OPTIMAL STOPPING PROBLEMS IN OPERATIONS MANAGEMENT

A DISSERTATION
SUBMITTED TO THE DEPARTMENT OF MANAGEMENT SCIENCE AND ENGINEERING
AND THE COMMITTEE ON GRADUATE STUDIES
OF STANFORD UNIVERSITY
IN PARTIAL FULFILLMENT OF THE REQUIREMENTS
FOR THE DEGREE OF
DOCTOR OF PHILOSOPHY

Sechan Oh
March 2010

This dissertation is online at: http://purl.stanford.edu/yx778cf5733

I certify that I have read this dissertation and that, in my opinion, it is fully adequate in scope and quality as a dissertation for the degree of Doctor of Philosophy.

Ali Ozer, Primary Adviser

I certify that I have read this dissertation and that, in my opinion, it is fully adequate in scope and quality as a dissertation for the degree of Doctor of Philosophy.

Warren Hausman, Co-Adviser

I certify that I have read this dissertation and that, in my opinion, it is fully adequate in scope and quality as a dissertation for the degree of Doctor of Philosophy.

Benjamin Van Roy

Approved for the Stanford University Committee on Graduate Studies.

Patricia J. Gumport, Vice Provost Graduate Education

This signature page was generated electronically upon submission of this dissertation in electronic format. An original signed hard copy of the signature page is on file in University Archives.

Abstract

Optimal stopping problems determine the time to terminate a process to maximize expected rewards. Such problems are pervasive in the areas of operations management, marketing, statistics, finance, and economics. This dissertation provides a method that characterizes the structure of the optimal stopping policy for a general class of optimal stopping problems. It also studies two important optimal stopping problems arising in Operations Management.

In the first part of the dissertation, we provide a method to characterize the structure of the optimal stopping policy for the class of discrete-time optimal stopping problems. Our method characterizes the structure of the optimal policy for some stopping problems for which conventional methods fail. Our method also simplifies the analysis of some existing results. Using the method, we determine sufficient conditions that yield threshold or control-band type optimal stopping policies. The results also help characterize parametric monotonicity of optimal thresholds and provide bounds for them.

In the second part of the dissertation, we first generalize the Martingale Model of Forecast Evolution to account for multiple forecasters who forecast demand for the same product. The result enables us to consistently model the evolution of forecasts generated by two forecasters who have asymmetric demand information. Using the forecast evolution model, we next study a supplier's problem of eliciting credible forecast information from a manufacturer when both parties obtain asymmetric demand information over multiple periods. For better capacity planning, the supplier designs and offers a screening contract that ensures the manufacturer's credible information sharing. By delaying to offer this incentive mechanism, the supplier can

obtain more information. This delay, however, may increase (resp., or decrease) the degree of information asymmetry between the two firms, resulting in a higher (resp., or lower) cost of screening. The delay may also increase capacity costs. Considering all such trade-offs, the supplier has to determine how to design a mechanism to elicit credible forecast information from the manufacturer and when to offer this incentive mechanism.

In the last part of the dissertation, we study a manufacturer's problem of determining the time to introduce a new product to the market. Conventionally, manufacturing firms determine the time to introduce a new product to the market long before launching the product. The timing decision involves considerable risk because manufacturing firms are uncertain about competing firms' market entry timing and the outcome of production process development activities at the time when they make the decision. As a solution for reducing such risk, we propose a dynamic market entry strategy under which the manufacturer makes decisions about market entry timing and process improvements in response to the evolution of uncertain factors. We show that the manufacturer can reduce profit variability and increase average profit by employing this dynamic strategy. Our study also characterizes the industry conditions under which the dynamic strategy is most effective.

Acknowledgements

I would like first thank my advisor Prof. Özalp Özer for his insightful advices throughout my doctoral studies. He has been a great mentor in every aspects of my graduate life, and this dissertation would have not been possible without his dedicated teaching.

I would also like to thank Prof. Warren H. Hausman for his great support. He has provided valuable personal advices during difficult times. His own doctoral research was the most important reference for the forecast evolution model that is presented in the third chapter of my dissertation.

I am also very grateful to Prof. Benjamin Van Roy, Prof. Sunil Kumar, and Prof. Thomas A. Weber for serving in the committee of my dissertation defense. They provided great suggestions to substantially improve this dissertation. They have also taught me courses such as dynamic programming, economic theory, and revenue management in the early stage of my graduate study. I also wish to acknowledge and thank Samsung Scholarship Foundation for financial support.

Graduate school would have been very difficult without my friends. I give special thanks to my MS&E and Korean friends who have filled my graduate life with joy. I was very fortunate to have met such brilliant and fun friends in the course of my graduate study.

Last but not least, I would like to thank my family for their unconditional love and support. All my accomplishments to this point are the outcome of my parents' sacrifice and dedication.

Contents

List of Tables

List of Figures

Chapter 1

Introduction

Optimal stopping problems determine the time to stop a process in order to maximize expected rewards. Such problems appear frequently in the areas of economics, finance, statistics, marketing and operations management. For example, a stock option holder faces the problem of determining the time to exercise the option in order to maximize the expected income. As another example, employers face the problem of determining the time to stop interviewing job candidates in order to hire the best candidate. This dissertation studies two important optimal stopping problems arising in Operations Management. It also provides a method that characterizes the structure of the optimal stopping policy for a general class of optimal stopping problems.

Optimal stopping problems often have simple threshold or control-band type optimal stopping policies. For example, for an American put option, a threshold policy under which the option holder exercises the stock option if the current stock price is below a certain threshold is optimal. Such structural properties of the optimal stopping policy are important for three reasons. First, knowing the structure of the optimal policy provides managerial insights. They provide actionable policies that a decision maker can follow to optimize her objective. Second, structural results also enable one to develop efficient numerical algorithms to solve optimal stopping problems. Finally, structural results are important when the optimal stopping problem is part of a higher-level and/or larger scale optimization problem. For these reasons, most research papers that study optimal stopping problems provide structural properties of the optimal stopping policy if such structures exist (see, e.g., Chen 1970, Yao and Zheng 1999b, Ben-Ameur et al. 2002, Alagoz et al. 2004).

In Chapter 2, we provide a method that characterizes the structure of the optimal stopping policy for the class of discrete-time optimal stopping problems. Our method is based on structural properties of the benefit function, which we define as the difference between the reward of continuing the process and the reward of stopping the process. To characterize its structure, we establish the benefit function's recursive relation with the one-step benefit function, which we define as the difference between the rewards of stopping at the next period and the current period. Using this recursive relation and the stochastic monotonicity of state-transition, we determine several sufficient conditions that yield threshold or control-band type optimal stopping policies. We show that our method can characterize the structure of the optimal policy of some stopping problems for which conventional methods fail. We also show that the method simplifies the analysis of some existing results.

Next, in Chapter 3, we study an optimal stopping problem faced by a supplier who invests in new capacity. For a timely delivery, the supplier has to secure component capacity prior to receiving a firm order from a product manufacturer. The supplier relies on the demand forecast for his capacity decision. However, the manufacturer often has other forward-looking information because of her superior relationship with or proximity to the market and expert opinion about her own product. To elicit the manufacturer's private information, the supplier needs to design and offer a screening contract. As the sales season approaches, both the supplier and the manufacturer can update their demand forecasts over time. Hence, by delaying to offer the screening contract, the supplier can reduce the demand uncertainty that he faces. However, delaying the capacity decision is not always beneficial for the supplier. For example, the delay may increase the degree of information asymmetry between the two firms if the manufacturer obtains more information than the supplier over time. Capacity costs may also increase as the supplier delays the capacity decision because of a tighter deadline for building capacity. By considering all such trade-offs, the supplier needs to determine when to offer a screening contract and how to design the screening contract to maximize his profit. In Chapter 3, we develop an optimal stopping problem to solve this problem.

The supplier's decision problem consists of two stages. The first stage is an optimal stopping problem that determines the optimal time to offer a screening contract, and the second stage is a mechanism design problem for forecast information sharing. Using the method that we develop in Chapter 2, we establish the optimality of a control band policy that prescribes when to offer an optimal incentive mechanism. Under this policy, the supplier offers a menu of contracts if the supplier's demand forecast falls within the control band. We also provide structural properties of the optimal screening contract and explicitly show how the optimal contract depends on the demand forecast and how the timing decision affects the mechanism design problem. Through numerical studies, we characterize the environment in which the supplier should offer the contract late or early. By comparing the profits of the dynamic strategy with those of a static one in which the supplier offers a contract in a fixed period, we show that the supplier can significantly improve his profit by optimally determining the time to offer a contract. However, the results also show that this dynamic strategy can reduce the total supply chain efficiency.

Modeling the aforementioned stopping problem requires a forecast evolution model that describes forecast sequences made by two decision makers. To develop such a model, we extend the Martingale Model of Forecast Evolution (MMFE) framework to the cases with multiple decision makers in Chapter 3. The MMFE is a general model that describes the evolution of forecasts arising from many statistical and judgment-based forecasting methods. Due to its descriptive power and generality, researchers have used the MMFE in many studies that involve dynamic forecast updates such as inventory control and production planning (e.g., Heath and Jackson 1994, Aviv 2001, Gallego and Özer 2001, Toktay and Wein 2001, Altug and Muharremoglu 2009, Iida and Zipkin 2009, Schoenmeyr and Graves 2009). Our extension enables the MMFE framework to model several plausible forecast evolution scenarios that involve multiple decision makers in a consistent way.

Finally, in Chapter 4, we study an optimal stopping problem faced by a manufacturer who introduces a new product to the market. When determining the timing for introducing the new product, the manufacturer takes into consideration the trade-off between the time-to-market and the completeness of the production processes. On the one hand, the manufacturer can attain a large market share by entering the market early. On the other hand, the manufacturer can improve the production process for the new product by investing more time in process design, which results in a reduction of production costs. However, the manufacturer are uncertain about both the timing of the competitors' market entry and the outcome of production process development activities. For this reason, the manufacturer needs to dynamically make the timing decision depending on the competitors' movements and the readiness of his own production process. We formulate the manufacturer's problem as an optimal stopping problem.

The manufacturer's decision process also consists of two stages. The first stage is an optimal stopping problem that determines the optimal time to introduce a new product to the market and optimal investment decisions to improve the production process. The second stage is a production and pricing decision problem that determines the production quantity and the sales prices for the new product. Using the method that we develop in Chapter 2, we establish the optimality of threshold-type market entry policies that prescribe the optimal time to introduce the new product. We also characterize structural properties of the optimal production and pricing decisions. By comparing to a static market entry decision, we show that the dynamic market entry decision yields a higher and less variable profit. Our study also characterizes when the value of the dynamic market entry is the greatest.

Chapter 2

Characterizing the Structure of the Optimal Stopping Policy

2.1. Introduction

Optimal stopping problems are determining the time to terminate a process to maximize expected rewards given the initial state of the process. Such problems appear frequently in the operations, marketing, finance and economics literature. Some examples are the problem of determining the time to exercise a stock option, to sell or purchase an asset, and to introduce a new product. Optimal stopping problems are rarely solvable in a closed form and they require computational methods. Hence, researchers often try to characterize the structure of the optimal stopping policy that determines when to stop the process based on the state of the process at each decision epoch. When possible, researchers also provide monotonicity results (comparative statics) regarding the optimal policy parameters. Such structural results enable actionable policies that a decision maker can follow to maximize rewards. They also help develop efficient numerical algorithms to solve the problem. This chapter provides a method to characterize the structure of an optimal stopping policy for the class of discrete-time optimal stopping problems. This chapter also determines sufficient conditions that yield simple threshold or control-band type stopping policies. These conditions are presented in eight propositions that can be used to characterize optimal stopping policy for various application areas.

Structural properties of the optimal stopping policy are helpful for three reasons. First, knowing the structure of the optimal policy provides managerial insights. They provide actionable policies that a decision maker can follow to optimize her objective. For example, exercising an American put option (stopping the process) is optimal when the current stock price (current state) is below a certain threshold. Another example is from Alagöz et al. (2007a,b) who establish an optimal organ-transplantation policy for a patient with end-stage liver disease. They show that given the current health condition of the patient, transplanting an organ is optimal if the quality[1] of the offered organ is above a certain level. Second, structural results also enable one to develop efficient numerical algorithms to solve optimal stopping problems (as in Yao and Zheng 1999b, Ben-Ameur et al. 2002, Wu and Fu 2003). For example, Wu and Fu (2003) first establish the optimality of a threshold-type exercise policy for an American-Asian option. Using this structure, they develop a computationally efficient simulation-based algorithm. Finally, structural results are important when the optimal stopping problem is part of a higher-level and/or larger scale optimization problem (Terwiesch and Xu 2004, Anily and Grosfeld-Nir 2006). For example, Terwiesch and Xu (2004) study a manufacturer's problem of pricing prototypes and the final product. To do so, they first model a customer's purchasing decision as an optimal stopping problem. They show that the customer's optimal purchase policy has a threshold structure. Using this result, the authors formulate the manufacturer's optimal pricing problem. For aforementioned reasons, several other papers also characterize structural properties of optimal stopping policies for various stopping problems (such as Chow et al. 1964, Cox et al. 1979, Chen et al. 1998, and Boyaci and Özer 2009).

The above observations motivated us to identify a universal method that can help determine the structure of optimal stopping policy for problems arising in various fields. The method also helps us to identify sufficient conditions that yield simple optimal policies. Before discussing our new method, we describe the two approaches currently used in the literature. The first approach is to verify whether the problem

[1]Quality of an organ is measured, for example, by the age of the donor (Alagoz et al. 2007b).

satisfies the monotone-case condition (Chow et al. 1971). This approach first determines the set of states at which the reward of immediate stopping is greater than the expected reward of stopping at the next period. The policy that stops the process if the current state is in this set is known as the one-step look-ahead policy. Chow et al. (1971) show that if this set is closed almost surely[2], then the one-step look-ahead policy is optimal and call such a case the monotone-case. Since the one-step look-ahead policy is easy to compute and implement, researchers often try to verify whether the problem satisfies the monotone-case condition (see, for example, Stadje 1991, Hui et al. 2008). However, most optimal stopping problems do not satisfy the monotone-case condition because the one-step look-ahead policy, a myopic policy, is not optimal in general. The second and the most common approach is based on the dynamic programming formulation of the optimal stopping problem. This approach determines structural properties of the *value function* of the dynamic program to characterize the optimal stopping policy (see, for example, Wu and Fu 2003, Babich and Sobel 2004, and Alagoz et al. 2007a). As we will illustrate in §2.3, these two approaches, although helpful, do not always yield the structure of the optimal stopping policy.

This chapter provides a different approach to characterize the structure of the optimal stopping policy. Our method is based on structural properties of the *benefit function*, which we define as the difference between the reward of continuing the process and the reward of stopping the process. To characterize its structure, we establish the benefit function's recursive relation with the *one-step benefit function*, which we define as the difference between the rewards of stopping at the next period and the current period. Next, we determine sufficient conditions that yield a threshold or control-band type optimal stopping policy using this recursive relation and the stochastic monotonicity of state-transition. We show that our method can characterize the structure of the optimal policy of some optimal stopping problems for which the two approaches discussed above fail to do so. We also show that the

[2]That is, when the current state is in this set, the future states will be in this set with probability one.

method simplifies the analysis of existing results. One can use the method to characterize the structure of the optimal policy before numerically solving the optimal stopping problem. The results also make the sufficient conditions that yield simple optimal stopping policies transparent and easier to determine.

For the analysis of our method, we use stochastic monotonicities of parameterized random variables. Stochastic monotonicities are widely used in the analysis of stochastic objective functions (Shaked and Shanthikumar 2007). For example, Athey (2000) uses them to characterize the structure of the objective functions that arise in economics. Smith and McCardle (2002) use them to characterize the structure of the *value function* of a Markov decision process (MDP). Even though optimal stopping problems are a sub-class of general MDPs, our method is not related to those in Smith and McCardle (2002). Our method for characterizing the structure of the optimal policy is based on the *benefit function*. The benefit function is the difference between the two reward functions: the expected reward of continuing the process and the reward of stopping the process. In contrast, the value function is the maximum of the two. Hence, the *benefit function* is different from the *value function* of MDPs. We show that the properties of the benefit function directly imply the structure of the optimal stopping policy when the structure of the value function and the result of Smith and McCardle (2002) do not. In addition, the benefit function and the value function have different recursions as we discuss in §2.2. Our approach is effective because there are only two actions for optimal stopping problems. Despite its effectiveness and simplicity, this method has been neglected (see §2.3.2 for some examples). Our research formalizes this method in a general framework.

There has also been extensive research on the monotonicity of optimal control. Several researchers have provided sufficient conditions that yield monotone optimal policies (Altman and Stidham 1995 and Glasserman and Yao 1994). However, the monotonicity of the optimal policies in this stream of research is based on the theory of ordered optimal solutions by Topkis (1978) and its generalizations. For example, Altman and Stidham (1995) show that a threshold policy is optimal for stationary Markov decision processes with two actions when the reward function is submodular in state and action and the state transition is stochastically monotone. We remark that *none* of the discrete-time optimal stopping problems satisfies Altman and Stidham (1995)'s assumptions set forth for binary Markov decision problems. Our method and sufficient conditions are based on the zero-crossing points of the benefit function and do not require the submodularity of reward functions. Hence, the results of this chapter can be applied to the problems that do not satisfy the sufficient conditions provided by the literature on monotone optimal control. We refer to Glasserman and Yao (1994) for additional references on this literature.

Due to the importance of stopping problems, extensive research has been done (Chow et al. 1971, Shiryaev 1978, Tsitsiklis and Van Roy 1999, Peskir and Shiryaev 2006). This line of research characterizes optimal stopping times and optimal rewards under various general assumptions. However, this literature has not focused on characterizing the structure of the optimal stopping policy. Our study, in contrast, focuses on a method to characterize the structure of optimal stopping policies and determines sufficient conditions to obtain them.

The rest of this chapter is organized as follows. In §2.2, we define the optimal stopping problem and propose the method to characterize the structure of the optimal stopping policy. In §2.3, we provide an example for which only our method can characterize the structure of the optimal policy and another example for which our method simplifies the analysis of an existing result. In §2.4, we provide sufficient conditions that yield a threshold or control-band type optimal stopping policy and provide monotonicity results and bounds for the optimal policy. In §2.5, we consider optimal stopping problems with additional decisions other than the stopping decision. In §2.6, we consider infinite-horizon optimal stopping problems. In §2.7, we provide example applications to facilitate the use of the proposed method. In §2.8, we conclude. The objective of this chapter is to introduce an easy and useful method to researchers in broad application areas. Hence, in Appendix A.1, we provide a theorem and examples that one can use to verify stochastic monotonicities of state transition models. Proofs not provided right after the propositions are deferred to Appendix A.2.

2.2. Optimal Stopping Problem and the Two-Step Method

Let $\{x_t | t = 1, 2, \ldots, T\}$ be a Markov process that evolves in a state space $X \subset \mathbb{R}^d$, defined on a probability space $(\Omega, \mathcal{F}, \mathcal{P})$. We denote the σ-algebra generated by $\{x_1, x_2, \ldots, x_t\}$ by $\mathcal{F}_t \subset \mathcal{F}$. A stopping time τ is a random variable that takes values in $\{1, 2, \ldots, T\}$ and satisfies $\{\omega \in \Omega | \tau(\omega) \leq t\} \in \mathcal{F}_t$ for all $t \leq T$. We denote the set of all such stopping times by U^T.

At each period $t \in \{1, 2, \ldots, T\}$, a decision maker observes the state x_t and decides whether to continue or stop a process. If the decision maker decides to continue, she attains a reward $C_t(x_t)$ and the state evolves. If she stops, then the decision maker attains a reward $S_t(x_t)$, and the problem is terminated. Without loss of generality, we assume that stopping is a forced decision at period T.[3] In addition, we assume that reward functions $C_t : \mathbb{R}^d \to \mathbb{R}$ and $S_t : \mathbb{R}^d \to \mathbb{R}$ are integrable, i.e., $E|S_t(x_t)| < \infty$ and $E|C_t(x_t)| < \infty$ for every $t < \infty$. The objective is to determine the optimal time to stop the process in order to maximize the total discounted rewards with a discount factor $\alpha \in (0, 1]$. This problem can be formulated as

$$V^*(x) \equiv \sup_{\tau \in U^T} E\left[\sum_{t=1}^{\tau-1} \alpha^{t-1} C_t(x_t) + \alpha^{\tau-1} S_\tau(x_\tau) \big| x_1 = x\right]. \tag{2.1}$$

The optimal value function $V^*(x)$ corresponds to the total expected reward when the optimal stopping time achieves the supremum in (2.1), and the initial state is x. Note that the optimal stopping time and the optimal value function are well-defined. However, this formulation does not help characterize an actionable policy that a decision maker can follow to maximize her expected reward. Hence, we provide a dynamic programming (DP) recursion that specifies an optimal action for each state at each period.

[3] In some problems, not stopping and receiving a reward of 0 at the end of the decision horizon is an option for the decision maker. For such problems, one can introduce a fictitious period $t = T + 1$ with $S_{T+1}(x_{T+1}) = 0$, and enforce the stopping at period $t = T + 1$. Therefore, the forced stopping assumption is not restrictive.

Let U_t^T be the set of stopping times that satisfy $\tau \in [t, T]$. Then, we define

$$V_t(x) \equiv \sup_{\tau \in U_t^T} E\left[\sum_{u=t}^{\tau-1} \alpha^{u-t} C_u(x_u) + \alpha^{\tau-t} S_\tau(x_\tau) \big| x_t = x\right],$$

which indicates the total expected rewards from period t when the process has not yet stopped and the current state is x. If the decision maker stops the process at period t, the reward is $S_t(x)$. If she continues, the expected reward is $C_t(x) + \alpha E[V_{t+1}(x_{t+1})|x_t = x]$. Therefore, $V_t(x)$ satisfies the following DP recursion (Theorem 3.2 of Chow et al. 1971):

$$V_t(x_t) = \max\{S_t(x_t), C_t(x_t) + \alpha E[V_{t+1}(\tilde{x}_{t+1}(x_t))]\}, \quad t < T, \tag{2.2}$$

where $V_T(x_T) = S_T(x_T)$. We denote $E[V_{t+1}(x_{t+1})|x_t]$ by $E[V_{t+1}(\tilde{x}_{t+1}(x_t))]$ to emphasize its functional dependence on x_t. Note that $V_t(x)$ is the *value function* of the DP recursion, and the *optimal* value function satisfies $V^*(x) = V_1(x)$. An optimal policy stops the process at period t if $S_t(x_t) \geq C_t(x_t) + \alpha E[V_{t+1}(\tilde{x}_{t+1}(x_t))]$.

To date, the most common approach for characterizing the structure of the optimal stopping policy is based on characterizing structural properties of the corresponding value function. However, this approach is not always useful in characterizing the structure of the optimal stopping policy. Note that the optimal stopping decision is based on the relative values of $S_t(x_t)$ and $C_t(x_t) + \alpha E[V_{t+1}(\tilde{x}_{t+1}(x_t))]$. For example, the information that $V_{t+1}(x_{t+1})$ is increasing in x_{t+1} is not useful in characterizing the structure of the optimal policy, when both $S_t(x_t)$ and $C_t(x_t) + \alpha E[V_{t+1}(\tilde{x}_{t+1}(x_t))]$ are increasing in x_t. In contrast, the structural properties of

$$B_t(x_t) \equiv \alpha E[V_{t+1}(\tilde{x}_{t+1}(x_t))] + C_t(x_t) - S_t(x_t)$$

help directly characterize the optimal stopping policy because the optimal policy at each period $t < T$ is to stop the process if $B_t(x_t) \leq 0$ and to continue, otherwise. We refer to this function as the *benefit function*. It is the expected benefit of delaying the stopping decision at period t. For optimal stopping problems, determining structural properties of the *benefit* function provides an easier way to characterize the optimal policy than determining structural properties of the *value* function. This approach is possible because optimal stopping problems have only two options to choose at each period: stop or continue.

We also define the *one-step look-ahead*, or in short, the *one-step benefit function*

$$M_t(x_t) \equiv \alpha E[S_{t+1}(\tilde{x}_{t+1}(x_t))] + C_t(x_t) - S_t(x_t),$$

which indicates the expected benefit of delaying the stopping decision at period t without considering the possible benefit of delaying the decision beyond period $t+1$. These two functions are closely related. The following recursion formalizes this relationship:

$$\begin{aligned} B_t(x_t) &= \alpha E[V_{t+1}(\tilde{x}_{t+1}(x_t))] + C_t(x_t) - S_t(x_t) \\ &= \alpha E[\max\{S_{t+1}(\tilde{x}_{t+1}(x_t)), B_{t+1}(\tilde{x}_{t+1}(x_t)) + S_{t+1}(\tilde{x}_{t+1}(x_t))\}] \\ &\quad + C_t(x_t) - S_t(x_t) \\ &= M_t(x_t) + \alpha E[\max\{0, B_{t+1}(\tilde{x}_{t+1}(x_t))\}], \ t < T-1, \qquad (2.3) \\ B_{T-1}(x_{T-1}) &= M_{T-1}(x_{T-1}). \end{aligned}$$

This recursive relationship highlights two important observations. First, structural properties of the *benefit function* is closely related to that of the *one-step* benefit function. Second, these properties also depend on the functional form of the state transition $\tilde{x}_{t+1}(x_t)$ and the corresponding transition probabilities. Naturally, establishing structural properties of $B_t(x_t)$ from this recursive relationship involves an inductive argument. Suppose $M_t(x_t)$ has a certain structural property for every t. Then $B_{T-1}(x_{T-1})$ inherits its structural property by definition. We will show that when the state transition $\tilde{x}_{t+1}(x_t)$ has an appropriate stochastic monotonicity, $B_t(x_t)$ also has the same property with $M_t(x_t)$ for all other periods. The two-step method is based on this idea. In particular,

First, we determine structural properties of $M_t(x_t)$ from $C_t(x_t)$, $S_t(x_t)$

and $\tilde{x}_{t+1}(x_t)$. Next, we verify whether $\tilde{x}_{t+1}(x_t)$ has the stochastic monotonicity that enables $B_t(x_t)$ to inherit the structural property of $M_t(x_t)$.

Then, structural properties of $B_t(x_t)$ directly imply the structure of the optimal stopping policy. For example, when $M_t(x_t)$ is increasing[4] in x_t, a stochastically increasing state transition $\tilde{x}_{t+1}(x_t)$ carries the increasing property to $B_t(x_t)$. Hence, a threshold policy is optimal. We provide such sufficient conditions in §2.4, which corresponds to the second step. We provide several examples to illustrate the first step in §2.7.

2.3. Two Example Applications

We introduce two examples that support the importance of our method discussed in the previous section. The first example illustrates that the conventional value-function-based approach for characterizing the structure of the optimal stopping policy does not always work. This example also does not satisfy the monotone-case condition. Hence, the conventional methods discussed previously fail to characterize the optimal policy for this first example. However, the proposed method successfully characterizes the structure of the optimal policy. The second example illustrates how the two-step method substantially simplifies the analysis of an existing result.

2.3.1 Time-to-Market Model

Consider a firm that decides when to introduce a new product. When the firm introduces the product earlier than competitors, it captures a larger market share. However, an early introduction results in high production costs and low profit margins due to low manufacturing yields. Hence, the firm needs to determine the optimal time to enter the market. Suppose that the total market demand D is deterministic. There are T periods in which the firm can introduce the new product. At each period $t \in \{1, 2, \ldots, T\}$, the firm decides whether to delay the market entry depending on the number of competitors who are already in the market, $x_t \in \{0, 1, \ldots\}$. Let $v(x_t)$

[4] We use the terms increasing and decreasing in the weak sense; i.e., increasing means non-decreasing.

be the market share of the firm when she enters the market after the x_tth competitor. It is decreasing concave in x_t. That is, as more competitors enter the market the firm loses more market share. Let p be the sales price of the product and c_t be the unit production cost when the firm enters the market at period t. The discounted profit margin $\alpha^{t-1}(p-c_t)$ increases with t due to higher manufacturing yields[5]. If the firm enters the market at period t, she attains a reward of $S_t(x_t) = v(x_t)(p-c_t)D$. If she delays the market entrance at period t, then ξ_t more competitors enter the market, and the state evolves as $\tilde{x}_{t+1}(x_t) = x_t + \xi_t$. The random variable ξ_t is independent of x_t.

This problem can be formulated as an optimal stopping problem. The value function is derived as $V_t(x) = \sup_{\tau \in U_t^T} E\left[\alpha^{\tau-1} S_\tau(x_\tau) \middle| x_t = x\right]$ and the benefit function is derived as $B_t(x_t) = \alpha E[V_{t+1}(\tilde{x}_{t+1}(x_t))] - S_t(x_t)$. The DP recursion is given by $V_t(x_t) = \max\{S_t(x_t), \alpha E[V_{t+1}(\tilde{x}_{t+1}(x_t))]\}$ with $V_T(x_T) = S_T(x_T)$. Figure 2.1 shows an example of $V_t(x_t)$, $B_t(x_t)$, and $M_t(x_t)$ for $t = 13$. The problem setting for this example is $T = 15$, $\alpha = 1$, $v(x) = 0.9 - 0.85e^{0.1(x-15)}$, $p = 5$, $c_t = 4 - 0.1t - 0.005t^2$, $D = 10$, and ξ_t is a Bernoulli r.v. with 0.7.

