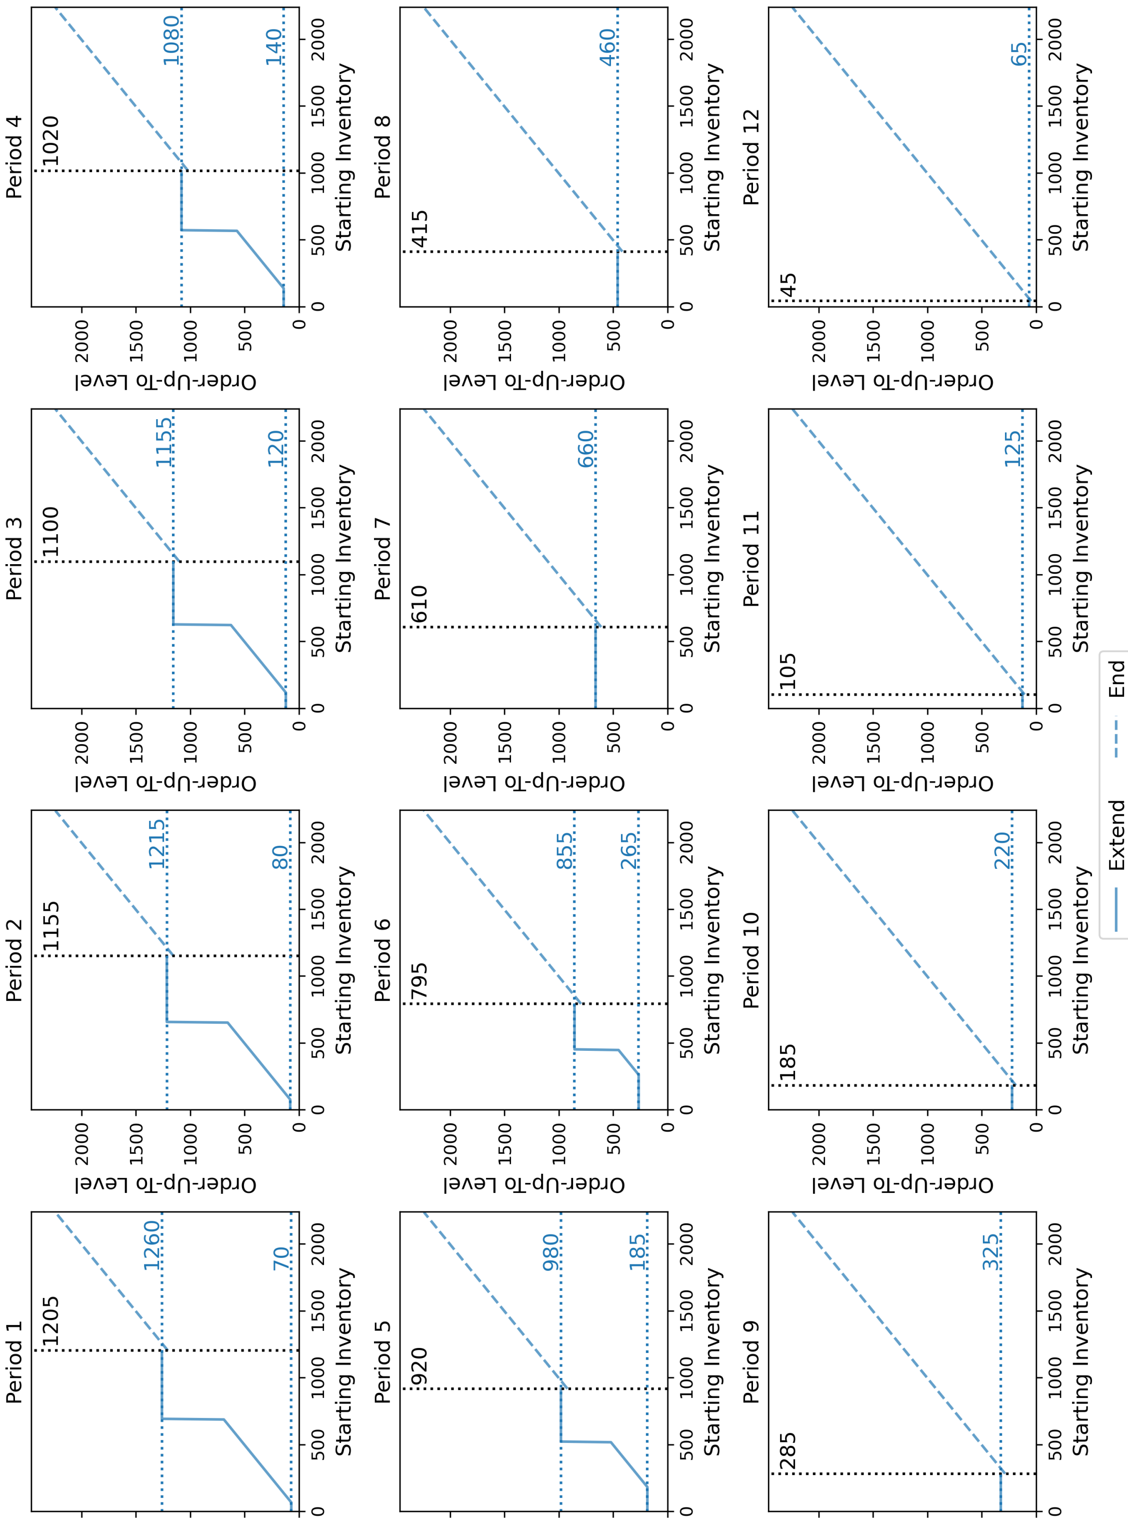

Figure 7.9: Comparison of Optimal and Heuristic Strategies - Heuristic Strategy - *Single Peak* Demand Pattern with $c_f = 8000$ and $c_s = 1000$

As shown in Table 7.7, the heuristic strategy has an identical *continuing buy* strategy through period six. The major differences are that in period seven the heuristic strategy calls for ordering up to 660 units (instead of the optimal 645) and in period eight ordering up to 460 units (instead of the optimal 450). Both of these differences are bolded in Table 7.7. The heuristic strategy thus reflects the assumption that no more orders may be placed after the *last time buy*, whereas the optimal strategy allows an extension of the capacity reservation and an additional order in the event of an unusually large demand immediately after the (anticipated) *last time buy*. The heuristic's resulting larger-than-optimal order-up-to level in period seven accounts for the heuristic optimality gap of 0.16%, which is relatively large among our problem instances.

Table 7.7: Comparison of the Optimal and Heuristic Strategies through Period 8 for Select Starting Inventory Levels for the Problem Instance with the *Single Peak* Demand Pattern, $c_f = 8000$, and $c_s = 1000$. These Starting Inventory Intervals are the Only Relevant Intervals If the Firm Begins Period 1 with No Inventory and Correctly Adheres to Either Strategy.

	Starting	**Order-Up-To Level**		
Period	**Inventory Level**	**Optimal**	**Heuristic**	**Strategy**
1	< 70	70	70	*Continuing Buy*
2	< 80	80	80	*Continuing Buy*
3	< 120	120	120	*Continuing Buy*
4	< 140	140	140	*Continuing Buy*
5	< 185	185	185	*Continuing Buy*
6	< 265	265	265	*Continuing Buy*
7	< 610	645	**660**	*Last Time Buy*
8	< 415	450	**460**	*Last Time Buy*
	≥ 415	n/a	n/a	*End Capacity Reservation*

More surprising is that the optimal pre-commit strategies are also very near optimal. For 35% of problem instances, the pre-commit strategy is the same as the optimal strategy. Such situations are more prevalent when the capacity reservation fee is extremely high (16,000 in our problem set) and neither the shortage costs nor the demand variability is high enough to make an investment in flexibility worthwhile.

We observed that the largest optimality gaps for the pre-commit variant occurred, not surprisingly, in cases with high capacity reservation fees and high shortage costs along with high demand variance early in the horizon. In such cases, one would like to avoid extending the capacity reservation beyond that determined in the pre-commit solution at nearly any cost—but the combination of high shortage costs and high demand variance makes the option to extend the capacity reservation for just one more period beyond that in the pre-commit solution extremely valuable in averting large quantities of high-cost shortages if unexpectedly high demand materializes.

One key reason why the optimal pre-commit strategies perform so well is that the *continuing-buy* order quantities are unconstrained and the system operates with order-up-to levels, so even if an extremely high or low demand observation occurs, the system effectively "regenerates" at the beginning of the next period. As such, the impact of any demand observation on the economic tradeoffs pertaining to the timing of the last time buy is limited. Furthermore, our assumption of lost sales combined with our use of relatively high service levels also dampens the effect of high demands on the timing of the last-time buy in the dynamic variant of the problem. Although the computing times to determine the optimal pre-commit solution are long, thus limiting its implementation in practice, a practical heuristic would entail determining the timing of the last time buy based on a deterministic analysis and using our heuristic policy to determine the order-up-to levels.

Table 7.8: Numerical Study - Summary of Results by Cost Parameter

(a) Shortage Cost

	Optimality Gap	
c_s	**Heuristic**	**Pre-commit**
200	0.01%	0.01%
400	0.02%	0.03%
600	0.03%	0.06%
800	0.04%	0.09%
1000	0.03%	0.09%

(b) Capacity Reservation Cost

	Optimality Gap	
c_f	**Heuristic**	**Pre-commit**
2000	0.04%	0.06%
4000	0.04%	0.07%
8000	0.01%	0.05%
16000	0.01%	0.05%

Tables 7.8a and 7.8b present the average optimality gap across problem instances with the same shortage cost or capacity reservation cost, respectively. As expected, larger optimality gaps are associated with larger shortage costs, especially for the pre-commit variant of the problem because the pre-commit variant provides no flexibility to adjust the inventory after the last time buy. On the other hand, smaller optimality gaps were observed in problem instances with larger capacity reservation costs, especially for the heuristic. This is to be expected, as the heuristic's two order-up-to levels (per period) are calculated by assuming the firm either will or will not continue the capacity reservation in the next period, and higher capacity reservation costs make it less likely the firm will choose to deviate from its plans.

Chapter 8

Dynamic Model with One-Time Buy-Back Option

We now consider a variant of the dynamic version of this problem in which we introduce the option of a one-time buy-back offer. Much of the setting and notation are the same as in the previous model, with the main difference being the addition of the option to make a one-time buy-back offer after the firm ends the capacity reservation.

We model two versions of the buy-back offer. In the first half of this chapter, we consider a version in which the buy-back quantity is a decision variable and the buy-back cost is a deterministic function of the selected quantity. Using this deterministic model, we develop the associated value and cost-to-go functions in Section 8.1, characterize the optimal solution in Section 8.2, modify our heuristic to account for the buy-back option in Section 8.3, and present the results of a numerical study in Section 8.4. Then, in Section 8.5, we introduce a version of the buy-back offer in which the buy-back quantity is a stochastic function of the buy-back price selected by the firm. We conclude with a numerical study of the stochastic buy-back model in Section 8.6.

We begin with the deterministic buy-back option. We represent the total cost to acquire, refurbish, and return B net units to inventory by $r(B)$. We assume $r(\cdot)$ is convex, as this reflects the increasing marginal cost of securing more units via buy-back. We also assume that $r(0) = 0$, as negligible cost and effort will yield essentially zero units. We assume that $r'(\cdot) \geq c_p$, as the customer's valuation of the item is tied to the retail price, which is often double the production cost or more. We also assume that $r'(0) < c_s$, which is necessary for the buy-back to be economically justified. If this assumption did not hold, then the optimal buy-back quantity in every situation would be zero units and the optimal buy-back decision would be trivial. We also assume that there exists some finite B such that $r'(B) > c_s$, which ensures that there is some buy-back quantity above which it is better to incur shortages than to buy back additional units. This is consistent with consumer behavior, as some customers may have high personal valuations for their items and will not be willing to part with them for any reasonable price. To represent the buy-back decision, we introduce a new binary decision variable z_n, which is equal to one if the buy-back is offered in period n, and zero otherwise. We also introduce a new binary state variable β_n, which is equal to one if the buy-back option is still available in the current period, and zero otherwise. All other variables and parameters are unchanged from the dynamic variant of the problem in Chapter 5.

In the remainder of the chapter, we present our formulation, analysis, and a numerical study. The formulation is complicated by the binary variables (α_n and β_n) in the state definition, which affect the set of decisions that can be made. For this reason, we divide our development of the value functions according to the values of these state variables, and for each such combination, we present the value function first for period N, and then for an arbitrary $n < N$. First, we consider the situation in which the capacity reservation has ended and the buy-back option has been exercised. We then consider the situation in which the capacity reservation has ended but the buy-back option is still available. Finally, we consider the setting in which the capacity reservation can still be extended and the buy-back option has yet to be exercised.

8.1 Model Formulation

As in previous chapters, the terminal value function is:

$$C_{N+1}(\cdot) = 0 \tag{8.1}$$

and therefore:

$$C^*_{N+1}(\cdot) = 0 \tag{8.2}$$

because there are no decisions to make at the end of the horizon. However, because the terminal value function does not affect the firm's decisions, we exclude it from the formulation in the remainder of this chapter.

After the capacity reservation has been terminated and the buy-back offer has been made, there are no more decisions to make and therefore the value function only depends on the state variables I_n, α_n, and β_n. We develop the value function by first considering the final period N and then recursively expressing the value function for period $n < N$. The various cost functions are analogous to those in Section 5.1 but now include the additional state variable β_n as well as the decision variables z_n and B. The value function in period N after the capacity reservation has been terminated, the buy-back offer has been made, and the starting inventory level is I_N, is composed entirely of the shortage costs:

$$C_N(I_N, \alpha_N = 0, \beta_N = 0) = c_s \int_{I_N}^{\infty} (x - I_N) f_N(x) dx \tag{8.3}$$

The value function in any arbitrary $n < N$ after the capacity reservation has been terminated, the buy-back offer has been made, and the starting inventory level is I_n, is the sum of the expected inventory holding costs, expected shortage costs, and the expected cost-to-go for the remaining periods:

$$\begin{aligned} C_n(I_n, \alpha_n = 0, \beta_n = 0) =& c_h \int_0^{I_n} (I_n - x) f_n(x) dx + c_s \int_{I_n}^{\infty} (x - I_n) f_n(x) dx \\ &+ \int_0^{I_n} C^*_{n+1}(I_n - x, \alpha_{n+1} = 0, \beta_{n+1} = 0) f_n(x) dx \\ &+ \int_{I_n}^{\infty} C^*_{n+1}(0, \alpha_{n+1} = 0, \beta_{n+1} = 0) f_n(x) dx \end{aligned} \quad (8.4)$$

We once again claim that, because the inventory holding and shortage costs are linear, the value function is also convex in the starting inventory level. We should also note that, because there are no decisions to make in this situation, the optimal expected cost-to-go is equivalent to the value function, and therefore for any period n:

$$C^*_n(I_n, \alpha_n = 0, \beta_n = 0) = C_n(I_n, \alpha_n = 0, \beta_n = 0) \quad (8.5)$$

We will now consider the value function when the firm has not yet made the buy-back offer, the definition of which depends on whether the firm is in the final period N or in any other period. When the buy-back option is available in period N, the firm may choose to buy back B units and pay the corresponding acquisition and refurbishment cost $r(B)$. It should be noted that, because this buy-back offer is a one-time option and the units become available to satisfy demand in the same period, the value function after the buy-back offer has been made is the same as if the buy-back offer had already been made in a previous period, and instead the firm has an additional B units in inventory at the start of period N at an additional cost of $r(B)$. Therefore, we can represent the value function in period N when the firm chooses to buy back B units as:

$$C_N(B, z_N = 1 | I_N, \alpha_N = 0, \beta_N = 1) = r(B) + C^*_N(I_N + B, \alpha_N = 0, \beta_N = 0) \quad (8.6)$$