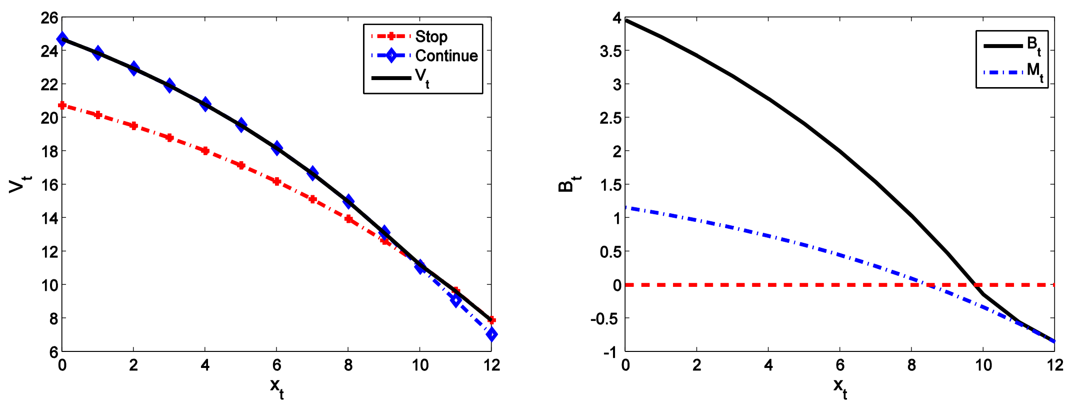

Figure 2.1: Value function and the benefit function of the time-to-market problem

First note from Figure 2.1 that $M_t(x)$ and $B_t(x)$ do not cross the zero line at the same point. The monotone-case condition of Chow et al. (1971) is also not satisfied.

[5]This example is a simplified version of a problem faced by Hitachi GST, a global provider of hard disk drives, as discussed in Özer and Uncu (2008).

Hence, the one-step look ahead policy is not optimal. Second, note that although the reward function $S_t(x_t)$ is decreasing concave in x_t for every t, the value function $V_t(x_t)$ is not concave in x_t. The decreasing property of $V_t(x_t)$ in x_t does not provide any information about the optimal stopping policy, either. The value function is the maximum of the reward of stopping and the reward of continuing, which are both decreasing in x_t. Hence, structural properties of the value function (and hence the methods of Smith and McCardle 2002) cannot characterize the structure of the optimal stopping policy in this case. However, by using the two-step method, we can easily verify the decreasing property of $B_t(x_t)$, which establishes the optimality of a threshold policy; i.e., the firm should enter the market at period t if $x_t \geq \overline{x}_t$ for a certain threshold $\overline{x}_t$. We provide the complete analysis in §2.7.

2.3.2 American-Asian Option

Our second example is the problem of pricing an American-Asian option studied in Ben-Ameur et al. (2002) and Wu and Fu (2003). A stock option is a financial derivative security that promises the option holder a payoff when the holder exercises the option. The payoff depends on the future prices of an underlying stock and the agreed upon strike price K. A holder of the American-Asian option can exercise the stock option at fixed periods $t \in \{1, 2, \ldots, T+1\}$.[6] The payoff depends on the average price of the underlying stock, where the average is taken for the stock prices at periods $1, 2, \ldots, t$. Let $x_{t,1}$ be the current stock price at period t and $x_{t,2}$ be the average stock price at period t. If the option holder exercises the option at period $t \leq T$, she receives a reward $S_t(x_t) = x_{t,2} - K$. If the option holder does not exercise the option at period t, the stock price changes as $\tilde{x}_{t+1,1}(x_t) = \xi x_{t,1}$, where ξ is a log-normal random variable. Accordingly, the average stock price is updated as $\tilde{x}_{t+1,2}(x_t) = \frac{t x_{t,2} + \xi x_{t,1}}{t+1}$. If the option is not exercised in any periods, then $S_{T+1}(\cdot) = 0$. By using the risk-neutral measure (Black and Scholes 1973, Harrison and Kreps 1979) of ξ, the option pricing problem can be formulated as an optimal stopping problem. The price of the option is the optimal value function of the stopping problem, and the optimal exercise policy

[6] Period $T+1$ is a fictitious period, which gives the option holder the right not to exercise the option at period T.

is the optimal stopping policy. The discount factor is $\alpha = e^{-r}$, where r is the risk-free interest rate. The value function is derived as $V_t(x) = \sup_{\tau \in U_t^T} E\left[\alpha^{\tau-t} S_\tau(x_\tau) \middle| x_t = x\right]$ and the benefit function is derived as $B_t(x_t) = \alpha E[V_{t+1}(\tilde{x}_{t+1}(x_t))] - S_t(x_t)$. The DP recursion is given by $V_t(x_t) = \max\{S_t(x_t), \alpha E[V_{t+1}(\tilde{x}_{t+1}(x_t))]\}$ with $V_{T+1}(x_{T+1}) = 0$.

Ben-Ameur et al. (2002) and Wu and Fu (2003) independently determine the structure of the optimal policy based on the properties of the *value* function. For example, Proposition 1 (in Ben-Ameur et al. 2002) verifies that $S_t(x_t)$ and $\alpha E[V_{t+1}(\tilde{x}_{t+1}(x_t))]$ are increasing and convex in $x_{t,1}$ and $x_{t,2}$. However, this information does not characterize the structure of the optimal policy. Instead, both of these papers establish the optimality of a state-dependent threshold policy by verifying that the increasing rate of $\alpha E[V_{t+1}(\tilde{x}_{t+1}(x_t))]$ in $x_{t,2}$ is less than or equal to 1. Although elegant, establishing the structure of the optimal policy using the value function requires a lengthy analysis that involves three propositions and one lemma (§4 in Ben-Ameur et al. 2002). However, as we will show in §2.7, the two-step method can be used to easily verify that the benefit function $B_t(x_t)$ is decreasing in $x_{t,2}$. This result directly implies the optimality of a *state-dependent* threshold policy. Under this policy, the option holder optimally exercises the option at period t if $x_{t,2} \geq \overline{x}_{t,2}(x_{t,1})$ for certain thresholds $\overline{x}_{t,2}(x_{t,1})$. We provide the complete analysis in §2.7.

2.4. Conditions that Imply the Structure of the Optimal Policy

We provide sufficient conditions on $M_t(x_t)$ and $\tilde{x}_{t+1}(x_t)$ that together imply the structure of the benefit function $B_t(x_t)$ and the optimal stopping policy. We also establish some monotonicity results and bounds for the optimal policy parameters. The result of this section corresponds to the second step of the two-step method.[7]

[7]We remark that the results of this section are not related to those in in Chow et al. (1971) and they do not imply the optimality of the one-step look-ahead policy. This policy calls for stopping the process when the reward of instant stopping is greater than the expected reward of stopping at the next period, i.e., when $M_t(x_t) \leq 0$. The one-step look-ahead policy is generally sub-optimal. Chow et al. (1971) provide sufficient conditions under which the one-step look-ahead policy is optimal and refer to it as the monotone-case. In the monotone-case, $\{x : M_t(x) \leq 0\} = \{x : B_t(x) \leq 0\}$. The

For the case of a multi-dimensional state space, we denote the ith element of the vector x_t by $x_{t,i}$ and the $d-1$ dimensional vector excluding the element $x_{t,i}$ from x_t by $x_{t,-i}$. Similarly, $x_{t,-(i,j)}$ denotes the $d-2$ dimensional vector excluding elements $x_{t,i}$ and $x_{t,j}$ from x_t. We write $(x'_{t,i}, x_{t,-i})$ for the state $(x_{t,1}, x_{t,2}, \ldots, x'_{t,i}, \ldots, x_{t,d})$ and write $(x'_{t,i}, x''_{t,j}, x_{t,-(i,j)})$ for the state $(x_{t,1}, x_{t,2}, \ldots, x'_{t,i}, \ldots, x''_{t,j}, \ldots, x_{t,d})$.

We also define the stopping set of period t as $\{x \in X : B_t(x) \leq 0\}$. It is the set of states for which the optimal policy stops the process at period t. For a single-dimensional state space problem, we define $\overline{x}_t \equiv \sup\{x \in X : B_t(x) \leq 0\}$ and $\underline{x}_t \equiv \inf\{x \in X : B_t(x) \leq 0\}$. Similarly, for a multi-dimensional state space problem, we define $\overline{x}_{t,i}(x_{t,-i}) \equiv \sup\{x_{t,i} : B_t(x_{t,i}, x_{t,-i}) \leq 0, x_t \in X\}$ and $\underline{x}_{t,i}(x_{t,-i}) \equiv \inf\{x_{t,i} : B_t(x_{t,i}, x_{t,-i}) \leq 0, x_t \in X\}$.

2.4.1 Univariate Benefit Functions

Consider optimal stopping problems with a single-dimensional state space and an increasing or decreasing one-step benefit function. Before stating the first proposition, we define the stochastic monotonicity necessary for the analysis. We follow the definitions in Shaked and Shanthikumar (2007) for all stochastic monotonicities.

Definition 2.1. *A set of random variables $\{\tilde{x}(\theta), \theta \in \mathbb{R}\}$ is stochastically increasing in θ if $E[u(\tilde{x}(\theta))]$ is increasing in θ for all increasing functions u.*

We note that many common Markov process models have this property. In Appendix A.1, we provide a theorem and examples that one can use to verify stochastic monotonicities of state transitions. Using this definition, we provide the first sufficient condition.

Proposition 2.1. *When $M_t(x_t)$ is increasing [resp., decreasing] in x_t, and $\tilde{x}_{t+1}(x_t)$ is stochastically increasing in x_t for every t, then the following statements are true for every t:*

1. *$B_t(x_t)$ is increasing [resp., decreasing] in x_t.*

conditions we provide in this section do not imply that these two sets are identical.

2. *A threshold policy that stops the process if $x_t \leq \overline{x}_t$ [resp., $x_t \geq \underline{x}_t$] is optimal.*

Proof. The proof is based on an induction argument. Consider the increasing one-step benefit function case. At period $t = T-1$, we have $B_{T-1}(x) = M_{T-1}(x)$. Hence, $B_t(x)$ is increasing in x for $t = T-1$. Next assume for the induction argument that $B_{t+1}(x_{t+1})$ is increasing in x_{t+1}. The composition of an increasing function and $\max\{0, x\}$ is also increasing, hence, $\max\{0, B_{t+1}(x)\}$ is an increasing function. Because the state transition $\tilde{x}_{t+1}(x_t)$ is stochastically increasing in x_t, $E[\max\{0, B_{t+1}(\tilde{x}_{t+1}(x_t))\}]$ is increasing in x_t. Because the increasing property is closed under summation, $B_t(x_t) = M_t(x_t) + \alpha E[\max\{0, B_{t+1}(\tilde{x}_{t+1}(x_t))\}]$ is increasing in x_t, which concludes the induction hypothesis and the proof of Part 1 for the increasing $M_t(x)$ case.

To prove Part 2, we define two sets $Q = \{x \in X : x \leq \overline{x}_t\}$ and $Q^* = \{x \in X : B_t(x) \leq 0\}$. We prove that $Q = Q^*$. When the state space is discrete or $B_t(x_t)$ is a continuous function, $\overline{x}_t$ satisfies $B_t(\overline{x}_t) = 0$. Then, every $x \in Q$ satisfies $B_t(x) \leq B_t(\overline{x}_t) = 0$ because $B_t(x_t)$ is increasing. Hence, $Q \subset Q^*$. Conversely, for every $x \in Q^*$, we have $x \leq \overline{x}_t$ by the definition of $\overline{x}_t$. Hence, $Q^* \subset Q$, which implies that $Q = Q^*$. Therefore, the optimal stopping policy stops the process at period t if $x_t \in Q$, i.e., if $x_t \leq \overline{x}_t$. Note that when the state space is continuous and $B_t(x_t)$ has a discontinuous point, it is possible that $B_t(\overline{x}_t) > 0$. In this case, the optimal policy stops the process if $x_t < \overline{x}_t$ instead of $x_t \leq \overline{x}_t$. However, the structural result does not change. Hence, we assume throughout this chapter that $B_t(x_t)$ is continuous when the state space is continuous. The decreasing case can be proved in a similar way. □

Recall that the recursive relationship in Equation (2.3) has $\max\{0, \cdot\}$. In general, $\max\{0, f(x)\}$ does not preserve structural properties of $f(x)$. Increasing, decreasing and convex properties are among the few properties that are preserved under this operation. Even concavity is not preserved. Other properties, such as subadditivity in x_t are preserved, but those properties are not useful in characterizing the structure of the optimal stopping policy. Note that the structure of the optimal policy does not depend on the maximum point of $B_t(x_t)$, but depends on the pattern of $B_t(x_t)$

crossing 0. Hence, we focus on increasing, decreasing, and convex properties.

Next, consider optimal stopping problems with a single-dimensional state space and a convex one-step benefit function. For the next result, we need a stochastically convex state transition.

Definition 2.2. *A set of random variables* $\{\tilde{x}(\theta), \theta \in \mathbb{R}\}$ *is stochastically convex in* θ *if* $E[u(\tilde{x}(\theta))]$ *is convex in* θ *for all convex functions* u*.*

Using this definition, we provide the second sufficient condition.

Proposition 2.2. *When* $M_t(x_t)$ *is convex in* x_t*, and* $\tilde{x}_{t+1}(x_t)$ *is stochastically convex in* x_t *for every* t*, then the following statements are true for every* t*:*

1. $B_t(x_t)$ *is convex in* x_t*.*
2. *A control-band policy that stops the process if* $x_t \in [\underline{x}_t, \overline{x}_t]$ *is optimal.*

2.4.2 Multivariate Benefit Functions with a Partially Dependent State Transition

Consider optimal stopping problems with a d-dimensional state space. We consider a partially dependent state transition in which there exists an element i such that the state transition $\tilde{x}_{t+1,-i}(x_t)$ is independent of $x_{t,i}$. In other words, state $x_{t,i}$ affects *only* the ith element of the next period's state. In this case, the state transition can be expressed as $\tilde{x}_{t+1}(x_t) = (\tilde{x}_{t+1,i}(x_t), \tilde{x}_{t+1,-i}(x_{t,-i}))$. Note that $\tilde{x}_{t+1,i}(x_t)$ can still depend on $x_{t,-i}$ in addition to depending $x_{t,i}$. Note also that a special case of the partially dependent state transition is the fully independent state transition case, in which the state transition can be expressed as $\tilde{x}_{t+1}(x_t) = (\tilde{x}_{t+1,1}(x_{t,1}), \tilde{x}_{t+1,2}(x_{t,2}), \ldots, \tilde{x}_{t+1,d}(x_{t,d}))$.

To better understand this case, recall the American-Asian option pricing problem discussed in §2.3.2. The state transition of the current stock price is $\tilde{x}_{t+1,1}(x_t) = \xi x_{t,1}$. This update is stochastically increasing in $x_{t,1}$ but is independent of $x_{t,2}$. However, the state transition of the average stock price $\tilde{x}_{t+1,2}(x_t) = \frac{t x_{t,2} + \xi x_{t,1}}{t+1}$ is stochastically increasing in both $x_{t,1}$ and $x_{t,2}$. Next, we provide sufficient conditions for the optimality of a state-dependent threshold or control-band policy.

Proposition 2.3. *When $M_t(x_t)$ is increasing [resp., decreasing] in $x_{t,i}$, $\tilde{x}_{t+1,i}(x_t)$ is stochastically increasing in $x_{t,i}$ and $\tilde{x}_{t+1,-i}(x_t)$ is independent of $x_{t,i}$ for every t, then the following statements are true for every t:*

1. *$B_t(x_t)$ is increasing [resp., decreasing] in $x_{t,i}$.*

2. *A state-dependent threshold policy that stops the process when $x_{t,i} \leq \overline{x}_{t,i}(x_{t,-i})$ [resp., $x_{t,i} \geq \underline{x}_{t,i}(x_{t,-i})$] is optimal for each $x_{t,-i}$.*

Proposition 2.4. *When $M_t(x_t)$ is convex in $x_{t,i}$, $\tilde{x}_{t+1,i}(x_t)$ is stochastically convex in $x_{t,i}$ and $\tilde{x}_{t+1,-i}(x_t)$ is independent of $x_{t,i}$ for every t, then the following statements are true for every t:*

1. *$B_t(x_t)$ is convex in $x_{t,i}$.*

2. *A state-dependent control-band policy that stops the process when $x_{t,i} \in [\underline{x}_{t,i}(x_{t,-i}), \overline{x}_{t,i}(x_{t,-i})]$ is optimal for each $x_{t,-i}$.*

2.4.3 Multivariate Benefit Functions with a Dependent State Transition

Elements of the state transition $\tilde{x}_{t+1}(x_t)$ are dependent on each other when the ith element of the next period state $\tilde{x}_{t+1,i}(x_t)$ depends on the ith element of the current state and also on the other elements of the current state. An interesting analysis can be applied to the case where the state transition $\tilde{x}_{t+1}(x_t)$ is stochastically increasing in x_t. Note that the state x_t has multiple dimensions in this case. Hence, we need a more general definition of the stochastically increasing property.

Definition 2.3. *A set of random vectors $\{\tilde{x}(\theta), \theta \in \mathbb{R}^d\}$ of dimension m is stochastically increasing in $\theta \in \mathbb{R}^d$ if $E[u(\tilde{x}(\theta))]$ is increasing for all increasing functions $u : \mathbb{R}^m \to \mathbb{R}$.*

Given this definition, we have the following result.

Proposition 2.5. *When $M_t(x_t)$ is increasing[resp., decreasing] in each $x_{t,i}$ and $\tilde{x}_{t+1}(x_t)$ is stochastically increasing in $x_t \in X$ for every t, then the following statements are true for every t.*

1. *$B_t(x_t)$ is increasing [resp., decreasing] in each $x_{t,i}$.*

2. *A state-dependent threshold policy that stops the process when $x_{t,i} \leq \overline{x}_{t,i}(x_{t,-i})$ [resp., $x_{t,i} \geq \underline{x}_{t,i}(x_{t,-i})$] is optimal for each i.*

At first sight, the conditions that the one-step benefit function is increasing in $x_{t,i}$ for all elements i and that the state transition is stochastic increasing appear to be somewhat restrictive. Consider, for example, a two-dimensional state space in which a higher value of $x_{t,2}$ leads to a lower value of $x_{t,1}$ in the next period. In this case, one can redefine the state variable as $y_{t,2} = -x_{t,1}$ and $y_{t,1} = x_{t,1}$, which makes the state transition $\tilde{y}_{t,1}(y_t)$ stochastically increasing in both $y_{t,1}$ and $y_{t,2}$. Similarly, if an initial formulation of $M_t(x_t)$ is increasing in $x_{t,1}$ and decreasing in $x_{t,2}$, a similar transformation makes the one-step benefit function $M_t(y_t)$ an increasing function. Therefore, the increasing benefit-function condition is not too restrictive. In general, one element of the current state may impact some elements of the next period state but not all of them. Suppose $\tilde{x}_{t+1,i}(x_t)$ is independent of $x_{t,j}$ for some $j \neq i$. Then, by definition $\tilde{x}_{t+1,i}(x_t)$ is stochastically increasing in $x_{t,j}$. Therefore, the stochastically increasing property of multi-dimensional state transition can also be satisfied easily.

Cases in which the state transition depends only on parts of the state space can be analyzed following the analyses given in this and the previous subsections. For example, consider the case in which the state transition of a three-dimensional state space $\tilde{x}_{t+1}(x_t)$ can be separated into $\tilde{x}_{t,(1,2)}(x_{t,(1,2)})$ and $\tilde{x}_{t,3}(x_{t,3})$, where the random transition $\tilde{x}_{t,(1,2)}(x_{t,(1,2)})$ is stochastically increasing in $x_{t,(1,2)}$. If the one-step benefit function $M_t(x_t)$ is increasing in both $x_{t,1}$ and $x_{t,2}$, we can apply a slight modification of Proposition 2.5 to $x_{t,(1,2)}$ for each fixed $x_{t,3}$. We omit the proposition and the analysis to avoid repetition.

2.4.4 Monotonicity Results and Bounds for Optimal Thresholds

We discuss two types of monotonicity results for the optimal thresholds. The first one is the parametric monotonicity of the state-dependent optimal thresholds. The second one is the time-monotonicity of optimal thresholds. We also provide bounds for the optimal thresholds. The result of this subsection is useful for developing efficient numerical algorithms. They also help characterize how policy parameters respond to the changes in the environment.

Proposition 2.6. *The following statements are true for every t:*

1. *If $B_t(x_t)$ is increasing [resp., decreasing] in both $x_{t,i}$ and $x_{t,j}$ for $i \neq j$, then $\overline{x}_{t,i}(x_{t,-i})$ [resp., $\underline{x}_{t,i}(x_{t,-i})$] is decreasing in $x_{t,j}$ and $\overline{x}_{t,j}(x_{t,-j})$ [resp., $\underline{x}_{t,j}(x_{t,-j})$] is also decreasing in $x_{t,i}$.*
2. *If $B_t(x_t)$ is increasing in $x_{t,i}$ and decreasing in $x_{t,j}$ for $i \neq j$, then $\overline{x}_{t,i}(x_{t,-i})$ is increasing in $x_{t,j}$ and $\underline{x}_{t,j}(x_{t,-j})$ is also increasing in $x_{t,i}$.*
3. *If $B_t(x_t)$ is increasing [resp., decreasing] in $x_{t,i}$ and convex in $x_{t,j}$ for $i \neq j$, then $\underline{x}_{t,j}(x_{t,-j})$ is increasing [resp., decreasing] in $x_{t,i}$ and $\overline{x}_{t,j}(x_{t,-j})$ is decreasing [resp., increasing] in $x_{t,i}$.*

Next we consider time-monotonicity of optimal thresholds in stationary optimal stopping problems. An optimal stopping problem is stationary if the Markov process x_t is time-homogeneous and the reward functions $C(x_t)$ and $S(x_t)$ are time-invariant. It is a well-known result that the value function $V_t(x)$ is decreasing in t for every x in such problems. See, for example, §4.4 of Bertsekas (2005). As t increases, the decision maker has less opportunity to delay the stopping decision, hence the value function $V_t(x)$ decreases in t. This property directly implies that $B_t(x)$ is decreasing in t, which in turn implies the following proposition.

Proposition 2.7. *For stationary optimal stopping problems, the following statements are true:*

1. *$\overline{x}_t$ is increasing in t and $\underline{x}_t$ is decreasing in t.*
2. *$\overline{x}_{t,i}(x_{t,-i})$ is increasing in t and $\underline{x}_{t,i}(x_{t,-i})$ is decreasing in t for every $x_{t,-i}$.*

Finally, we provide bounds for the optimal thresholds. By definition, $B_t(x_t) \geq M_t(x_t)$ for every x_t. Hence, if the one-step look-ahead policy continues the process at period t with x_t, i.e., $M_t(x_t) > 0$, then the optimal policy also continues the process at period t, i.e, $B_t(x_t) > 0$. This property implies the following proposition.

Proposition 2.8. *The optimal thresholds have the following bounds:*

1. $\overline{x}_t \leq \sup\{x \in X : M_t(x) \leq 0\}$ *and* $\underline{x}_t \geq \inf\{x \in X : M_t(x) \leq 0\}$.
2. $\overline{x}_{t,i}(x_{t,-i}) \leq \sup\{x_{t,i} : M_t(x_{t,i}, x_{t,-i}) \leq 0, x_t \in X\}$ *and* $\underline{x}_{t,i}(x_{t,-i}) \geq \inf\{x_{t,i} : M_t(x_{t,i}, x_{t,-i}) \leq 0, x_t \in X\}$.

Note that determining the x that satisfies $\sup\{x \in X : M_t(x) \leq 0\}$ is simple and does not involve recursive computation. Often it can be derived in a closed form. Hence, these bounds together with the monotonicity results considerably help reduce the computational time required to determine the optimal thresholds and resulting expected profit by reducing the search region. They also provide qualitative understanding of a decision process modeled as an optimal stopping problem.

2.5. Optimal Stopping Problems with Additional Decisions

The general theory of optimal stopping has focused primarily on problems in which stopping time is the only decision to make. Yet, it is possible to have optimal stopping problems with additional decisions. In particular, at each decision period $t \in \{1, 2, \ldots, T\}$, a decision maker observes the state $x_t \in X$ of a process and decides whether to stop or continue the process. When stopping the process, the decision maker attains a reward of $S_t(x_t)$ and the process is terminated. When continuing the process, the decision maker takes an action $a_t \in A_t$ and attains a reward of $C_t(a_t, x_t)$. Then, the state evolves. The state transition depends on both x_t and a_t, and we denote it by $\tilde{x}_{t+1}(a_t, x_t)$. The action set A_t is independent of the state. Let π be a policy that specifies both the action to take for every t and every x_t and the stopping time τ that satisfy $\tau \leq T$. Let Π^T be the set of all admissible policies. Then, the decision maker's optimal stopping problem is to determine the optimal action to take and the optimal time stop the process in order to maximize the total discounted rewards. We can formulate this problem as

$$V^*(x) \equiv \sup_{\pi \in \Pi^T} E\left[\sum_{t=1}^{\tau-1} \alpha^{t-1} C_t(a_t, x_t) + \alpha^{\tau-1} S_\tau(x_\tau) \middle| x_1 = x\right].$$

The following DP specifies the optimal action for each state at each period.

$$V_t(x_t) \quad = \quad \max\{S_t(x_t), \sup_{a_t \in A_t} [C_t(a_t, x_t) + \alpha E[V_{t+1}(\tilde{x}_{t+1}(a_t, x_t))]]\}, \quad t < T, \tag{2.4}$$

where $V_T(x_T) = S_T(x_T)$. The optimal value function satisfies $V^*(x) = V_1(x)$. An optimal policy stops at period t if $S_t(x_t) \geq \sup_{a_t \in A_t} [C_t(a_t, x_t) + \alpha E[V_{t+1}(\tilde{x}_{t+1}(a_t, x_t))]]$. The optimal action when continuing the process at period t is the maximizer of the function inside $\sup[\cdot]$. Note that the optimal action depends on the state. Then, we define the one-step benefit function and the benefit function as

$$\begin{aligned} M_t(a_t, x_t) &\equiv \alpha E[S_{t+1}(\tilde{x}_{t+1}(a_t, x_t))] + C_t(a_t, x_t) - S_t(x_t), \\ B_t(a_t, x_t) &\equiv \alpha E[V_{t+1}(\tilde{x}_{t+1}(a_t, x_t))] + C_t(a_t, x_t) - S_t(x_t). \end{aligned}$$

Additionally, we define the maximal benefit function as $\overline{B}_t(x_t) \equiv \sup_{a_t \in A_t} B_t(a_t, x_t)$. We have:

$$\begin{aligned} & B_t(a_t, x_t) = \alpha E[V_{t+1}(\tilde{x}_{t+1}(a_t, x_t))] + C_t(a_t, x_t) - S_t(x_t) \\ & = \alpha E\left[\max\left\{0, \overline{B}_{t+1}(\tilde{x}_{t+1}(a_t, x_t))\right\} + S_{t+1}(\tilde{x}_{t+1}(a_t, x_t))\right] + C_t(a_t, x_t) - S_t(x_t) \\ & = M_t(a_t, x_t) + \alpha E\left[\max\left\{0, \overline{B}_{t+1}(\tilde{x}_{t+1}(a_t, x_t))\right\}\right], \quad t < T - 1, \qquad (2.5) \\ & B_{T-1}(a_{T-1}, x_{T-1}) = M_{T-1}(a_{T-1}, x_{T-1}). \end{aligned}$$

Note that the optimal policy stops the process at period t when $\overline{B}_t(x_t) \leq 0$. Hence, we need to determine structural properties of $\overline{B}_t(x_t)$ to characterize the structure of the optimal stopping policy. However, even when $B_t(a_t, x_t)$ has a certain structural property in x_t for each fixed a_t, $\overline{B}_t(x_t)$ is not guaranteed to have the same property, because the optimal action depends on x_t. Fortunately, increasing, decreasing and convex properties are preserved under maximization. Following the two-step method, we provide the following sufficient conditions for the case of a single-dimensional state space. We define $\overline{x}_t \equiv \sup\{x \in X : \overline{B}_t(x) \leq 0\}$ and $\underline{x}_t \equiv \inf\{x \in X : \overline{B}_t(x) \leq 0\}$.