On the other hand, if the firm chooses to not offer the buy-back in period N, it instead satisfies demand in the final period to the extent possible with the available inventory. The firm only incurs expected shortage costs, as all inventory is disposed at the end of the horizon. Thus, the value function associated with not offering the buy-back in period N can be written as:

$$C_N(z_N = 0 | I_N, \alpha_N = 0, \beta_N = 1) = c_s \int_{I_N}^{\infty} (x - I_N) f_N(x) dx \quad (8.7)$$

The expressions in (8.6) and (8.7) differ in the value assigned to the z_N variable. The firm faces the decision in period N of which option to take. The value function in period N then becomes:

$$\begin{aligned} C_N(B, z_N | I_N, \alpha_N = 0, \beta_N = 1) =& z_N \left[r(B) + C^*_N(I_N + B, \alpha_N = 0, \beta_N = 0) \right] \\ &+ (1 - z_N) \left[c_s \int_{I_N}^{\infty} (x - I_N) f_N(x) dx \right] \end{aligned} \quad (8.8)$$

We can now turn to an arbitrary period $n < N$. By the same argument we used to develop (8.6), we can write the value function in an arbitrary period $n < N$ as a function of the number of buy-back units as follows:

$$C_n(B, z_n = 1 | I_n, \alpha_n = 0, \beta_n = 1) = r(B) + C_n^*(I_n + B, \alpha_n = 0, \beta_n = 0) \tag{8.9}$$

On the other hand, if the firm chooses to forgo the buy-back option in an arbitrary period $n < N$, the firm faces the expected inventory holding and shortage costs from satisfying demand to the extent possible with the available inventory. The firm also retains the option to offer the buy-back in a future period, which is reflected in the value of β_{n+1} in the cost-to-go function for the next period, $n + 1$. Therefore, we can represent the value function in an arbitrary period $n < N$ when the firm chooses to not offer a buy-back in the current period as:

$$\begin{aligned} &C_n(z_n = 0 | I_n, \alpha_n = 0, \beta_n = 1) \\ &\qquad = \int_0^{I_n} \left[c_h(I_n - x) + C_{n+1}^*(I_n - x, \alpha_{n+1} = 0, \beta_{n+1} = 1)\right] f_n(x) dx \\ &\qquad + \int_{I_n}^{\infty} \left[c_s(x - I_n) + C_{n+1}^*(0, \alpha_{n+1} = 0, \beta_{n+1} = 1)\right] f_n(x) dx \end{aligned} \tag{8.10}$$

Given the expressions (8.9) and (8.10), we can once again express the value function that accounts for the choice of whether to offer the buy-back now as:

$$\begin{aligned} C_n(B, z_n | I_n, \alpha_n = 0, \beta_n = 1) &= z_n \left[r(B) + C_n^*(I_n + B, \alpha_n = 0, \beta_n = 0) \right] \\ &+ (1 - z_n) \left[\int_0^{I_n} \left[c_h(I_n - x) + C_{n+1}^*(I_n - x, \alpha_{n+1} = 0, \beta_{n+1} = 1)\right] f_n(x) dx \right. \\ &\left. + \int_{I_n}^{\infty} \left[c_s(x - I_n) + C_{n+1}^*(0, \alpha_{n+1} = 0, \beta_{n+1} = 1)\right] f_n(x) dx \right] \end{aligned} \tag{8.11}$$

With the value function in period N defined by (8.8) and in an arbitrary period $n < N$ defined by (8.11), we can now express the optimal expected cost-to-go in any period n with the buy-back option available as:

$$C_n^*(I_n, \alpha_n = 0, \beta_n = 1) = \min_{z_n \in \{0,1\}, 0 \le B \le M z_n} C_n(B, z_n | I_n, \alpha_n = 0, \beta_n = 1) \tag{8.12}$$

In this optimization, the constraints reflect the requirement that the buy-back quantity can only be positive if the buy-back offer is made in period n, where once again M is a very large positive number.

We now turn to the value function when the firm still has the option to extend the capacity reservation, starting with the final period N and then considering the case of any arbitrary period $n < N$. If the capacity reservation was extended in period $N - 1$, the firm faces the decision in period N to either extend the reservation for another period or terminate the capacity reservation permanently. Suppose that the firm chooses not to extend the reservation. It must then try to satisfy demand with the starting inventory I_N in conjunction with potential units obtained from a future buy-back offer, which is still available. Therefore, the value function associated with terminating the capacity reservation in period N is equivalent to the cost-to-go of having the option to make the buy-back offer, i.e.,

$$C_N(y_N = 0|I_N, \alpha_N = 1, \beta_N = 1) = C_N^*(I_N, \alpha_N = 0, \beta_N = 1) \tag{8.13}$$

Alternatively, if the firm chooses to extend the reservation in period N, it incurs the capacity reservation fee c_f and must choose an order-up-to level S_N and pay for any units ordered. It must also pay the expected shortage costs associated with the S_N units of inventory that are now available. Note that we do not allow the firm to both order units and offer the buy-back in the same period, and therefore, the value function of extending the reservation in the final period N and ordering up to S_N units is:

$$C_N(S_N, y_N = 1|I_N, \alpha_N = 1, \beta_N = 1) = c_f + c_p(S_N - I_N) + c_s \int_{S_N}^{\infty} (x - S_N) f_N(x) dx \tag{8.14}$$

with the constraint $S_N \geq I_N$. Note that the expressions in (8.13) and (8.14) are distinguished by the value of the binary decision variable y_N. As before, we write the value function in period N when the capacity reservation was extended in period $N-1$ as:

$$\begin{aligned} C_N(S_N, y_N|I_N, \alpha_N = 1, \beta_N = 1) =& y_N \left[c_f + c_p(S_N - I_N) + c_s \int_{S_N}^{\infty} (x - S_N) f_N(x) dx \right] \\ &+ (1 - y_N) C_N^*(I_N, \alpha_N = 0, \beta_N = 1) \end{aligned} \tag{8.15}$$

with the constraint $I_N \leq S_N \leq I_N + My_N$.

With the value function defined for the final period N, we can now turn to an arbitrary period $n < N$. By the same argument we used to develop (8.13), we can write the expected cost-to-go of terminating the capacity reservation in any arbitrary period $n < N$ as:

$$C_n(y_n = 0|I_n, \alpha_n = 1, \beta_n = 1) = C_n^*(I_n, \alpha_n = 0, \beta_n = 1) \tag{8.16}$$

On the other hand, the value function of extending the reservation in an arbitrary period $n < N$ and ordering up to S_n units consists of the capacity reservation fee, variable procurement costs, the expected inventory holding and shortage costs in period n, and the cost-to-go of entering the next period with the remaining inventory and the option of extending the capacity reservation still available. This value function can be expressed as:

$$\begin{aligned} C_n(S_n, y_n = 1|I_n, \alpha_n = 1, \beta_n = 1) &= c_f + c_p(S_n - I_n) \\ &+ \int_0^{S_n} \left[c_h(S_n - x) + C_{n+1}^*(S_n - x, \alpha_{n+1} = 1, \beta_{n+1} = 1) \right] f_n(x) dx \\ &+ \int_{S_n}^{\infty} \left[c_s(x - S_n) + C_{n+1}^*(0, \alpha_{n+1} = 1, \beta_{n+1} = 1) \right] f_n(x) dx \end{aligned} \tag{8.17}$$

with the constraint that $S_n \geq I_n$. As before, we write the overall value function in an arbitrary period $n < N$ when the capacity reservation was extended in period $n-1$ as:

$$\begin{aligned}
C_n(S_n, y_n | I_n, \alpha_n = 1, \beta_n = 1) = y_n \Big[& c_f + c_p(S_n - I_n) \\
& + \int_0^{S_n} \big[c_h(S_n - x) + C^*_{n+1}(S_n - x, \alpha_{n+1} = 1, \beta_{n+1} = 1)\big] f_n(x)dx \\
& + \int_{S_n}^{\infty} \big[c_s(x - S_n) + C^*_{n+1}(0, \alpha_{n+1} = 1, \beta_{n+1} = 1)\big] f_n(x)dx \Big] \\
& + (1 - y_n) C^*_n(I_n, \alpha_n = 0, \beta_n = 1) \qquad (8.18)
\end{aligned}$$

with the constraint $I_n \leq S_n \leq I_n + My_n$. With the value function in period N defined by (8.15) and in any arbitrary period $n < N$ defined by (8.18), we can now express the optimal expected cost-to-go in any period $n \in \{1, ..., N\}$ when the capacity reservation can be extended as:

$$C^*_n(I_n, \alpha_n = 1, \beta_n = 1) = \min_{y_n \in \{0,1\}, I_n \leq S_n \leq I_n + My_n} C_n(S_n, y_n | I_n, \alpha_n = 1, \beta_n = 1) \qquad (8.19)$$

The constraints in the above minimization reflect the requirement that units can only be ordered in the current period if the capacity reservation is extended.

With the formulation complete, we now turn to some properties of the optimal buy-back policy.

8.2 Properties of the Optimal Buy-Back Policy

We now investigate properties of the buy-back value function. We should first note that, although the state variables differ between the dynamic variant in Chapter 5 and the model presented in this chapter, if no further decisions can be made, the expected cost-to-go functions are identical. Therefore, for an arbitrary n, $C^*_n(I_n, \alpha_n = 0, \beta_n = 0)$ as defined in equation (8.5) is convex in the starting inventory level I_n.

We would like to determine the optimal buy-back quantity B if the firm decides to offer a buy-back in the current period, n. The value functions associated with offering the buy-back in either period N, as shown in equation (8.6), or any arbitrary period $n < N$, as shown in equation (8.9), are both composed of the buy-back cost function and the expected cost-to-go after there are no decisions remaining. Therefore, the following argument applies equally to both. Consider the function $C_n(B, z_n = 1 | I_n, \alpha_n = 0, \beta_n = 1)$ expressed in (8.9) and repeated below:

$$C_n(B, z_n = 1 | I_n, \alpha_n = 0, \beta_n = 1) = r(B) + C^*_n(I_n + B, \alpha_n = 0, \beta_n = 0)$$

Proposition 8.1. *Given that the firm has decided to offer the buy-back in an arbitrary period n, the buy-back value function $C_n(B, z_n = 1 | I_n, \alpha_n = 0, \beta_n = 1)$ as defined in (8.9) is jointly convex in (I_n, B).*

Proof. This is a function in two variables, I_n and B, because all other state and decision variables have been fixed. By assumption, the first term $r(\cdot)$ is convex. Similarly, we have shown that the second term is convex in its first argument. Because B and $I_n + B$ are both affine maps, both terms in (8.9) are also convex and, therefore because the sum of convex functions is convex, the overall expression is jointly convex in (I_n, B). □

One consequence of the joint convexity of (8.9) is that a value of B that minimizes (8.9) for a given starting inventory level I_n can be found easily. Let $B_n^*(I_n)$ be an optimal buy-back quantity that minimizes (8.9) for a given period n and starting inventory level I_n with the constraint that $B \geq 0$. Because (8.9) is convex in (I_n, B), $B \geq 0$ is a convex non-empty set, and (8.9) is bounded below, we know that:

$$C_n(B_n^*(I_n), z_n = 1 | I_n, \alpha_n = 0, \beta_n = 1) = \min_{B \geq 0} C_n(B, z_n = 1 | I_n, \alpha_n = 0, \beta_n = 1) \tag{8.20}$$

is also convex in I_n. This allows us to represent the choice in period n between offering the buy-back or not in the current period as the minimum of two functions. Therefore the decision problem expressed in (8.12) can be written as:

$$\begin{aligned} C_n^*(I_n, \alpha_n = 0, \beta_n = 1) = \min \Big\{ & C_n(B_n^*(I_n), z_n = 1 | I_n, \alpha_n = 0, \beta_n = 1), \\ & C_n(z_n = 0 | I_n, \alpha_n = 0, \beta_n = 1) \Big\} \end{aligned} \tag{8.21}$$

We now turn to some properties of the buy-back value function that enable us to characterize the optimal number of buy-back units to acquire in terms of the starting inventory I_n.

Lemma 8.2. *Given that the firm chooses to offer the buy-back in an arbitrary period n, the minimum of the buy-back value function* (8.9) *occurs at a starting inventory level for which the corresponding optimal buy-back quantity is zero.*

We present a formal proof of Lemma 8.2 in Appendix A but provide an informal outline of the proof here. If the minimum of the function in (8.9) occurs at some starting inventory level with a positive buy-back quantity, that implies the expected cost-to-go is improved with more inventory. Therefore, if that inventory level could be achieved for free (without the buy-back cost), the cost-to-go can only be smaller, which contradicts the assumption of the starting inventory level being the minimizer.

Theorem 8.3. *Given that the firm chooses to offer the buy-back in an arbitrary period n, there is a unique starting inventory threshold $\hat{I}_n$ above which the optimal buy-back quantity is zero, i.e., $B_n^*(I) = 0$ for all I such that $I \geq \hat{I}_n$.*

We prove this by using the convexity of the cost-to-go function. For any starting inventory level with an optimal buy-back quantity of zero, the marginal buy-back cost must be larger than the marginal savings of having an additional unit of inventory. By the convexity of the cost-to-go, any larger inventory value must have an even smaller marginal savings from having an additional unit, and therefore the buy-back is not beneficial for that higher starting inventory level either.

Proof. Consider some starting inventory level $\hat{I}$ with the corresponding optimal buy-back quantity $B_n^*(\hat{I}) = 0$. An optimal buy-back quantity of zero implies that $\frac{d}{dB}r(B)|_{B=0} \geq -\frac{d}{dB}C_n^*(\hat{I}+B, \alpha_n = 0, \beta_n = 0)|_{B=0}$ (the cost of adding an incremental unit via buy-back is greater than the savings from having that incremental unit available). Now consider another starting inventory level $\tilde{I} > \hat{I}$ with the corresponding optimal buy-back quantity $B_n^*(\tilde{I}) = \tilde{B} > 0$. An optimal positive buy-back quantity implies that $\frac{d}{dB}r(B)|_{B=0} < -\frac{d}{dB}C_n^*(\tilde{I} + B, \alpha_n = 0, \beta_n = 0)|_{B=0}$. This, in turn, implies that $\frac{d}{dB}C_n^*(\hat{I} + B, \alpha_n = 0, \beta_n = 0)|_{B=0} > \frac{d}{dB}C_n^*(\tilde{I}+B, \alpha_n = 0, \beta_n = 0)|_{B=0}$, which contradicts the convexity of $C_n^*(I+B, \alpha_n = 0, \beta_n = 0)$. Therefore, for any $\hat{I}$ such that $B_n^*(\hat{I}) = 0$, all $I_n \geq \hat{I}$ also have $B_n^*(I_n) = 0$. This means that there is at most a single starting inventory threshold above which the optimal buy-back quantity is zero units. □

Theorem 8.3 enables us to characterize the intersection between the value functions associated with buying back now and not buying back now, as outlined in (8.21). For any starting inventory level and period combination, the expected cost of offering a buy-back with an optimal buy-back quantity of zero will never be less than the expected cost of not buying back now. Therefore, we know that, for any arbitrary period n, any starting inventory level for which it is more expensive to not buy back now must occur at a point where the optimal buy-back quantity is strictly positive. This knowledge enables us to more quickly find the intersection point through numerical methods, and therefore also the inventory threshold below which it is optimal to offer the buy-back in the current period.

Once again, the non-convexity of the general DP value functions makes it necessary to find the optimal strategy numerically. However, before we present the results of our numerical study, we will first introduce an extension of our heuristic for the dynamic model with a one-time buy-back option. Such a heuristic could be used if solutions need to be found quickly.