Proposition 2.9. *When $M_t(a_t, x_t)$ is increasing [resp., decreasing] in x_t, and $\tilde{x}_{t+1}(a_t, x_t)$ is stochastically increasing in $x_t \in X$ for every a_t and every t, then the following statements are true for every t:*

1. *$\overline{B}_t(x_t)$ is increasing [resp., decreasing] in x_t.*
2. *A threshold policy that stops the process if $x_t \leq \overline{x}_t$ [resp., $x_t \geq \underline{x}_t$] is optimal.*

Proof. We prove the first part; then the second part follows Proposition 2.1 Part 2. The proof is based on an induction argument. Consider the increasing one-step benefit function case. At period $t = T-1$, we have $B_{T-1}(a_t, x) = M_{T-1}(a_t, x)$. Let $a_t^*(x) = \arg\max_{a_t \in A_t} B_t(a_t, x)$. For any $x^1 \leq x^2$, we have $\overline{B}_{T-1}(x^1) = B_{T-1}(a_{T-1}^*(x^1), x^1) \leq B_{T-1}(a_{T-1}^*(x^1), x^2) \leq B_{T-1}(a_{T-1}^*(x^2), x^2)$, where the first inequality is from the fact that $B_{T-1}(a, x)$ is increasing in x, and the second inequality is by the definition of $a_{T-1}^*(x)$. Next assume for the induction argument that $\overline{B}_{t+1}(x_{t+1})$ is increasing in x_{t+1}. The composition of an increasing function and $\max\{0, x\}$ is also increasing, hence, $\max\{0, \overline{B}_{t+1}(x)\}$ is an increasing function of x. Because the state transition $\tilde{x}_{t+1}(a_t, x_t)$ is stochastically increasing in x_t, $E[\max\{0, \overline{B}_{t+1}(\tilde{x}_{t+1}(a_t, x_t))\}]$ is increasing in x_t for each a_t. Because the increasing property is preserved under summation, the benefit function $B_t(a_t, x_t) = M_t(a_t, x_t) + \alpha E[\max\{0, \overline{B}_{t+1}(a_t, \tilde{x}_{t+1}(x_t))\}]$ is increasing in x_t for each a_t. By applying the same argument that we applied on $\overline{B}_{T-1}(x)$, we can verify that $\overline{B}_t(x_t) = \sup_{a_t \in A_t} B_t(a_t, x_t)$ is increasing in x_t, which concludes the induction hypothesis and the proof of the proposition. □

Proposition 2.10. *When $M_t(a_t, x_t)$ is convex in x_t, and $\tilde{x}_{t+1}(a_t, x_t)$ is stochastically convex in $x_t \in X$ for every a_t and every t, then the following statements are true for every t:*

1. *$\overline{B}_t(x_t)$ is convex in x_t.*
2. *A control-band policy that stops the process if $x_t \in [\underline{x}_t, \overline{x}_t]$ is optimal.*

We can derive a similar result for the case of the multi-dimensional state space as in §2.4. We omit the discussion here to avoid repetition.

2.6. Infinite-Horizon Optimal Stopping Problems

Here, we show that the results of this chapter can be applied to infinite-horizon optimal stopping problems. First, we define the infinite-horizon optimal stopping problem. As before, $\{x_t|t = 1, 2, \dots\}$ is a Markov process that evolves in a state space $X \subset \mathbb{R}^d$. A stopping time τ is a random variable that takes values in $\{1, 2, \dots, \infty\}$ and satisfies $\{\omega \in \Omega|\tau(\omega) \leq t\} \in \mathcal{F}_t$ for all finite t. We denote the set of all such stopping times by U. In the infinite-horizon case, we limit our interest to problems with a time-homogeneous Markov process and reward functions that are integrable. Unlike in the finite-horizon case, the stopping time can have an infinite value; hence, we need to agree on $S(x_\tau)$ for $\tau = \infty$. Clearly, if $\lim_{t\to\infty} S(x_t)$ exists, then it is natural to set $S(x_\tau)$ to this value. Otherwise, we can set $S(x_\tau)$ to be a fixed value, such as 0, or set $S(x_\tau) = \limsup_{t\to\infty} S(x_t)$, but this choice depends on the specific problem setup. Our results are valid regardless of this choice. As before, the decision maker's optimal stopping problem is the problem of determining the optimal time to stop the process in order to maximize the total discounted rewards. We can formulate this problem as

$$V^*(x) \equiv \sup_{\tau \in U} E\left[\sum_{t=1}^{\tau-1} \alpha^{t-1} C(x_t) + \alpha^{\tau-1} S(x_\tau)\big|x_1 = x\right]. \tag{2.6}$$

Unlike in the finite-horizon case, a DP recursion is not possible because the backward recursion requires a well-defined final period. To apply the two-step method, we consider a T-period optimal stopping problem that has the same Markov process and the reward functions as the infinite-horizon problem. We use the notation $(\cdot|T)$ to emphasize that the functions we consider are the finite horizon counter parts of the infinite horizon problem. For example, we denote the optimal value function of the T-period problem by $V^*(x|T)$. Because the infinite-horizon problem can be seen as a finite-horizon problem with a long decision horizon, one can conjecture that the optimal value function of the infinite horizon problem is the limit of the sequence of optimal value functions of finite-horizon problems, i.e., $V^*(x) = \lim_{T\to\infty} V^*(x|T)$. In many cases, the optimal value function also satisfies the Bellman equation $V^*(x) = \max\{S(x), C(x) + \alpha E[V^*(\tilde{x}(x))]\}$, where $\tilde{x}(x)$ denotes the one-step state transition, and a stationary policy that stops the process if $S(x) \geq C(x) + \alpha E[V^*(\tilde{x}(x))]$ is optimal. Although these properties do not always hold, researchers have found several sufficient conditions that guarantee these properties. For example, optimal stopping problems with negative reward functions satisfy these properties as noted in §3.1 of Bertsekas (2007). Shiryaev (1978) also provides several such conditions. The two-step method can be applied to all infinite-horizon problems with these properties. Therefore, instead of providing a sufficient condition, we directly assume the following properties for all infinite-horizon problems we consider:

Assumption 2.1. *1. $V^*(x) = \max\{S(x), C(x)+\alpha E[V^*(\tilde{x}(x))]\}$, and a stationary policy that stops the process if $S(x) \geq C(x) + \alpha E[V^*(\tilde{x}(x))]$ is optimal.*

2. $\lim_{T\to\infty} V^*(x|T) = V^*(x)$.

As before, we define the one-step benefit function and the benefit function as $M(x) \equiv \alpha E[S(\tilde{x}(x))] + C(x) - S(x)$ and $B^*(x) \equiv \alpha E[V^*(\tilde{x}(x))] + C(x) - S(x)$. From Part 1 of Assumption 2.1, a stationary policy that stops the process if $B^*(x) \leq 0$ is optimal. Then, the following proposition constructs the relationship between $B^*(x)$ and $B_1(x|T)$ and the optimal thresholds. Before stating the proposition, we define $\overline{x} \equiv \sup\{x \in X : B^*(x) \leq 0\}$ and $\underline{x} \equiv \inf\{x \in X : B^*(x) \leq 0\}$ for the case of a single-dimensional state space and define $\overline{x}_i(x_{-i}) \equiv \sup\{x_i : B^*(x_i, x_{-i}) \leq 0, x \in X\}$ and $\underline{x}_i(x_{-i}) \equiv \inf\{x_i : B^*(x_i, x_{-i}) \leq 0, x \in X\}$ for the case of a multi-dimensional state space. Similarly, we denote the thresholds of the T-period problem by $\overline{x}_{t|T}$, $\underline{x}_{t|T}$, $\overline{x}_{t,i|T}(x_{-i})$ and $\underline{x}_{t,i|T}(x_{-i})$.

Proposition 2.11. *For infinite-horizon problems that satisfy Assumption 2.1, the following statements are true:*

1. $B_1(x|T) \uparrow B^*(x)$ as $T \to \infty$.

2. $\overline{x}_{1|T} \downarrow \overline{x}$, $\underline{x}_{1|T} \uparrow \underline{x}$, $\overline{x}_{1,i|T}(x_{-i}) \downarrow \overline{x}_i(x_{-i})$ and $\underline{x}_{1,i|T}(x_{-i}) \uparrow \underline{x}_i(x_{-i})$ as $T \to \infty$.

Increasing, decreasing and convex properties are preserved under limit. Hence, when $B_1(x|T)$ has any of these structural properties for every T, $B^*(x)$ also has the same structural property. Therefore, we can apply the results of previous sections to infinite-horizon problems as follows: First, fix a finite time horizon T. Next, determine a structural property of $B_1(x|T)$ using the results of previous sections. Verify these properties are preserved under limit. If so, the structure of the optimal policy follows from the finite horizon counter part. Consider, for example, the case in which the state space is single-dimensional, $M(x)$ is increasing in x and $\tilde{x}(x)$ is stochastically increasing in x. Proposition 2.1 implies that $B_1(x|T)$ is increasing in x for every finite T. Then, from Proposition 2.11, $B^*(x) = \lim_{T\to\infty} B_1(x|T)$, which implies that $B^*(x)$ is also increasing in x. When $B^*(x)$ is increasing in x, it is optimal to stop the process if $x \le \overline{x}$. Hence, a threshold policy that stops the process if $x \le \overline{x}$ is optimal. We can apply similar arguments to other cases.

Part 2 of Proposition 2.11 is useful when running the value iteration algorithm (Bertsekas 2005 Section 7) to solve the infinite horizon problem. The value iteration algorithm is based on Part 2 of Assumption 2.1. It recursively computes $V^*(x|T) = V_1(x|T)$, which converges to $V^*(x)$ as $T \to \infty$.[8] From the time-homogeneity of the state transition and the reward functions, we have $V_t(x|T) = V_{t+1}(x|T+1)$, which implies $V_1(x|T+1) = \max\{S(x), C(x) + \alpha E[V_2(\tilde{x}(x)|T+1)]\} = \max\{S(x), C(x) + \alpha E[V_1(\tilde{x}(x)|T)]\}$. Hence, the value iteration recursion is as follows:

$$V_1(x|T+1) \quad = \quad \max\{S(x), C(x) + \alpha E[V_1(\tilde{x}(x)|T)]\}, \tag{2.7}$$

where $V_1(x|1) = S(x)$. Then, consider the case in which $B_1(x|T)$ is increasing in x. From the proposition, we have $\overline{x}_{1|T} \ge \overline{x}_{1|T+1}$. By the definition of $\overline{x}_{1|T+1}$, we have $S(x) < C(x) + \alpha E[V_2(\tilde{x}(x)|T+1)]$ for every $x > \overline{x}_{1|T+1}$. Because $\overline{x}_{1|T} \ge \overline{x}_{1|T+1}$ and $V_2(x|T+1) = V_1(x|T)$, we have $S(x) < C(x) + \alpha E[V_1(\tilde{x}(x)|T)]$ for every $x > \overline{x}_{1|T}$. Hence, we need not compute the value of $S(x)$ in (2.7) for $x > \overline{x}_{1|T}$. One can apply similar arguments to the other cases.

[8] We do not need to compute $B_1(x|T)$ to determine the optimal policy. The benefit function is an auxiliary function to determine the *structure* of the optimal policy.

2.7. Example Applications

We apply the two-step method to some optimal stopping problems from the literature including those discussed in §2.3. The objective of this section is three-fold: (1) We illustrate how to use the two-step method to obtain structural results. (2) We show that this method can be used for a variety of applications that arise in fields such as Finance, Marketing and Operations. Hence, the approach can be used for broad application areas. (3) We illustrate that the method makes it easy and transparent to characterize structural results. This transparency also allows one to obtain some results that were not reported in the original papers.

In what follows, we do not state all the assumptions, model elements and results from each paper but instead focus on the basic model and the optimal policy. For each example, we first show how to obtain structural properties of the one-step benefit function. Next, we characterize the structure of the optimal policy using the results of §2.4-2.6. We refer the reader to Appendix A.1 for the stochastic monotonicities of state transitions.

2.7.1 Time-to-Market Model

Recall the time-to-market model from §2.3.1. The reward function and the state transition of this problem have the following properties:

1. $S_t(x_t) = v(x_t)(p - c_t)D$.
2. $C_t(x_t) = 0$.
3. $\tilde{x}_{t+1}(x_t) = x_t + \xi_t$, where $\xi_t \geq 0$ is independent of x_t.

The one-step benefit function can be expressed as

$$\begin{aligned} M_t(x_t) &= \alpha E[S_{t+1}(\tilde{x}_{t+1}(x_t))] - S_t(x_t) \\ &= \alpha E[v(x_t + \xi_t) - v(x_t)](p - c_{t+1})D + v(x_t)(\alpha(p - c_{t+1}) - p + c_t)D. \end{aligned}$$

The first term is decreasing in x_t because $E[v(x_t + \xi_t) - v(x_t)]$ is decreasing in x_t due to concavity of $v(x)$. The second term is also decreasing in x_t because $v(x_t)$ is decreasing in x_t and $p - c_t \leq \alpha(p - c_{t+1})$. Therefore, $M_t(x_t)$ is decreasing in x_t. Because the state transition is stochastically increasing, a threshold policy is optimal from Proposition 2.1. Under this policy, the firm should enter the market at period t if $x_t \geq \overline{x}_t$.

2.7.2 Option Pricing Problems

We discuss two option pricing problems. First recall the American-Asian option pricing problem from §2.3.2. The reward function and the state transition are as follows:

1. $S_t(x_t) = x_{t,2} - K$ for $t \leq T$ and $S_{T+1}(x) = 0$.
2. $C_t(x_t) = 0$.
3. $\tilde{x}_{t+1,1}(x_t) = \xi x_{t,1}$ and $\tilde{x}_{t+1,2}(x_t) = \frac{tx_{t,2}+\xi x_{t,1}}{t+1}$, where ξ is a log-normal random variable and is independent of x_t.

Discounting factor is $\alpha = e^{-r}$, where r is the risk-free rate. Note that period $T+1$ is a fictitious period to apply the forced stopping restriction at the last period. The one-step benefit function is

$$\begin{aligned}
M_t(x_t) &= \alpha E[\frac{tx_{t,2} + \xi x_{t,1}}{t+1} - K] - (x_{t,2} - K) \\
&= \left(\frac{\alpha t - (t+1)}{t+1} x_{t,2} + \frac{\alpha E[\xi]}{t+1} x_{t,1} + (1-\alpha)K\right), \quad t < T, \\
M_T(x_T) &= -(x_{T,2} - K).
\end{aligned}$$

For every t, $M_t(x_t)$ is decreasing in $x_{t,2}$. The state transition $\tilde{x}_{t+1,2}(x_t)$ is stochastically increasing in $x_{t,2}$, and $\tilde{x}_{t+1,1}(x_t)$ is independent of $x_{t,2}$. Therefore, Proposition 2.3 establishes the optimality of the state-dependent threshold policy that exercises the option if $x_{t,2} \geq \overline{x}_{t,2}(x_{t,1})$.

Next we consider an American call option pricing problem. An option holder can exercise the stock option at periods $t = 1, 2, \ldots, T$ at a fixed strike price K. Let x_t be the price of the underlying stock at period t. If the option holder exercises the option at period t, she receives a reward of $x_t - K$. It is a well-known result that early exercise is never optimal for this option. We can prove this fact by using the two-step method. To do so, we use the binomial valuation approach from Cox et al. (1979). The reward function and the state transition are as follows:

1. $S_t(x) = x - K$ for $t \leq T$ and $S_{T+1}(x) = 0$.
2. $C_t(x) = 0$.
3. $\tilde{x}_{t+1}(x_t) = \xi x_t$, where $\xi = u$ with probability $p = \frac{r-d}{u-d}$ and $\xi = d$ with $1 - p$.

Note that the probability measure of ξ is the risk-neutral measure of Harrison and Kreps (1979), where r is the risk-free rate. The discount factor is $\alpha = \frac{1}{r}$. For $t < T-1$, the one-step benefit function is

$$M_t(x_t) = \alpha\left[(ux_t - K)p + (dx_t - K)(1-p)\right] - (x_t - K) = (1-\alpha)K,$$

which is always strictly positive. Because $B_t(x_t) \geq M_t(x_t) > 0$ for every x_t, stopping is never optimal at period $t < T - 1$.

Early exercising can be optimal if the underlying stock pays a dividend during the life of the call option. Suppose that the stock pays a dividend of δx_m at period $m < T$ with certain values of δ and m. the reward functions are identical to the no-dividend case, but the state transition is $\tilde{x}_{t+1}(x_t) = \xi x_t$ for $t \neq m-1$ and $\tilde{x}_{t+1}(x_t) = \xi(1-\delta)x_t$ for $t = m - 1$. Hence, the one-step benefit function is identical to the no dividend case when $t \neq m - 1$. When $t = m - 1$,

$$M_t(x_t) = \alpha\left[(u(1-\delta)x_t - K)p + (d(1-\delta)x_t - K)(1-p)\right] - (x_t - K) = (1-\alpha)K - \delta x_t,$$

which is a strictly decreasing function of x_t. From Proposition 2.1, exercising the option is optimal if $x_t \geq \overline{x}_t$ for a threshold $\overline{x}_t$. The structure of the optimal exercise policy of the American option is also studied in Chen (1970) and Chapter 1 of Ross (1983) under different modeling assumptions.

2.7.3 Dynamic Quality Control Problem

Consider the dynamic quality control problem studied in Chen et al. (1998) and Yao and Zheng (1999b). The objective is to determine the optimal time to stop an inspection process to minimize the total inspection and warranty costs. For a batch of T units, period t indicates that exactly t units have been inspected and repaired if defective. The state variable x_t denotes the total number of defective units among t units inspected. Each unit in a batch is defective with probability Θ, which is an unknown random variable, and can be estimated with the observation x_t and a prior distribution f_Θ. Let $\Theta_t(x_t)$ be the best estimate of Θ at period t, with the observation x_t. The decision maker incurs an inspection cost c_i per each unit inspected and a repair cost c_r per each unit repaired. When the decision is to stop the inspection at period t and the defective rate is θ, the expected warranty cost is $\phi(t, \theta)$. We note that the expected cost depends on the state x_t through an intermediate random variable $\Theta_t(x_t)$. The properties of the profit functions and state transition are as follows:

1. $S_t(x_t) = E[\phi(t, \Theta_t(x_t))]$, where $\phi(t, \theta)$ is K-submodular[9] in (t, θ) with $K = c_r$.
2. $C_t(x_t) = c_i + c_r E[\Theta_t(x_t)]$.
3. $\tilde{x}_{t+1}(x_t) = x_t + D(x_t)$, where $D(x_t)$ is a Bernoulli random variable with parameter $E[\Theta_t(x_t)]$.

The objective is to minimize the total cost instead of maximizing the total profit, hence $S_t(x_t)$ and $C_t(x_t)$ are defined as cost functions. The random variable $\Theta_t(x_t)$ is stochastically increasing in x_t, so is $\tilde{x}_{t+1}(x_t)$. In addition $\Theta_{t+1}(\tilde{x}_{t+1}(x_t))$ has the same distribution with $\Theta_t(x_t)$. Then, the one-step penalty of delaying can be derived as

$$\begin{aligned} M_t(x_t) &= E[S_{t+1}(\tilde{x}_{t+1}(x_t))] + C_t(x_t) - S_t(x_t) \\ &= E[\phi(t+1, \Theta_{t+1}(\tilde{x}_{t+1}(x_t)))] + E[c_i + c_r\Theta_t(x_t)] - E[\phi(t, \Theta_t(x_t))] \\ &= E[\phi(t+1, \Theta_t(x_t)) - \phi(t, \Theta_t(x_t)) + c_i + c_r\Theta_t(x_t)]. \end{aligned}$$

[9] Chen et al. (1998) define a function $\phi(t, \theta)$ as K-submodular in (t, θ) if $[\phi(t+1, \theta) + \phi(t, \theta')] - [\phi(t+1, \theta') + \phi(t, \theta)] \geq K(\theta' - \theta)$ for all t and $\theta' \geq \theta$.

The one-step penalty function $M_t(x_t)$ is decreasing in x_t because $\phi(t+1,\theta)-\phi(t,\theta)+c_r\theta$ is decreasing in θ from K-submodularity, and $\Theta_t(x_t)$ is stochastically increasing in x_t. Hence, from Proposition 2.1, a threshold policy is optimal.

Yao and Zheng (1999a) also study a two-stage dynamic quality control problem. In this problem, each unit in a batch can possibly have two types of defects. The batch is processed in two inspection stages. Type 1 defects can be inspected only in the first-stage, and type 2 defects can be inspected only in the second-stage. The second stage problem is a good example of a two-dimensional optimal stopping problem. In their second stage problem, the number of units with type 1 defects, $x_{t,1}$, and the number of units of type 2 defects, $x_{t,2}$, are the state variables. The one-step benefit function $M_t(x_{t,1}, x_{t,2})$ is increasing in both $x_{t,1}$ and $x_{t,2}$, and the state transition is stochastically increasing in each variable. From Proposition 2.5, a state-dependent threshold policy that stops the process if $x_{t,1} \le \overline{x}_{t,1}(x_{t,2})$ or $x_{t,2} \le \overline{x}_{t,2}(x_{t,1})$, is optimal. The two-step method enables us to determine a new result that is not reported in Yao and Zheng (1999a). Because $B_t(x_t)$ is increasing in both $x_{t,1}$ and $x_{t,2}$, from Proposition 2.6, we have:

Proposition 2.12. *The threshold $\overline{x}_{t,1}(x_{t,2})$ is decreasing in $x_{t,2}$ and $\overline{x}_{t,2}(x_{t,1})$ is decreasing in $x_{t,1}$.*

2.7.4 Organ Transplantation Problem

Alagoz et al. (2007b) study the organ transplantation decision problem faced by patients with end-stage liver disease. In this problem, a patient has to decide whether to accept an allocated organ at each period t or to wait for another organ. The state $x_{t,1} \in \{1, \ldots, H+1\}$ indicates the condition of the patient, and the state $x_{t,2} \in \{1, \ldots, L+1\}$ indicates the quality of the organ allocated at period t. Higher values of $x_{t,1}$ and $x_{t,2}$ indicate worse conditions. The reward $S(x_t)$ that the patient receives when he accepts the organ at period t is decreasing in both $x_{t,1}$ and $x_{t,2}$. When the patients waits for another organ, there is a continuing reward $C(x_t)$, which is decreasing in $x_{t,1}$ and independent of $x_{t,2}$. The state of the organ that will be allocated in the next period $\tilde{x}_{t+1,2}(x_t)$ is independent of the state of current allocated organ, $x_{t,2}$, but depends on the current health condition of the patient, $x_{t,1}$. The authors consider an infinite-horizon problem with $\alpha < 1$. Alagoz et al. (2007b) provide qualitative properties of the reward function and the state transition as follows:

1. $S(x_t)$ is decreasing in both $x_{t,1}$ and $x_{t,2}$
2. $C(x_t)$ is decreasing in $x_{t,1}$ and independent of $x_{t,2}$
3. $\tilde{x}_{t+1}(x_t)$ is independent of $x_{t,2}$.

To apply the two-step method, we first consider the T-period problem that has the same Markov process and the reward functions as the infinite-horizon problem. Because $\tilde{x}_{t+1}(x_t)$ is independent of $x_{t,2}$, the one-step benefit function $M_t(x_t|T) = \alpha E[S(\tilde{x}_{t+1}(x_t))] + C(x_t) - S(x_t)$ is increasing in $x_{t,2}$ for each fixed $x_{t,1}$. From the same reason, $\alpha E[\max\{0, B_{t+1}(\tilde{x}_{t+1}(x_t)|T)\}]$ is independent of $x_{t,2}$. Hence, $B_t(x_t|T) = M_t(x_t|T) + \alpha E[\max\{0, B_{t+1}(\tilde{x}_{t+1}(x_t)|T)\}]$ is increasing in $x_{t,2}$. Finally, from Proposition 2.11, $B^*(x)$ is increasing in $x_{t,2}$, which establishes the optimality of the state-dependent threshold policy that stops the process at period t if $x_{t,2} \leq \overline{x}_2(x_{t,1})$.

2.7.5 The Secretary Problem

Consider an extension of the classical secretary problem studied in Chow et al. (1964) and Freeman (1983). A decision maker interviews T candidates for a single position in a random order. The decision maker can accept only the current candidate he is interviewing. The utility of accepting ith best candidate among T candidates is $T-i$. At period t, let state x_t denote the relative rank of the current candidate among all t candidates the decision maker has interviewed. When the decision maker accepts the x_tth best candidate among t, the expected utility is $T - \frac{T+1}{t+1}x_t$. The properties of the profit functions and state transition for this problem are as follows:

1. $S_t(x_t) = T - \frac{T+1}{t+1}x_t$.
2. $C_t(x_t) = 0$.
3. $\tilde{x}_{t+1}(x_t)$ is uniformly distributed from 1 to $t+1$ and independent of x_t.

Because x_{t+1} is independent of x_t, the one-step benefit function $M(x_t) = \alpha E[S(\tilde{x}_{t+1}(x_t))] - S(x_t)$ is increasing in x_t. From the same reason, $\alpha E[\max\{0, B_{t+1}(\tilde{x}_{t+1}(x_t))\}]$ is independent of x_t. Therefore, $B_t(x_t) = M_t(x_t) + \alpha E[\max\{0, B_{t+1}(\tilde{x}_{t+1}(x_t))\}]$ is increasing in x_t, which establishes the optimality of the state-dependent threshold policy that accepts a candidate at period t if $x_t \leq \overline{x}_t$.

2.8. Conclusion

This section has proposed a two-step method to characterize the structure of the optimal stopping policy for a general class of optimal stopping problems. The method first characterizes structural properties of the one-step benefit function, which depends on the reward functions and the state transition over two periods. The method then verifies the stochastic monotonicity of the state transition that enables the benefit function to inherit the structural property of the *one-step* benefit function, which characterizes the structure of the optimal stopping policy. We have also considered optimal stopping problems with additional decisions other than the stopping decision. By applying the proposed method to examples from several fields, we have illustrated how to use the method; why it is needed to obtain structural results where other conventional methods fail; and how it simplifies the analysis. Our hope is that the method and the propositions will also help researchers to easily determine specific conditions needed and the resulting optimal policy structure for their optimal stopping problems appearing in broad application areas.

Chapter 3

Mechanism Design for Capacity Planning under Dynamic Evolution of Asymmetric Demand Forecasts

3.1. Introduction

This chapter studies a supplier's problem of eliciting credible forecast information from a manufacturer when both parties obtain forecast information over time. The supplier relies on the demand forecast for his capacity decision. Both parties obtain information and update their forecasts over time. However, the manufacturer often has other forward-looking information because of her superior relationship with or proximity to the market and expert opinion about her own product. Hence, firms have asymmetric information, which changes over time. In such a dynamic environment, what is the optimal mechanism/contract that maximizes supplier's profit by enabling credible forecast information sharing? What is the right time for the supplier to offer this mechanism? Does time play an important role? If so, how? This chapter addresses these questions. In doing so, this chapter also characterizes the supplier's optimal mechanism/contract and shows how this mechanism changes over time as a function of forecast updates. This chapter also rigorously models the information evolution process for multiple decision makers who forecast the same object. To do so, this chapter generalizes the *Martingale Model of Forecast Evolution (MMFE)* to account for multiple decision makers who forecast demand for the product.

The MMFE successfully describes the evolution of forecasts arising from many statistical and judgment-based forecasting methods. There are two variants of the MMFE: the multiplicative and the additive models. Both variants frequently appear in the literature. Hausman (1969) first develops the multiplicative MMFE and verifies the model using actual data from several independent forecast-revision processes (agricultural supply forecasts, financial forecasts, and forecasts of seasonal sales of apparel). Hausman and Peterson (1972) extend the model to a multiple products system. Graves et al. (1986) develop the additive MMFE for a single-product system. Heath and Jackson (1994) generalize both variants of MMFE by allowing the correlation of demands of different time periods. Sapra and Jackson (2009) develop the continuous-time analog of the MMFE. Due to its descriptive power and generality, researchers have used the MMFE in several studies that involve dynamic forecast updates, for example, to develop effective production, inventory, capacity management methods that are responsive to forecast updates (Heath and Jackson 1994, Graves et al. 1998, Gallego and Özer 2001, Toktay and Wein 2001, Iida and Zipkin 2006, Altug and Muharremoglu 2009, Schoenmeyr and Graves 2009). This model is also used to understand and quantify the value of information sharing (Chen and Lee 2009 and Iida and Zipkin 2009) and collaborative forecasting (Aviv 2001, 2002, 2007).

This chapter provides a framework to model evolutions of forecasts for multiple decision makers in a consistent and rigorous manner. With the exception of Aviv (2001) and Iida and Zipkin (2009), the literature focuses on forecast models from a single forecaster's perspective. As pointed out by Aviv, the consideration of multiple forecasters and ensuring consistency among forecast evolution across multiple forecasters is an important and nontrivial process. We provide a framework that ensures consistency, and that can be used to model several plausible forecast evolution scenarios, such as collaborative forecasting and delayed information among forecasters. Consider, for example, multiple forecasters who have different capability and speed in learning information about the demand for a product. These forecasters would have asymmetric demand information that would change over time. Depending on how much information each party obtains, information asymmetry among forecasters may get larger or smaller. We model this scenario using our framework and refer to it as the *Martingale Model of Asymmetric Forecast Evolution (MMAFE)*. During the last decade, Operations Management research has been focusing increasingly on problems with multiple decision makers who may employ different forecasters to obtain information about the same object (such as demand). Our hope is that the framework provided in this chapter will also enable researchers to consider and revisit, for example, the performance of supply chain contracts in dynamic environments. In this chapter, we provide one such study.

Using the MMAFE for multiple decision makers, we revisit the incentive problem observed in forecast information sharing when firms need to share this information for better planning. In particular, we study a supplier's capacity planning problem. For a timely delivery, the supplier has to secure component capacity prior to receiving a firm order from a product manufacturer. The supplier makes the capacity decision based on demand forecasts. The manufacturer possesses more demand information than the supplier due to, for example, her proximity to the market and her expert opinion about the final product. Hence, the supplier can better plan for capacity by using the manufacturer's forecast information. However, without a proper incentive mechanism, when asked for this information, the manufacturer may inflate her forecast so that the supplier secures more capacity. Being aware of this situation, the supplier may not find the information *credible*. This interaction may result in forecast manipulation[1] that reduces both parties' expected profits.