8.3 Extension of the Heuristic for the Dynamic Model with Buy-Back

The heuristic presented in Chapter 6 can be adapted easily to the dynamic model with a one-time buy-back option. In our heuristic, we find two near-optimal, inventory-state-independent order-up-to levels for each period by assuming that the firm either will or will not extend the capacity reservation in the next period. We called these two values the *continuing-buy* and *last-time-buy* order-up-to levels, respectively. We use the same assumptions and nomenclature as we did in the heuristic for the problem variant without a buy-back.

First, let us assume that the firm knows (with certainty) that it will extend the capacity reservation in the next period. Because the buy-back offer cannot be made in the same period in which the capacity reservation is extended, the firm also knows that the next period will not be the buy-back offer period. Therefore, the choice of order-up-to level in the current period is not affected by the buy-back decision, as there will be at least one more opportunity to order from the contract manufacturer before having to consider the buy-back offer. For this reason, we can use the same newsvendor-style analysis as detailed in Section 6.1 for estimating the *continuing-buy* order-up-to level, S_n^c.

However, if we assume the firm knows (with certainty) that it will not extend the capacity reservation in the next period, then the optimal policies for the dynamic model with a one-time buy-back option and the dynamic model without a buy-back option may differ because the former must take into account the opportunity to make a buy-back offer in future periods. In Section 6.2 we presented a method for finding the *last-time-buy* order-up-to level based on the properties of the expected cost-to-go function after the capacity reservation has ended. In the dynamic model without a buy-back option, this expected cost-to-go function was convex, as there were no further decisions and therefore the expected cost-to-go was composed only of inventory holding and shortage costs. However, in the dynamic model with a buy-back option, the firm still has a decision to make after the capacity reservation has ended: the choice of when (if at all) to offer the buy-back and how many units to buy back at that time. As we are not able to characterize the optimal buy-back policy analytically, in this version of the heuristic we will determine the *last-time-buy* order-up-to level, S_n^ℓ, numerically.

As in our other heuristic, we assume that the firm will not extend the capacity reservation after ordering up to the *last-time-buy* order-up-to level, S_n^ℓ. In order to find a starting-inventory-independent order-up-to level, we also assume that the starting inventory level is less than the to-be-selected last-time-buy order-up-to level. Because we are using the marginal cost of ordering when determining S_n^ℓ, it is sufficient to restrict our attention to a starting inventory level of zero, as the same S_n^ℓ will be optimal for all starting inventory levels between zero and S_n^ℓ. With both of these assumptions in place, the *last-time-buy* value function in an arbitrary period n is:

$$C_n(S_n|0, \alpha_n = 1, \beta_n = 1) = c_f + c_p(S_n - 0) + C_n^*(S_n, \alpha_n = 0, \beta_n = 1) \tag{8.22}$$

and we define S_n^ℓ as the value of S_n that minimizes (8.22) in period n.

Finding S_n^ℓ requires knowing the expected cost-to-go function when the buy-back option is available, $C_n^*(I_n, \alpha_n = 0, \beta_n = 1)$, and therefore our heuristic requires solving the dynamic program for all $(I_n, \alpha_n = 0, \beta_n = 1)$ and $(I_n, \alpha_n = 0, \beta_n = 0)$ states. However, because the buy-back is a one-time option, this requires only a fraction of the computing time required to solve the dynamic program for the entire state space.

After finding the candidate order-up-to levels S_n^c and S_n^ℓ, we can use the heuristic to find an approximate solution. We do so by recursively solving the simplified dynamic program using the following heuristic expected cost-to-go function defined for all periods $1 \leq n \leq N$

and starting inventory levels I_n, in much the same way as in the heuristic presented earlier:

$$C_n^*(I_n,\alpha_n = 1, \beta_n = 1) = \min_{y_n \in \{0,1\}, I_n \le S_n \le I_n + My_n, S_n \in \{I_n, S_n^c, S_n^\ell\}} C_n(S_n, y_n | I_n, \alpha_n = 1, \beta_n = 1) \tag{8.23}$$

where $C_n(S_n, y_n|I_n, \alpha_n = 1, \beta_n = 1)$ is defined as in equation (8.18). Once again, it is necessary to discretize the demand and inventory levels to a consistent level of granularity in order to make the problem tractable.

8.4 Numerical Study

In this section, we present the results of a numerical study for the dynamic model with a one-time, deterministic buy-back option. Our goals for the study are: (1) to explore how the introduction of the buy-back option changes the optimal strategy, and (2) to evaluate the performance of the heuristic relative to that of the optimal strategy. We utilized the same 100 problem instances as in Chapter 7 and discretized the demand distribution and the set of options for inventory levels and order-up-to levels in the same way as in Chapter 7. We direct the reader to Section 7.2 for a detailed overview of the problem instances as well as our implementation process.

For this numerical study, we assume that the total buy-back cost as a function of the number of buy-back units B is as follows:

$$r(B) = \frac{B^2}{20} + c_p B \tag{8.24}$$

This function meets our requirements that $r(\cdot)$ is convex, $r(0) = 0$, and $r'(\cdot) \ge c_p$. Dividing both sides by B yields the per-unit buy-back cost as a function of the buy-back quantity:

$$\frac{r(B)}{B} = \frac{B}{20} + c_p \tag{8.25}$$

This function is linear, meaning that the number of customers who respond to the buy-back (in this deterministic model) is a linearly increasing function of the per-unit buy-back offer value. In this case, for every additional unit acquired via the buy-back, the per-unit buy-back cost increases by 0.05. This interpretation will be relevant when we introduce the stochastic version of the buy-back in Section 8.5.

In Figure 8.1 we plot the cost to acquire various quantities of units via the buy-back or paying a capacity reservation cost (in this case 8000) and ordering from the contract manufacturer. Note that, due to the capacity reservation cost, for order quantities of a few hundred units or less, it is less expensive to acquire the units via the buy-back. However, as the buy-back is a one-time option and the capacity reservation can be extended in future periods if necessary, the optimal strategy hinges on more than the variable costs alone.

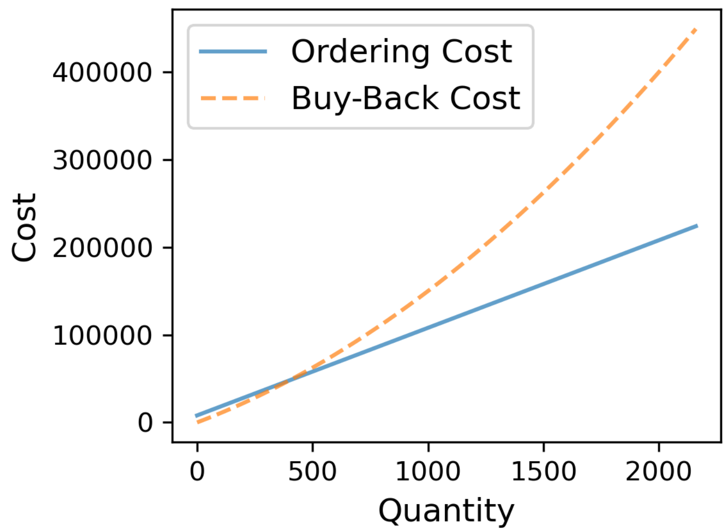

Figure 8.1: Comparison of the Buy-Back Cost and the Cost to Extend the Capacity Reservation and Order when $c_f = 8000$

We chose this particular buy-back cost function so that the intersection point plotted in Figure 8.1 would occur at a value of a few hundred units (the exact point depends on the choice of c_f). A few hundred units are equivalent to a few periods of expected demand, and therefore the buy-back is complementary to, but not a full replacement for, ordering a last time buy from the contract manufacturer.

In this numerical study, we implemented both the optimal and heuristic algorithms. CPU times are approximately 28 minutes to find the optimal dynamic solution and approximately 3 minutes to find the heuristic solution. The increase in the computing times for the heuristic relative to those in Chapter 7 is due to the need to partially solve the dynamic program in order to find the heuristic order-up-to levels. For the full results of the numerical study for all 100 problem instances, see Tables B.6 through B.10 in Appendix B.

Characteristics of the Optimal Strategy

In this numerical study, we observed that the optimal strategy for the vast majority of problem instances has the same structure as the common optimal strategy first described in Section 7.3, with up to three threshold values that demarcate intervals of starting inventory in each period, each corresponding to exactly one of the *continuing buy*, *extend only*, *last time buy*, or *end capacity reservation* strategies. Of the 1200 problem-instance-period combinations, only one has an optimal strategy that deviates from this basic structure. As such, in this subsection, we will primarily focus on how the optimal strategies for the dynamic model with the buy-back option differ from the optimal strategies for the dynamic model without the buy-back option for the same base problem instances.

We now consider an example from the numerical study. This problem instance was selected because of clear differences between the optimal strategies with and without the buy-back option, but it is representative of the results of our numerical study overall. In Figure 8.2, we plot the optimal order-up-to levels in the first period for the dynamic model with the one-time buy-back option and the dynamic model without the buy-back option for the problem instance with the *Double Peak* demand pattern, a capacity reservation cost of 8000, and a shortage cost of 600. The most prominent differences between the optimal strategies for the buy-back and non-buy-back models are the optimal *last-time-buy* order-up-to levels. As expected, the presence of the buy-back option leads the firm to acquire fewer units from the contract manufacturer with the *last time buy*. On the other hand, the *continuing-buy* order-up-to levels are unchanged at 70. Another major difference in the optimal strategies is the significantly smaller *end-capacity-reservation* threshold (shown by the vertical dotted line) for the model with the buy-back option, as the firm can now more easily take the risk of ending the capacity reservation with less inventory on hand due to the option to acquire units via the buy-back in the future.

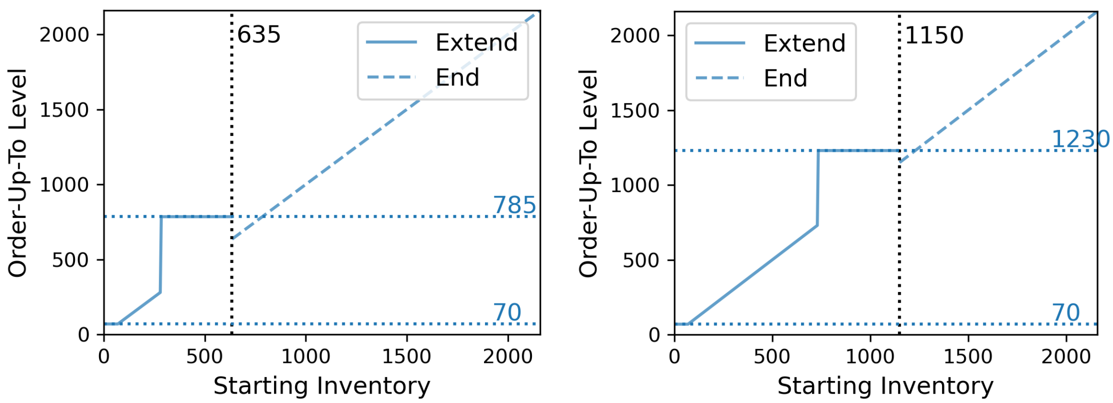

(a) With Buy-Back Option

(b) Without Buy-Back Option

Figure 8.2: Comparison of the Optimal Order-Up-To Levels in Period 1 With or Without the Buy-Back Option for the *Double Peak* Demand Pattern with $c_f = 8000$ and $c_s = 600$

The optimal strategies for both variants differ significantly in later periods, as well. In Figure 8.3 we present a comparison of the optimal strategies with or without the buy-back option for the same problem instance as the one on which Figure 8.2 is based. Two differences are readily apparent: (1) the *end-capacity-reservation* thresholds are much lower for the variant with a buy-back option, and (2) the optimal strategy including the buy-back has an anticipated *last time buy* in period four, rather than in period five. Combined, all of the differences induced by the buy-back option are responsible for an 8.7% reduction in the expected cost over the full horizon for this problem instance when the firm begins the first period with no starting inventory.

Another notable feature of the optimal strategy with the buy-back option (as detailed in

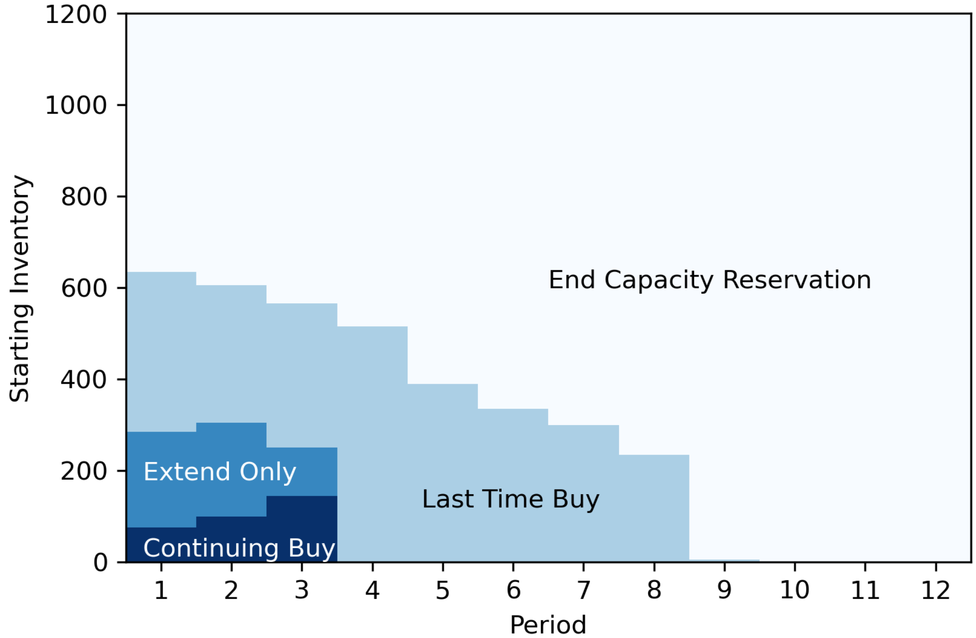

(a) With Buy-Back Option

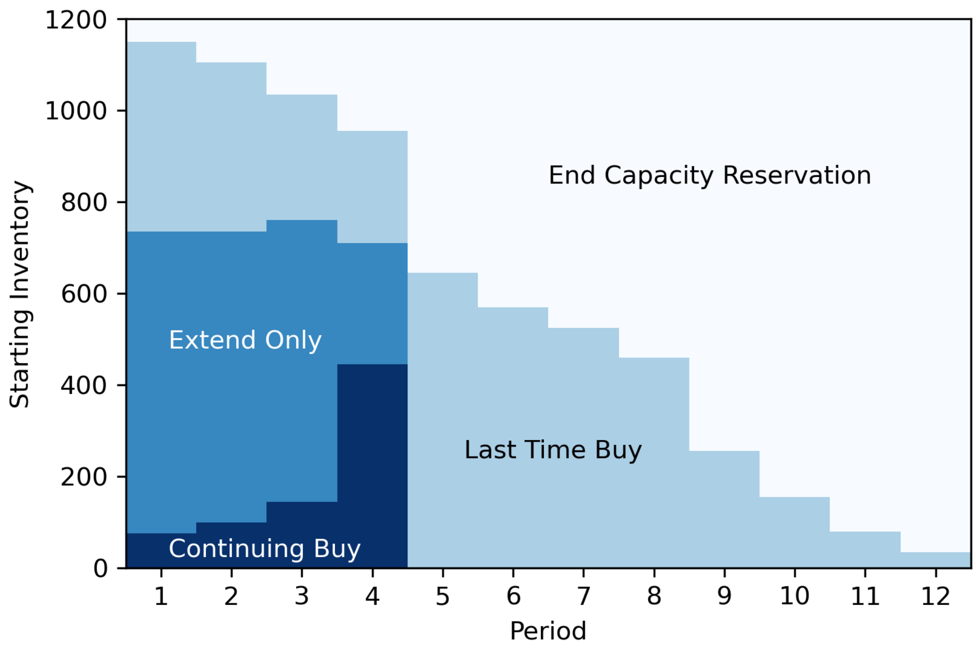

(b) Without Buy-Back Option

Figure 8.3: Comparison of the Optimal Strategies With or Without the Buy-Back Option for the *Double Peak* Demand Pattern with $c_f = 8000$ and $c_s = 600$

Figure 8.3a) is that, late in the time horizon, it is optimal to end the capacity reservation for any starting inventory level—even zero units. In our previous numerical study, this strategy was optimal only when the capacity reservation costs were high, the expected remaining demand was low, and the shortage costs were low. As Figure 8.3b shows, if the firm finds itself in period 10 with low starting inventory and the opportunity to extend the capacity reservation, it would do so if no buy-back were available. In the same situation in the dynamic model with a buy-back option, purchasing from the contract manufacturer has been totally supplanted by the buy-back offer.

We now briefly discuss the optimal buy-back strategy after the capacity reservation has been terminated. In Figure 8.4 we plot the optimal buy-back-up-to levels in period four for the same problem instance that provides the basis for Figures 8.2 and 8.3. The buy-back order-up-to levels are qualitatively different than those for the *last time buy*: rather than the commonly-observed constant *last-time-buy* order-up-to levels over an interval of starting inventory levels, the buy-back-up-to levels increase slightly—but less than one-for-one—with the starting inventory level. This partially mitigates the effect of the convex increasing buy-back costs by dampening the order *quantities* to a greater extent as the starting inventory level declines.

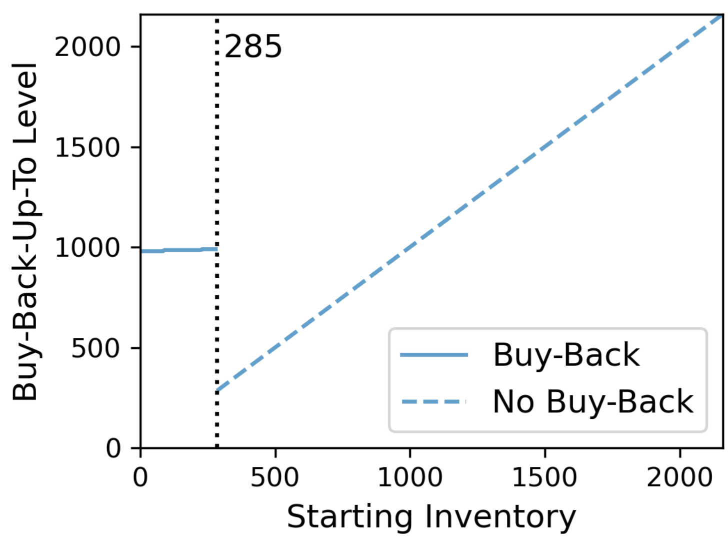

Figure 8.4: Optimal Buy-Back-Up-To Strategy in Period 4 for the *Double Peak* Demand Pattern with $c_f = 8000$ and $c_s = 600$

As on the order-up-to level figures, we provide a vertical dotted line to show the starting inventory threshold below which it is optimal to make a buy-back offer in the current period. In Figure 8.4, this threshold (285) is much less than the optimal buy-back-up-to level (approximately 1000) and slightly less than the expected demand in period four for the *Double Peak* demand pattern (300). Due to the one-time nature of the buy-back offer, the firm is willing to tolerate a relatively high number of expected shortages in exchange for retaining the buy-back option for a future period. As a result, the firm will not make the buy-back offer unless it is likely to need a large number of units in the current period, because it does not have to forgo the buy-back opportunity if it does not exercise it now.

The ability to forgo the buy-back in the current period without any fixed cost contrasts with the sequential capacity requirement, in which retaining the option to produce in future periods requires paying the capacity reservation fee. The buy-back option thus enables the firm to place the *last time buy* earlier in the horizon, as it can instead expect to satisfy some late-stage demand via the buy-back. As a result, we expect that the buy-back option will allow the firm to meaningfully lower its expected costs relative to the dynamic variant without a buy-back option.

Comparison with the Dynamic Model without a Buy-Back Option

The option to make a one-time buy-back offer results in a significant improvement of the expected cost of the optimal strategy. On average for our set of problem instances, the expected total cost when the firm begins the first period with no starting inventory is 7% lower for the optimal strategy for the dynamic model with a buy-back option than for the optimal strategy for the dynamic model without a buy-back option. We did not observe a relationship between the demand pattern or the shortage cost and the magnitude of the benefit, but the capacity reservation cost has a strong impact on the magnitude of the benefit. Table 8.1 presents the expected total cost benefit averaged across problem instances with the same capacity reservation cost. The table includes the percentage benefit, absolute benefit, as well as the ratio of the benefit to the capacity reservation cost. The average benefit is about 1.5 to 2 times the single-period capacity reservation cost, with the cost savings arising from a combination of fewer capacity reservation fees and reduced inventory holding costs.

Table 8.1: Expected Total Cost Benefit - by Capacity Reservation Cost

	Expected Total Cost		
c_f	**% Benefit**	**Absolute Benefit**	**Benefit/c_f**
2000	1.75%	2635	1.3
4000	4.00%	6777	1.7
8000	8.67%	17077	2.1
16000	13.60%	30135	1.9

Most interestingly, the buy-back option benefit (7%) is *much* larger than the optimality gap between the dynamic model and the pre-commit model (0.06%) studied in Chapter 7. We should note that the comparison between the dynamic and pre-commit models was based on the assumption that the per-period capacity reservation costs were identical. However, in reality, the contract manufacturer would impose larger charges in the form of higher capacity reservation fees, higher variable costs, or an additional fixed cost, for the flexibility of the dynamic capacity reservation contract. The large cost improvement due to the buy-back and the minimal savings from the dynamic model versus the pre-commit model imply that the firm would be better off pursuing a buy-back option as opposed to securing a period-to-period capacity reservation contract with the contract manufacturer. The differences in overall costs that we have observed here are consequences of the sequential capacity reservation cost structure, and they highlight the potentially critical role of the buy-back option or other alternative source of units that *does not require* a large one-time investment or the continuation of fee payments just to maintain the option to order.

We note that the exact magnitudes of the various cost differences will depend upon the cost parameters, but the insights from our numerical study support our expectation that the directional effects will generalize to other sets of cost parameters.

Performance of the Heuristic

The heuristic produced near-optimal results for our set of 100 problem instances assuming beginning-of-horizon inventory levels of zero, with an average optimality gap of 0.01%. (We omit the adjusted optimality gap used in Chapter 7 as expected variable procurement costs differ for the buy-back and no-buy-back options.) Demand patterns were the only factor to have a noticeable effect on the magnitudes of optimality gaps, with gaps being effectively zero except for the *Double Peak* demand pattern, for which the average optimality gap was 0.05%. These results parallel those for the dynamic model without a buy-back.

The heuristic performs quite well in an absolute sense, but its relative performance improves further with the introduction of the buy-back option. The main reason is that the *last-time-buy* order-up-to level in the heuristic is determined under the assumption that the capacity reservation will not be extended in the next period. The existence of the buy-back option makes it even more likely that it is either not necessary, or not economical, to extend the capacity reservation contract beyond the heuristically-determined period in which the (anticipated) *last time buy* occurs.

8.5 Generalization with Stochastic Buy-Back Yield

We also generalize the model in the previous section to allow uncertainty in the number of customers responding to a buy-back offer at any given buy-back price. In this generalization, the firm selects a single buy-back price p_b, which generates a random buy-back quantity $B(p_b)$.

The value function associated with making a buy-back offer at a buy-back price p_b in an arbitrary period n with a starting inventory level of I_n is:

$$C_n(p_b, z_n = 1 | I_n, \alpha_n = 0, \beta_n = 1) = \int_0^{B_{max}} \left[p_b x + C_n^*(I_n + x, \alpha_n = 0, \beta_n = 0) \right] f_{B(p_b)}(x) dx \tag{8.26}$$

where $f_{B(p_b)}(\cdot)$ is the pdf of the random variable $B(p_b)$ and B_{max} is the maximum possible number of buy-back units that can be acquired (e.g., the total number of units owned by the firm's customers). We continue to assume that the refurbishment yield is deterministic, so for simplicity, we define $B(p_b)$ as the quantity *after* refurbishment. We also have the expected cost-to-go function:

$$C_n^*(I_n, \alpha_n = 0, \beta_n = 1) = \min_{z_n \in \{0,1\}, c_p \le p_b \le c_p + M z_n} C_n(p_b, z_n | I_n, \alpha_n = 0, \beta_n = 1) \tag{8.27}$$

where the constraints reflect the requirement that the buy-back price must be larger than the per-unit procurement cost c_p, where M is a very large positive number. This assumption is the same as in our model with a deterministic buy-back option.

In the simplest form, the random variable $B(p_b)$ could be from a family of distributions that depends on p_b. For instance, if the firm believes that the mean customer response to an increase in the buy-back price is roughly linear and the coefficient of variation remains constant across a range of mean values, $B(p_b)$ could be defined as a Normal random variable with mean $(p_b - c_p)\mu$ and standard deviation $(p_b - c_p)\sigma$, where μ and σ are constants. As with the deterministic buy-back case, the parameters are shifted by c_p to ensure that the buy-back yield is zero if $p_b = c_p$. For many consumer electronics devices, especially higher-end devices, some customers trade in their devices and receive trade-in value toward their new devices when replacing them. Therefore, it is reasonable to expect the firm to be able to construct the distribution $B(p_b)$ for a range of buy-back prices based on the firm's own experience with buy-back or trade-in offers, or based on knowledge of industry trends.

8.6 Numerical Study with Stochastic Buy-Back Yield

In this section, we present the results of a numerical study for the dynamic model with a one-time buy-back option, in which the yield from the buy-back offer is a stochastic function of the buy-back offer price. Our goal for this study is to explore how the change from a deterministic to a stochastic response to the buy-back offer affects both the expected cost and the optimal strategy of the firm. Once again, we utilized the same 100 problem instances as in Chapter 7 using the same discretization process. We direct the reader to Section 7.2 for a detailed overview of the problem instances considered as well as our implementation process.

We model the random buy-back quantity $B(p_b)$ as a (discretized) Normal random variable with a mean of $\mu = 20(p_b - c_p)$ and a standard deviation of $\sigma = 4(p_b - c_p)$. In the deterministic case, for every additional unit acquired via the buy-back, the per-unit buy-back cost increases by 0.05 (see equation (8.25)), which is reflected in the definition of μ and facilitates a comparison between the stochastic and deterministic cases. Our choice of σ implies a coefficient of variation equal to 0.2. This provides for a sufficiently high degree of uncertainty that we anticipate observing some changes in optimal strategies. In practice, the firm would estimate the buy-back response based on historical trade-in offers and market conditions.

We utilized values of p_b that are multiples of 0.25 because this leads to random variables $B(p_b)$ for which, after small downward adjustments of the standard deviations, our discretization results in buy-back quantities that are multiples of five units. We implemented a discretization process that is the same as that for our discretized Normal demands, except the discrete points were distributed from three standard deviations below the mean to three standard deviations above the mean (rather than four) in one standard deviation increments (rather than one-half).

For this numerical study, we only found the optimal strategy for the dynamic program, which required approximately 45 minutes of computing time for each problem instance. Our heuristic could be applied to the stochastic yield case as well, but in view of the near-optimal performance of the heuristic discussed in Section 8.4 we elected to focus on a comparison of optimal strategies for the deterministic and stochastic versions of the buy-back model.

On average, for our set of 100 problem instances assuming beginning-of-horizon inventory levels of zero, the expected total cost increased by 1.5% in the stochastic problem relative to that of the deterministic problem. On average, problem instances with larger cost parameter values had larger cost increases between the stochastic problem and the corresponding deterministic problem, as expected. For the full results of the numerical study for all 100 problem instances, see Tables B.11 through B.15 in Appendix B. We next present a representative example from the numerical study to illustrate how the optimal strategy changes with the introduction of a stochastic buy-back.

Characteristics of the Optimal Strategy

We now consider the same problem instance that was the basis for Figure 8.4, as it is once again representative of the results of our numerical study. In Figure 8.5a we plot the optimal buy-back-up-to level in period four. The result is remarkably similar to that for the deterministic version of the problem instance, which we have repeated here in Figure 8.5b for convenience. In this problem instance, the stochastic buy-back model results in slightly lower buy-back-up-to levels and a slightly higher inventory threshold below which the buy-back is initiated relative to those for the deterministic case. The results shown here are representative of the relatively small changes in optimal buy-back strategy due to the change to a stochastic yield.

The sawtooth pattern in the optimal buy-back-up-to levels in Figure 8.5a is an artifact of our discretization process. In order to discretize to multiples of five units, we elected to adjust the standard deviations (recall that we used a constant coefficient of variation of 0.2) down to the nearest multiple of five. Due to this adjustment, neighboring sets of buy-back prices correspond to random variables with the same standard deviation (but different means). For example, in our numerical study, the set of buy-back prices 105, 105.25, 105.5, 105.75, and 106 correspond to Normal random variables with means of 100, 105, 110, 115, and 120, respectively. If we followed the standard deviation formula mentioned at the beginning of Section 8.6 exactly, the corresponding standard deviations would be 20, 21, 22, 23, and 24, respectively. However, this would violate our desired discretization to multiples of five

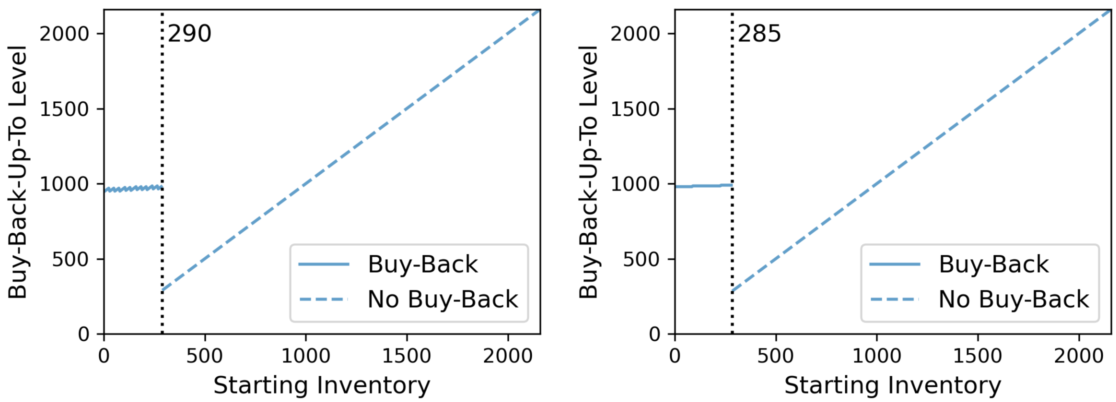

(a) Stochastic Yield (b) Deterministic Yield

Figure 8.5: Optimal Buy-Back-Up-To Strategy with Stochastic and Deterministic Buy-Back Yield in Period 4 for the *Double Peak* Demand Pattern with $c_f = 8000$ and $c_s = 600$

units, so therefore the standard deviations were each rounded down to 20. Because of this rounding, some of the buy-back prices correspond to random variables with coefficients of variation which are slightly lower than the desired 0.2. As a result, the optimal buy-back prices are almost always those that correspond to random variables with the largest mean for any fixed standard deviation, as these have slightly lower coefficients of variation than their neighbors. We believe the discretization that we have used is accurate enough to allow for comparisons between the deterministic and stochastic versions of the buy-back models, which is our primary interest.