Forecast manipulation has been widely observed in industry and its adverse effects are also well-documented (see, for example, Files 2001 and Clark 2007). To address this problem, researchers have recently started to provide some analytical remedies. For example, Cachon and Lariviere (2001) provide some properties of incentive mechanisms, and Özer and Wei (2006) design explicit contracts to ensure credible forecast information sharing. The literature that provides analytical remedies for credible

[1]This manipulation could be either due to the manufacturer who inflates her forecast report or the supplier who "corrects" the forecast information thinking that the manufacturer inflated her forecast.

forecast or cost information sharing in a *strategic* setting is relatively recent (e.g, Corbett and de Groote 2000, Ha 2001, Lovejoy 2006). Chen (2003) provides a review of this literature in strategic supply chain settings. To date, the mechanism design literature and the supply chain contracting literature have primarily focused on *static* environments and designed incentive mechanisms for them. For example, the supplier (principal) can design and offer a menu of contracts to the manufacturer (agent) and screen her private forecast information while maximizing his objective function. Özer and Wei (2006) show that such a mechanism results in a capacity reservation contract where the manufacturer pays a fee to reserve capacity (instead of directly sharing her private forecast). This literature assumes that the information is static and not updated over time. It also assumes that information asymmetry between the principal and the agent is static. The MMAFE framework help us to address this forecast sharing problem when both parties update their forecasts over time.

The supplier[2] and the manufacturer can obtain demand information and update their forecasts as the sales season approaches. Hence, by delaying to offer an incentive mechanism, the supplier (and the manufacturer) can obtain more information, which reduces demand uncertainty and increases expected profits. This delay, however, may increase (resp., or decrease) the degree of information asymmetry between the two decision makers, resulting in higher (resp., or lower) cost of screening. The capacity cost may also increase as the supplier delays the capacity decision because of tighter deadline for building capacity. Considering all such trade-offs, the supplier has to determine (i) when to offer an incentive mechanism/contract to elicit credible forecast information and (ii) how to design the mechanism so as to maximize his profit while ensuring that the manufacturer participates and credibly reveals her forecast.

The aforementioned timing, contract and capacity decisions are closely linked because the optimal contract and the capacity decision depend on the supplier's forecast information and the information asymmetry, both of which change over time. Hence, we formulate this problem as a two-stage closely-embedded stochastic decision process. The first-stage determines the optimal time to offer a capacity reservation

[2] To be consistent with the literature, we use the supplier-manufacturer narrative, which could also be described as a supplier-retailer interaction.

contract. The solution of this problem depends on the solution of the second-stage problem, which is about designing an optimal mechanism for capacity planning. Similarly, the solution of the second stage depends on the timing decision obtained from the first-stage. We establish the optimality of a control band policy that prescribes when to offer an optimal incentive mechanism. Under this policy, the supplier offers a menu of contracts if the supplier's demand forecast falls within the control band. We also establish properties of this optimal stopping policy. Next we provide structural properties of the optimal incentive mechanism (which is interpreted as a capacity reservation contract) and explicitly show how the optimal mechanism depends on the demand forecast and how the timing decision affects the mechanism design problem. Using this framework, we also solve the problem from a centralized decision maker's perspective. In this case, the decision maker has access to all relevant information and forecast updates and determines the optimal time to decide and build capacity. Through numerical studies, we characterize the environment in which the supplier should offer the contract late or early. By comparing the profits of the dynamic strategy with those of a static one in which the supplier offers a contract in a fixed period, we show that the supplier can significantly improve his profit by optimally determining the time to offer a contract. However, the results also show that this dynamic strategy can reduce the total supply chain efficiency.

The literature on mechanism design for a dynamic framework is sparse. Plambeck and Zenios (2000), Zhang and Zenios (2008), Lutze and Özer (2008) and Akan et al. (2009) are among the few exceptions. These authors study a principal's problem of designing a long-term contract when the agent takes actions over multiple periods. The principal in their model offers a contract at the beginning of the planning horizon. In contrast, the principal in our model dynamically determines the time to design and offer a contract and the agent takes a single action. Note that when such a dynamic system is managed by a centralized decision maker (i.e., the integrated firm solution), the problem reduces to determining when the firm should make a decision given information updates. The centralized decision maker does not face the problem of information asymmetry. For example, Boyaci and Özer (2009) study a supplier's problem of determining when to stop advance selling under demand information updates. Ulu and Smith (2009), and Wang and Tomlin (2009) also consider the optimal timing decisions under multiple information updates. Although it is not a central part of our study, the present chapter also determines the integrated firm's optimal time to determine capacity when the firm obtains multiple forecast updates over a capacity planning horizon.

The rest of the chapter is organized as follows. In §3.2, we develop the MMFE for multiple decision makers and introduce the MMAFE. In §3.3 and §3.4, we describe the basic elements of our model and provide a formulation to solve the problem. In §3.5, we provide structural properties of the optimal stopping policy. We also characterize an optimal contract that enables credible forecast information sharing. In §3.6, we provide the integrated firm solution. In §3.7, we present numerical studies. In §3.8, we discuss extensions. In §3.9, we conclude. We defer all proofs to the Appendix.

3.2. The Martingale Model of Forecast Evolution for Multiple Decision Makers

This section develops the MMFE for multiple decision makers who forecast demand for the same product. When several decision makers forecast demand for the same product, their forecast revisions are likely to be positively correlated. The correlation may occur inter-temporally, because the decision makers may obtain demand information with a time delay. For example, the supplier and the retailer of a product can use past sales data to update demand forecasts, but the supplier may obtain this information later than the retailer (Lee et al. 2000). The forecasting model for multiple decision makers should successfully describe all possible interactions between decision makers' forecast sequences. The forecast sequences should also be consistent in such a way that they converge to the same value, i.e., the actual demand. Hence, we aim to develop a descriptive framework that characterizes the dynamics of general forecasting processes across multiple decision makers while being consistent.

In §3.2.1, we discuss the general MMFE for multiple decision makers and provide its properties. These properties help us to better understand the two important variants of the model: multiplicative and additive forecast revisions. When the size of forecast revisions are small compared to the size of the forecast, the difference between the two models is negligible. However, the forecasts are often made long before the beginning of the sales season; hence, forecast revisions are likely to be large. Similarly, several researchers point out that the multiplicative model fit empirical data better than the additive model does (see, for example, Hausman 1969, Heath and Jackson 1994, Chod and Rudi 2006 and Wang and Tomlin 2009). Therefore, in §3.2.2, we focus on the multiplicative case, and defer the development of the additive MMFE for multiple decision makers to the Appendix. In §§3.2.3, 3.2.4, we apply the model to some forecasting scenarios discussed in the literature that involve multiple decision makers.

3.2.1 The General MMFE

Consider N periods during which each decision maker independently forecasts demand for a single product. We denote the sales season by period $N+1$. Demand for the product is X_{N+1}, which is a random variable prior to the sales season. At the beginning of each period $n \in \{1, \dots N\}$, demand information available to decision maker i is given by the set $\mathcal{F}_n^i$, which is a σ-field. The demand forecast of decision maker i at the beginning of period n is $X_n^i \equiv E[X_{N+1}|\mathcal{F}_n^i]$; i.e., the expected demand given information $\mathcal{F}_n^i$. We denote the differences between subsequent forecasts by $\Delta_n^i \equiv X_{n+1}^i - X_n^i$. In Appendix A, we provide a glossary of notation for an easy reference.

Definition 3.1. *The forecast evolution X_n^i constructed by $(X_{N+1}, \mathcal{F}_n^i)$ is an MMFE if it satisfies the properties (a) X_{N+1} is square-integrable, (b) $\mathcal{F}_n^i \subseteq \mathcal{F}_{n+1}^i$ for every n, and (c) $\sigma(X_{N+1}) \subseteq \mathcal{F}_{N+1}^i$.*

Condition (a) is required to define the Martingale differences and is identical to the condition that X_{N+1} has a finite variance. Condition (b) implies that the decision makers do not lose information over time. Condition (c) implies that demand is revealed to decision maker i during the sales period. From this definition, we have the following theorem.

Theorem 3.1. *If X_n^i constructed by $(X_{N+1}, \mathcal{F}_n^i)$ is an MMFE, then we have the following properties for every n:*

(a) X_n^i is a Martingale adapted to $\mathcal{F}_n^i$.

(b) $E[X_{N+1}|\mathcal{F}_n^i] = E[X_{N+1}|X_n^i] = X_n^i$.

(c) $E[X_{n+l}^i|\mathcal{F}_n^i] = E[X_{n+l}^i|X_n^i] = X_n^i$ for every $l \geq 0$.

(d) $E[\Delta_n^i] = 0$ and Δ_l^i is uncorrelated with $\mathcal{F}_n^i$ for every $l \geq n$.

Theorem 3.1 is first discussed in Heath and Jackson (1994) without a proof. We provide a formal proof in the Appendix. Part (a) verifies that the MMFE is indeed a Martingale. Part (b) implies that in forecasting X_{N+1}, the value of X_n^i is sufficient information for decision maker i. Part (c) implies that the forecast is unbiased. Finally, Part (d) implies that X_n^i is indeed the best forecast for decision maker i. Note that if Δ_n^i is correlated with any past information, the decision maker can use the correlation to improve the forecast.[3]

3.2.2 The Multiplicative MMFE

We denote the ratio of successive forecasts by $\delta_n^i \equiv X_{n+1}^i / X_n^i$ for $n < N$ and $\delta_N^i \equiv X_{N+1}/X_N^i$ for each $i \in \{s, m\}$. We consider only two decision makers for the sake of brevity and without loss of generality. We refer to these decision makers as the supplier (s) and the manufacturer (m) because we apply the model to a forecast sharing problem between these two decision makers in §3.3.

The classical MMFE assumes that the multiplicative forecast update for each decision maker, i.e., δ_n^i, is log-normally distributed for every n. Part (c) and (d) of

[3]From the tower property of conditional expectation, $E[X_{N+1}|\mathcal{F}_n^i] = E[X_{N+1}|X_n^i] = E[E[X_{N+1}|\mathcal{F}_{n+1}^i]|X_n^i] = E[X_{n+1}^i|X_n^i] = X_n^i + E[\Delta_n^i|X_n^i]$. If Δ_n^i is correlated to X_n^i, the value of $E[\Delta_n^i|X_n^i]$ may not be 0. In this case, decision maker i's best demand forecast at period t would be $X_n^i + E[\Delta_n^i|X_n^i]$ rather than X_n^i.

Theorem 3.1 imply that δ_n^i is independent of X_n^i and has a mean value of 1. Hence, the initial forecast X_1^i and the variances of $\log(\delta_n^i)$ fully characterize the evolution of X_n^i for *each* decision maker i[4]. However, the variance of $\log(\delta_n^i)$ is *not* sufficient to characterize the interaction between the two forecast processes X_n^s and X_n^m.

One may determine the correlation coefficient between $\log(\delta_{n_s}^s)$ and $\log(\delta_{n_m}^m)$ for every n_s and n_m to characterize the interaction between the two forecast sequences. However, this approach can lead to inconsistency. We provide one such example. Suppose that the correlation coefficient between $\log(\delta_1^s)$ and $\log(\delta_1^m)$ is 1 and the correlation coefficient between $\log(\delta_2^s)$ and $\log(\delta_1^m)$ is also 1. Then, by obtaining the value of δ_1^s, the supplier obtains the full information of δ_1^m, which also contains the full information of δ_2^s. Hence, the property that δ_1^s and δ_2^s are independent does not hold.

We propose a different approach to model the interaction between X_n^s and X_n^m, which does not suffer from any inconsistency such as the one discussed above. Decision makers update their forecasts by obtaining information about events that affect demand. Following Hausman (1969), suppose there are in total K such events and let e_j be the random variable that models the impact of event j. According to the theory of proportional effect (Aitchison and Brown 1957), the change in the forecast by each event is proportional to the size of the current forecast. In other words, after obtaining the information of event j, decision maker i updates the forecast from X_n^i to $X_n^i e_j$. Following this explanation, we first express demand by $X_{N+1} = \prod_{j=1}^{K} e_j$. Next, we divide the set of all events into $(N+1) \times (N+1)$ sets by the time at which the information is obtained by each decision maker. More specifically, we define E_{n_s,n_m} as the set of events whose information is obtained by the supplier during period n_s and by the manufacturer during period n_m.

We define $\delta_{n_s,n_m} \equiv \prod_{j \in E_{n_s,n_m}} e_j$, which indicates the total information obtained by the supplier at period n_s and by the manufacturer at period n_m. We assume that each δ_{n_s,n_m} is log-normally distributed and has a mean value of 1 except $\delta_{0,0}$[5]. When

[4]The assumption that $E[\delta_n^i] = 1$ for a log-normal random variable δ_n^i implies $E[\log(\delta_n^i)] = -Var(\log(\delta_n^i))/2$. Therefore, the variance of $\log(\delta_n^i)$ is sufficient to characterize δ_n^i.

[5]By taking the logarithm of δ_{n_s,n_m}, we get $\log(\delta_{n_s,n_m}) = \sum_{j \in E_{n_s,n_m}} \log(e_j)$. When the number of events in E_{n_s,n_m} becomes large, $\sum_{j \in E_{n_s,n_m}} \log(e_j)$ will be asymptotically normal from the central

E_{n_s,n_m} is an empty set, i.e., when no information is obtained by the supplier at period n_s and by the manufacturer at period n_m, then $\delta_{n_s,n_m} = 1$. Note that by construction a distinct piece of information is contained in one event set, hence δ_{n_s,n_m} forms an independent set of random variables.

Given this construction, we can express demand as $X_{N+1} = \prod_{n_s=0}^{N} \prod_{n_m=0}^{N} \delta_{n_s,n_m}$. The supplier's information set at the beginning of period n is

$$\mathcal{F}_n^s \equiv \sigma([\delta_{0,0}, \ldots, \delta_{0,N}], \ldots, [\delta_{n-1,0}, \ldots, \delta_{n-1,N}]).$$

Then, the supplier's demand forecast is $X_n^s = E[X_{N+1}|\mathcal{F}_n^s] = \prod_{n_s=0}^{n-1} \prod_{n_m=0}^{N} \delta_{n_s,n_m}$, and the ratio of successive forecasts is $\delta_n^s = \prod_{n_m=0}^{N} \delta_{n,n_m}$. Because the multiplication of log-normal random variables is also a log-normal random variable, δ_n^s is also log-normally distributed. Therefore, from the supplier's perspective, the forecast evolution is consistent with the classical MMFE. The manufacturer's forecast can be expressed in a similar way. Figure 3.1 illustrates the information structure of the MMFE for two decision makers. During each period, the supplier obtains all information given in the row corresponding to that period, whereas the manufacturer obtains all information given in the corresponding column.

Figure 3.1: Information Structure of the MMFE

n	0		1		$\cdots$		N	(s)
0	$\delta_{0,0}$	$\times$	$\delta_{0,1}$	$\times$	$\cdots$	$\times$	$\delta_{0,N}$	X_1^s
	$\times$		$\times$		$\times$		$\times$	$\times$
1	$\delta_{1,0}$	$\times$	$\delta_{1,1}$	$\times$	$\cdots$	$\times$	$\delta_{1,N}$	δ_1^s
	$\times$		$\times$		$\times$		$\times$	$\times$
$\vdots$	$\vdots$	$\times$	$\vdots$	$\times$	$\ddots$	$\times$	$\vdots$	$\vdots$
	$\times$		$\times$		$\times$		$\times$	$\times$
N	$\delta_{N,0}$	$\times$	$\delta_{N,1}$	$\times$	$\cdots$	$\times$	$\delta_{N,N}$	δ_N^s
(m)	X_1^m	$\times$	δ_1^m	$\times$	$\ldots$	$\times$	δ_N^m	X_{N+1}

From this construction, we can fully characterize the evolution of X_n^s and X_n^m by

limit theorem. Because both decision makers have the information $\delta_{0,0}$ before the beginning of the forecast horizon, we assume that $\delta_{0,0}$ is a deterministic value. When $E[\delta_{n_s,n_m}] \neq 1$ for some (n_s, n_m), we can push this information to $\delta_{0,0}$ and normalize δ_{n_s,n_m} by $\delta_{n_s,n_m}/E[\delta_{n_s,n_m}]$, hence the assumption $E[\delta_{n_s,n_m}] = 1$ is without loss of generality.

determining the value of $\delta_{0,0}$ and the variances of $\log(\delta_{n_s,n_m})$.

Note that the demand and the forecast revisions have the following relationship; $X_{N+1} = X_1^i \prod_{k=1}^N \delta_k^i$ for each decision maker i. At the beginning of period n, decision maker i has the information $X_1^i, \delta_1^i, \ldots, \delta_{n-1}^i$, but does not have the information $\delta_n^i, \delta_{n+1}^i, \ldots, \delta_N^i$. Hence, the multiplication $\prod_{k=n}^N \delta_k^i$ represents the demand uncertainty faced by decision maker i at the beginning of period n.

3.2.3 Collaborative Forecasting, Delayed Information and Information Sharing

Consider the collaborative forecasting process discussed in Aviv (2001, 2002, and 2007). When the two decision makers collaborate to forecast demand, they share all available information. Based on Definition 3.1, we define the collaborative information set as $\mathcal{F}_n^{cf} \equiv \mathcal{F}_n^s \cup \mathcal{F}_n^m$. Because the union of two σ-fields is also a σ-field, $\mathcal{F}_n^{cf}$ is a well-defined information set. Then, the collaborative forecast (CF) of the two decision makers is $X_n^{cf} \equiv E[X_{N+1}|\mathcal{F}_n^{cf}]$. Next we derive the most important property of the CF.

Theorem 3.2. *The CF has a smaller mean-squared-error than the forecast of a single decision maker, i.e.,* $E[(X_{N+1} - X_n^{cf})^2] \leq E[(X_{N+1} - X_n^i)^2]$ *for every* $i \in \{s, m\}$.

This result states that two decision makers who collaborate can predict demand more accurately.

For the case of multiplicative MMFE, the collaborative information set $\mathcal{F}_n^{cf}$ includes all δ_{n_s,n_m} such that $n_s \leq n$ or $n_m \leq n$. Hence, the initial forecast is $X_0^{cf} = \delta_{0,0} \prod_{n_s=1}^N \delta_{n_s,0} \prod_{n_m=1}^N \delta_{0,n_m}$, and the ratio of successive forecasts are

$$\delta_n^{cf} = \delta_{n,n} \prod_{n_s=n+1}^{N} \delta_{n_s,n} \prod_{n_m=n+1}^{N} \delta_{n,n_m}.$$

Figure 3.2(a) illustrates the information structure available to decision makers under a collaborative forecasting scheme. Because each δ_n^{cf} is a log-normal random variable with the mean value of 1, the collaborative forecast is also a multiplicative MMFE.

Figure 3.2: Information Structure of the multiplicative MMFE

(a) Collaborative Forecast (b) Asymmetric Forecast Evolution

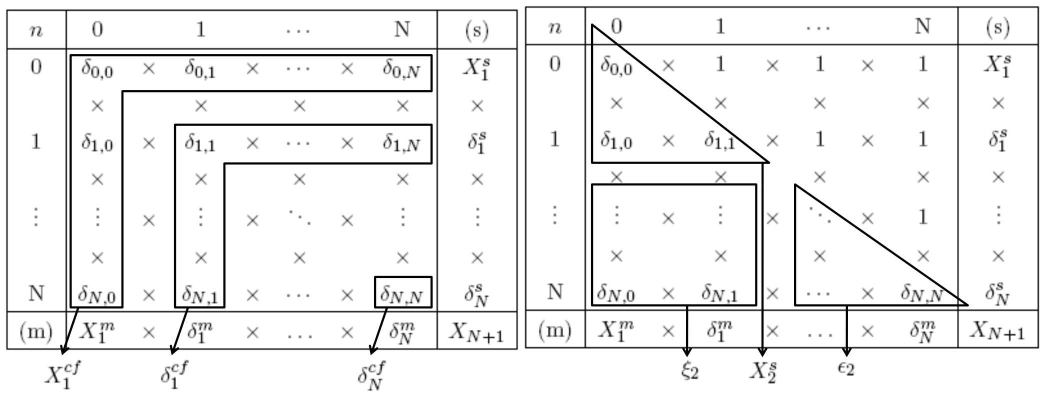

Aviv (2001) also uses the MMFE to describe the forecast sequences of the two decision makers. He first models the forecast sequence of each decision maker as an MMFE with the initial forecast of $X_1^i = \mu\delta_0^i$ and the multiplicative forecast revisions of δ_n^i. Then, he assumes that $Var(\log(\delta_n^i)) = (\eta^i\sigma_n)^2$ for every $n = 0, \ldots, N-1$, and $Var(\log(\delta_N^i)) = \sigma^2 - \sum_{n=0}^{N-1}(\eta^i\sigma_n)^2$. The value of σ represents the degree of total demand uncertainty, and the value of η^i represents the forecasting power of decision maker i. He models the interaction between the two forecast sequences by assuming that the correlation coefficient between $\log(\delta_{n_s}^s)$ and $\log(\delta_{n_m}^m)$ is ρ for $n_s = n_m$, and 0 for $n_s \neq n_m$. In other words, the forecast revisions of two decision makers are correlated, but not inter-temporally[6].

In the MMFE for multiple decision makers, we construct a single demand model, which automatically constructs the forecast sequence of each decision maker. By construction, our model does not suffer from inconsistency. In contrast, Aviv (2001) constructs the forecast sequence of each decision maker separately and then models the interaction between them. This approach may lead to inconsistency. For example, when $\eta^s = \eta^m = 1$ and $\sigma_{N-1} = \sigma$, both decision makers obtain all demand information during period $N-1$, hence the correlation coefficient ρ should be exactly 1 for the demand models to be consistent. For this reason, Aviv (2001) provides a sufficient

[6] Iida and Zipkin (2009) also use the MMFE to describe the forecast sequences of two decision makers in a similar way to Aviv (2001).

condition on (ρ, η^s, η^m) that guarantees consistency.

As mentioned above, Aviv (2001) assumes the inter-temporal independence between the two decision makers' forecast revisions. However, some important forecasting scenarios do not follow this assumption. We discuss two such examples (delayed information and asymmetric forecast evolution) shortly. The MMFE for multiple decision makers covers the general case including that of Aviv (2001). Here, we describe his model with our framework. The inter-temporal independence means that $\delta_{n_s,n_m} = 1$ unless $n_s = n_m$ or $n_i = N$ for $i \in \{s, m\}$. Hence, the information obtained by the two decision makers at period n consists of three parts, $\delta_{n,n}$, $\delta_{n,N}$ and $\delta_{N,n}$. They correspond to the three corners of a single block in Figure 3.2(a). The correlated part of δ_n^s and δ_n^m corresponds to $\delta_{n,n}$, hence we can set $Var(\log(\delta_{n,n})) = \rho\eta^s\eta^m\sigma_n^2$. Similarly, the uncorrelated parts of δ_n^s and δ_n^m correspond to $\delta_{n,N}$ and $\delta_{N,n}$, hence we can set $Var(\log(\delta_{n,N})) = (\eta^s\sigma_n)^2 - \rho\eta^s\eta^m\sigma_n^2$ and $Var(\log(\delta_{N,n})) = (\eta^m\sigma_n)^2 - \rho\eta^s\eta^m\sigma_n^2$.

Next consider the delayed information scenario discussed by Chen (1999). In this case, the manufacturer observes demand of a product at each period and makes a replenishment order to the supplier. The supplier of the product receives the manufacturer's order with the delay of l periods. The decision makers use the demand history to update forecasts. Therefore, the information sets of the two decision makers are identical with l periods of delay. In other words, we have $\mathcal{F}_{n+l}^s = \mathcal{F}_n^m$ for every n. Then, the supplier and the manufacturer have the same sequence of forecasts with a delay of l periods. In the multiplicative MMFE, the delayed information can be represented by $\delta_{n,k} = 1$ for every $k + l \neq n$.

3.2.4 Asymmetric Forecast Evolution and Information Sharing

Consider the scenario in which the manufacturer has all information that the supplier has at each period; i.e., $\mathcal{F}_n^s \subseteq \mathcal{F}_n^m$ for every n. Both decision makers obtain new demand information from a third-party research firm (such as Gartner; see http://www.gartner.com) over time, but the manufacturer has additional information because of her proximity to the market and the insider information about her own product. In other words, the manufacturer has *private* demand information that is not available to the other decision maker. We refer to this forecast evolution model as the *Martingale Model of Asymmetric Forecast Evolution* (MMAFE). We denote the difference between the two decision makers' forecasts by $A_n \equiv X_n^m - X_n^s$. Then, we have the following properties of the MMAFE:

Theorem 3.3. *If* (X_n^s, X_n^m) *constructed by* $(X_{N+1}, \mathcal{F}_n^s, \mathcal{F}_n^m)$ *is an MMAFE, then we have the following properties for every* n*:*

(a) $E[X_{n+l}^m|\mathcal{F}_n^s] = E[X_{n+l}^m|X_n^s] = X_n^s$ *for every* $l \geq 0$*.*

(b) $E[X_{N+1}|X_n^s, A_n] = E[X_{N+1}|F_n^m] = X_n^m$*.*

(c) $E[A_n] = 0$ *and* A_n *is uncorrelated with* $\mathcal{F}_n^s$*.*

Part (a) implies that the supplier's estimate of the manufacturer's forecast is the same as his own forecast. Part (b) implies that by knowing the value of A_n, the supplier can make the best forecast. Part (c) implies that A_n is uncorrelated with the supplier's information set, $\mathcal{F}_n^s$.

For the case of multiplicative MMFE, we can model the asymmetric information scenario by setting $\delta_{n_s,n_m} = 1$ for every $n_m > n_s$. In other words, the supplier obtains no information earlier than the manufacturer. Hence, the information obtained by the supplier at period n consists of $\delta_{n,0}, \delta_{n,1}, \ldots, \delta_{n,n}$, where each δ_{n,n_m} has already been obtained or is being obtained at the same time by the manufacturer. We refer this case by the *multiplicative Martingale Model of Asymmetric Forecast Evolution* (m-MMAFE) and we will use this model in the second part of the chapter. The information structure of the m-MMAFE is provided in Figure 3.2(b). We refer the reader to this figure for better understanding of the following discussion.

The manufacturer's private demand information represents the information asymmetry between the two decision makers. The manufacturer's demand uncertainty is also the demand uncertainty faced by the system. Recall that the multiplication of $\delta_n^m, \delta_{n+1}^m, \ldots, \delta_N^m$ represents the demand uncertainty faced by the manufacturer at the beginning of period n, and we denote it by $\epsilon_n \equiv \prod_{k=n}^N \delta_k^m$. From the manufacturer's perspective, demand is $X_{N+1} = X_n^m \epsilon_n$, where X_n^m is her current forecast, which is deterministically known to her. The remaining market uncertainty ϵ_n is resolved over periods n to N as the manufacturer obtains information, i.e. the forecast updates.

In contrast, the demand uncertainty faced by the supplier at the beginning of period n is $\prod_{k=n}^{N} \delta_k^s$. The manufacturer has already obtained part of this information. To distinguish the known part, we rewrite

$$\begin{aligned} \prod_{k=n}^{N} \delta_k^s &= \prod_{k=n}^{N} \left(\prod_{n_m=0}^{n-1} \delta_{k,n_m} \prod_{n_m=n}^{N} \delta_{k,n_m} \right) \\ &= \underbrace{\left(\prod_{k=n}^{N} \prod_{n_m=0}^{n-1} \delta_{k,n_m} \right)}_{\xi_n} \underbrace{\left(\prod_{k=n}^{N} \prod_{n_m=n}^{N} \delta_{k,n_m} \right)}_{\epsilon_n} . \end{aligned} \tag{3.1}$$

The first part of Equation (3.1) represents the demand information that is already obtained by the manufacturer. The second part represents the demand information that is not yet obtained by the manufacturer. Because $\delta_{n_s,n_m} = 1$ for $n_m > n_s$, the second part of (3.1) is

$$\prod_{k=n}^{N} \prod_{n_m=n}^{N} \delta_{k,n_m} = \prod_{n_m=n}^{N} \prod_{k=n}^{N} \delta_{k,n_m} = \prod_{n_m=n}^{N} \prod_{k=0}^{N} \delta_{k,n_m} = \prod_{n_m=n}^{N} \delta_{n_m}^m,$$

which is equal to ϵ_n. The first part of (3.1) is the manufacturer's private information, and we denote it by ξ_n. Then, demand can be represented as $X_{N+1} = X_n^s \xi_n \epsilon_n$. From the supplier's perspective, X_n^s is deterministic, ξ_n and ϵ_n are uncertain. By construction, X_n^s, ξ_n and ϵ_n are independent. Note also that $X_n^m = X_n^s \xi_n$. The supplier obtains only part of the information of ϵ_n and ξ_n during period n and he obtains the full information of ϵ_n and ξ_n over periods n to N.