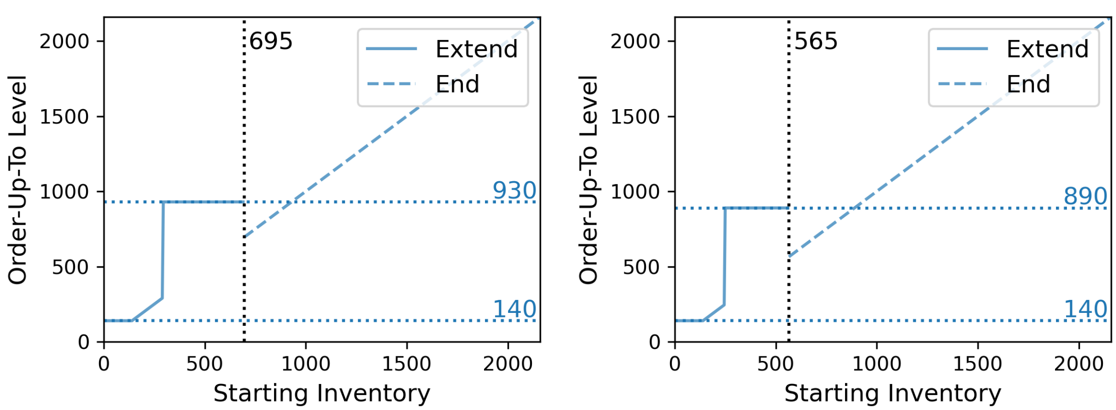

(a) Stochastic Yield (b) Deterministic Yield

Figure 8.6: Optimal Order-Up-To Strategy with Stochastic and Deterministic Buy-Back Yield in Period 3 for the *Double Peak* Demand Pattern with $c_f = 8000$ and $c_s = 600$

We are also interested in how the stochastic buy-back yield may change the optimal strategy for the last time buy decision(s). In Figure 8.6 we plot the optimal order-up-to levels for the cases of stochastic and deterministic buy-back yields, respectively, for the previous period (period three) of the same problem instance. It is evident that the relatively small changes in the optimal buy-back-up-to strategy lead to moderately larger changes in the optimal order-up-to strategy. In the stochastic yield case, the optimal strategy calls for ordering up to a slightly higher *last-time-buy* order-up-to level (930 vs 890) as well as a considerably higher *end-capacity-reservation* threshold (695 vs 565). However, the *continuing-buy* order-up-to level is not affected by the change to a stochastic buy-back yield, as expected.

In this problem instance, the change to a stochastic buy-back did not result in very large changes to the *structure* of the optimal strategy. However, the changes to the *values* of the decision variables and *thresholds* in the optimal strategy impact the expected cost. The uncertainty of the stochastic buy-back causes the firm to act conservatively and select smaller buy-back quantities, as evidenced by the slightly lower buy-back-up-to levels in Figure 8.5a. As a result, the firm instead chooses to order more units from the contract manufacturer (as shown by the larger *last-time-buy* order-up-to level in Figure 8.6a) and to extend the capacity reservation for larger levels of starting inventory (as shown by the larger *end-capacity-reservation* threshold in Figure 8.6a) in order to compensate for the lower buy-back quantity. As a result, in this problem instance, the stochastic buy-back model has 1.7% higher expected costs, which is slightly above the average across all problem instances.

Overall, for our problem instances, the buy-back option lowers the expected cost of the optimal solution considerably, especially in problem instances with large capacity reservation costs. This result holds for both the deterministic and stochastic buy-back models. Therefore, our results suggest that consumer electronics firms should consider pursuing a buy-back option as a part of their last time buy strategy when they are subject to a contract manufacturer's sequential capacity reservation requirement.

The strong performance of the buy-back option raises the question of why they are not more common in practice. The answer may be partly that selling refurbished or "refreshed" devices, especially for higher-end products like smartphones, is a large market in and of itself, and therefore using these devices to satisfy warranty claims may not be the most profitable use of them. Additionally, the firm may find it is competing with other companies to acquire used items. For instance, in the case of smartphones, it is common for carriers (such as AT&T) to make their own trade-in offers, especially to entice customers to switch away from other carriers. The carriers' offers may be more appealing than those of the original manufacturer, as the revenue from the accompanying long-term service contract can offset the higher cost of a more generous trade-in offer. A more realistic model would require taking into account the other competing trade-in or buy-back offers available to customers.

Aside from competing offers, the effectiveness of the buy-back option depends heavily on a number of other factors that are not included in our model. Customers may respond to past buy-back offers by not upgrading to the next-generation product at launch and instead waiting for a buy-back offer. This behavioral response from customers would need to be accounted for in the cost of the buy-back offer. Additionally, with large buy-back volumes, operational efficiencies in the collection process or the capacity of the refurbishing facility may become relevant factors. Fully modeling the buy-back option would require incorporating more detail and nuances to capture these effects and is a potential direction for future research.

Chapter 9

Conclusion

Our research was motivated by a problem facing modern consumer electronics firms, which increasingly design their own products but outsource production to contract manufacturers. These firms' products are often multi-generational, with short product life cycles that necessitate an eventual end-of-production decision for each product generation. To answer the questions of *when* to end the production of a particular product and *how much* to order from the contract manufacturer at that time, we developed a new variant of the last time buy problem.

To capture the economics of capacity reservations at major contract manufacturers, we introduced the notion of a *sequential capacity reservation* fee into a firm's procurement planning problem. In this realistic cost structure, which, to the best of our knowledge, has not been studied in the literature, the contract manufacturer requires the firm to pay a per-period capacity reservation fee if the firm would like to retain the option to order again in the subsequent period. In contrast to the more standard setup cost often included in other inventory models, this fee must be paid even if the firm does not order any units in the current period. However, if the firm does not pay the capacity reservation fee, the contract manufacturer will terminate the production of the product and shift the associated resources to a different product, potentially one sold by a competitor. This cost structure reflects the relatively high opportunity cost of capacity at the contract manufacturer.

As noted above, we address the firm's problem of determining whether to extend the capacity reservation in each period (if the choice has not already been made) and how much to order in each period. The goal is to minimize the expected total cost of capacity reservation fees, variable production costs, inventory holding costs, and shortage costs. We analyzed the resulting dynamic programming problem and found that the structure of the optimal policy is quite different than if the more standard fixed cost per period without the sequential requirement were imposed. The dynamic programming value functions are not guaranteed to be convex, but we have found them to be sufficiently well-behaved (e.g., a small number of local minima) that the optimal strategy can be computed numerically for any problem instance. We show that, if the firm extends the capacity reservation in the current period, the optimal order-up-to policy is characterized by a set of starting inventory intervals, each corresponding to a different order-up-to level. There may be uncountably many of these intervals and therefore the optimal strategy can be quite complex. However, in our numerical study, we find that in most cases the optimal strategy in any given period consists of *continuing buy*, *extend only*, *last time buy*, and *end capacity reservation* policy regions.

This strong pattern that we observed in the optimal strategies in our numerical study inspired our heuristic, which entails finding inventory-state-independent *continuing-buy* and *last-time-buy* order-up-to levels by assuming that the firm either extends or ends the capacity reservation in the next period, respectively. Our heuristic found near-optimal solutions for the dynamic model in a fraction of the time required to optimally solve the dynamic program.

We also found that the expected cost of the optimal dynamic strategy was almost identical to that of the optimal strategy for the pre-commit variant of our problem, in which the firm must commit at the start of the problem horizon to a specific last-time-buy period. This suggests that the flexibility afforded to the firm in the dynamic model is not very valuable even when the capacity reservation fees are assumed to be the same for the dynamic and pre-commit settings. Furthermore, in practice, it is reasonable to expect that the contract manufacturer would impose larger charges (either fixed or variable) in exchange for the period-to-period flexibility of our dynamic model. The findings from our numerical study indicate that a firm should not be willing to pay much—if anything at all—for a high degree of contract flexibility if the supplier enforces a sequential requirement on capacity reservation fees.

We also introduced an extension of our model, in which the firm has the one-time option to buy back units from its customers after ending the capacity reservation. This model is applicable when the late-stage demand is for units to satisfy warranty claims. In this setting, the firm would buy back units with the intent of refurbishing them to serve as warranty replacement units. We initially assumed that the number of buy-back units is a deterministic function of the buy-back price offered to the firm's customers. In contrast to the negligible cost savings from the flexibility provided by the dynamic model versus the pre-commit model, we found that the optimal strategy with the buy-back option led to considerably lower costs, especially in problem instances with large capacity reservation costs. This result also held for a generalization of the model in which the number of buy-back units was a stochastic function of the buy-back price offered. To summarize, the results of our numerical study suggest that, in a sequential capacity reservation setting, the buy-back option is more valuable than the flexibility afforded by a period-to-period capacity reservation contract whether the buy-back yields are deterministic or stochastic.

There are several avenues for future research. First, we have assumed that demand forecasts in the form of statistical distributions are available at the beginning of the planning horizon and are not updated. More research is needed to account for commonly-used forecast updating methods and possibly to devise specially-designed forecasting and forecast updating methods for settings in which sequential capacity reservation requirements exist. Second, with the growth of contract manufacturing, it has become increasingly more common for explicit or implicit fixed-charge capacity reservations to be imposed, and due to the economics and/or logistics involved in changing production lines from producing one product to another, capacity reservations are often required to be sequential. More research is needed to develop optimal and near-optimal procurement algorithms for such settings. Finally, we have considered only fixed-charge sequential capacity reservations, but a supplier may charge both a reservation fee for the option to use the capacity and a separate setup cost for a positive order quantity. This would lead to a much more complex procurement optimization problem, but our work provides a first step toward solving it.

Bibliography

[1] Sina Behfard et al. “Last Time Buy and repair decisions for fast moving parts.” *International Journal of Production Economics* 197 (2018), pp. 158–173.

[2] Sina Behfard et al. “Last time buy and repair decisions for spare parts.” *European Journal of Operational Research* 244.2 (2015), pp. 498–510.

[3] Marco Bijvank and Iris FA Vis. “Lost-sales inventory theory: A review.” *European Journal of Operational Research* 215.1 (2011), pp. 1–13.

[4] Youssef Boulaksil, Jan C Fransoo, and Tarkan Tan. “Capacity reservation and utilization for a manufacturer with uncertain capacity and demand.” *OR Spectrum* 39.3 (2017), pp. 689–709.

[5] James R Bradley and Héctor H Guerrero. “Lifetime buy decisions with multiple obsolete parts.” *Production and Operations Management* 18.1 (2009), pp. 114–126.

[6] James R Bradley and Héctor H Guerrero. “Product design for life-cycle mismatch.” *Production and Operations Management* 17.5 (2008), pp. 497–512.

[7] Gérard P Cachon. “Supply chain coordination with contracts.” *Handbooks in Operations Research and Management Science* 11 (2003), pp. 227–339.

[8] Kyle D Cattani and Gilvan C Souza. “Good buy? Delaying end-of-life purchases.” *European Journal of Operational Research* 146.1 (2003), pp. 216–228.

[9] Feng Cheng and Suresh P Sethi. “Optimality of state-dependent (s, S) policies in inventory models with Markov-modulated demand and lost sales.” *Production and Operations Management* 8.2 (1999), pp. 183–192.

[10] Dwayne Cole, Burak Kazaz, and Scott Webster. “Final purchase and trade-in decisions in response to a component phase-out announcement: a deterministic analysis.” *International Journal of Production Research* 54.5 (2016), pp. 1257–1272.

[11] Dwayne Cole, Burak Kazaz, and Scott Webster. “Satisfying Warranty Claims on Obsolete Products.” *Trends and Research in the Decision Sciences* 1 (2015), pp. 227–247.

[12] European Commission. *Consumer guarantees.* [Online; accessed on November 2, 2022]. 2022. URL: `https://europa.eu/youreurope/business/dealing-with-customers/consumer-contracts-guarantees/consumer-guarantees/index_en.htm`.

[13] Intel Corporation. *Annual Report*. [Online; accessed on November 20, 2022]. 2021. URL: https://www.intc.com/filings-reports/all-sec-filings/content/0000050863-22-000007/0000050863-22-000007.

[14] Daniel Costa and Edward A Silver. "Exact and approximate algorithms for the multi-period procurement problem where dedicated supplier capacity can be reserved." *OR Spektrum* 18.4 (1996), pp. 197–207.

[15] Leonard Fortuin. "Reduction of the all-time requirement for spare parts." *International Journal of Operations & Production Management* 2.1 (1981), pp. 29–37.

[16] Leonard Fortuin. "The all-time requirement of spare parts for service after sales—theoretical analysis and practical results." *International Journal of Operations & Production Management* 1.1 (1980), pp. 59–70.

[17] Fortune. *Global 500: Hon Hai Precision Industry*. [Online; accessed on November 20, 2022]. 2022. URL: https://fortune.com/company/hon-hai-precision-industry/global500/.

[18] JBG Frenk, Sonya Javadi, and Semih O Sezer. "An optimal stopping approach for the end-of-life inventory problem." *Mathematical Methods of Operations Research* 90.3 (2019), pp. 329–363.

[19] JBG Frenk, Canan Pehlivan, and Semih O Sezer. "Order and exit decisions under non-increasing price curves for products with short life cycles." *Mathematical Methods of Operations Research* 90.3 (2019), pp. 365–397.

[20] JBG Frenk et al. "An exact static solution approach for the service parts end-of-life inventory problem." *European Journal of Operational Research* 272.2 (2019), pp. 496–504.

[21] Matthew Galloway. *Finite discrete approximation to the normal distribution*. [Online; accessed on August 31, 2022]. 2017. URL: https://math.stackexchange.com/questions/920252/finite-discrete-approximation-to-the-normal-distribution.

[22] JHJ Geurts and JMC Moonen. "On the robustness of 'insurance type'spares provisioning strategies." *Journal of the Operational Research Society* 43.1 (1992), pp. 43–51.

[23] Matthieu van der Heijden and Bermawi P Iskandar. "Last time buy decisions for products sold under warranty." *European Journal of Operational Research* 224.2 (2013), pp. 302–312.

[24] Mansik Hur, Burcu B Keskin, and Charles P Schmidt. "End-of-life inventory control of aircraft spare parts under performance based logistics." *International Journal of Production Economics* 204 (2018), pp. 186–203.

[25] Karl Inderfurth and Rainer Kleber. "An advanced heuristic for multiple-option spare parts procurement after end-of-production." *Production and Operations Management* 22.1 (2013), pp. 54–70.

[26] Karl Inderfurth and Kampan Mukherjee. "Decision support for spare parts acquisition in post product life cycle." *Central European Journal of Operations Research* 16.1 (2008), pp. 17–42.

[27] Uday S Karmarkar, Sham Kekre, and Sunder Kekre. "The dynamic lot-sizing problem with startup and reservation costs." *Operations Research* 35.3 (1987), pp. 389–398.

[28] Rainer Kleber and Karl Inderfurth. "A heuristic approach for inventory control of spare parts after end-of-production." In: *Logistikmanagement 2007*. Ed. by Andreas Otto and Robert Obermaier. Logistikmanagement. DUV, 2007, pp. 185–200.

[29] Rainer Kleber, Tobias Schulz, and Guido Voigt. "Dynamic buy-back for product recovery in end-of-life spare parts procurement." *International Journal of Production Research* 50.6 (2012), pp. 1476–1488.

[30] Nicholas W Leifker, Philip C Jones, and Timothy J Lowe. "A continuous-time examination of end-of-life parts acquisition with limited customer information." *The Engineering Economist* 57.4 (2012), pp. 284–301.

[31] Nicholas W Leifker, Philip C Jones, and Timothy J Lowe. "Determining optimal order amount for end-of-life parts acquisition with possibility of contract extension." *The Engineering Economist* 59.4 (2014), pp. 259–281.

[32] Jianbin Li et al. "Supply chain coordination through capacity reservation contract and quantity flexibility contract." *Omega* 99 (2021), p. 102195.