For notational simplicity, we denote the standard deviation of $\log(Z)$ of a log-normal random variable Z as σ_Z throughout this chapter. The value of σ_{ϵ_n} represents the degree of demand uncertainty of the system at period n, and the value of σ_{ξ_n} represents the degree of information asymmetry between the supplier and the manufacturer at period n. By construction, σ_{ϵ_n} always decreases in n. In contrast, σ_{ξ_n} can either increase or decrease in n depending on the values of $\sigma_{\delta_n^s}$ and $\sigma_{\delta_n^m}$. When the supplier obtains more information than the manufacturer during period n, i.e., $\sigma_{\delta_n^s} > \sigma_{\delta_n^m}$, we have $\sigma_{\xi_{n+1}} < \sigma_{\xi_n}$, and vice versa.

3.3. Determining the Optimal Time to Offer an Optimal Mechanism

Consider a supplier who sells a key component to a manufacturer at a wholesale price. Because of the long leadtime required in building capacity, the supplier determines the production capacity before receiving a firm order from the manufacturer. The supplier has some flexibility on when to decide and build capacity but delaying the production capacity decision may require the supplier to incur expediting costs or overtime labor. We refer to the time window during which the supplier sets capacity as the capacity planning horizon. Both decision makers are uncertain about the product demand during the capacity planning horizon. Hence, the supplier relies on the demand forecast for his capacity decision. During the capacity planning horizon, the supplier and the manufacturer may obtain new demand information and update their forecasts. The new information includes changes in the general economic conditions, competitor's sales data, and market research done by a third-party research firm, but the two decision makers may also obtain different information. The manufacturer, for example, often has other forward-looking information because of her superior relationship with or proximity to the market and expert opinion about her own product. Such a forecast process for two decision makers; i.e, (X_n^s, X_n^m) can be cast as an m-MMAFE as defined in §3.2.4.

The supplier can use the manufacturer's private information for better capacity planning. However, when asked for this information, the manufacturer has an incentive to inflate her forecast so that the supplier secures more capacity. Credible information sharing requires an appropriate mechanism. This incentive problem has recently been studied by various researchers. (see, for example, Cachon and Lariviere 2001, Özer and Wei 2006). Being aware of this incentive, the supplier can design and offer a mechanism to screen the manufacturer's forecast information truthfully. The screening contract in this case can be interpreted as a capacity reservation contract. This interpretation will become clearer once we solve for the supplier's optimal screening contract. In contrast to the previous literature on credible forecast information sharing, here we study the timing of when the supplier should offer the contract and its connection with the optimal screening contract. By delaying the capacity decision, the supplier can reduce the demand uncertainty that he faces. At the same time, the manufacturer also obtain more demand information. Depending on the information that the two decision makers obtain, the degree of information asymmetry can either increase or decrease over time. Even when it increases, the supplier might be better off by delaying because he can use the manufacturer's more-accurate demand forecast for better capacity planning. This delay may change the capacity cost due to the aforementioned reasons. Considering all these trade-offs, the supplier needs to address the two questions: (1) When is the optimal time to offer a capacity reservation contract to the manufacturer? and (2) What is the optimal capacity reservation contract that maximizes the profit? Because the optimal capacity reservation contract depends on the time when the supplier offers the contract and the demand forecasts at that time, these two questions are closely coupled.

The sequence of events is as follows. At the beginning of period $n \in \{1, 2, \ldots, N\}$ of the capacity planning horizon, the supplier has demand forecast, X_n^s, and decides whether to offer a capacity reservation contract to the manufacturer. If he does not offer a contract, the supplier obtains the forecast update δ_n^s and the manufacturer obtains a forecast update δ_n^m during period n. Otherwise, the supplier offers a menu of contracts $\{K(\xi_n), P(\xi_n)\}$. Both capacity $K(\cdot)$ and corresponding payment $P(\cdot)$ are functions of the manufacturer's private information ξ_n[7]. Given this menu, the manufacturer chooses a particular contract $(K(\check{\xi}), P(\check{\xi}))$ to maximize her profit if her expected profit from the contract is larger than her reservation profit $\underline{\pi}^m$. By doing so, she announces her forecast information to be $\check{\xi}$, which could differ from her true forecast information. The supplier commits to not offering another contract if an offer

[7] Recall from Section 3.2.4 that demand and forecasts have the following relationship; $X_{N+1} = X_n^m \epsilon_n = X_n^s \xi_n \epsilon_n$, where $\epsilon_n \equiv \prod_{k=n}^N \delta_k^m$ represents the demand uncertainty of the system and $\xi_n \equiv \prod_{k=n}^N \delta_k^s / \epsilon_n$ represents the manufacturer's private information.

is declined, and thus the manufacturer chooses a contract and commits to it when the supplier offers a menu of contracts for the first time[8]. The supplier builds $K(\check{\xi})$ units of capacity, and charges $P(\check{\xi})$. The supplier builds capacity at a unit cost of c_n and a fixed cost of C_n. The unit capacity cost c_n and the fixed capacity cost C_n change over time[9]. During the sales period $N+1$, the manufacturer observes demand, X_{N+1}, and places an order at a unit wholesale price of w. The supplier produces the component at a unit production cost of c and fulfils the order to the extent possible given the reserved capacity. Then, the manufacturer sells the final product to the market at a unit retail price of r.[10] Unmet demand is lost, and unsold components have zero salvage value. We denote the c.d.f. and the p.d.f. of ϵ_n by $G_n(.)$ and $g_n(.)$ and those of ξ_n by $F_n(.)$ and $f_n(.)$. We assume that $r > w > c + c_n$ holds for every n. Because of the long leadtime for capacity construction, the capacity decision is irreversible. Therefore, the supplier has a single opportunity to offer a contract to the manufacturer, and we do not consider the possibility of renegotiation after the capacity decision is made.

3.4. Formulation

To solve the aforementioned problem, we formulate a two-stage dynamic program. The first-stage problem is an optimal stopping problem to determine the time to offer an optimal menu of contracts. The second-stage problem solves the mechanism design problem. These two stages are nested optimization problems; i.e., the solution of one stage depends on the solution of the other.

[8] Alternatively, the manufacturer is not aware of the supplier's search of the time to offer contracts, and hence she does not defer accepting a contract to later time periods.

[9] Although it is natural to assume that c_n and C_n are increasing in n, we do not make this assumption to provide a more general result.

[10] The manufacturer may carry out some value added operations that cost, say, m per unit. She sells at a unit price $r' > 0$. Her effective sales price is $r = r' - m$. Hence, without loss of generality, we assume $m = 0$.

3.4.1 The First Stage Problem

At the beginning of period n, the supplier's demand forecast is X_n^s. Given this information, the supplier decides whether to offer a menu of contracts or to delay this offer to the next period;

$$u_n(X_n^s) = \begin{cases} u^s, & \text{offer a contract,} \\ u^d, & \text{delay to offer a contract,} \end{cases}$$

for $n < N$. If the supplier offers the menu of contracts at the beginning of period n, then the state is updated to indicate that the process has already been stopped. To do so, we define the termination state t. If the supplier delays to offer the menu of contracts to the next period, he obtains demand information δ_n^s during period n, and updates his forecast as $X_{n+1}^s = X_n^s \delta_n^s$. Hence, the state transition is

$$X_{n+1}^s = \begin{cases} t, & \text{if } X_n^s = t, \text{ or } X_n^s \neq t \text{ and } u_n(X_n^s) = u^s, \\ X_n^s \delta_n^s, & \text{otherwise.} \end{cases}$$

We denote the supplier's expected profit when he offers the optimal menu of contracts at period n by $\pi_n(X_n^s)$. This profit depends on the optimal mechanism offered in the second stage. We will explicitly define this function in the next section. The reward function $h_n(X_n^s)$ is given by

$$h_n(X_n^s) = \begin{cases} \pi_n(X_n^s), & \text{if } X_n^s \neq t \text{ and } u_n(X_n^s) = u^s, \\ 0, & \text{if } X_n^s = t \text{ or } u_n(X_n^s) = u^d. \end{cases}$$

Let $\mathcal{P} = \{u_1(X_1^s), \ldots, u_N(X_N^s)\}$ represent a policy that determines when to offer a contract. Then, the optimal stopping problem is given as

$$\max_{\mathcal{P}} E\left[\sum_{n=1}^{N} h_n(X_n^s)\right],$$

where the maximization is taken over all admissible policies. The following dynamic programming algorithm provides the solution to this problem:

$$V_N(X_n^s) = \begin{cases} \pi_N(X_n^s), & \text{if } X_n^s \neq t, \\ 0, & \text{if } X_n^s = t, \end{cases}$$

$$V_n(X_n^s) = \begin{cases} \max\{\pi_n(X_n^s), E[V_{n+1}(X_{n+1}^s)|X_n^s]\}, & \text{if } X_n^s \neq t \text{ for } n < N, \\ 0, & \text{if } X_n^s = t \text{ for } n < N. \end{cases}$$

Notice that it is optimal for the supplier to offer a capacity reservation contract at the beginning of period n when $\pi_n(X_n^s) \geq E[V_{n+1}(X_{n+1}^s)|X_n^s]$; otherwise, it is optimal to delay to offer a contract.

3.4.2 The Second Stage Problem

Suppose the supplier offers a menu of contracts $\{K(\xi), P(\xi)\}$ at the beginning of period n. The manufacturer chooses a specific contract $(K(\check{\xi}), P(\check{\xi}))$ that maximizes her expected profit while implying that her forecast information is $\check{\xi}$. Recall that this forecast information could be different from her true information. Hence, by choosing a contract, the manufacturer defines the supplier's expected profit, her expected profit, and the total supply chain's expected profit. As a function of the manufacturer's private forecast information ξ_n, these profits are defined as follows:

$$\begin{aligned} \Pi_n^s(K(\check{\xi}), P(\check{\xi}), \xi_n, X_n^s) &\equiv (w-c)E_{\epsilon_n}[\min(X_n^s\xi_n\epsilon_n, K(\check{\xi}))] + P(\check{\xi}) - (c_nK(\check{\xi}) + C_n), \\ \Pi_n^m(K(\check{\xi}), P(\check{\xi}), \xi_n, X_n^s) &\equiv (r-w)E_{\epsilon_n}[\min(X_n^s\xi_n\epsilon_n, K(\check{\xi}))] - P(\check{\xi}), \\ \Pi_n^{tot}(K(\check{\xi}), \xi_n, X_n^s) &\equiv (r-c)E_{\epsilon_n}[\min(X_n^s\xi_n\epsilon_n, K(\check{\xi}))] - (c_nK(\check{\xi}) + C_n). \end{aligned} \tag{3.2}$$

The supplier's objective is to design a menu of contracts that maximizes his expected profit among all possible contracts. From the *revelation principle* (Myerson 1982), the supplier can limit the search for the optimal menu of contracts to the class of incentive-compatible, direct-revelation contracts. Under this type of contract, the manufacturer credibly reports her private information ξ_n. In this case, the supplier can screen the true value of ξ_n from the manufacturer's contract choice. Hence, the supplier's expected profit from offering an incentive-compatible, direct-revelation contract $(K(\xi), P(\xi))$ is $\Pi_n^s(K(\xi_n), P(\xi_n), \xi_n, X_n^s)$. To identify an optimal menu of contracts, the supplier solves

$$\begin{aligned} \pi_n(X_n^s) \quad = \quad & \max_{K(.),P(.)} E_{\xi_n}[\Pi_n^s(K(\xi_n), P(\xi_n), \xi_n, X_n^s)] \text{ subject to} \\ & (IC) : \Pi_n^m(K(\xi), P(\xi), \xi, X_n^s) \geq \Pi_n^m(K(\check{\xi}), P(\check{\xi}), \xi, X_n^s) \text{ for every } \check{\xi} \neq \xi \\ & (PC) : \Pi_n^m(K(\xi), P(\xi), \xi, X_n^s) \geq \underline{\pi}^m \text{ for every } \xi. \end{aligned} \tag{3.3}$$

The first constraint is the incentive compatibility constraint (IC), which ensures the manufacturer's credible information sharing. The second constraint is the participation constraint (PC), which ensures the manufacturer's participation on the contract. Note that the solution of this problem is a function of the supplier's forecast information, which depends on his timing decision.

3.5. Analysis

To solve the first-stage optimal stopping problem, we need the optimal profit obtained in the second-stage, which is $\pi_n(X_n^s)$. Hence, we solve the mechanism design problem first and embed its solution to the optimal stopping problem to determine the optimal time to offer the optimal menu of contracts.

3.5.1 Optimal Capacity Reservation Contract

Once the supplier decides to offer a capacity reservation contract at period n, he determines the optimal contract given his forecast information X_n^s, market uncertainty ϵ_n, and the belief about the manufacturer's private forecast information ξ_n. When the forecast model follows an m-MMAFE, the random variables ϵ_n and ξ_n are log-normally distributed. The result of this subsection also holds for other random variables. Hence, we do not assume that ϵ_n and ξ_n are log-normally distributed in this subsection. We denote the optimal menu of contracts that solves the optimization problem in (3.3) by $\{K_n^{dc}(\xi), P_n^{dc}(\xi)\}$, where $\epsilon_n \in [\underline{\epsilon}_n, \overline{\epsilon}_n]$ and $\xi_n \in [\underline{\xi}_n, \overline{\xi}_n]$.

Lemma 3.1. *(a) A menu of contracts* $\{K(\xi), P(\xi)\}$ *is feasible for (3.3) if and only if it satisfies the following conditions for all* $\xi \in [\underline{\xi}_n, \overline{\xi}_n]$*:*

$$(i) \quad \Pi_n^m(K(\xi), P(\xi), \xi, X_n^s) = \underline{\pi}^m + \int_{\underline{\xi}_n}^{\xi} \left[(r-w) X_n^s \int_{\underline{\epsilon}_n}^{\frac{K(x)}{x X_n^s}} y g_n(y) dy \right] dx,$$

(ii) $K(\xi)$ *is increasing in* ξ*.*

(b) The optimization problem (3.3) has the following equivalent formulation:

$$\begin{aligned} \pi_n(X_n^s) \quad &= \quad \max_{K(.)} E_{\xi_n} \Big[(r-c) E_{\epsilon_n} [\min(X_n^s \xi_n \epsilon_n, K(\xi_n))] - (c_n K(\xi_n) + C_n) \\ &\quad - \frac{1 - F_n(\xi_n)}{f_n(\xi_n)} (r-w) X_n^s \int_{\underline{\epsilon}_n}^{\frac{K(\xi_n)}{\xi_n X_n^s}} y g_n(y) dy \Big] - \underline{\pi}^m \qquad (3.4) \\ s.t. \quad &\quad K(\xi) \text{ is increasing.} \end{aligned}$$

After determining the optimal capacity reservation function $K_n^{dc}(.)$ from (3.4), we derive the corresponding payment as:

$$\begin{aligned} P_n^{dc}(\xi) \quad &= \quad (r-w) E_{\epsilon_n} [\min(X_n^s \xi \epsilon_n, K_n^{dc}(\xi))] \\ &\quad - \int_{\underline{\xi}_n}^{\xi} \left[(r-w) X_n^s \int_{\underline{\epsilon}_n}^{\frac{K_n^{dc}(x)}{x X_n^s}} y g_n(y) dy \right] dx - \underline{\pi}^m. \qquad (3.5) \end{aligned}$$

Note that the optimal menu of contracts depends on the forecast, X_n^s. This dependency does not result in the analytical complexity for the structural properties of the optimal mechanism and the optimal stopping policy. However, to *numerically* solve the optimal stopping problem, one needs to determine the optimal capacity reservation contract for every X_n^s to derive $\pi_n(X_n^s)$. To reduce the computational complexity, we introduce a normalized version of (3.4), which we define as follows:

$$\hat{\pi}_n \equiv \max_{K(.)} E_{\xi_n}\Big[(r-c)E_{\epsilon_n}[\min(\xi_n\epsilon_n, K(\xi_n))] - c_n K(\xi_n) - \frac{1-F_n(\xi_n)}{f_n(\xi_n)}(r-w)\int_{\underline{\epsilon}_n}^{\frac{K(\xi_n)}{\xi_n}} y g_n(y)dy\Big]$$
$$s.t. \quad K(\xi) \text{is increasing.} \tag{3.6}$$

We denote the optimal solution of (3.6) by $\hat{K}_n^{dc}(.)$, and we define

$$\hat{P}_n^{dc}(\xi) \equiv (r-w)E_{\epsilon_n}[\min(\xi\epsilon_n, \hat{K}_n^{dc}(\xi))] - \int_{\underline{\xi}_n}^{\xi}\left[(r-w)\int_{\underline{\epsilon}_n}^{\frac{\hat{K}_n^{dc}(x)}{x}} y g_n(y)dy\right]dx.$$

Then, we have the following properties of the optimal contract and the expected profit:

Theorem 3.4. *The following statements are true for all n:*

(a) $K_n^{dc}(\xi) = X_n^s \hat{K}_n^{dc}(\xi)$.

(b) $P_n^{dc}(\xi) = X_n^s \hat{P}_n^{dc}(\xi) - \underline{\pi}^m$.

(c) $\pi_n(X_n^s) = X_n^s \hat{\pi}_n - C_n - \underline{\pi}^m$.

These properties and results show that the optimal menu of contracts is state-dependent. The menu depends on the supplier's forecast information at the time the menu is offered to the manufacturer. The multiplicative nature of the results is due to the multiplicative nature of forecast updates. If we ignore the fixed parts $\underline{\pi}^m$ and C_n, the forecast X_n^s scales the problem (3.4). However, the supplier needs to guarantee $\underline{\pi}^m$ to ensure the manufacturer's participation regardless of X_n^s. Hence, in Part (a) and (b) of the theorem, the optimal capacity reservation and prices are linearly increasing in X_n^s, where the capacity prices have the y-intercept of $-\underline{\pi}^m$. Note that the unit capacity reservation price is $\frac{P_n^{dc}(\xi)}{K_n^{dc}(\xi)} = \frac{\hat{P}_n^{dc}(\xi)}{\hat{K}_n^{dc}(\xi)} - \frac{\underline{\pi}^m}{X_n^s \hat{K}_n^{dc}(\xi)}$. Without the reservation profit, the supplier would charge a unit capacity reservation price of $\hat{P}_n^{dc}(\xi)/\hat{K}_n^{dc}(\xi)$ to the manufacturer for all X_n^s. However, to ensure the participation of the manufacturer, the supplier discounts $\frac{\pi^m}{X_n^s \hat{K}_n^{dc}(\xi)}$ from the unit reservation price. The supplier also incurs the fixed capacity cost C_n regardless of X_n^s. Hence, in Part (c), the supplier's optimal expected profit is linearly increasing in X_n^s with the y-intercept of $-C_n - \underline{\pi}^m$. This result also implies that one needs to solve only N optimization problems to derive $\pi_n(X_n^s)$ for every decision period n to numerically solve the optimal stopping problem.

Next, we discuss how to solve (3.6), which is a function-space optimization problem. Without the constraint that $K(\xi)$ is increasing, we can derive the optimal function $K(.)$ by choosing the value K that maximizes the inner function of (3.6) for each ξ. If the optimal $K(.)$ of the unconstrained problem is increasing in ξ, then it is the optimal solution of (3.6). We prove that this approach works when ϵ_n and ξ_n have increasing generalized failure rates (IGFR)[11]. Note that the log-normal random variables have IGFRs, hence the following result holds when the demand forecast is an m-MMAFE.

Theorem 3.5. *If ϵ_n and ξ_n have IGFRs, then the following properties hold:*

(a) $\hat{K}_n^{dc}(\xi)$ is the unique solution of the first-order condition

$$(r-c)(1-G_n(\frac{K}{\xi})) - c_n - \frac{1-F_n(\xi)}{\xi f_n(\xi)}(r-w)\frac{K}{\xi}g_n(\frac{K}{\xi}) = 0. \tag{3.7}$$

(b) Both $K_n^{dc}(\xi)$ and $P_n^{dc}(\xi)$ are increasing in ξ.

(c) Both $\Pi_n^s(K_n^{dc}(\xi), P_n^{dc}(\xi), \xi, X_n^s)$ and $\Pi_n^m(K_n^{dc}(\xi), P_n^{dc}(\xi), \xi, X_n^s)$ are increasing in ξ.

(d) $P_n^{dc}(K_n^{dc})$ is an increasing concave function of K_n^{dc}.

Together with Theorem 3.4 and Equation (3.5), Part (a) fully characterizes the optimal menu of contracts. Part (b) implies that the manufacturer reserves more capacity and pays more when her forecast is larger. We can represent the manufacturer's

[11] The generalized failure rate of a random variable is defined as $\frac{xf(x)}{1-F(x)}$, where $f(x)$ and $F(x)$ are the p.d.f. and the c.d.f. of the random variable. All random variables with increasing failure rates have IGFRs, and other common classes of random variables have IGFRs, including Log-normal, Gamma and Weibull.

expected profit under the optimal capacity reservation contract as

$$\Pi_n^m(K_n^{dc}(\xi), P_n^{dc}(\xi), \xi, X_n^s) = \underline{\pi}^m + (\Pi_n^m(K_n^{dc}(\xi), P_n^{dc}(\xi), \xi, X_n^s) - \underline{\pi}^m).$$

The first term is the manufacturer's reservation profit, and the second term is the informational rent of the manufacturer. This rent is the price that the supplier has to pay to learn the manufacturer's private information. When the manufacturer has no private information, i.e., $\sigma_{\xi_n} = 0$, the informational rent is 0. Then, Part (c) implies that the manufacturer's informational rent increases in ξ. Because both $K_n^{dc}(\xi)$ and $P_n^{dc}(\xi)$ are increasing in ξ, we can map the capacity K_n^{dc} to the price P_n^{dc} directly without ξ. This mapping is denoted by $P_n^{dc}(K_n^{dc})$. For this reason, Özer and Wei (2006) referred to this contract as the capacity reservation contract[12]. To reserve K units of capacity, the manufacturer has to pay $P_n^{dc}(K)$ to the supplier. The capacity reservation contract, $P_n^{dc}(K)$, is a concave increasing function as shown in Part (d). The concavity implies that the marginal reservation price is decreasing in the capacity level, hence the supplier provides more incentive to a manufacturer who has a higher forecast.

3.5.2 Optimal Time to Offer the Contract

The previous section reveals that primarily three factors affect the timing decision: cost of demand uncertainty, cost of asymmetric information (i.e., cost of screening), and cost of capacity. When the supplier offers the contract at period n, his resulting expected profit is given in Theorem 3.4(c). This profit consists of the variable part $X_n^s \hat{\pi}_n$ and the fixed part $-C_n - \underline{\pi}^m$. The normalized profit $\hat{\pi}_n$ in Equation (3.6) is determined by demand uncertainty ϵ_n, information asymmetry ξ_n, and the unit capacity cost c_n. The impact of these factors on the supplier's profit is proportional to his demand forecast. In contrast, the impact of the fixed capacity cost C_n is independent of the demand forecast. The first two terms in Equation (3.6) represent

[12]We remark that Özer and Wei (2006) studied the static mechanism design problem for the additive uncertainty. This chapter considers the dynamic version of this problem for the multiplicative uncertainty that is resolved through multiple forecast updates. In the Appendix, we provide the dynamic version of the problem for the additive case.

the total supply chain profit, which is increasing with reduced demand uncertainty. The last term can be interpreted as the cost of screening[13]. If the supplier offers the menu of contracts in the next period instead of the current period, his profit increases when the benefit of reduced demand uncertainty outweighs the change in the cost of asymmetric information and the change in the capacity cost. The optimal stopping formulation considers these trade-offs in determining the optimal time to offer the menu of contracts. The following theorem characterizes the optimal stopping policy.

Theorem 3.6. *A control band policy that offers a capacity reservation contract at period n if $X_n^s \in [L_n, U_n]$ is optimal, and the optimal thresholds are given as $U_n \equiv \sup\{X_n^s : \pi_n(X_n^s) \geq E[V_{n+1}(X_{n+1}^s)|X_n^s]\}$, and $L_n \equiv \inf\{X_n^s : \pi_n(X_n^s) \geq E[V_{n+1}(X_{n+1}^s)|X_n^s]\}$.*

Theorem 3.6 establishes the optimality of a control band policy. Under this policy, the supplier offers the capacity reservation contract at period n if the current forecast falls within the control band, i.e., $X_n^s \in [L_n, U_n]$. When the demand forecast is larger than U_n, the benefit of reducing demand uncertainty and asymmetric information is significant and outweighs the cost of delaying. Recall that demand uncertainty is multiplicative. Hence, the supplier can significantly increase his profit by reducing the degree of demand uncertainty when the forecast is large. In contrast, the potential increase in the fixed capacity cost is constant regardless of the forecast level. In this case, delaying to offer the contract is optimal when $X_n^s > U_n$. When the demand forecast is smaller than L_n, the cost of delaying to offer the contract (e.g., possible increase in the unit capacity cost) is negligible. In contrast, the potential decrease in the fixed capacity cost is again constant regardless of the forecast level. In this case, delaying to offer the contract is optimal. The following set of results further characterizes the optimal stopping policy and sheds some light onto these trade-offs.

Theorem 3.7. *The following statements are true for all n:*

(a) When $C_{n+1} > C_n$ for all n, the lower threshold, L_n, is 0 for all n. Hence, an

[13] Without this last term, the supplier's objective would be the same as that of a centralized system. This term is the information rent that the supplier (and also the total supply chain) has to forgo to enable credible information sharing.

upper threshold policy that offers the capacity reservation contract at period n if $X_n^s \leq U_n$ is optimal.

(b) *Let $n^* \equiv \arg\max_n \hat{\pi}_n$. When $C_{n+1} = C_n$ for all n, the lower and upper thresholds satisfy that $L_n = U_n = 0$ for $n \neq n^*$, and that $L_n = 0$ and $U_n = \infty$ for $n = n^*$. Hence, a state-independent stopping policy that offers the capacity reservation contract at period n^* is optimal.*

When the supplier always incurs more fixed capacity costs by delaying the capacity decision, he should optimally offer the contract when his demand forecast is small. As the forecast level converges to 0, the benefit of reduced demand uncertainty and information asymmetry vanishes, whereas the loss in the fixed capacity cost remains constant. Hence, the supplier should delay to offer the contract only when the demand forecast is large, which implies that the lower threshold is 0 and that an upper threshold policy is optimal. When the fixed capacity cost is constant over time, the optimal time to offer the contract is fully determined by the trade-offs among the demand uncertainty, information asymmetry, and the unit capacity cost, whose impacts on the supplier's profit are all proportional to the forecast level. Hence, if the expected profit of offering the contract at period n is greater than the expected profit of offering the contract at period m, then it is true regardless of the forecast level. Thus, the optimal policy is to always stop at the optimal stopping period, n^*, at which the normalized expected profit is the greatest.

3.6. Centralized Supply Chain

In a centralized system, a single decision maker determines both decisions: (1) the optimal time to decide and build capacity and (2) the optimal capacity level that maximizes the total supply chain profit. Note that the demand forecast for the centralized decision maker is X_n^m because the centralized decision maker would have access to all information; i.e., private information concept is irrelevant in this case. At the beginning of each period n, the centralized decision maker decides whether to set the capacity level or delay the decision to the next period. This problem can be formulated as a two-stage dynamic program similar to that in §3.4. However, in this case, the second-stage is a simple newsvendor-type optimization problem instead of a mechanism design problem.

First we consider the second-stage problem. The centralized decision maker determines the optimal capacity level that maximizes the total supply chain expected profit as:

$$\pi_n^{cs}(X_n^m) \equiv \max_K (r-c)E_{\epsilon_n}[\min(X_n^m\epsilon_n, K)] - (c_nK + C_n).$$

We define the normalized expected profit $\hat{\pi}_n^{cs} \equiv \max_K(r-c)E_{\epsilon_n}[\min(\epsilon_n, K)] - c_nK$, and then we can represent the optimal expected profit as $\pi_n^{cs}(X_n^m) = X_n^m\hat{\pi}_n^{cs} - C_n$. The optimal capacity level is given as $K_n^{cs}(X_n^m) \equiv X_n^m G_n^{-1}(\frac{r-c-c_n}{r-c})$.

Next we consider the first-stage optimal stopping problem. The formulation is similar to that of the decentralized supplier's problem in §3.4.1, hence we focus on the differences between the two problems. In the centralized case, the state variable is given by X_n^m instead of X_n^s. If the centralized decision maker delays the capacity decision, she obtains the forecast update δ_n^m and updates the forecast as $X_{n+1}^m = X_n^m\delta_n^m$. If she stops and sets the capacity, the reward of stopping at period n is $\pi_n^{cs}(X_n^m)$. Hence, to solve the problem for a centralized system, we replace X_n^s with X_n^m and $\pi_n(X_n^s)$ with $\pi_n^{cs}(X_n^m)$ in the formulation of §3.4.1. We denote the value-to-go function of the centralized case by $V_n^{cs}(X_n^m)$. Then, we have the following results:

Theorem 3.8. *For the centralized supply chain, the following statements are true for all n:*

(a) *A control band policy that determines the capacity level at period n if $X_n^m \in [L_n^{cs}, U_n^{cs}]$ is optimal, and the optimal thresholds are given as $U_n^{cs} \equiv \sup\{X_n^m : \pi_n^{cs}(X_n^m) \geq E[V_{n+1}^{cs}(X_{n+1}^m)|X_n^m]\}$, and*
$L_n^{cs} \equiv \inf\{X_n^m : \pi_n^{cs}(X_n^m) \geq E[V_{n+1}^{cs}(X_{n+1}^m)|X_n^m]\}$.

(b) *When $C_{n+1} > C_n$ for all n, an upper threshold policy that determines the capacity level at period n if $X_n^m \leq U_n^{cs}$ is optimal.*

(c) *When $C_{n+1} = C_n$ for all n, a state-independent policy that determines the* *capacity level at period n^{cs} is optimal, where the optimal stopping period is given as $n^{cs} \equiv \arg\min_n \hat{\pi}_n^{cs}$.*

(d) The capacity level of the centralized supply chain is equal to or greater than the capacity level of the decentralized supply chain, i.e., $K_n^{cs}(\xi) \geq K_n^{dc}(\xi)$.

The centralized decision maker's optimal stopping policy has the same structure as the supplier's policy discussed in the previous section. Note, however, that this theorem does not imply that the optimal stopping thresholds for the centralized and decentralized systems are the same. The decentralized supplier may determine the capacity level at a different time than the centralized decision maker, which is one source of channel inefficiency. Part (d) of the theorem also shows that the decentralized supplier builds less capacity than the centralized decision maker, which is another source of channel inefficiency. We quantify these differences through a numerical study.

3.7. Numerical Study

The purpose of this section is four-fold. First we illustrate some of our results through numerical examples, such as the form of optimal capacity reservation contract. Second, we compare the capacity levels and profits of the centralized supply chain with those of decentralized supply chain. This comparison enables us to quantify the cost of suboptimal timing and capacity decisions on profits; i.e., inefficiency due to decentralization. We also compare the optimal thresholds of the decentralized and centralized systems. These comparisons characterize the environments in which the supplier's self-interested strategy deviates largely from the socially optimal action. Third, we quantify the impact of the three factors - market uncertainty, information asymmetry and capacity cost - on the expected profits as well as the optimal stopping thresholds. These comparisons help us to characterize when the supplier should offer the contract early or late. Finally, we compare our model with a static model in which the supplier offers the menu of contracts at a fixed time to quantify the value of determining the optimal time to offer the capacity reservation contract.

3.7.1 Optimal Capacity Reservation Contract

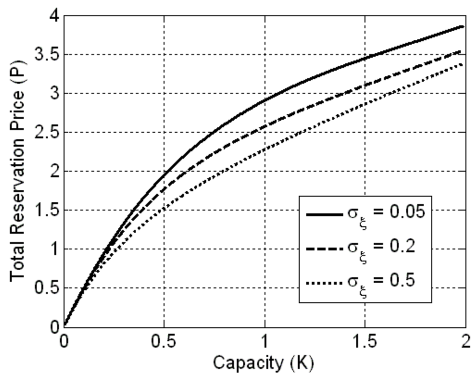

(a) Optimal Menu of Contract

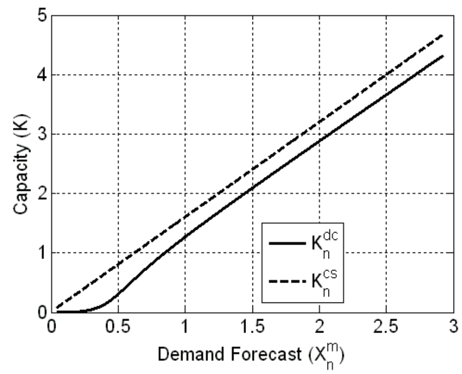

(b) Capacity Levels for Centralized and Decentralized Systems

Figure 3.3: Optimal Capacity Reservation Contract

Figure 3.3(a) illustrates the optimal capacity reservation contracts for three different levels of information asymmetry, which is measured by the standard deviation of $\log(\xi_n)$, i.e., σ_{ξ_n}. The supply chain setting for this test is $\sigma_{\xi_n} \in \{0.05, 0.2, 0.5\}$, $r = 15$, $c = 3$, $w = 10$, $c_n = 2$, $C_n = 0$, $\underline{\pi}^m = 0$, $\sigma_{\epsilon_n} = 1$, and $X_n^s = 1$. As shown in Theorem 3.5 Part (e), the optimal capacity reservation contract $P_n^{dc}(K)$ is an increasing concave function of K. This figure also shows that the supplier charges less to reserve capacity as the degree of information asymmetry, σ_{ξ_n} increases. In other words, the supplier's cost of screening the manufacturer's private information decreases with the degree of information asymmetry. Figure 3.3(b) illustrates the capacity level determined by the optimal capacity reservation contract and the optimal capacity level of the centralized system when $\sigma_{\xi_n} = 0.5$. In both cases, the decision makers builds more capacity as the forecast increases, but the supplier always builds less capacity than a centralized decision maker.

When the supplier delays to offer a capacity reservation contract, the expected profits of each decision maker change due to the changes in the three values: the degree of the demand uncertainty σ_{ϵ_n}, the degree of information asymmetry σ_{ξ_n}, and the capacity cost (c_n, C_n). Table 3.1 shows the expected profits and the total supply chain profit for various values.

Table 3.1: Expected Profits When the Supplier Offers the Contract at Period n

σ_{ϵ_n}	Π_n^s	Π_n^m	Π_n^{tot}	σ_{ξ_n}	Π_n^s	Π_n^m	Π_n^{tot}	c_n	Π_n^s	Π_n^m	Π_n^{tot}
1.00	35.74	19.49	55.22	0.80	33.21	21.16	54.37	2.0	35.74	19.49	55.22
0.95	37.68	19.78	57.46	0.75	33.54	20.94	54.48	2.3	32.26	18.64	50.90
0.90	39.63	20.06	59.69	0.70	33.90	20.70	54.60	2.6	29.07	17.88	46.96
0.85	41.59	20.32	61.91	0.65	34.29	20.44	54.73	2.9	26.14	17.20	43.34
0.80	43.56	20.55	64.11	0.60	34.72	20.15	54.88	3.2	23.42	16.58	40.00
0.75	45.54	20.75	66.28	0.55	35.20	19.84	55.04	3.5	20.90	16.01	36.91
0.70	47.52	20.91	68.43	0.50	35.74	19.49	55.22	3.8	18.54	15.49	34.03
$c_n = 2,\ \sigma_{\xi_n} = 0.5$				$c_n = 2,\ \sigma_{\epsilon_n} = 1$				$\sigma_{\epsilon_n} = 1,\ \sigma_{\xi_n} = 0.5$			
$r = 15,\ c = 3,\ w = 10,\ C_n = 0,\ \underline{\pi}^m = 10,\ X_n^s = 10$											

The first section of Table 3.1 shows that the expected profits of both decision makers increase as the demand uncertainty decreases. The second section of Table 3.1 shows that the supplier's profit and the total supply chain profit increase but the manufacturer's expected profit decreases as the degree of information asymmetry decreases. In other words, the manufacturer's informational rent decreases with the degree of information asymmetry. The third section of Table 3.1 shows that the increased capacity cost reduces the expected profits of both decision makers.

3.7.2 Optimal Time to Offer the Contract

The basic supply chain setting for this and the next subsection is $N = 6$, $r = 15$, $c = 3$, $w = 10$, $\underline{\pi}^m = 10$, $c_n = 2$, and $C_n = 20+(n-1)\Delta_C$. The term Δ_C indicates the rate of increase in the fixed cost of capacity. We set the parameters for the forecast evolution as $\sigma_{\delta_N^m} = 1$, $\sigma_{\delta_N^s} = \sqrt{1.6}$, $\sigma_{\delta_n^m} = \sigma^m$ and $\sigma_{\delta_n^s} = \sigma^s$ for $n = 1, 2, \ldots, N-1$. The value of σ^i quantifies the amount of information that decision maker i obtains at each period of the capacity planning horizon, and the value of $\sigma_{\delta_N^i}$ indicates the degree of residual demand uncertainty of decision maker i at the end of the capacity planning horizon. When $\sigma^m < \sigma^s$, the supplier obtains more information than the manufacturer during the capacity planning horizon. We set the values of σ^s and σ^m such that the forecast model follows an m-MMAFE, i.e., we maintain $(n-1)(\sigma^s)^2+\sigma^2_{\delta_N^s} \geq (n-1)(\sigma^m)^2+\sigma^2_{\delta_N^m}$

for every $n = 1, 2, \ldots, N$.

Table 3.2 shows the optimal thresholds U_n of the decentralized supplier and the optimal thresholds U_n^{cs} of the centralized decision maker for several values of σ^s, σ^m, and Δ_C. Recall that when C_n is increasing in n, the optimal stopping policy is an upper threshold policy. Hence, U_n is sufficient to describe the optimal stopping decision. The difference between U_n and U_n^{cs} indicates how much the decentralized supplier's strategy deviates from the socially optimal action[14]. When the optimal

Table 3.2: Optimal Thresholds of Decentralized and Centralized Supply Chains

	$(\sigma^m)^2$	$(\sigma^s)^2$	Δ_C	U_1	U_3	U_5	U_1^{cs}	U_3^{cs}	U_5^{cs}
(a)	0.1	0.1	8	39.49	41.12	44.44	30.13	31.49	34.30
(b)	0.1	0.2	8	31.39	33.25	36.83	30.13	31.49	34.30
(c)	0.2	0.1	8	31.44	26.97	26.11	15.38	16.03	17.57
(d)	0.1	0.1	4	19.75	20.56	22.22	15.07	15.75	17.15

threshold of period n is large, the optimal policy delays to offer a contract at period n only when the forecast is large. Therefore, high threshold levels indicate that the decision maker tends to offer a contract early and low threshold levels indicate the opposite.

First we focus on the optimal thresholds of the decentralized supplier. Case (a) of Table 3.2 is the basis of our experiment. In case (b), the supplier obtains more information than he does in case (a). Therefore, the supplier waits longer in case (b) than in case (a) in order to obtain more demand information. In case (c), the manufacturer obtains more information than she does in case (a), but the supplier's obtains the same amount of information in both cases. Even in this case, the supplier waits longer than in case (a) so that the manufacturer obtains more demand information. It is a surprising result because the delay of offering the contract increases the degree of information asymmetry in case (c). Therefore, the increase of the degree of information asymmetry is not necessarily bad for the supplier if by offering the contract later, the supplier can reduce the demand uncertainty through screening the manufacturer who obtains more information in the current period. Finally, the result

[14] Even though the stopping decision of the decentralized supplier is made based on X_n^s and the centralized decision maker on X_n^m, the forecasts X_n^s and X_n^m are the same in expectation.

of case (d) implies that the supplier offers a contract early when the capacity cost increases rapidly. To summarize, the supplier waits longer when the decision makers obtain more demand information over the capacity planning horizon and when the capacity cost increase less rapidly.

Next, we focus on the difference between U_n and U_n^{cs}. The centralized system of case (b) is identical to that of case (a) because in both cases the centralized system has the same information. In case (a), the information asymmetry between the two decision makers remains constant over time. In this case, the decentralized supplier stops the process earlier than the centralized decision maker, which implies that the supplier's self-interested strategy keeps the system from waiting until the socially optimal time and hence reduces channel efficiency. In case (b), the information asymmetry between the two decision makers decreases over time. Therefore, the decentralized supplier waits longer because he can reduce the degree of demand uncertainty and information asymmetry at the same time. Even though the supplier's optimal policy is close to the socially optimal policy, this results not from altruism but from the supplier's self-interest. In case (c), the information asymmetry between the two decision makers increases over time. Therefore, the supplier offers a contract much earlier than the socially optimal point to prevent the information asymmetry from growing too much. Finally, case (d) shows that both the decentralized supplier and the centralized decision maker commit early when the capacity cost increases rapidly. From this result, we can conclude that the supplier offers a contract earlier than the socially optimal time when the manufacturer obtains much information.

3.7.3 Value of Determining the Optimal Time

In this subsection, we assess the value of determining the optimal time to offer a capacity reservation contract. Clearly, the supplier is better off by optimizing the time to offer a contract. However, it is unclear whether or not this strategy benefits the manufacturer and the total supply chain. To clarify this point, we compare the profits from our model with those from a static model in which the supplier offers a contract at a fixed period regardless of the forecast level. We consider N static models, which offer the contract at each fixed period $n \in \{1, 2, \ldots, N\}$, and report the average gain and the lowest gain in the expected profits. We denote the percentage improvement in the supplier's expected profit of using the dynamic model compared to profit from the static model that offers the contract at period n by

$$I_n^s \equiv \frac{V_1(X_1^s) - E[\pi_n(X_n^s)|X_1^s]}{E[\pi_n(X_n^s)|X_1^s]} \times 100\%.$$

Similarly, we denote the percentage improvement in the manufacturer's profit and the total supply chain profit by I_n^m and I_n^{tot}. For each $j \in \{s, m, tot\}$, we define $I_{ave}^j \equiv \frac{\sum_{n=1}^N I_n^j}{N}$ and $I_{min}^j \equiv \min\{I_1^j, \ldots, I_N^j\}$.

Table 3.3 shows these values for several supply chain settings. The basic setting is the same as that we use in Table 3.2. The improvement in the supplier's profit is

Table 3.3: Percentage Improvements in Expected Profits

	$(\sigma^m)^2$	$(\sigma^s)^2$	Δ_C	X_1^s	I_{ave}^s	I_{min}^s	I_{ave}^m	I_{min}^m	I_{ave}^{tot}	I_{min}^{tot}
(a)	0.1	0.1	8	41.12	4.21	0.63	1.10	-5.43	3.12	2.66
(b)	0.1	0.1	8	45.00	3.94	2.01	2.15	-4.55	3.32	1.82
(c)	0.2	0.1	8	26.97	14.85	0.00	-22.55	-38.00	-1.02	-5.70
(d)	0.2	0.1	8	35.00	7.57	3.23	15.45	-9.18	9.31	-0.37
(e)	0.1	0.2	8	33.25	8.27	1.19	0.49	-2.09	5.31	2.15
(f)	0.1	0.2	8	37.00	7.39	3.25	0.89	-1.76	5.07	3.64
(g)	0.1	0.1	4	20.56	5.79	0.85	0.88	-4.61	3.55	3.03
(h)	0.1	0.1	4	23.00	5.09	3.03	2.01	-3.74	3.80	1.84

always positive, and we observe $I_{min}^s = 0$ only in case (c). In this case, the initial forecast X_1^s is below the optimal threshold U_1^s. Therefore, the supplier stops at the initial period, and the expected profit from the dynamic strategy is the same as that from the static model that stops at period 1. The average improvements in the supplier's profit range from 3.94% to 14.85% in our experiment. Therefore, we can conclude that the supplier can significantly improve his expected profit by optimally determining the time to offer a contract. However, this strategy may not lead to an improvement in the total supply chain profit. In cases (c) and (d), I_{min}^{tot} is negative, and in case (c), even I_{ave}^{tot} is negative. As discussed in the previous subsection, the supplier in cases (c) and (d) offers the contract very early to prevent the manufacturer from obtaining too much information. This example verifies that the mechanism designer's self-interested strategy to determine the optimal time can decrease the total supply chain profit. The impact of this strategy is even worse on the manufacturer. In Table 3.3, I^m_{min} is always negative, and I^m_{ave} is much smaller than I^s_{ave} except for case (d).

3.8. Extensions and Generalizations

In this section, we discuss some extensions and generations of our model. We discuss how the structural properties of the capacity planning problem would change in each of the generalized models. We also discuss the MMFE for more than two decision makers.

3.8.1 Endogenous Wholesale Price

In previous sections, we have assumed that the wholesale price is determined before the beginning of the capacity planning horizon. In this subsection, we consider the case in which the supplier determines the wholesale price w when he offers the screening contract. In this case, the supplier includes $w(\xi)$ in his menu of contracts, i.e., he offers $\{K(\xi), P(\xi), w(\xi)\}$ to the manufacturer. As before, the manufacturer chooses a specific contract $(K(\check{\xi}), P(\check{\xi}), w(\check{\xi}))$ that maximizes her expected profit given this menu. We denote the optimal menu of contracts at period n by $\{K^{dc}_n(\xi), P^{dc}_n(\xi), w^{dc}_n(\xi)\}$.

Theorem 3.9. *The optimal menu of contracts is given as* $w^{dc}_n(\xi) = r$, $K^{dc}_n(\xi) = X^s_n \xi G^{-1}_n(\frac{r-c-c_n}{r-c})$, *and* $P^{dc}_n(\xi) = -\underline{\pi}^m$. *The supplier's optimal expected profit is given as* $\pi^s_n(X^s_n) = X^s_n \hat{\pi}^{cs}_n - C_n - \underline{\pi}_n$.

It is important to note that the manufacturer's unit profit margin is 0 under the optimal menu of contracts because $w = r$. This result implies that with the ability to determine the wholesale price w, the supplier obtains the manufacturer's business in exchange for the manufacturer's reservation profit. Hence, the capacity level determined by the optimal contract is the same as the optimal capacity level of the centralized supply chain, and the supplier's optimal expected profit has the same structure as in the original problem except that $\hat{\pi}_n^s$ is now replaced by $\hat{\pi}_n^{cs}$. Because the structure of the optimal stopping policy does not depend on the value of $\hat{\pi}_n^s$ in the original problem, Theorems 3.6 and 3.7 hold for the endogenous wholesale price case.

3.8.2 Forecast Update Costs

Previously, the decision makers obtain new demand information over time without incurring an additional cost specific to each update (as in Özer et al. 2007 and all forecast-related papers referenced in this chapter except for Taylor and Xiao 2009, and Ulu and Smith 2009). Firms update forecasts as they get closer to the sales period because they obtain information about, for example, the overall economy, consumer tastes or past sales data. Often forecast related costs are sunk because firms invest in forecasting upfront regardless of whether they obtain information updates. Hence, most managerial decisions and related literature treats these costs as exogenous to the decision problem. However, there may be cases in which firms may incur some fixed cost to obtain more information (as in Taylor and Xiao 2009, and Ulu and Smith 2009). Let κ_n^s be the cost to obtain the demand information δ_n^s for the supplier, and let κ_n^m be the fixed cost to obtain the demand information δ_n^m for the manufacturer. Because the manufacturer incurs extra costs to update demand information, the manufacturer's reservation profit is now $\underline{\pi}_n^m = \underline{\pi}_1^m + \sum_{k=1}^{n-1} \kappa_k^m$. In this case, the fixed part of the expected profit $\pi_n(X_n^s)$ becomes $-C_n - \underline{\pi}_n^m$ instead of $-C_n - \underline{\pi}^m$. In addition, the reward function in §3.4.1 is now

$$h_n(X_n^s) = \begin{cases} \pi_n(X_n^s), & \text{if } X_n^s \neq t \text{ and } u_n(X_n^s) = u^s, \\ -\kappa_n^s, & \text{if } X_n^s = t \text{ or } u_n(X_n^s) = u^d. \end{cases}$$

The optimal stopping policy is also a control-band policy. However, now the fixed loss of delaying to offer a contract at period n is $\kappa_n^s + (C_{n+1} + \underline{\pi}_{n+1}^m - C_n - \underline{\pi}_n^m) = \kappa_n^s + \kappa_n^m + (C_{n+1} - C_n)$. Because $\kappa_n^s + \kappa_n^m + (C_{n+1} - C_n) \geq C_{n+1} - C_n \geq 0$, the optimal stopping policy is still the upper threshold policy when $C_{n+1} > C_n$ for all n. However, even when $C_{n+1} = C_n$, the fixed loss of delaying is non-negative, hence Part (b) of Theorem 3.7 does not hold any more.

3.8.3 State Dependent Reservation Profit

In some cases, the manufacturer's reservation profit is proportionally increasing in her demand forecast X_n^m, i.e., the reservation profit is given as $X_n^m \underline{\pi}^m = X_n^s \xi_n \underline{\pi}^m$. In this case, the mechanism design problem (3.3) now has the following participation constraint: $PC\ :\ \Pi_n^m(K(\xi), P(\xi), \xi, X_n^s) \geq X_n^s \xi \underline{\pi}^m$ for every ξ. In this case, the structure of the optimal stopping policy remains the same. To prove this result, we define $\hat{K}(.) = X_n^s K(.)$ and $\hat{P}(.) = X_n^s P(.)$. Then, the supplier's and the manufacturer's profit in Equation (3.2) can be expressed as $X_n^s\{\Pi_n^s(\hat{K}(\check{\xi}), \hat{P}(\check{\xi}), \xi, 1) + C_n\} - C_n$ and $X_n^s \Pi_n^m(\hat{K}(\check{\xi}), \hat{P}(\check{\xi}), \xi, 1)$, respectively. Using these properties, the supplier's optimization problem can be reformulated as

$$\begin{aligned} \pi_n(X_n^s) &= X_n^s \left(\max_{\hat{K}(.), \hat{P}(.)} \left\{ E_{\xi_n}[\Pi_n^s(\hat{K}(\xi_n), \hat{P}(\xi_n), \xi_n, 1)] + C_n \right\} \right) - C_n \qquad (3.8) \\ \text{s.t.} \quad & (IC) : \Pi_n^m(\hat{K}(\xi), \hat{P}(\xi), \xi, 1) \geq \Pi_n^m(\hat{K}(\check{\xi}), \hat{P}(\check{\xi}), \xi, 1) \text{ for every } \ \check{\xi} \neq \xi \\ & (PC) : \Pi_n^m(\hat{K}(\xi), \hat{P}(\xi), \xi, 1) \geq \xi \underline{\pi}^m \text{ for every } \ \xi. \end{aligned}$$

Because all constraints and the term inside of $\{.\}$ in (3.8) are independent of X_n^s, we have $\pi_n(X_n^s) = X_n^s \hat{\pi}_n - C_n$, where $\hat{\pi}_n \equiv \pi_n(1) + C_n$. Hence, following the proofs of Theorems 3.6 and 3.7, one can verify that the structure of the optimal stopping policy remains the same.

Although the optimal stopping policy remains the same, the optimal mechanism cannot be derived in a simple form any more. The key property that establishes Lemma (3.1) is the fact that the participation constraint is binding only at $\underline{\xi}_n$. When the reservation profit is increasing in ξ, this property does not hold in general, and thus we cannot obtain a simpler formulation for the mechanism design problem as in Lemma (3.1).

3.8.4 Dynamic Mechanism Design under Non-Commitment

Previously, we have assumed that the supplier commits to not offering another menu of contracts if the manufacturer declines an offer. This assumption has been commonly made in the auction literature (McAfee and McMillan 1987), and similar take-it-or-leave-it strategies are prevalent in various markets. Riley and Zeckhauser (1983) proved that such a commitment is indeed the best strategy for the seller of a product who faces a single buyer. However, in some cases, the agent (manufacturer) may not find the principal's (supplier's) commitment of not offering another contract credible. For example, consider the case in which the supplier and the manufacturer would never be involved in the same business after the current business, and the manufacturer is the only one who purchases the component that the supplier currently provides. In this case, if the manufacturer declines an offer from the supplier, the supplier has an incentive to provide a new menu of contracts instead of committing to not offering another menu. Such cases are called non-commitment in the economics literature (Laffont and Tirole 1988, Skreta 2006).

Deriving an optimal mechanism under non-commitment in a dynamic framework is a very challenging problem. Solving our capacity planning problem under non-commitment is especially challenging for three reasons. First, it has been shown by Laffont and Tirole (1988) that the revelation principle is not valid when the principal offers multiple short-term contracts over time under non-commitment except for the very last period. Hence, the supplier's mechanism design problem cannot be derived in a simple form as in Equation (3.3) except for period N. Second, the supplier can acquire some information about the manufacturer's private information by observing the manufacturer's refusal of offered contracts. Hence, the supplier needs to design the optimal mechanism by taking this information into account. The probability distribution function of the manufacturer's private information given such information is not that of a log-normal random variable any more. It is important to recall that the IGFR property of log-normal random variables is critical for deriving the simple optimality condition given in Theorem 3.5. For this reason, we cannot easily derive the optimal mechanism even for period N at which the revelation principle is valid. Third, the manufacturer's private information evolves over time in our model. Because of the complexity of fully deriving an optimal mechanism for a two-period problem with static private information, Laffont and Tirole (1988) provide some important properties of special classes of equilibria instead of characterizing the optimal contract. Although our problem setting is different from that of Laffont and Tirole (1988), the dynamic nature of the private information makes our problem more complex than theirs.

3.8.5 MMFE for $J > 2$ decision makers

When we construct the MMFE for two decision makers, we divide the total information into $(N+1)^2$ groups by the time that each decision maker obtains the information. Similarly, for $J > 2$ decision makers, we divide the information sets into $(N+1)^J$ groups such that each $\delta_{n_1,\ldots,n_J}$ represents demand information that is obtained by J decision makers at the specified period. By determining the standard deviations of each $\log(\delta_{n_1,\ldots,n_J})$, one can similarly describe the forecast evolution for J decision makers in a consistent manner.

3.9. Conclusion

In this chapter, we have provided a framework to generalize the MMFE for multiple decision makers who forecast demand for the same product. We have shown that this framework is consistent and can be used to model several forecast evolution scenarios when multiple decision makers employ different forecasters. We model the scenario in which forecasters have asymmetric demand information that changes over time and referred to this model as the *Martingale Model of Asymmetric Forecast Evolution*. Using this model, we have studied a supplier's problem of determining the optimal time to offer a capacity reservation contract to a manufacturer. We have characterized structural properties of the optimal time to offer the contract, and the optimal capacity reservation contract. We have also established how these decisions are linked. We have provided managerial insights through numerical studies. For example, we have shown that the supplier can significantly improve his profit by optimally determining the time to offer a contract. Even though our focus has been on the capacity reservation contract, we have also discussed how the framework can be used to determine the optimal time to design a mechanism for other problems with asymmetric information and dynamic information updates. For example, the seller of a product can delay to design a selling mechanism to update his knowledge of the customers' valuation on the product by observing more early sales data. The consumer of an airline ticket may delay to bid for a name-your-own-price ticket to update her valuation of the product as the traveling date approaches (Courty and Li 2000 and Akan et al. 2009).

Chapter 4

A Dynamic Strategy to Optimize Market Entry Timing and Process Improvement Decisions

4.1. Introduction

Manufacturing firms determine the timing for introducing a new product to the market in the presence of two major uncertainties. First, manufacturing firms are uncertain about competing firms' market entry timing. Entering the market later than competitors results in a drastic reduction of profit in a highly competitive environment. For example, Kumar and McCaffrey (2003) estimate the penalty of being late to market by one quarter in the hard disk drive industry at $106 million, which corresponds to fifty percent of gross profit. Second, manufacturing firms are also uncertain about whether they can complete the development of the new production process by the product launch time because the outcomes of manufacturing process development activities are often highly unpredictable. Product launch with an ill-prepared production process significantly reduces profit. In 2005, Microsoft launched the Xbox360 one year ahead of the competing game consoles' market entry. The early market entry resulted in a huge number of failing units, which cost Microsoft $1.15 billion for repairs (Taub 2007). In the presence of such uncertainties, optimizing the market entry timing is a challenging problem. The market entry decision also involves considerable risk, because the decision is irreversible. This chapter seeks for a solution for optimizing the market entry timing in consideration of risk.

Conventionally, manufacturing firms determine the market entry timing long before launching the product. When determining the market entry timing, firms take into consideration the trade-off between the time-to-market and the completeness of the production processes. On the one hand, manufacturing firms can attain a large market share by entering the market early. On the other hand, manufacturing firms can improve the production process for the new product by investing more time in process design, which results in a reduction of production costs. When determining the market entry timing, manufacturing firms are uncertain about the timing of the competitors' market entry. Hence, they often make the timing decision based on their estimation of the competitors' market entry timing. The decision may have a poor outcome if the actual timing of the competitors' market entry deviates significantly from the estimation. In addition, manufacturing firms frequently encounter failures of production process development activities for a new product (Pisano 1996), i.e., uncertainties reside in the learning activities for process design. Hence, without effectively adjusting the market entry timing depending on the outcome of process development activities, manufacturing firms may introduce a new product with an ill-understood production process. As a remedy for these failures, manufacturing firms can adopt a strategy that determines the market entry timing dynamically in response to the evolution of uncertain factors.

For the dynamic market entry strategy to be effective, the coordination between the market entry timing and the development of the production process is important. Consider, for example, the case in which competitors of a manufacturer have introduced their products earlier than expected. In this case, the manufacturer would want to accelerate his market entry in order to avoid a significant loss in the market share. However, to complete the new production process by the accelerated market entry timing, the manufacturer has to invest more capital in process design to expedite the learning rate. As another example, consider the case in which the manufacturer has encountered several failures of manufacturing process development projects during the process design. In this case, the manufacturer may suffer from a low production yield and a large production cost unless he delays the market entry timing. Without flexible management of process development, the dynamic market entry strategy cannot be effective.

The production and the pricing decisions for a new product also have to be coordinated with the market entry decision. For example, if a manufacturing firm introduces a new product much earlier than competitors, the firm can set a high monopolist price for the new product. The firm also needs to produce a large number of products to fulfil high demand. On the other hand, when manufacturing firms determine the market entry timing, they take into consideration the production and pricing decisions for the new product. In other words, manufacturing firms need to jointly optimize the prior-market-entry decisions, i.e., market entry timing and process improvement decisions, and post-market-entry decisions, i.e., production and pricing decisions, for a new product.

In this chapter, we consider the problem of a manufacturer who employs a dynamic strategy to optimize the decisions about market entry timing and process improvement. The manufacturer faces a two-stage stochastic decision process, which consists of the *process design stage* and the *production and sales stage*. During the process design stage, the manufacturer first determines whether to continue process design or stop it. When the manufacturer continues process design, he makes investment decisions to improve the production process for the new product. However, the size of the potential market for the new product decreases as the manufacturer delays market entry. The decision process proceeds to the production and sales stage when the manufacturer stops process design. In this stage, the manufacturer determines the production quantity and the sales prices for the new product in the presence of demand uncertainty.