[33] Boy Lüthje. "Electronics contract manufacturing: global production and the international division of labor in the age of the internet." *Industry and Innovation* 9.3 (2002), pp. 227–247.

[34] Carme Martínez-Costa et al. "A review of mathematical programming models for strategic capacity planning in manufacturing." *International Journal of Production Economics* 153 (2014), pp. 66–85.

[35] Microsoft. *Annual Report 2021*. [Online; accessed on November 20, 2022]. 2021. URL: `https://www.microsoft.com/investor/reports/ar21/index.html`.

[36] John R Moore Jr. "Forecasting and scheduling for past-model replacement parts." *Management Science* 18.4, Part I (1971), B-200–B-213.

[37] Emin Ozyoruk, Nesim Kohen Erkip, and Çağın Ararat. "End-of-life inventory management problem: Results and insights." *International Journal of Production Economics* 243 (2022), p. 108313.

[38] Sung Il Park and Jong Soo Kim. "A mathematical model for a capacity reservation contract." *Applied Mathematical Modelling* 38.5-6 (2014), pp. 1866–1880.

[39] Morteza Pourakbar and Rommert Dekker. "Customer differentiated end-of-life inventory problem." *European Journal of Operational Research* 222.1 (2012), pp. 44–53.

[40] Morteza Pourakbar, JBG Frenk, and Rommert Dekker. "End-of-Life Inventory Decisions for Consumer Electronics Service Parts." *Production and Operations Management* 21.5 (2012), pp. 889–906.

[41] Morteza Pourakbar, Erwin van der Laan, and Rommert Dekker. "End-of-Life Inventory Problem with Phaseout Returns." *Production and Operations Management* 23.9 (2014), pp. 1561–1576.

[42] Anyan Qi, Hyun-Soo Ahn, and Amitabh Sinha. "To share or not to share? capacity reservation in a shared supplier." *Production and Operations Management* 28.11 (2019), pp. 2823–2840.

[43] Sanmina. *Full Fiscal 2022 Financial Results*. [Online; accessed on November 20, 2022]. 2022. URL: `https://ir.sanmina.com/news/news-details/2022/SANMINA-REPORTS-FOURTH-QUARTER-AND-FULL-FISCAL-2022-FINANCIAL-RESULTS/default.aspx`.

[44] Dogan A Serel. "Capacity reservation under supply uncertainty." *Computers & Operations Research* 34.4 (2007), pp. 1192–1220.

[45] Yuelin Shen and Sean P Willems. "Modeling sourcing strategies to mitigate part obsolescence." *European Journal of Operational Research* 236.2 (2014), pp. 522–533.

[46] Zhenyang Shi and Shaoxuan Liu. "Optimal inventory control and design refresh selection in managing part obsolescence." *European Journal of Operational Research* 287.1 (2020), pp. 133–144.

[47] Tim John Sturgeon. *Turn-key Production Networks: Industry Organization, Economic Development, and the Globalization of Electronics Contract Manufacturing*. PhD Dissertation. University of California, Berkeley, 1999.

[48] Ruud H Teunter and Leonard Fortuin. "End-of-life service." *International Journal of Production Economics* 59.1-3 (1999), pp. 487–497.

[49] Ruud H Teunter and Willem K Klein Haneveld. "Inventory control of service parts in the final phase." *European Journal of Operational Research* 137.3 (2002), pp. 497–511.

[50] Ruud H Teunter and Willem K Klein Haneveld. "The 'final order' problem." *European Journal of Operational Research* 107.1 (1998), pp. 35–44.

[51] John PJ Van Kooten and Tarkan Tan. "The final order problem for repairable spare parts under condemnation." *Journal of the Operational Research Society* 60.10 (2009), pp. 1449–1461.

[52] Arthur F Veinott Jr. "On the Opimality of (s,S) Inventory Policies: New Conditions and a New Proof." *SIAM Journal on Applied Mathematics* 14.5 (1966), pp. 1067–1083.

[53] Arthur F Veinott Jr and Harvey M Wagner. "Computing optimal (s, S) inventory policies." *Management Science* 11.5 (1965), pp. 525–552.

[54] Harvey M Wagner and Thomson M Whitin. "Dynamic version of the economic lot size model." *Management Science* 5.1 (1958), pp. 89–96.

[55] S David Wu, Murat Erkoc, and Suleyman Karabuk. "Managing capacity in the high-tech industry: A review of literature." *The Engineering Economist* 50.2 (2005), pp. 125–158.

[56] Yanyi Xu, Arnab Bisi, and Maqbool Dada. "New structural properties of (s, S) policies for inventory models with lost sales." *Operations Research Letters* 38.5 (2010), pp. 441–449.

Appendix A

Extended Proofs

In this appendix, we present proofs that were omitted from the main body of the dissertation.

A.1 Extended Proof for Chapter 4

Proof of Proposition 4.1

Proposition 4.1. *An optimal constrained order-up-to $S_n^*(I_n)$ exists that minimizes the value function $C_n(S_n|I_n, \boldsymbol{y})$ for arbitrary n and I_n.*

Proof. A global minimum (and minimizer) exists if the value function is continuous, is bounded below by zero, and increases to infinity as the order-up-to level increases to infinity. We next prove that the value function satisfies all three conditions. Recall that the function $C_n(S_n|I_n, \mathbf{y})$ as defined in (4.5) is defined on the left-bounded and right-unbounded interval $[I_n, \infty)$. It is obvious that this function is continuous and that it is bounded below by zero because each term is non-negative. Because the function is left-bounded, it only remains for us to prove that the function goes to ∞ as S_n goes to ∞. Note that the term $c_p(S_n - I_n)$ goes to ∞ as S_n goes to ∞ and all other terms are non-negative. Therefore, the function must attain a minimum value somewhere on its domain. □

A.2 Extended Proof for Chapter 5

Proof of Proposition 5.2

Proposition 5.2. *For an arbitrary period n and starting inventory level I_n, there exists at least one optimal constrained order-up-to level $S_n^*(I_n)$ that minimizes the value function $C_n(S_n, y_n = 1|I_n, \alpha_n = 1)$.*

Proof. A global minimum (and minimizer) exists if the value function is continuous, is bounded below by zero, and increases to infinity as the order-up-to level increases to infinity. We next prove that the value function satisfies all three conditions. Recall that the value function $C_n(S_n, y_n = 1|I_n, \alpha_n = 1)$ is defined on the left-bounded and right-unbounded interval $[I_n, \infty)$. It is obvious that this function is continuous and that it is bounded below by zero because each term is non-negative. Because the function is left-bounded, it only remains for us to prove that the function goes to ∞ as S_n goes to ∞. Note that the term $c_p(S_n - I_n)$ goes to ∞ as S_n goes to ∞ and all other terms are non-negative. Therefore, the function must attain a minimum value somewhere on its domain. □

A.3 Extended Proof for Chapter 8

Proof of Lemma 8.2

Lemma 8.2. *Given that the firm chooses to offer the buy-back in an arbitrary period n, the minimum of the buy-back value function* (8.9) *occurs at a starting inventory level for which the corresponding optimal buy-back quantity is zero.*

Proof. Consider the buy-back value function expressed in (8.9) and repeated below:

$$C_n(B, z_n = 1|I_n, \alpha_n = 0, \beta_n = 1) = r(B) + C_n^*(I_n + B, \alpha_n = 0, \beta_n = 0) \tag{A.1}$$

In this value function, I_n is the starting inventory level in period n, while B is a quantity of inventory that must be acquired via the buy-back. Let $\hat{I} = I_n + B$ and consider

$$C_n(0, z_n = 1|\hat{I}, \alpha_n = 0, \beta_n = 1) = r(0) + C_n^*(\hat{I}, \alpha_n = 0, \beta_n = 0) \tag{A.2}$$

The right-hand side of (A.2) is less than the right-hand side of (A.1) because the second terms on the right-hand sides of both expressions are equal and $r(B) > r(0)$ for any $B > 0$. Therefore (8.9) achieves its minimum at $B = 0$ for some I_n.

□

Appendix B

Complete Results of the Numerical Studies

This appendix contains additional details on the implementation and results of our numerical studies. Section B.1 covers the numerical study presented in Chapter 7, beginning with an explanation of the back-of-the-envelope calculations we performed to inform our choice of parameters and concluding with a summary of results for all 100 problem instances. Section B.2 contains a summary of the results for the numerical study of the dynamic variant with the deterministic buy-back option included in Section 8.4. Finally, Section B.3 contains a summary of the results for the numerical study of the dynamic variant with stochastic buy-back option included in Section 8.6.

B.1 Complete Results of the Dynamic Variant Numerical Study

Back-of-the-Envelope Calculation for Estimating the Timing of the Last Time Buy

When designing the numerical study, we wanted to ensure that we chose a range of parameter values that leads to interesting and varied optimal solutions. To that end, we made use of a simple back-of-the-envelope calculation to estimate the timing of the last time buy to help us ensure that the solutions for our problem set would have a range of last-time-buy periods. To do so, we assume monthly demand is deterministic and constant at the mean over the horizon. This allows us to ignore shortage costs and ignore the variable costs of procurement. We can then express the total capacity reservation and inventory holding costs if we place the last time buy in period n as:

$$C(n) = nc_f + \frac{(N-n)(N-n+1)}{2} c_h D \tag{B.1}$$

where the first term represents the capacity reservation costs and the second term is the total inventory holding cost incurred to satisfy all demand in the final $N - n$ periods. This expression is convex in n, so we can easily find the optimal timing of the last time buy in this deterministic case for the values for c_f, c_s, and D that we selected for our numerical study. Although the optimal timing of the last time buy in the stochastic case will likely differ, the deterministic approximation allows us to predict the optimal timing of the last time buy with sufficient accuracy to ensure that our cost parameters lead to a range of values for the timing of the last time buy. This prediction is confirmed by our numerical study. When we consider the deterministic approximation with a shortage cost of 200 and a per-period demand of 100 for capacity reservation costs of 2000, 4000, 8000, and 16,000, we estimate the last time buy will be placed in periods 11, 9, 6, and 1, respectively. This result is quite similar to the optimal solution to the pre-commit variant for the same cost parameters and *Flat* demand pattern, for which the last time buy periods are 11, 9, 5, and 1 for the same capacity reservation costs, respectively. This similarity gives us confidence that the back-of-the-envelope calculation accurately estimates the timing of the last time buy. Furthermore, the fact that the back-of-the-envelope calculation finds last time buy timing estimates that are dispersed along the time horizon provides justification for our choice of cost parameters.

Summary of Results for All 100 Problem Instances

We now provide a summary of the results for all 100 problem instances in our Chapter 7 numerical study. In Tables B.1 through B.5 we report the results separated into five tables, one for each demand pattern. Within each table, the results are reported for all three problem variants without the buy-back option: the dynamic variant (DP), the heuristic for the dynamic variant (H), and the pre-commit variant (PC). We provide the expected total cost from all three solution methods assuming the beginning-of-horizon inventory level is zero units. We also report the computing time (in seconds) necessary for the completion of each algorithm. We provide the optimal last-time-buy period for only the pre-commit variant, as the last-time-buy period is determined dynamically in the other two problem variants and thus depends upon the trajectory of demand observations. Finally, we report two different measures of the optimality gap relative to the optimal solution of the dynamic variant: (i) the percentage gap considering all costs; and (ii) an estimated percentage gap considering controllable costs, where controllable costs do not include the variable costs to cover expected demand. These are referred to as the *Optimality Gap* and the *Adj. Optimality Gap*, respectively. To assist in interpreting the results, we have color-coded the optimality gap columns, using red to indicate a larger optimality gap. We summarize patterns that we observed in these results in Chapter 7.

Table B.1: Numerical Study Results - *Flat* Demand Pattern

Parameters			Expected Total Cost			LTB Period	Computing Time			Optimality Gap		Adj. Optimality Gap	
Demand	c_f	c_s	DP	H	PC	PC	DP	H	PC	H	PC	H	PC
Flat	2000	200	149506	149553	149586	11	1446	6	9549	0.03%	0.05%	0.16%	0.27%
Flat	2000	400	151970	152141	152261	11	1460	6	9614	0.11%	0.19%	0.54%	0.91%
Flat	2000	600	152855	153050	153051	12	1463	6	9616	0.13%	0.13%	0.59%	0.60%
Flat	2000	800	153503	153688	153689	12	1449	6	9534	0.12%	0.12%	0.55%	0.56%
Flat	2000	1000	154081	154204	154295	12	1447	6	9530	0.08%	0.14%	0.36%	0.63%
Flat	4000	200	169371	169371	169371	9	1447	6	9550	0.00%	0.00%	0.00%	0.00%
Flat	4000	400	173178	173178	173206	10	1461	6	9598	0.00%	0.02%	0.00%	0.05%
Flat	4000	600	174423	174434	174501	10	1451	6	9561	0.01%	0.04%	0.02%	0.14%
Flat	4000	800	175275	175431	175491	10	1450	6	9548	0.09%	0.12%	0.28%	0.39%
Flat	4000	1000	176001	176020	176173	10	1454	6	9541	0.01%	0.10%	0.03%	0.31%
Flat	8000	200	197076	197076	197076	5	1389	6	8929	0.00%	0.00%	0.00%	0.00%
Flat	8000	400	204121	204121	204121	6	1346	6	9017	0.00%	0.00%	0.00%	0.00%
Flat	8000	600	206751	206751	206754	6	1379	6	8880	0.00%	0.00%	0.00%	0.00%
Flat	8000	800	208435	208435	208446	7	1379	6	8945	0.00%	0.01%	0.00%	0.01%
Flat	8000	1000	209558	209558	209570	7	1373	6	8884	0.00%	0.01%	0.00%	0.01%
Flat	16000	200	209495	209495	209495	1	1397	6	8953	0.00%	0.00%	0.00%	0.00%
Flat	16000	400	226062	226062	226062	1	1392	6	8937	0.00%	0.00%	0.00%	0.00%
Flat	16000	600	231091	231091	231091	1	1374	6	8914	0.00%	0.00%	0.00%	0.00%
Flat	16000	800	234147	234147	234147	1	1387	6	8963	0.00%	0.00%	0.00%	0.00%
Flat	16000	1000	236303	236303	236303	1	1376	6	8848	0.00%	0.00%	0.00%	0.00%

Table B.2: Numerical Study Results - *Linear Decline* Demand Pattern

Parameters			Expected Total Cost			LTB Period	Computing Time			Optimality Gap		Adj. Optimality Gap	
Demand	c_f	c_s	**DP**	**H**	**PC**	**PC**	**DP**	**H**	**PC**	**H**	**PC**	**H**	**PC**
Linear Decline	2000	200	147298	147298	147298	9	1436	6	9344	0.00%	0.00%	0.00%	0.00%
Linear Decline	2000	400	149686	149686	149700	10	1451	6	9445	0.00%	0.01%	0.00%	0.05%
Linear Decline	2000	600	150485	150491	150524	10	1431	6	9321	0.00%	0.03%	0.02%	0.13%
Linear Decline	2000	800	151089	151167	151197	10	1454	6	9413	0.05%	0.07%	0.25%	0.35%
Linear Decline	2000	1000	151629	151629	151715	10	1447	6	9354	0.00%	0.06%	0.00%	0.27%
Linear Decline	4000	200	163840	163840	163840	7	1453	6	9429	0.00%	0.00%	0.00%	0.00%
Linear Decline	4000	400	167511	167513	167598	8	1434	6	9323	0.00%	0.05%	0.00%	0.18%
Linear Decline	4000	600	168705	168779	168871	9	1450	6	9421	0.04%	0.10%	0.15%	0.34%
Linear Decline	4000	800	169548	169598	169604	9	1438	6	9341	0.03%	0.03%	0.10%	0.11%
Linear Decline	4000	1000	170201	170201	170228	9	1438	6	9354	0.00%	0.02%	0.00%	0.05%
Linear Decline	8000	200	184487	184487	184487	3	1460	6	9389	0.00%	0.00%	0.00%	0.00%
Linear Decline	8000	400	192122	192122	192123	5	1468	6	9456	0.00%	0.00%	0.00%	0.00%
Linear Decline	8000	600	194604	194604	194610	5	1469	6	9468	0.00%	0.00%	0.00%	0.01%
Linear Decline	8000	800	196203	196207	196219	5	1455	6	9386	0.00%	0.01%	0.01%	0.02%
Linear Decline	8000	1000	197382	197382	197417	5	1465	6	9464	0.00%	0.02%	0.00%	0.04%
Linear Decline	16000	200	195513	195513	195513	1	1468	6	9446	0.00%	0.00%	0.00%	0.00%
Linear Decline	16000	400	209998	209998	209998	1	1450	6	9384	0.00%	0.00%	0.00%	0.00%
Linear Decline	16000	600	215331	215331	215333	1	1445	6	9385	0.00%	0.00%	0.00%	0.00%
Linear Decline	16000	800	218580	218580	218591	1	1465	6	9467	0.00%	0.00%	0.00%	0.01%
Linear Decline	16000	1000	220886	220886	220903	1	1452	6	9384	0.00%	0.01%	0.00%	0.02%

Table B.3: Numerical Study Results - *Exponential Decline* Demand Pattern

Parameters			Expected Total Cost			LTB Period	Computing Time			Optimality Gap		Adj. Optimality Gap	
Demand	c_f	c_s	**DP**	**H**	**PC**	**PC**	**DP**	**H**	**PC**	**H**	**PC**	**H**	**PC**
Exponential Decline	2000	200	147385	147385	147385	9	1542	6	10029	0.00%	0.00%	0.00%	0.00%
Exponential Decline	2000	400	150010	150012	150024	10	1543	6	10061	0.00%	0.01%	0.01%	0.05%
Exponential Decline	2000	600	150825	150831	150865	10	1549	6	10117	0.00%	0.03%	0.02%	0.13%
Exponential Decline	2000	800	151443	151521	151551	10	1534	6	9996	0.05%	0.07%	0.25%	0.34%
Exponential Decline	2000	1000	151998	152007	152084	10	1539	6	10034	0.01%	0.06%	0.03%	0.27%
Exponential Decline	4000	200	161409	161409	161409	5	1535	6	9989	0.00%	0.00%	0.00%	0.00%
Exponential Decline	4000	400	165481	165482	165481	4	1555	6	10061	0.00%	0.00%	0.00%	0.00%
Exponential Decline	4000	600	166989	166989	166990	4	1539	6	10018	0.00%	0.00%	0.00%	0.00%
Exponential Decline	4000	800	168023	168023	168029	7	1559	6	10114	0.00%	0.00%	0.00%	0.01%
Exponential Decline	4000	1000	168776	168776	168782	7	1545	6	10049	0.00%	0.00%	0.00%	0.01%
Exponential Decline	8000	200	174460	174460	174460	2	1509	6	10157	0.00%	0.00%	0.00%	0.00%
Exponential Decline	8000	400	183880	183881	183911	3	1505	6	10112	0.00%	0.02%	0.00%	0.05%
Exponential Decline	8000	600	186637	186637	186663	4	1500	6	10086	0.00%	0.01%	0.00%	0.04%
Exponential Decline	8000	800	188167	188167	188216	4	1513	6	10170	0.00%	0.03%	0.00%	0.07%
Exponential Decline	8000	1000	189330	189345	189393	4	1519	6	10161	0.01%	0.03%	0.02%	0.09%
Exponential Decline	16000	200	183132	183132	183132	1	1528	6	10147	0.00%	0.00%	0.00%	0.00%
Exponential Decline	16000	400	199642	199643	199726	1	1526	6	10147	0.00%	0.04%	0.00%	0.11%
Exponential Decline	16000	600	205804	205866	206244	1	1518	6	10095	0.03%	0.21%	0.07%	0.51%
Exponential Decline	16000	800	209361	209561	210244	1	1519	6	10092	0.10%	0.42%	0.22%	0.99%
Exponential Decline	16000	1000	211466	211719	212388	2	1528	6	10119	0.12%	0.44%	0.28%	1.01%

Table B.4: Numerical Study Results - *Single Peak* Demand Pattern

Parameters			Expected Total Cost			LTB Period	Computing Time			Optimality Gap		Adj. Optimality Gap	
Demand	c_f	c_s	**DP**	**H**	**PC**	**PC**	**DP**	**H**	**PC**	**H**	**PC**	**H**	**PC**
Single Peak	2000	200	148268	148311	148358	10	1621	6	10551	0.03%	0.06%	0.15%	0.32%
Single Peak	2000	400	150694	150791	150813	11	1610	6	10518	0.06%	0.08%	0.31%	0.39%
Single Peak	2000	600	151437	151521	151557	11	1603	6	10479	0.06%	0.08%	0.27%	0.38%
Single Peak	2000	800	152012	152020	152114	11	1625	7	10564	0.01%	0.07%	0.03%	0.32%
Single Peak	2000	1000	152509	152509	152671	11	1598	6	10512	0.00%	0.11%	0.00%	0.50%
Single Peak	4000	200	166118	166118	166132	8	1604	6	10463	0.00%	0.01%	0.00%	0.03%
Single Peak	4000	400	170014	170046	170097	9	1616	6	10573	0.02%	0.05%	0.06%	0.16%
Single Peak	4000	600	171296	171349	171464	9	1614	7	10581	0.03%	0.10%	0.10%	0.33%
Single Peak	4000	800	172180	172260	172440	9	1596	7	10492	0.05%	0.15%	0.15%	0.50%
Single Peak	4000	1000	172912	173030	173233	9	1613	7	10562	0.07%	0.19%	0.22%	0.61%
Single Peak	8000	200	193532	193532	193532	6	1628	6	10310	0.00%	0.00%	0.00%	0.00%
Single Peak	8000	400	201083	201088	201266	7	1635	7	10329	0.00%	0.09%	0.01%	0.23%
Single Peak	8000	600	203485	203586	203956	7	1632	7	10397	0.05%	0.23%	0.12%	0.56%
Single Peak	8000	800	205035	205265	205717	7	1646	6	10368	0.11%	0.33%	0.27%	0.80%
Single Peak	8000	1000	206198	206533	206968	8	1634	7	10363	0.16%	0.37%	0.39%	0.89%
Single Peak	16000	200	212274	212274	212274	1	1634	6	10365	0.00%	0.00%	0.00%	0.00%
Single Peak	16000	400	228568	228568	228568	1	1647	7	10369	0.00%	0.00%	0.00%	0.00%
Single Peak	16000	600	234527	234527	234527	1	1638	7	10372	0.00%	0.00%	0.00%	0.00%
Single Peak	16000	800	238178	238178	238178	1	1635	6	10298	0.00%	0.00%	0.00%	0.00%
Single Peak	16000	1000	240770	240770	240770	1	1647	7	10381	0.00%	0.00%	0.00%	0.00%

Table B.5: Numerical Study Results - *Double Peak* Demand Pattern

Parameters			Expected Total Cost			LTB Period	Computing Time			Optimality Gap		Adj. Optimality Gap	
Demand	c_f	c_s	**DP**	**H**	**PC**	**PC**	**DP**	**H**	**PC**	**H**	**PC**	**H**	**PC**
Double Peak	2000	200	147772	147772	147776	10	1095	4	7151	0.00%	0.00%	0.00%	0.02%
Double Peak	2000	400	150318	150318	150333	10	1483	6	9700	0.00%	0.01%	0.00%	0.05%
Double Peak	2000	600	151407	151596	151449	10	1476	6	9644	0.13%	0.03%	0.60%	0.13%
Double Peak	2000	800	152168	152247	152247	10	1465	6	9561	0.05%	0.05%	0.25%	0.25%
Double Peak	2000	1000	152684	152693	152771	10	1470	6	9606	0.01%	0.06%	0.03%	0.27%
Double Peak	4000	200	165890	166005	166040	9	1471	6	9570	0.07%	0.09%	0.25%	0.33%
Double Peak	4000	400	169198	169656	169316	9	1469	6	9586	0.27%	0.07%	0.93%	0.24%
Double Peak	4000	600	170476	170743	170750	9	1472	6	9563	0.16%	0.16%	0.53%	0.54%
Double Peak	4000	800	171390	171514	171767	9	1473	7	9570	0.07%	0.22%	0.24%	0.73%
Double Peak	4000	1000	171997	172179	172448	9	1486	6	9624	0.11%	0.26%	0.35%	0.87%
Double Peak	8000	200	190624	190624	190641	4	1484	6	9169	0.00%	0.01%	0.00%	0.02%
Double Peak	8000	400	199055	199055	199055	5	1480	6	9224	0.00%	0.00%	0.00%	0.00%
Double Peak	8000	600	202241	202241	202241	5	1479	6	9226	0.00%	0.00%	0.00%	0.00%
Double Peak	8000	800	204299	204299	204299	5	1462	6	9238	0.00%	0.00%	0.00%	0.00%
Double Peak	8000	1000	205826	205826	205826	5	1461	6	9277	0.00%	0.00%	0.00%	0.00%
Double Peak	16000	200	206725	206725	206725	1	1467	6	9203	0.00%	0.00%	0.00%	0.00%
Double Peak	16000	400	224447	224447	224447	1	1462	6	9216	0.00%	0.00%	0.00%	0.00%
Double Peak	16000	600	231159	231159	231159	1	1465	6	9204	0.00%	0.00%	0.00%	0.00%
Double Peak	16000	800	235280	235280	235280	1	1466	6	9222	0.00%	0.00%	0.00%	0.00%
Double Peak	16000	1000	238208	238208	238208	1	1467	6	9185	0.00%	0.00%	0.00%	0.00%

B.2 Complete Results of the Numerical Study for the Dynamic Variant with Deterministic Buy-Back

This appendix provides a summary of the results for all 100 problem instances in our Section 8.4 numerical study. In Tables B.6 through B.10 we report the results separated into five tables, one for each demand pattern. Within each table, the results are reported for both solution methods: the dynamic variant with buy-back (DP w/ BB) and the heuristic (H). We provide the expected total cost from both solution methods assuming the beginning-of-horizon inventory level is zero units. We also report the computing time (in seconds) necessary for the completion of each algorithm. The computing time is affected by the number of economically-sensible buy-back quantities for the given parameters. In particular, we only need to consider buy-back quantities with marginal costs less than c_s, so the viable range of buy-back quantities is smallest for the lowest value of $c_s = 200$. However, as this is true for both the dynamic and heuristic algorithms, the computing time comparison between them is fair for any given value of c_s. We report the optimality gap of the heuristic solution relative to the optimal solution of the dynamic variant. (We omit the adjusted optimality gap used in Chapter 7 as expected variable procurement costs differ for the buy-back and no-buy-back options.) We also report the expected total cost for the dynamic variant without a buy-back option (DP w/o BB) from the numerical study covered in Chapter 7. This enables us to calculate the expected benefit of the buy-back option, which we present as a percentage benefit, an absolute benefit, and as the ratio of the benefit to the capacity reservation cost. To assist in interpreting the results, we have color-coded the optimality gap column (using darker shades of red to indicate a larger optimality gap) and the expected benefit columns (using darker shades of green to indicate a larger expected benefit). We summarize patterns that we observed in these results in Section 8.4.

Table B.6: Numerical Study Results with Deterministic Buy-Back - *Flat* Demand Pattern

Parameters			Expected Total Cost			Computing Time		Optimality Gap	Expected Benefit of Buy-Back Option		
Demand	c_f	c_s	DP w/ BB	H	DP w/o BB	DP	H	Heuristic	% Benefit	Benefit	Benefit/c_f
Flat	2000	200	147706	147706	149506	1565	88	0.00%	1.22%	1800	0.9
Flat	2000	400	149957	149957	151970	1677	185	0.00%	1.34%	2012	1.0
Flat	2000	600	150803	150803	152855	1669	185	0.00%	1.36%	2052	1.0
Flat	2000	800	151427	151427	153503	1678	184	0.00%	1.37%	2076	1.0
Flat	2000	1000	152045	152045	154081	1683	183	0.00%	1.34%	2036	1.0
Flat	4000	200	164195	164195	169371	1571	90	0.00%	3.15%	5176	1.3
Flat	4000	400	167427	167427	173178	1668	183	0.00%	3.44%	5751	1.4
Flat	4000	600	168698	168698	174423	1666	184	0.00%	3.39%	5725	1.4
Flat	4000	800	169557	169557	175275	1667	183	0.00%	3.37%	5718	1.4
Flat	4000	1000	170232	170232	176001	1667	183	0.00%	3.39%	5770	1.4
Flat	8000	200	179265	179265	197076	1527	86	0.00%	9.94%	17811	2.2
Flat	8000	400	185833	185833	204121	1638	180	0.00%	9.84%	18288	2.3
Flat	8000	600	188066	188066	206751	1636	183	0.00%	9.94%	18685	2.3
Flat	8000	800	189439	189439	208435	1627	181	0.00%	10.03%	18996	2.4
Flat	8000	1000	190432	190432	209558	1642	180	0.00%	10.04%	19126	2.4
Flat	16000	200	187265	187265	209495	1540	88	0.00%	11.87%	22229	1.4
Flat	16000	400	194155	194155	226062	1630	182	0.00%	16.43%	31908	2.0
Flat	16000	600	196598	196598	231091	1648	179	0.00%	17.54%	34492	2.2
Flat	16000	800	198072	198072	234147	1641	179	0.00%	18.21%	36075	2.3
Flat	16000	1000	199143	199143	236303	1645	178	0.00%	18.66%	37160	2.3

Table B.7: Numerical Study Results with Deterministic Buy-Back - *Linear Decline* Demand Pattern

Parameters			**Expected Total Cost**			**Computing Time**		**Optimality Gap**	**Expected Benefit of Buy-Back Option**		
Demand	c_f	c_s	**DP w/ BB**	**H**	**DP w/o BB**	**DP**	**H**	**Heuristic**	**% Benefit**	**Benefit**	**Benefit/c_f**
Linear Decline	2000	200	145021	145021	147298	1519	88	0.00%	1.55%	2277	1.1
Linear Decline	2000	400	147170	147170	149686	1630	180	0.00%	1.68%	2516	1.3
Linear Decline	2000	600	147922	147922	150485	1609	182	0.00%	1.70%	2563	1.3
Linear Decline	2000	800	148486	148486	151089	1605	180	0.00%	1.72%	2603	1.3
Linear Decline	2000	1000	148988	148988	151629	1616	178	0.00%	1.74%	2642	1.3
Linear Decline	4000	200	157771	157771	163840	1523	86	0.00%	3.70%	6068	1.5
Linear Decline	4000	400	160573	160573	167511	1614	184	0.00%	4.14%	6937	1.7
Linear Decline	4000	600	161530	161530	168705	1603	182	0.00%	4.25%	7175	1.8
Linear Decline	4000	800	162160	162160	169548	1616	180	0.00%	4.36%	7388	1.8
Linear Decline	4000	1000	162689	162689	170201	1619	180	0.00%	4.41%	7512	1.9
Linear Decline	8000	200	170295	170295	184487	1591	85	0.00%	7.69%	14192	1.8
Linear Decline	8000	400	174897	174897	192122	1673	182	0.00%	8.97%	17226	2.2
Linear Decline	8000	600	176455	176455	194604	1704	179	0.00%	9.33%	18149	2.3
Linear Decline	8000	800	177394	177394	196203	1681	181	0.00%	9.59%	18809	2.4
Linear Decline	8000	1000	178051	178051	197382	1697	181	0.00%	9.79%	19331	2.4
Linear Decline	16000	200	178295	178295	195513	1569	85	0.00%	8.81%	17218	1.1
Linear Decline	16000	400	183433	183433	209998	1682	182	0.00%	12.65%	26565	1.7
Linear Decline	16000	600	185234	185234	215331	1698	180	0.00%	13.98%	30097	1.9
Linear Decline	16000	800	186298	186298	218580	1689	183	0.00%	14.77%	32282	2.0
Linear Decline	16000	1000	187027	187027	220886	1606	188	0.00%	15.33%	33859	2.1