We establish the optimality of several threshold-type market entry policies that prescribe the optimal time to introduce the new product to the market. Under a threshold-type market entry policy, the manufacturer stops process design if the state of the problem exceeds a threshold and continues it otherwise. For example, we prove the optimality of a knowledge-level-based threshold policy under which the manufacturer enters the market if the level of cumulative knowledge regarding the production process exceeds a certain threshold. We also characterize monotonicities of the optimal thresholds. We next provide properties of the optimal production and pricing decisions. To do so, we derive the solution of the second-stage problem as functions of the outcome of the first-stage problem. These functions enable us to quantify the trade-off between the time-to-market and process improvements in the process design stage.

Using our modeling framework, we develop two measures that assess the value of the dynamic strategy. The measures are based on comparisons of the dynamic strategy to a conventional static strategy. The first measure evaluates the profit gain that the dynamic strategy provides, and the second measure evaluates the reduction in the variability of profit achieved by the dynamic strategy. Our numerical study shows that the manufacturer can increase profit, while reducing the variability of profit by employing the dynamic strategy. Our numerical study also characterizes industry characteristics under which the dynamic strategy is the most effective.

The trade-off between the time-to-market and the completeness of development in new product introduction has been studied extensively. By developing quantitative models that assess this trade-off, researchers have determined optimal market entry decisions under various industry characteristics (e.g., Kalish and Lilien 1986, Cohen et al. 1996, Bayus 1997, Bayus et al. 1997, and Ülkü et al. 2005). Researchers have also studied market entry decisions in game theoretic frameworks (e.g., Klastorin and Tsai (2004) and Savin and Terwiesch (2005)). However, the existing research on time-to-market decisions has focused on static market entry strategies. One exceptional example is the work by Özer and Uncu (2008) who studied a supplier's problem of dynamically determining the time to apply for product qualification. Our study is different from Özer and Uncu (2008) in four dimensions. First, we consider a general manufacturing firm which does not face a qualification procedure, whereas they consider a supplier who have to pass qualification tests to sell their products. Second, we take into consideration the uncertainties in the development of the production process. The dynamic strategy we propose enables manufacturing firms to respond to the realization of unexpected events driven by such uncertainties. Third, in addition to the market entry decision, we consider the decisions about process improvements. As mentioned above, the coordination between the market entry decision and process improvement decisions are crucial for the effectiveness of the dynamic strategy. Fourth, we investigate the risk reduction benefit of the dynamic strategy, which has been neglected by Özer and Uncu (2008). We refer the reader to Krishnan and Ulrich (2001) and Shane and Ulrich (2004) for further references regarding time-to-market decisions.

A group of researchers have addressed the problem of managing process improvement decisions during a product introduction stage (e.g., Terwiesch and Bohn 2001, Terwiesch and Xu 2004, and Carrillo and Franza 2006). In contrast to the existing studies on this literature, we develop a stochastic learning model that incorporates the uncertainties in the outcome of process improvement activities. Fine and Porteus (1989) have also developed a stochastic learning model for process improvements, but they consider process improvements in the middle of a product's life cycle. In our model, process improvement decisions have to be jointly optimized with the market entry decision, whereas no market entry decision needs to be made in Fine and Porteus (1989)'s model.

The rest of the chapter is organized as follows. In §4.2, we introduce our model and notation. In §4.3, we formulate a two-stage dynamic program to solve the problem. In §4.4, we provide structural properties of the optimal market entry policy. We also characterize the optimal production and pricing decisions. In §4.5, we present the results of our numerical study and generate managerial insights regarding the dynamic strategy. In §4.6, we conclude and provide possible future research directions. All proofs are in the Appendix.

4.2. Model

We consider a manufacturer who dynamically determines the timing for introducing a new product to the market. Prior to introducing the product, the manufacturer can improve the production process for the new product by investing in learning activities such as adjustments of the process recipe, development of faster inspection methods, and reduction of defect rates (Terwiesch and Bohn 2001)[1]. After launching the product, the manufacturer makes production and pricing decisions for the new product in the presence of demand uncertainty. We model the manufacturer's problem as a two-stage stochastic decision process. The first stage is the *process design stage* during which the manufacturer dynamically determines (i) learning activities to improve the production process and (ii) the time to stop process design and introduce the product. The second stage is the *production and sales stage*, during which the manufacturer determines (i) the production quantity and (ii) the sales prices for the new product. All revenues and costs are discounted by $\alpha \in (0, 1]$.

The process design stage consists of T periods, indexed from 1 to T. Let x_t be the manufacturer's level of cumulative knowledge related to the production process at the beginning of period t of the process design stage. At the beginning of each period t, the manufacturer first determines whether to stop process design or continue it. If the manufacturer decides to continue process design, then he chooses a learning activity i_t from the set of available options I_t and invests in it. If the manufacturer invests in option i_t, the manufacturer's knowledge level increases[2] by a random amount $k_t(i_t)$, which is independent of x_t[3]. Learning activities incur investment costs, which we denote by $c_i(i_t)$.

As the manufacturer continues process design, the expected size of the potential market for the new product decreases for two reasons. First, the delay in market entry increases the chance of competitors' launching competing products earlier than the manufacturer, which reduces the manufacturer's market share (Savin and Terwiesch 2005, Özer and Uncu 2008). Second, by delaying the market entry, the manufacturer loses a fixed amount of time for selling the new product because the life cycle (or the time window) of one generation of products is often determined exogenously (Cohen et al. 1996, Kumar and McCaffrey 2003, Klastorin and Tsai 2004). We model the dynamics of the market potential following Kalyanaram and Krishnan (1997)

[1] Such learning activities are also called learning-before-doing (Pisano 1996, Carrillo and Gaimon 2000, Terwiesch and Xu 2004).

[2] We use the terms increasing and decreasing in the weak sense; i.e., increasing means non-decreasing.

[3] Similar additive learning models have been used by Fine and Porteus (1989), Cohen et al. (1996), and Terwiesch and Xu (2004).

and Ülkü et al. (2005). To incorporate uncertainties in the dynamics of the market potential, we extend Ülkü et al. (2005)'s model. Let s_t be the size of the potential market for the new product when the manufacturer introduces the product at period t. If the manufacturer continues process design, then the market potential decreases by a random amount $w_t(s_t)$, which is independent of $k_t(i_t)$. The random variable $w_t(s_t)$ models uncertainties residing in external factors such as competing firms' market entry timing. We assume that $s_{t+1} = s_t - w_t(s_t)$ is stochastically increasing[4] in s_t. In other words, if the market potential at period t is larger, then the market potential at period $t+1$ is also larger. We do not make any additional assumptions on the dynamics of the market potential, and thus this model can describe various plausible cases of market potential changes including that of Ülkü et al. (2005).

When the manufacturer stops process design, the problem proceeds to the production and sales stage, which consists of a regular sales period r and a salvage period s. Demand during each period $n \in \{r, s\}$ depends on the market potential s_t and price p_n, and has the following form:

$$D_n(s_t, p_n) = s_t A_n p_n^{-b}.$$

The parameter $b > 1$ is the price elasticity of demand, and A_n is a random variable that models demand uncertainty. We denote the support of A_n by $[\underline{A}_n, \overline{A}_n]$ for each $n \in \{r, s\}$. At the beginning of the regular sales period, the manufacturer first determines the number of products to produce, Q, and then determines the regular sales price, p_r. During the regular sales period, random demand $D_r(s_t, p_r)$ is realized, and the problem proceeds to the salvage period. At the beginning of the salvage period, the manufacturer determines the salvage price, p_s, to sell remaining products from the regular sales period. During the salvage period, random demand $D_s(s_t, p_s)$ is realized, and unsold products have no zero value after the salvage period. Revenue from the salvage period is discounted by $\beta \in (0, 1]$.

The unit production cost for the new product is determined by the manufacturer's knowledge level x_t at the time when he stops process design. If the manufacturer stops

[4] A parameterized random variable $x(\theta)$ is stochastically increasing [resp., decreasing] in θ if $E[u(x(\theta))]$ is increasing in θ for all increasing [resp., decreasing] functions u.

process design with knowledge level x_t, the unit production cost of the new product is $c_p(x_t) \equiv \delta_0 + \delta_1 e^{-\gamma x_t}$, which consists of the irreducible cost δ_0 and the reducible cost $\delta_1 e^{-\gamma x_t}$ (Savin and Terwiesch 2005). The parameter γ indicates the rate of return of knowledge. The unit production cost decreases in x_t, but the marginal benefit becomes smaller as the knowledge level becomes higher. There is well-documented evidence for diminishing returns in learning (see, for example, Zangwill and Kantor 1998 and Laprè et al. 2000). Because the learning activities we consider are not intended to improve the product but to improve the production process, the attributes and quality of the new product are independent of the knowledge level. Hence, the knowledge level does not directly affect demand for the new product. However, it has an indirect impact on demand because the manufacturer's pricing decision depends on the unit production cost, which in turn depends on the knowledge level.

The sequence of events is as follows. At the beginning of period $t \in \{1, 2, \ldots, T\}$ of the process design stage, the manufacturer first determines whether to continue process design or stop it, depending on the market potential, s_t, and the knowledge level, x_t. The manufacturer only knows the distribution of $w_t(s_t)$ and $k_t(i_t)$ at this point. If the manufacturer decides to continue process design, he makes an investment decision $i_t \in I_t$ to improve the production process, which incurs an investment cost $c_i(i_t)$. At the end of period t, $k_t(i_t)$ and $w_t(s_t)$ are realized, and the states are updated as $x_{t+1} = x_t + k_t(i_t)$ and $s_{t+1} = s_t - w_t(s_t)$. Then, the problem proceeds to period $t + 1$. If the manufacturer decides to stop process design at the beginning of period t, the problem proceeds to the production and sales stage. The manufacturer has to stop process design at the beginning of period $T + 1$ (or at the end of period T). Before the beginning of the regular sales period of the production and sales stage, the manufacturer produces Q units of products at the unit production cost of $c_p(x_t)$. The manufacturer then determines the regular sales price, p_r. During the regular sales period, A_r is realized, and the manufacturer collects revenue $p_r \min\{Q, s_t A_r p_r^{-b}\}$ and loses unmet demand $(Q - s_t A_r p_r^{-b})^+$. Then, the problem proceeds to the salvage period. At the beginning of the salvage period, the manufacturer determines the salvage price p_s to sell remaining products. During the salvage period, A_s is realized, and the manufacturer collects revenue. Appendix A provides a glossary of notation for easy reference.

4.3. Formulation

This section describes a two-stage dynamic program that we have formulated to solve the manufacturer's problem. The first-stage problem is an optimal stopping problem with additional decisions to determine the time to introduce the new product and investment decisions for process improvement. The second-stage problem is a pricing and production decision problem. The solution of each stage depends on the solution of the other stage, i.e., the two stages are nested.

4.3.1 The First-Stage Problem

At the beginning of period $t \in \{1, \ldots, T\}$, the manufacturer's knowledge level is x_t and the market potential is s_t. Given this information, the manufacturer decides whether to continue process design or to stop it;

$$u_t = \begin{cases} u^s, & \text{stop process design} \\ u^c, & \text{continue process design.} \end{cases}$$

If the manufacturer stops process design at period t, then the state is updated to indicate that the process has already been stopped. To do so, we define the terminal state $s_t = S$. If the manufacturer decides to continue process design at period t, he invests in a learning activity $i_t \in I_t$ to improve the production process. Then, the knowledge level is updated as

$$x_{t+1} = x_t + k_t(i_t),$$

and the market potential is updated as

$$s_{t+1} = \begin{cases} S, & \text{if } s_t = S\text{, or } s_t \neq S \text{ and } u_t = u^s \\ s_t - w_t(s_t), & \text{otherwise.} \end{cases}$$

The manufacturer has to stop process design at the beginning of period $T+1$ if he has not done so before. We denote the manufacturer's optimal expected profit when he stops process design at period t by $\Pi_t(s_t, x_t)$. We explicitly define this profit function in the next subsection. The reward that the manufacturer attains at period $t \leq T$ is given as

$$g_t(s_t, x_t, i_t) = \begin{cases} -c_i(i_t), & \text{if } s_t \neq S \text{ and } u_t = u^c, \\ \Pi_t(s_t, x_t), & \text{if } s_t \neq S \text{ and } u_t = u^s, \\ 0, & \text{otherwise,} \end{cases}$$

and the reward at period $t = T+1$ is given as

$$g_{T+1}(s_{T+1}, x_{T+1}) = \begin{cases} \Pi_{T+1}(s_{T+1}, x_{T+1}), & \text{if } s_{T+1} \neq S \\ 0, & \text{otherwise.} \end{cases}$$

Let $\mathcal{P} = \{u_1(s_1, x_1), i_1(s_1, x_1), \ldots, u_T(s_T, x_T), i_T(s_T, x_T)\}$ indicate a policy that determines investment decisions and the stopping decision. Then, the manufacturer's first-stage problem is given as

$$\max_{\mathcal{P}} E\left[\sum_{t=1}^{T} \alpha^{t-1} g_t(s_t, x_t, i_t) + \alpha^T g_{T+1}(s_{T+1}, x_{T+1})\right], \tag{4.1}$$

where the maximum is taken for all admissible policies.

We can use the following dynamic programming algorithm to solve this problem:

$$V_{T+1}(s_{T+1}, x_{T+1}) = \Pi_{T+1}(s_{T+1}, x_{T+1})$$
$$V_t(s_t, x_t) = \max\left\{\Pi_t(s_t, x_t), \max_{i_t \in I_t}\left(-c_i(i_t) + \alpha E\left[V_{t+1}(s_t - w_t(s_t), x_t + k_t(i_t))\right]\right)\right\}.$$

It is optimal to stop process design at period t if

$$\Pi_t(s_t, x_t) \geq \max_{i_t \in I_t}\left(-c_i(i_t) + \alpha E\left[V_{t+1}(s_t - w_t(s_t), x_t + k_t(i_t))\right]\right), \tag{4.2}$$

and it is optimal otherwise to continue process design by investing in

$$i_t^*(s_t, x_t) \equiv \arg\max_{i_t \in I_t}\left[-c_i(i_t) + \alpha E[V_{t+1}(s_t - w_t(s_t), x_t + k_t(i_t))]\right]. \tag{4.3}$$

4.3.2 The Second-Stage Problem

The second stage problem has two decision points: the beginning of the regular sales period and the beginning of the salvage period. We formulate the second-stage problem by a backward induction. Suppose that the manufacturer has Q_s units of unsold products before the salvage period. The manufacturer determines the salvage price to maximize the expected revenue as

$$J_s(s_t, Q_s) \equiv \max_{p_s} p_s E[\min\{Q_s, s_t A_s p_s^{-b}\}]. \tag{4.4}$$

Next we suppose that the manufacturer has Q units of products before the regular sales period. The manufacturer determines the regular sales price to maximize the revenue-to-go as

$$J_r(s_t, Q) \equiv \max_{p_r} p_r E[\min\{Q, s_t A_r p_r^{-b}\}] + \beta E[J_s(s_t, [Q - s_t A_r p_r^{-b}]^+)]. \tag{4.5}$$

Finally, the optimal production quantity is determined to maximize the total profit as

$$\Pi_t(s_t, x_t) = \max_{Q} J_r(s_t, Q) - c_p(x_t)Q. \tag{4.6}$$

4.4. Analysis

To solve the first-stage problem, we need the profit function $\Pi_t(s_t, x_t)$ from the second-stage problem. Hence, we solve the second-stage problem first, and then embed the solution in the first-stage problem.

4.4.1 Optimal Production and Pricing Decisions

We first derive the optimal salvage price, $p_s^*(s_t, Q_s)$ as a function of the market potential s_t and the number of remaining products Q_s. To do so, we transform the optimization problem (4.4) into an optimization problem that is independent of states by changing the decision variable. We define $z_s \equiv \frac{Q_s}{s_t p_s^{-b}}$, which is called the stocking factor (Petruzzi and Dada 1999). The stocking factor indicates the number of standard deviations that stocking quantity deviates from the expected demand. Hence, a higher stocking level results in a higher demand fill-rate. We refer the reader to Petruzzi and Dada (1999) and Monahan et al. (2004) for further discussions of the stocking factor. Then, by replacing p_s by $\left(\frac{s_t z_s}{Q_s}\right)^{1/b}$ in (4.4), we can derive

$$\begin{aligned} J_s(s_t, Q_s) &= \max_{p_s} p_s E[\min\{Q_s, s_t A_s p_s^{-b}\}] \\ &= s_t \left(\frac{Q_s}{s_t}\right)^{1-1/b} \max_{z_s} E[\min\{z_s^{1/b}, A_s z_s^{-1+1/b}\}], \end{aligned} \tag{4.7}$$

which decouples the states s_t and Q_s from the optimization problem. We denote the optimal solution and the resulting value of the decoupled problem by $z_s^* \equiv \arg\max_{z_s} E[\min\{z_s^{1/b}, A_s z_s^{-1+1/b}\}]$ and $J_s^* \equiv E[\min\{(z_s^*)^{1/b}, A_s(z_s^*)^{-1+1/b}\}]$. Then, the following theorem characterizes the optimal salvage price and the optimal revenue-to-go function.

Theorem 4.1. *The optimal stocking factor always satisfies $z_s^* \in [\underline{A}_s, \overline{A}_s]$. If, in addition, A_s has an increasing generalized failure rate (IGFR)*[5]*, z_s^* is the unique solution of the equation*

$$zProb(A_s > z) = (1 - \frac{1}{b})E[\min\{z, A_s\}]. \tag{4.8}$$

Then, the optimal salvage price is given as $p_s^*(s_t, Q_s) = \left(\frac{s_t z_s^*}{Q_s}\right)^{1/b}$, and the optimal revenue-to-go function is given as $J_s(s_t, Q_s) = s_t \left(\frac{Q_s}{s_t}\right)^{1-1/b} J_s^*$.

The optimality of z_s^* that satisfies the equation (4.8) is shown by Monahan et al. (2004).

Theorem 4.1 shows that the optimal salvage price increases in the market potential and decreases in the number of remaining products. As the size of the potential

[5]The generalized failure rate of a random variable is defined as $\frac{xf(x)}{1-F(x)}$, where $f(x)$ and $F(x)$ are the p.d.f. and the c.d.f. of the random variable. All random variables with increasing failure rates have IGFRs, and other common classes of random variables have IGFRs, including Log-normal, Gamma and Weibull.

market becomes larger, i.e., as s_t becomes larger, the manufacturer charges a larger salvage price to maximize revenue. However, the salvage price decreases as the number of remaining products increases because the manufacture would want to salvage most remaining products during the salvage period. The theorem also shows that the optimal revenue-to-go function increases in both the market potential and the number of remaining products.

We next derive the optimal regular sales price, $p_r^*(s_t, x_t)$, and the optimal production quantity, $Q^*(s_t, x_t)$, as functions of the market potential s_t and the knowledge level x_t. To do so, we first transform the optimization problem (4.5) into a state-independent optimization problem as before. We define the stocking factor of the regular sales period as $z_r \equiv \frac{Q}{s_t p_r^{-b}}$. Then, by replacing p_r by $\left(\frac{s_t z_r}{Q}\right)^{1/b}$ in (4.5), we can derive

$$\begin{aligned} J_r(s_t, Q) &= \max_{p_r} \left\{ p_r E[\min\{Q, s_t A_r p_r^{-b}\}] + \beta E[J_s(s_t, [Q - s_t A_r p_r^{-b}]^+)] \right\} \\ &= \max_{p_r} \left\{ p_r E[\min\{Q, s_t A_r p_r^{-b}\}] + \beta E\left[s_t \left(\frac{[Q - s_t A_r p_r^{-b}]^+}{s_t} \right)^{1-1/b} J_s^* \right] \right\} \\ &= s_t \left(\frac{Q}{s_t} \right)^{1-1/b} \max_{z_r} \left\{ E[\min\{z^{1/b}, A_r z^{-1+1/b}\}] + \beta J_s^* z^{-1+1/b} E[((z - A_r)^+)^{1-1/b}] \right\}, \end{aligned}$$

which again decouples the states from the optimization problem. For notational simplicity, we define $f(z) \equiv E[\min\{z^{1/b}, A_r z^{-1+1/b}\}] + A_s z^{-1+1/b} E[((z - A_r)^+)^{1-1/b}]$, and denote the optimal solution and the resulting value of the decoupled problem by $z_r^* \equiv \arg\max_z f(z)$ and $J_r^* \equiv f(z_r^*)$.

Theorem 4.2. *The optimal stocking factor z_r^* is either the unique solution of the equation*

$$\beta J_s^* E[A_r (z - A_r)^{-1/b}] = E[A_r] \tag{4.9}$$

or $\arg\max_{z \in [\underline{A}_r, \overline{A}_r]} f(z)$. *Then, the optimal revenue-to-go function is give as $J_r(s_t, Q) = s_t J_r^* \left(\frac{Q}{s_t}\right)^{1-1/b}$, and the optimal regular sales price and the optimal production quantity* *are given as*

$$p_r^*(s_t, x_t) = \frac{(b-1)J_r^*}{bc_p(x_t)(z_r^*)^{1/b}}$$
$$Q^*(s_t, x_t) = s_t \left(\frac{(b-1)J_r^*}{bc_p(x_t)}\right)^b.$$

Finally, the optimal expected profit that the manufacturer can attain by introducing the product at period t is

$$\Pi_t(s_t, x_t) = s_t c_p(x_t)^{1-b}\pi^*, \tag{4.10}$$

where $\pi^* \equiv \frac{1}{b-1}\left(\frac{(b-1)J_r^*}{b}\right)^b$.

The optimal stocking factor for the regular sales period is either in the set $[\underline{A}_r, \overline{A}_r]$ or greater than $\overline{A}_r$. Because $f(z)$ is not a unimodal function in general, we cannot derive the optimal z_r^* from a simple equation as before. However, for $z > \overline{A}_r$, this function is quasi-concave and the maximum is at the point that satisfies the equation (4.9). Hence, to determine the optimal stocking factor z_r^*, one needs to evaluate $f(z)$ for every $z \in [\underline{A}_r, \overline{A}_r]$ and for the single point that satisfies the equation (4.9). Although numerically determining such z_r^* is computationally difficult, the optimal stocking factor depends on neither s_t nor Q. Hence, by solving only one optimization problem, we can derive the optimal production quantity, regular sales price, and the revenue-to-go function for all states.

Theorem 4.2 shows that the optimal production quantity is linearly increasing in the market potential and also increasing in the knowledge level. A higher knowledge level means a lower unit production cost. Hence, this result implies that the manufacturer produces more products when the unit production cost is smaller. The optimal regular sales price is increasing in the knowledge level and is independent of the market potential. It is important to note that the market potential affects neither the price elasticity of demands nor the coefficient of variation of demands. The market potential only scales the size of the second-stage problem. Hence, the optimal regular price is independent of it. Finally, the optimal expected profit is increasing in the market potential and the knowledge level. When the manufacturer continues process design, the market potential and the knowledge level change in opposite directions, i.e., there is a trade-off between the market potential and the knowledge level.

4.4.2 Optimal Market Entry Policy

When determining whether to continue process design or to stop it, the manufacturer takes three factors into account: investment cost, the profit gain achieved by improved knowledge, and the profit loss incurred by decrease in the market potential. Among the three factors, investment cost depends neither on the current knowledge level nor on the current market potential, whereas the other two factors depend on both. For example, when the current knowledge level is very high, additional knowledge does not improve the manufacturer's profit. The benefit of improved knowledge is also negligible when the size of the potential market is small. By investigating such dependencies, we establish the optimality of threshold-type optimal market entry policies. A threshold-type market entry policy, for example, has the following form: *stop* process design if the current knowledge level exceeds a certain threshold and *continue* otherwise. Such threshold-type control policies are easy to implement and are also useful for numerically solving the optimal stopping problem.

Before discussing optimal market entry policies, we define three functions and four thresholds that facilitate the analysis of the optimal market entry policy. We first define the *one-step benefit function* as

$$M_t(s_t, x_t, i_t) \equiv -c_i(i_t) + \alpha E[\Pi_{t+1}(s_t - w_t(s_t), x_t + k_t(i_t))] - \Pi_t(s_t, x_t),$$

which indicates the myopic benefit of investing in learning activity i_t at period t. The one-step benefit function assesses the benefit of investing in i_t by taking only one period into account. Similarly, we define the *benefit function* as

$$B_t(s_t, x_t, i_t) \equiv -c_i(i_t) + \alpha E[V_{t+1}(s_t - w_t(s_t), x_t + k_t(i_t))] - \Pi_t(s_t, x_t),$$

which indicates the true benefit of investing in learning activity i_t at period t. We also define the *maximal benefit function* as $\overline{B}_t(s_t, x_t) \equiv B_t(s_t, x_t, i_t^*(s_t, x_t))$, which indicates the benefit of continuing process design at period t.[6] Then, a market entry policy that stops process design if $\overline{B}_t(s_t, x_t) \leq 0$ is optimal. The set of states at which stopping process design is optimal is given as $\{(s_t, x_t) : \overline{B}_t(s_t, x_t) \leq 0\}$. We denote the boundaries of this set by $\underline{x}_t(s_t) \equiv \inf\{x : \overline{B}_t(s_t, x) \leq 0\}$, $\overline{x}_t(s_t) \equiv \sup\{x : \overline{B}_t(s_t, x) \leq 0\}$, $\underline{s}_t(x_t) \equiv \inf\{s : \overline{B}_t(s, x_t) \leq 0\}$, and $\overline{s}_t(x_t) \equiv \sup\{s : \overline{B}_t(s, x_t) \leq 0\}$, which are the optimal thresholds for the market entry policies that we derive shortly.

First, we establish the optimality of a knowledge-level-based *lower* threshold policy.

Theorem 4.3. *Let $\hat{t}$ be the first period at which $x_t \geq -\frac{1}{\gamma}\log(\frac{\delta_0}{(b-1)\delta_1})$ holds for every realization of x_t. Then, (i) $\overline{B}_t(s_t, x_t)$ is decreasing in x_t, and (ii) a knowledge-level-based lower threshold policy that stops process design if $x_t \geq \underline{x}_t(s_t)$ is optimal for every $t \geq \hat{t}$.*

The manufacturer's expected profit, $\Pi_t(s_t, x_t)$, is increasing in x_t with an increasing return for small values of x_t, and is increasing in x_t with a decreasing return for large values of x_t.[7] In other words, the manufacturer's expected profit increases rapidly as the knowledge level increases from a small value, but the benefit of additional knowledge diminishes as the knowledge is accumulated. At periods $t \geq \hat{t}$, the knowledge level is sufficiently high such that additional knowledge does not improve the manufacturer's profit much. In contrast, investment cost and the market potential loss are independent of the knowledge level. In this case, the benefit of continuing process design, i.e., $\overline{B}_t(s_t, x_t)$, decreases as the manufacturer's knowledge level increases, and thus the manufacturer should stop process design if the knowledge level exceeds a certain threshold.

For periods earlier than period $\hat{t}$, the structure of the optimal market entry policy is generally unknown. However, the knowledge-level-based lower threshold policy is likely to be optimal for those periods as well. If the knowledge level is low at early

[6] Note from (4.3) that $\arg\max_{i_t \in I_t} B_t(s_t, x_t, i_t) = \arg\max_{i_t \in I_t} -c_i(i_t) + \alpha E[V_{t+1}(s_t - w_t(s_t), x_t + k_t(i_t))] = i_t^*(s_t, x_t)$.

[7] $\Pi_t(s_t, x_t)$ is proportional to $1/c_p(x_t)^{b-1}$, which is convex increasing in x_t for $x_t < -\frac{1}{\gamma}\log(\frac{\delta_0}{(b-1)\delta_1})$ and concavely increasing in x_t for $x_t \geq -\frac{1}{\gamma}\log(\frac{\delta_0}{(b-1)\delta_1})$

periods, the manufacturer would want to continue process design for multiple periods to sufficiently reduce the unit production cost. On the other hand, if the knowledge level is high at early periods, the benefit of continuing process design decreases in x_t for the same reason as in periods $t \geq \hat{t}$. In such a case, the knowledge-level-based lower threshold policy is optimal.

The optimal market entry policy may have a reversed structure if the manufacturer does not have enough time to sufficiently improve the production process. If the manufacturer's knowledge level can never reach a certain point, the profit gain that additional knowledge provides always increases as the knowledge level increases. In this case, continuing process design is more beneficial when the knowledge level is higher, and thus a knowledge-level-based *upper* threshold policy that stops process design if $x_t \leq \overline{x}_t(s_t)$ is optimal. In §4.5.1, we provide one extreme example under which such a policy is optimal. However, this case can hardly arise in practice. Manufacturing firms usually allocate enough time for process design to sufficiently reduce the unit production cost.

We next establish the optimality of a market-potential-based *upper* threshold policy.

Theorem 4.4. *Suppose that the condition*

$$\frac{dE[w_t(s_t)]}{ds_t} \leq 1 - \min_{i_t \in I_t} \frac{c_p(x_t)^{1-b}}{\alpha E[c_p(x_t + k_t(i_t))^{1-b}]} \tag{4.11}$$

holds for every t. Then, (i) $\overline{B}_t(s_t, x_t)$ is increasing in s_t, and (ii) a market-potential-based upper threshold policy that stops process design if $s_t \leq \overline{s}_t(x_t)$ is optimal for every t.