Table B.8: Numerical Study Results with Deterministic Buy-Back - *Exponential Decline* Demand Pattern

Parameters			Expected Total Cost			Computing Time		Optimality Gap	Expected Benefit of Buy-Back Option		
Demand	c_f	c_s	**DP w/ BB**	**H**	**DP w/o BB**	**DP**	**H**	**Heuristic**	**% Benefit**	**Benefit**	**Benefit/c_f**
Exponential Decline	2000	200	144140	144140	147385	1617	87	0.00%	2.20%	3245	1.6
Exponential Decline	2000	400	146730	146732	150010	1733	189	0.00%	2.19%	3279	1.6
Exponential Decline	2000	600	147561	147561	150825	1711	190	0.00%	2.16%	3265	1.6
Exponential Decline	2000	800	148165	148165	151443	1719	193	0.00%	2.16%	3278	1.6
Exponential Decline	2000	1000	148708	148708	151998	1732	191	0.00%	2.16%	3290	1.6
Exponential Decline	4000	200	153280	153280	161409	1619	87	0.00%	5.04%	8129	2.0
Exponential Decline	4000	400	156368	156369	165481	1734	189	0.00%	5.51%	9113	2.3
Exponential Decline	4000	600	157490	157490	166989	1738	191	0.00%	5.69%	9499	2.4
Exponential Decline	4000	800	158248	158248	168023	1742	189	0.00%	5.82%	9775	2.4
Exponential Decline	4000	1000	158828	158828	168776	1733	191	0.00%	5.89%	9948	2.5
Exponential Decline	8000	200	162001	162001	174460	1638	93	0.00%	7.14%	12459	1.6
Exponential Decline	8000	400	165886	165886	183880	1776	201	0.00%	9.79%	17995	2.2
Exponential Decline	8000	600	167274	167274	186637	1313	160	0.00%	10.37%	19363	2.4
Exponential Decline	8000	800	168105	168105	188167	1760	201	0.00%	10.66%	20062	2.5
Exponential Decline	8000	1000	168674	168674	189330	1764	201	0.00%	10.91%	20656	2.6
Exponential Decline	16000	200	170001	170001	183132	1656	95	0.00%	7.17%	13130	0.8
Exponential Decline	16000	400	173890	173890	199642	1759	200	0.00%	12.90%	25752	1.6
Exponential Decline	16000	600	175280	175280	205804	1762	200	0.00%	14.83%	30524	1.9
Exponential Decline	16000	800	176111	176111	209361	1770	199	0.00%	15.88%	33249	2.1
Exponential Decline	16000	1000	176681	176681	211466	1776	199	0.00%	16.45%	34785	2.2

Table B.9: Numerical Study Results with Deterministic Buy-Back - *Single Peak* Demand Pattern

Parameters			**Expected Total Cost**			**Computing Time**		**Optimality Gap**	**Expected Benefit of Buy-Back Option**		
Demand	c_f	c_s	**DP w/ BB**	**H**	**DP w/o BB**	**DP**	**H**	**Heuristic**	**% Benefit**	**Benefit**	**Benefit/c_f**
Single Peak	2000	200	145968	145968	148268	1751	95	0.00%	1.55%	2300	1.2
Single Peak	2000	400	148038	148038	150694	1870	209	0.00%	1.76%	2656	1.3
Single Peak	2000	600	148736	148736	151437	1877	207	0.00%	1.78%	2701	1.4
Single Peak	2000	800	149276	149276	152012	1877	207	0.00%	1.80%	2736	1.4
Single Peak	2000	1000	149788	149788	152509	1891	209	0.00%	1.78%	2721	1.4
Single Peak	4000	200	161033	161033	166118	1765	96	0.00%	3.06%	5085	1.3
Single Peak	4000	400	163682	163682	170014	1897	208	0.00%	3.72%	6333	1.6
Single Peak	4000	600	164665	164665	171296	1880	206	0.00%	3.87%	6631	1.7
Single Peak	4000	800	165336	165336	172180	1868	206	0.00%	3.98%	6844	1.7
Single Peak	4000	1000	165884	165884	172912	1892	207	0.00%	4.06%	7027	1.8
Single Peak	8000	200	183030	183030	193532	1762	91	0.00%	5.43%	10502	1.3
Single Peak	8000	400	187562	187562	201083	1871	209	0.00%	6.72%	13521	1.7
Single Peak	8000	600	188909	188909	203485	1879	210	0.00%	7.16%	14576	1.8
Single Peak	8000	800	189753	189753	205035	1892	208	0.00%	7.45%	15282	1.9
Single Peak	8000	1000	190369	190369	206198	1882	210	0.00%	7.68%	15829	2.0
Single Peak	16000	200	191030	191030	212274	1761	92	0.00%	10.01%	21244	1.3
Single Peak	16000	400	199675	199675	228568	1868	209	0.00%	12.64%	28893	1.8
Single Peak	16000	600	202331	202331	234527	1874	206	0.00%	13.73%	32196	2.0
Single Peak	16000	800	203891	203891	238178	1878	206	0.00%	14.40%	34287	2.1
Single Peak	16000	1000	204949	204949	240770	1873	205	0.00%	14.88%	35821	2.2

Table B.10: Numerical Study Results with Deterministic Buy-Back - *Double Peak* Demand Pattern

Parameters			Expected Total Cost			Computing Time		Optimality Gap	Expected Benefit of Buy-Back Option		
Demand	c_f	c_s	**DP w/ BB**	**H**	**DP w/o BB**	**DP**	**H**	**Heuristic**	**% Benefit**	**Benefit**	**Benefit/c_f**
Double Peak	2000	200	145339	145399	147772	1599	88	0.04%	1.65%	2433	1.2
Double Peak	2000	400	147599	148124	150318	1690	184	0.36%	1.81%	2719	1.4
Double Peak	2000	600	148591	148870	151407	1707	184	0.19%	1.86%	2815	1.4
Double Peak	2000	800	149275	149351	152168	1698	183	0.05%	1.90%	2893	1.4
Double Peak	2000	1000	149714	149807	152684	1702	182	0.06%	1.94%	2970	1.5
Double Peak	4000	200	161364	161365	165890	1586	88	0.00%	2.73%	4525	1.1
Double Peak	4000	400	163741	164094	169198	1714	181	0.22%	3.22%	5456	1.4
Double Peak	4000	600	164764	164873	170476	1704	184	0.07%	3.35%	5712	1.4
Double Peak	4000	800	165404	165404	171390	1689	183	0.00%	3.49%	5985	1.5
Double Peak	4000	1000	165867	165868	171997	1694	184	0.00%	3.56%	6130	1.5
Double Peak	8000	200	176554	176554	190624	1639	94	0.00%	7.38%	14069	1.8
Double Peak	8000	400	183065	183065	199055	1737	199	0.00%	8.03%	15991	2.0
Double Peak	8000	600	184624	184624	202241	1740	197	0.00%	8.71%	17617	2.2
Double Peak	8000	800	185541	185541	204299	1745	199	0.00%	9.18%	18758	2.3
Double Peak	8000	1000	186200	186209	205826	1749	199	0.00%	9.54%	19626	2.5
Double Peak	16000	200	184554	184554	206725	1644	93	0.00%	10.72%	22171	1.4
Double Peak	16000	400	193444	193444	224447	1736	198	0.00%	13.81%	31003	1.9
Double Peak	16000	600	196844	196844	231159	1739	197	0.00%	14.85%	34316	2.1
Double Peak	16000	800	198940	198940	235280	1748	199	0.00%	15.45%	36339	2.3
Double Peak	16000	1000	200439	200439	238208	1758	198	0.00%	15.86%	37769	2.4

B.3 Complete Results of the Numerical Study for the Dynamic Variant with Stochastic Buy-Back

This appendix provides a summary of the results for all 100 problem instances in our Section 8.6 numerical study. In Tables B.11 through B.15 we report the results separated into five tables, one for each demand pattern. We report the expected total cost from the dynamic solution method assuming the beginning-of-horizon inventory level is zero units. We also report the computing time (in seconds) necessary for the completion of the algorithm. We also report the expected total cost for the deterministic buy-back variant (DP w/ Det BB) from the numerical study in Section 8.4 to enable easy comparison with the stochastic buy-back variant (DP w/ Stoch BB). Finally, we report the expected cost increase due to the stochasticity of the buy-back yields as a percentage of the expected cost of the deterministic model. To assist in interpreting the results, we have color-coded the cost increase column, using darker shades of red to indicate a larger percentage cost increase. We summarize patterns that we observed in these results in Section 8.6.

Table B.11: Numerical Study Results with Stochastic Buy-Back - *Flat* Demand Pattern

Parameters			**Expected Total Cost**		**Expected Cost**
Demand	c_f	c_s	**DP w/ Stoch BB**	**DP w/ Det BB**	**Increase (%)**
Flat	2000	200	147942	147706	0.16%
Flat	2000	400	150393	149957	0.29%
Flat	2000	600	151308	150803	0.34%
Flat	2000	800	152033	151427	0.40%
Flat	2000	1000	152589	152045	0.36%
Flat	4000	200	165121	164195	0.56%
Flat	4000	400	168724	167427	0.77%
Flat	4000	600	169862	168698	0.69%
Flat	4000	800	170708	169557	0.68%
Flat	4000	1000	171329	170232	0.64%
Flat	8000	200	180855	179265	0.89%
Flat	8000	400	190835	185833	2.69%
Flat	8000	600	193780	188066	3.04%
Flat	8000	800	195273	189439	3.08%
Flat	8000	1000	196346	190432	3.11%
Flat	16000	200	188855	187265	0.85%
Flat	16000	400	200361	194155	3.20%
Flat	16000	600	204492	196598	4.02%
Flat	16000	800	206876	198072	4.44%
Flat	16000	1000	208506	199143	4.70%

Table B.12: Numerical Study Results with Stochastic Buy-Back - *Linear Decline* Demand Pattern

Parameters			**Expected Total Cost**		**Expected Cost**
Demand	c_f	c_s	**DP w/ Stoch BB**	**DP w/ Det BB**	**Increase (%)**
Linear Decline	2000	200	145424	145021	0.28%
Linear Decline	2000	400	147792	147170	0.42%
Linear Decline	2000	600	148452	147922	0.36%
Linear Decline	2000	800	149009	148486	0.35%
Linear Decline	2000	1000	149467	148988	0.32%
Linear Decline	4000	200	158879	157771	0.70%
Linear Decline	4000	400	162240	160573	1.04%
Linear Decline	4000	600	163404	161530	1.16%
Linear Decline	4000	800	164181	162160	1.25%
Linear Decline	4000	1000	164789	162689	1.29%
Linear Decline	8000	200	171752	170295	0.86%
Linear Decline	8000	400	179035	174897	2.37%
Linear Decline	8000	600	181096	176455	2.63%
Linear Decline	8000	800	182313	177394	2.77%
Linear Decline	8000	1000	183185	178051	2.88%
Linear Decline	16000	200	179752	178295	0.82%
Linear Decline	16000	400	188347	183433	2.68%
Linear Decline	16000	600	191274	185234	3.26%
Linear Decline	16000	800	192937	186298	3.56%
Linear Decline	16000	1000	194069	187027	3.77%

Table B.13: Numerical Study Results with Stochastic Buy-Back - *Exponential Decline* Demand Pattern

Parameters			**Expected Total Cost**		**Expected Cost**
Demand	c_f	c_s	**DP w/ Stoch BB**	**DP w/ Det BB**	**Increase (%)**
Exponential Decline	2000	200	144858	144140	0.50%
Exponential Decline	2000	400	147587	146730	0.58%
Exponential Decline	2000	600	148345	147561	0.53%
Exponential Decline	2000	800	148960	148165	0.54%
Exponential Decline	2000	1000	149455	148708	0.50%
Exponential Decline	4000	200	154315	153280	0.68%
Exponential Decline	4000	400	158570	156368	1.41%
Exponential Decline	4000	600	160094	157490	1.65%
Exponential Decline	4000	800	161081	158248	1.79%
Exponential Decline	4000	1000	161790	158828	1.86%
Exponential Decline	8000	200	163010	162001	0.62%
Exponential Decline	8000	400	169388	165886	2.11%
Exponential Decline	8000	600	171509	167274	2.53%
Exponential Decline	8000	800	172833	168105	2.81%
Exponential Decline	8000	1000	173779	168674	3.03%
Exponential Decline	16000	200	171010	170001	0.59%
Exponential Decline	16000	400	177770	173890	2.23%
Exponential Decline	16000	600	180324	175280	2.88%
Exponential Decline	16000	800	181884	176111	3.28%
Exponential Decline	16000	1000	182963	176681	3.56%

Table B.14: Numerical Study Results with Stochastic Buy-Back - *Single Peak* Demand Pattern

Parameters			**Expected Total Cost**		**Expected Cost**
Demand	c_f	c_s	**DP w/ Stoch BB**	**DP w/ Det BB**	**Increase (%)**
Single Peak	2000	200	146264	145968	0.20%
Single Peak	2000	400	148486	148038	0.30%
Single Peak	2000	600	149256	148736	0.35%
Single Peak	2000	800	149844	149276	0.38%
Single Peak	2000	1000	150355	149788	0.38%
Single Peak	4000	200	161685	161033	0.41%
Single Peak	4000	400	164827	163682	0.70%
Single Peak	4000	600	165996	164665	0.81%
Single Peak	4000	800	166674	165336	0.81%
Single Peak	4000	1000	167220	165884	0.81%
Single Peak	8000	200	184598	183030	0.86%
Single Peak	8000	400	189777	187562	1.18%
Single Peak	8000	600	191531	188909	1.39%
Single Peak	8000	800	192620	189753	1.51%
Single Peak	8000	1000	193410	190369	1.60%
Single Peak	16000	200	192598	191030	0.82%
Single Peak	16000	400	204837	199675	2.59%
Single Peak	16000	600	208480	202331	3.04%
Single Peak	16000	800	210497	203891	3.24%
Single Peak	16000	1000	211847	204949	3.37%

Table B.15: Numerical Study Results with Stochastic Buy-Back - *Double Peak* Demand Pattern

Parameters			**Expected Total Cost**		**Expected Cost**
Demand	c_f	c_s	**DP w/ Stoch BB**	**DP w/ Det BB**	**Increase (%)**
Double Peak	2000	200	145856	145339	0.36%
Double Peak	2000	400	148253	147599	0.44%
Double Peak	2000	600	149244	148591	0.44%
Double Peak	2000	800	149987	149275	0.48%
Double Peak	2000	1000	150419	149714	0.47%
Double Peak	4000	200	161951	161364	0.36%
Double Peak	4000	400	164758	163741	0.62%
Double Peak	4000	600	165941	164764	0.71%
Double Peak	4000	800	166738	165404	0.81%
Double Peak	4000	1000	167237	165867	0.83%
Double Peak	8000	200	178303	176554	0.99%
Double Peak	8000	400	185741	183065	1.46%
Double Peak	8000	600	187729	184624	1.68%
Double Peak	8000	800	188959	185541	1.84%
Double Peak	8000	1000	189832	186200	1.95%
Double Peak	16000	200	186303	184554	0.95%
Double Peak	16000	400	199976	193444	3.38%
Double Peak	16000	600	203725	196844	3.50%
Double Peak	16000	800	205674	198940	3.38%
Double Peak	16000	1000	206989	200439	3.27%