The left-hand-side of (4.11) indicates the sensitivity of the market potential reduction in the current market potential, and the right-hand-side of (4.11) indicates the minimum relative profit improvement achieved by additional knowledge. Hence, the condition (4.11) holds, for example, when investing in learning activities sufficiently improves the production process. The condition also holds when the history of the market potential does not provide much information about future changes in the market potential, i.e., when $w_t(s_t)$ is insensitive to s_t. We have mentioned above that three factors affect the manufacturer's market entry decision: investment cost, the profit gain achieved by improved knowledge, and the profit loss incurred by decrease in the market potential. Among them, investment cost is independent of the market potential, and the profit gain achieved by improved knowledge is proportionally increasing in the current market potential. The profit loss incurred by market potential reduction can either increase or decrease in the current market potential depending on the sensitivity of $w_t(s_t)$ in s_t. When the change in the market potential is sufficiently insensitive to the current market potential, i.e., when $\frac{dE[w_t(s_t)]}{ds_t}$ is sufficiently small, or when additional knowledge substantially improves profit, i.e., when $1-\min_{i_t\in I_t}\frac{c_p(x_t)^{1-b}}{\alpha E[c_p(x_t+k_t(i_t))^{1-b}]}$ is large, the benefit of continuing process design increases as the market potential increases, and thus the market-potential-based upper threshold policy is optimal.

When the large value of the current market potential signals a big market potential loss for the upcoming time period, a reversed market-potential-based threshold policy is optimal.

Theorem 4.5. *Suppose that the condition*

$$\frac{dE[w_t(s_t)]}{ds_t} \geq 1 - \max_{i_t\in I_t}\frac{c_p(x_t)^{1-b}}{\alpha E[c_p(x_t+k_t(i_t))^{1-b}]} \tag{4.12}$$

holds for every t. Then, (i) $\overline{B}_t(s_t, x_t)$ is increasing in s_t, and (ii) a market-potential-based lower threshold policy that stops process design if $s_t \geq \underline{s}_t(x_t)$ is optimal for every t.

If the expected market potential loss, $E[w_t(s_t)]$, increases in the current market potential more rapidly than the relative profit improvement, $1-\frac{c_p(x_t)^{1-b}}{\alpha E[c_p(x_t+k_t(i_t))^{1-b}]}$, the benefit of investing in learning activities decreases as the current market potential increases. In this case, the manufacturer should stop process design if the market potential is larger than a certain threshold. In §4.5.1, we provide an example under which the market-potential-based *lower* threshold policy is optimal, but the condition (4.12) rarely holds in general.

Theorem 4.6. *(a) If $\overline{B}_t(s_t, x_t)$ is decreasing [resp., increasing] in x_t and increasing [resp., decreasing] in s_t, then $\underline{x}_t(s_t)$ [resp., $\overline{x}_t(s_t)$] is increasing in s_t and $\overline{s}_t(x_t)$ [resp., $\underline{s}_t(x_t)$] is increasing in x_t.*
(b) If $\overline{B}_t(s_t, x_t)$ is increasing [resp., decreasing] in both x_t and s_t, then $\overline{x}_t(s_t)$ [resp., $\underline{x}_t(s_t)$] is decreasing in s_t and $\overline{s}_t(x_t)$ [resp., $\underline{s}_t(x_t)$] is decreasing in x_t.

Consider the case in which $\overline{B}_t(s_t, x_t)$ is decreasing in x_t and increasing in s_t in Theorem 4.6(a). The increasing property of $\underline{x}_t(s_t)$ implies that when the market potential is larger, continuing process design is optimal for a larger region of knowledge levels. Similarly, the increasing property of $\overline{s}_t(x_t)$ implies that when the knowledge level is higher, continuing process design is optimal for a smaller region of the market potential. Such monotonicities of the optimal thresholds help us to better understand how the optimal decision responds to the changes in the environment and are also useful for numerically solving the problem. The other cases of Theorem 4.6 can also be explained in a similar way.

4.5. Numerical Study

This section presents the result of our numerical studies. The purpose of this section is three-fold. First, we illustrate some examples of the optimal market entry and process improvement decisions. The examples further generate managerial insights regarding the market entry and process improvement decisions. Second, we develop two measures that quantify the value of the dynamic strategy. The two measures respectively assess the profit improvement benefit and the risk reduction benefit of the dynamic strategy. Third, by evaluating the two measures under various conditions, we characterize industry characteristics under which the dynamic strategy is the most effective.

For our numerical study, we model the dynamics of process improvement, market potential changes, and demand uncertainty as follows. At each period $t \in \{1, 2, \ldots, T\}$ of the process design stage, the manufacturer has two investment options for process improvement: regular learning, r, and expedited learning, e. Regular learning incurs investment costs of $c_i(r)$ and increases the manufacturer's knowledge level by $k_t(r)$, which is uniformly distributed from $\underline{k}$ to $\overline{k}$. Expedited learning incurs investment costs of $c_i(e)$, and increases the knowledge level by $k_t(e)$, which is uniformly distributed from $2\underline{k}$ to $2\overline{k}$, i.e., expedited learning doubles the learning rate of regular learning at additional costs of $c_i(e) - c_i(r)$. The market potential decreases by w_t at each period, where w_t is normally distributed and independent of s_t. We denote the coefficient of variation of w_t by cv_w. The demand uncertainty A_r is also normally distributed, and we denote the coefficient of variation of A_r by cv_A. We assume A_s is deterministic in the numerical studies.

The base numerical setting for our studies is as follows: $T = 4$, $\alpha = 1$, $\beta = 1$, $\underline{k} = 3$, $\overline{k} = 13$, $c_i(r) = 0$, $c_i(e) = 4$, $cv_w = 0.2$, $E[A_r] = 1$, $cv_A = 0.2$, $A_s = 0.1$, $b = 1.5$, $\delta_0 = 1$, $\delta_1 = 10$, and $\gamma = 0.075$. We illustrate the values of $E[s_t]$ and $E[w_t]$ in Figure 4.1. The investment cost for regular learning is sunk, i.e., $c_i(r) = 0$.

Figure 4.1: Expected Values of Market Potential

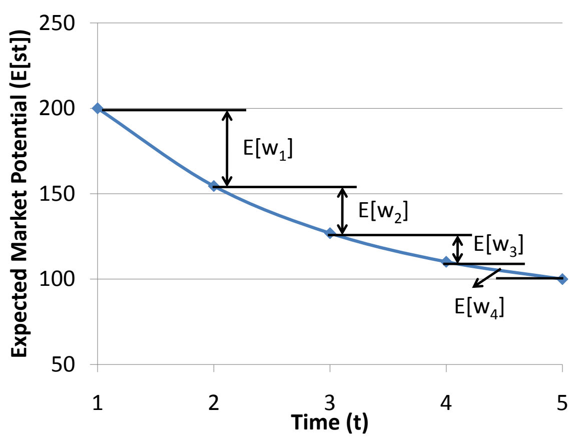

4.5.1 Optimal Market Entry and Process Improvement Decisions

Figure 4.2 illustrates the optimal market entry and process improvement decisions for periods 2 and 3 under the base numerical setting. Because the base numerical setting satisfies the sufficient condition of Theorem 4.4, the market-potential-based upper threshold policy is optimal for both periods, i.e., stopping process design is

Figure 4.2: Optimal Investment and Stopping Policy of Base Numerical Setting

(a) Period 2 (b) Period 3

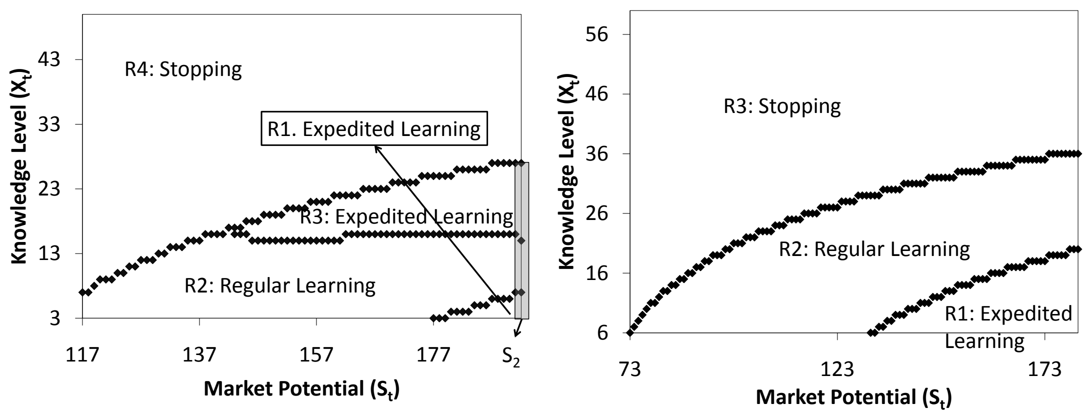

optimal if the market potential is smaller than the optimal threshold. In the base setting, a larger value of the current market potential does not imply a bigger loss in the market potential, and thus the benefit of continuing process design increases as the current market potential increases.

Although some possible realizations of x_2 and x_3 do not satisfy the condition $x_t \geq -\frac{1}{\gamma}\log(\frac{\delta_0}{(b-1)\delta_1})$ in Theorem 4.3, Figure 4.2 shows that the knowledge-level-based lower threshold policy is optimal for both periods, i.e., stopping process design is optimal if the current knowledge level exceeds the optimal threshold. As we discussed above, the knowledge-level-based market entry policy can have a reversed structure only when the manufacturer does not have enough time to sufficiently improve the production process. To verify this argument, we examine the optimal market entry decision when the manufacturer has only one period to improve the production process, i.e., when $T = 1$, and report the result in Figure 4.3. The other numerical settings are the same as in the base numerical setting except $b = 2.5$, $c_i(r) = 2$, $c_i(e) = 8$, and $E[w_1] = 20$. In this case, the manufacturer cannot increase his knowledge level to the point above which benefit of additional knowledge diminishes as the knowledge level increases, because he has only one period to improve the process. Except for in such an extreme case, the knowledge-level-based lower threshold policy is optimal for

Figure 4.3: Knowledge-Level-Based Upper Threshold Policy

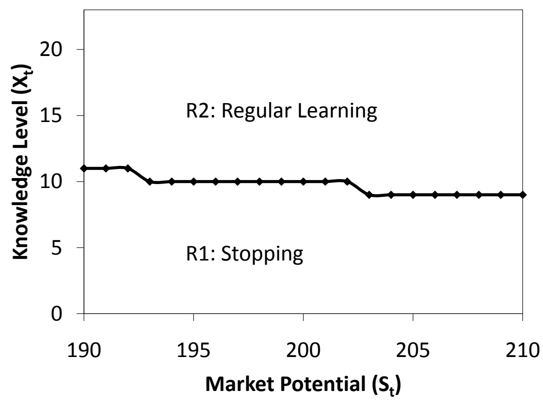

every period. We have tested 18 variations of the base numerical setting[8], and the knowledge-level-based lower threshold turns out to be optimal for every period in all tests.

When a large value of the current market potential signals a big market potential loss for the upcoming time period, the market-potential-based market entry policy can also have a reversed structure. We consider the case in which $T = 1$, $b = 2$, $w_1(s_1) = s_1 - 120$, and other settings are the same as in the base numerical setting. In this case, the market potential of period 2 is always $s_2 = s_1 - w_1(s_1) = 120$ regardless of the market potential of period 1, which implies that the manufacturer completely loses the benefit of a large market potential by delaying the market entry. Hence, continuing process design is less beneficial when the market potential is larger. In this case, the market-potential-based lower threshold policy is optimal as shown in Figure 4.4.

We next discuss the optimal process improvement decisions in Figure 4.2. The optimal investment policy at period 3 is to invest in expedited learning when the knowledge level is low (Region 1) and invest in regular learning when the knowledge level is high (Region 2). Intuitively, a low knowledge level may significantly delay the market entry unless the manufacturer expedite the learning rate, and thus the manufacturer should invest more capital in learning when the knowledge level is lower.

[8]We refer the reader to §4.5.3 for the details of the test settings.

Figure 4.4: Market-Potential-Based Lower Threshold Policy

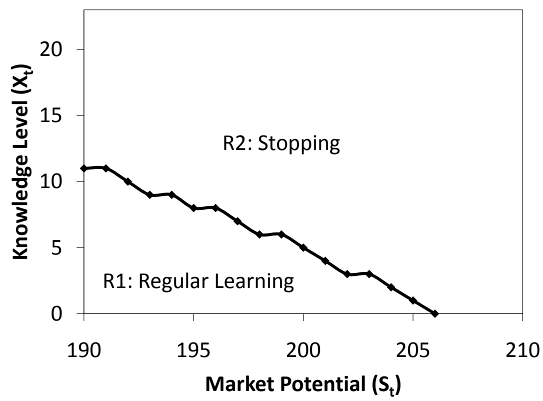

However, the optimal process improvement decision of period 2 shows a different structure. In Figure 4.2(a), investing in expedited learning is optimal not only when the knowledge level is very low (Region 1), but also when the knowledge level is reasonably high (Region 3). To explain the reasoning behind this structure, we measure the probability that the manufacturer stops process design at period 3 for different states of period 2. For each (s_2, x_2) in the set $\Pi \equiv \{(s_2, x_2) : s_2 = 192, x_2 \in [3, 27]\}$,[9] we first calculate the probability distribution of (s_3, x_2) when the manufacturer follows the optimal process improvement decision at period 2. Using this probability distribution, we next calculate $Prob(\overline{B}_3(s_3, x_3) \leq 0 | s_2, x_2)$, which indicates the probability that the manufacturer stops process design at period 3 for each given (s_2, x_2). Figure 4.5 illustrates this probability for every $(s_2, x_2) \in \Pi$ as a function of x_2. For all states across Regions 1 and 2, the probability of stopping at period 3 is almost 0, which means that the manufacturer will introduce the product either at period 4 or 5. In contrast, when the knowledge level falls in Region 3, the manufacturer introduces the product at period 3 with a substantial probability. When the knowledge level is sufficiently high, the manufacturer can effectively accelerate the market entry timing by expediting the learning rate accordingly. This result shows the importance of coordinating process improvement decisions with the market entry decision. By dynamically adjusting the learning rate depending on states, the manufacturer can

[9]The set Π corresponds to the gray box at the bottom right corner of Figure 4.2(a).

Figure 4.5: Probability of Stopping at Period 3

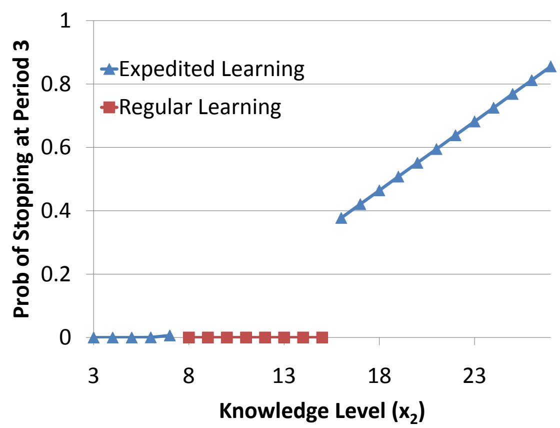

successfully control the market entry timing.

4.5.2 Measuring the Value of the Dynamic Strategy

In this subsection, we propose two measures that assess the value of the dynamic strategy. Compared to a static strategy under which the manufacturer determines market entry timing and process improvement decisions upfront, the dynamic strategy increases the manufacturer's expected profit. In addition, the dynamic strategy also reduces the variability of profit by enabling the manufacturer to respond to an unexpectedly low knowledge level or unusual changes in market potential. The two measures quantify these benefits respectively.

We first explicitly define the static strategy using our modeling framework. We define period 0 as the period at which the manufacturer makes the market entry timing and process improvement decisions. At this time, the manufacturer is uncertain about the initial knowledge level x_1 and the initial market potential s_1. However, the manufacturer has the information that the initial knowledge level x_1 is uniformly distributed from 0 to 23 and the initial market potential is normally distributed with a mean value of 200 and the coefficient of variation of 0.5. The market entry timing and process improvement decisions under the static strategy are state-independent, i.e., the decisions are static. For example, the manufacturer may decide to invest in regular learning at period 1 and enter the market at period 2 regardless of the states. Among all such static decisions, the manufacturer chooses the optimal static decision that maximizes his expected profit. Mathematically, this problem is identical to the stochastic optimization problem (4.1) except that now the maximum is taken over all static policies.

It is important to note that the optimal policy under the static strategy is an admissible policy for the original problem (4.1). Hence, the dynamic strategy always yields a higher expected profit than the static strategy. Let W_d be the profit of the dynamic strategy, and W_s be the profit of the static strategy. Then, the percentage difference in the expected profit, $I \equiv \frac{E[W_d]-E[W_s]}{E[W_s]} \times 100\%$, measures the profit improvement benefit of the dynamic strategy. By enabling the manufacturer to respond to the realization of uncertain states, the dynamic strategy also reduces the variability of profit. The percentage difference in the coefficient of variation of profit, $R \equiv \frac{\sqrt{Var(W_s)}/E[W_s]-\sqrt{Var(W_d)}/E[W_d]}{\sqrt{Var(W_s)}/E[W_s]} \times 100\%$, measures this risk reduction benefit of the dynamic strategy. The variability of profit is a widely used measure of risks involved in managerial decisions (e.g., Martinez-de Albeniz and Simchi-Levi 2003)[10].

Table 4.1 shows the expected profits, coefficient of variations of profits, I, and R under the base numerical setting. We have computed $Var(W_d)$ and $Var(W_s)$ using 40,000 independent samples of W_d and W_s. Under the base numerical setting, the

Table 4.1: Expected Value and Coefficient of Variation of Profits

$E[W_d]$	$E[W_s]$	$I(\%)$	$\frac{\sqrt{Var(W_d)}}{E[W_d]}$	$\frac{\sqrt{Var(W_s)}}{E[W_s]}$	$R(\%)$
34.44	33.50	2.836	0.2061	0.2288	9.911

expected profit of the dynamic strategy is 2.836% larger than that of the static strategy. In addition, the profit of the dynamic strategy has a 9.911% smaller coefficient of variation than that of the static strategy. In other words, the dynamic strategy yields a higher and less variable profit for the manufacturer.

[10]For fair comparison, we use the coefficient of variation of profit instead of the standard deviation of profit, because the dynamic strategy and the static strategy yield different expected profits.

4.5.3 Effectiveness of the Dynamic Strategy under Various Industrial Conditions

We next evaluate the two measures, I and R, under various numerical settings. The objective of this test to identify when the dynamic strategy becomes the most effective. In particular, we examine the impact of the six factors illustrated in Table 4.2 on the value of the dynamic strategy. For each factor, we first explain the test setting and

Table 4.2: Key Factors that Determine the Value of the Dynamic Strategy

	Factors of Interest
Process Characteristics	uncertainty in learning
	cost of expedited learning
	reducible unit production cost
Market Characteristics	uncertainty in market potential changes
	demand uncertainty
	size of the salvage market

then report I and R that we have computed via simulation.

Uncertainty in learning: We first evaluate the impact of the uncertainty in learning, i.e., the impact of the variability of $k_t(i_t)$. To change the variability of $k_t(i_t)$, we change the difference $\overline{k} - \underline{k}$ while holding the mean of $\underline{k}$ and $\overline{k}$ to be as in the base numerical setting. The larger difference between $\overline{k}$ and $\underline{k}$ means a more variable outcome of the learning activities. Because we hold the average of $\underline{k}$ and $\overline{k}$ to be constant, the average process improvement rate remains constant for all test settings.

Figure 4.6 illustrates I and R for several values of $\overline{k} - \underline{k}$. As the degree of the uncertainty in learning increases, both I and R increase rapidly. This result implies that the dynamic strategy is effective when the outcome of process improvement activities is highly uncertain. For example, the dynamic strategy is promising for industries in which failures of manufacturing process development projects are pervasive. The dynamic strategy enables manufacturing firms to effectively adjust the market entry timing depending on the outcome of R&D activities for process improvement.

Figure 4.6: Impact of Uncertainties in Learning

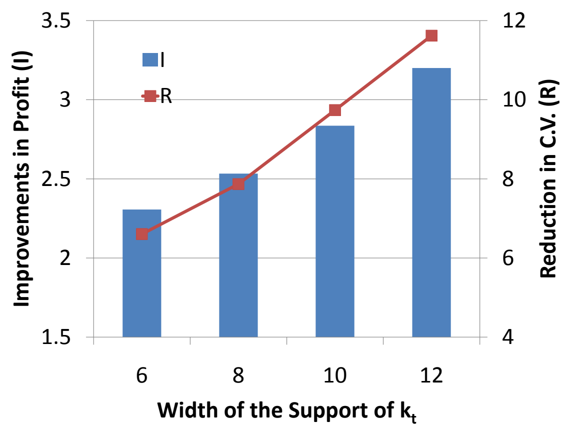

Reducible unit production cost: We next consider the impact of the reducible unit production cost, δ_1. Recall that the unit production cost, $c_p(x_t)$, consists of the irreducible cost, δ_0, and the reducible cost, $\delta_1 e^{-\gamma x_t}$. As the manufacturer's knowledge level increases, the reducible production cost decreases, whereas the irreducible cost remains constant regardless of the knowledge level. Hence, process improvement decisions are more important when the reducible cost is larger. For this reason, the dynamic strategy, which optimally controls process improvement decisions, becomes more effective as δ_1 becomes larger as shown in Figure 4.7. When the knowledge

Figure 4.7: Impact of Reducible Unit Production Cost

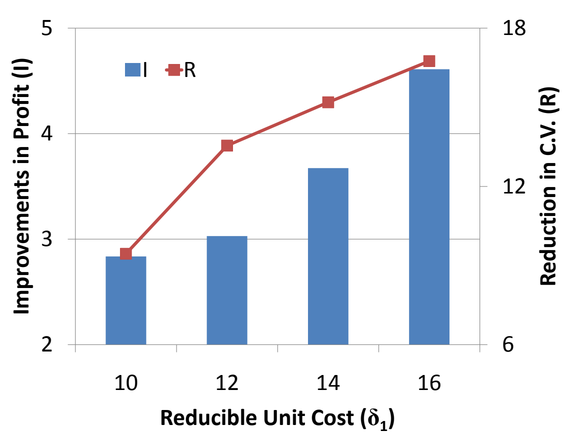

regarding the production process is the key determinant of the production cost such as in the hard disk drive industry (Bohn and Terwiesch 1999), the dynamic strategy can significantly improve the firm's profit.

Cost of expedited learning: We next evaluate the impact of the cost of expedited learning, $c_i(e)$. Recall from §4.5 that the regular learning cost is sunk, i.e., $c_i(r) = 0$. Hence, $c_i(e)$ indicates the additional cost that the manufacturer has to pay to expedite the learning rate. In Figure 4.8, we illustrate the value of the dynamic strategy for several values of $c_i(e)$. The dynamic strategy provides the flexibility of determining

Figure 4.8: Impact of the Cost of Expedited Learning

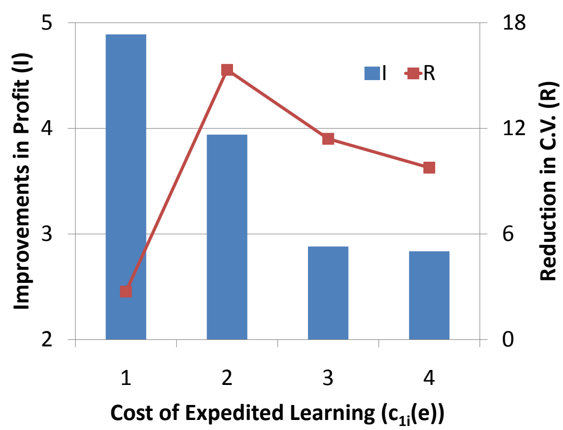

the learning rate depending on the realization of uncertain states. However, when expedited learning incurs a substantially higher cost than regular learning, i.e., when $c_i(e)$ is large, the manufacturer does not often invest in expedited learning, and thus the value of the flexible management is low. From this reason, the profit improvement achieved by the dynamic strategy decreases as the cost of expedited learning increases.

In Figure 4.8, the variability reduction benefit shows a non-monotonic pattern as the cost increases from 1 to 2. When $c_i(e) = 1$, the cost difference between regular learning and expedited learning is small, and hence the manufacturer who employs the static strategy invests in expedited learning at period 1, whereas he invests in regular learning at period 1 when $c_i(e) = 2, 3, 4$. For this reason, when $c_i(e) = 1$, the manufacturer enters the market earlier than the other cases, which substantially reduces the variability of profit. Such drastic changes in the optimal static policy can yield a non-monotonic pattern of R as in Figure 4.8, but the profit improvement benefit, I, always shows a monotonic pattern in all our numerical studies. Hence, we focus on I when discussing the impact of environments on the value of the dynamic strategy.

Uncertainty in market potential changes: We next examine the impact of the uncertainty in market potential changes, i.e., the impact of the variability of w_t. We fix the expected value $E[w_t]$ as in the base setting and change the coefficient of variation of w_t, i.e., cv_w, for this test. Figure 4.9 illustrates I and R for several values of cv_w. As the coefficient of variation increases, the dynamic strategy becomes more

Figure 4.9: Impact of Uncertainties in Market Potential Change

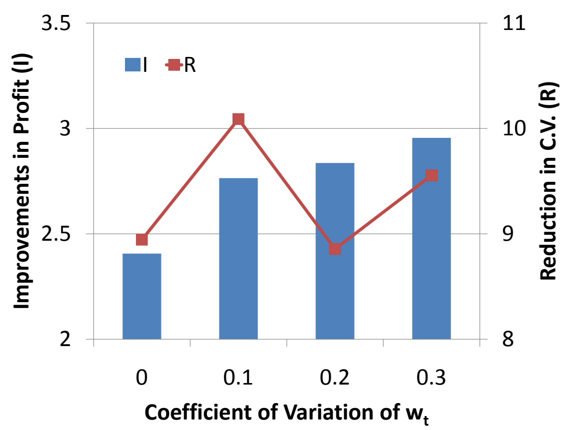

effective. The market potential changes are highly uncertain when, for example, competitors' strategies are difficult to anticipate or economic conditions are volatile. In such cases, the market entry timing decision that has been made before the realization of uncertain states may force the manufacturer to wait until he loses a great amount of market potential. In contrast, the market entry decision that has been made based on the realization of states enables the manufacturer to respond to such events. Hence, the value of the dynamic strategy is greater when the changes in the market potential are more uncertain.

Demand uncertainty: We next evaluate the impact of demand uncertainty, i.e., the impact of cv_A. In Figure 4.10, the value of the dynamic strategy increases as the

Figure 4.10: Impact of Demand Uncertainty

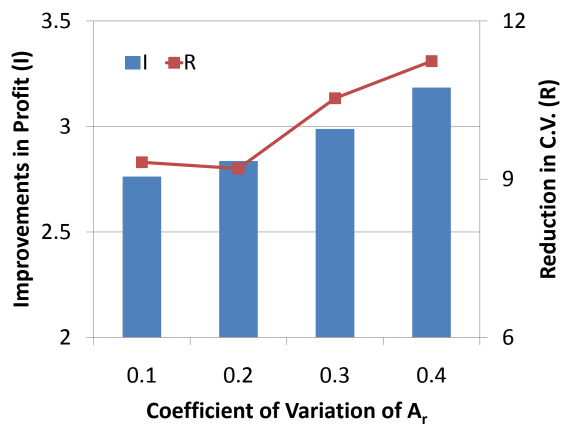

degree of demand uncertainty increases. The expected profit that the manufacturer can attain during the production and sales stage is smaller when the demand is more uncertain. On the other hand, the optimal control of pre-introduction decisions - market entry and process improvement decisions - is more critical when the expected profit during the production and sales season is smaller. Hence, the value of the dynamic strategy increases as the demand becomes more uncertain.

Size of the salvage market: Finally, we consider the impact of the size of the salvage market, A_s. Figure 4.11 illustrates the value of the dynamic strategy for several values of A_s. The manufacturer's expected profit during the production and sales stage increases as the size of the salvage market increases, which implies that the size of the salvage market and the degree of demand uncertainty have opposite effects on the value of the dynamic strategy. From this reason, both I and R decrease as the size of the salvage market increases.

Figure 4.11: Impact of the Size of the Salvage Market

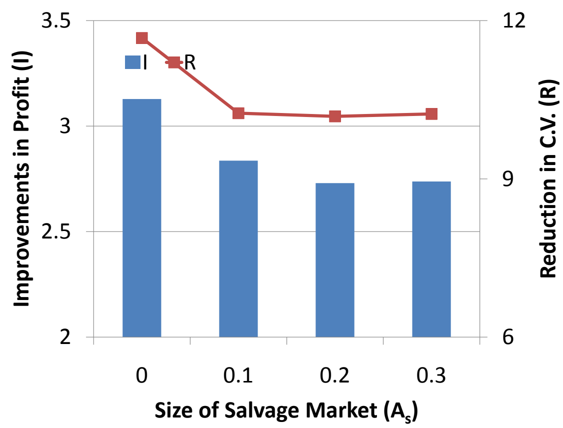

4.6. Conclusion and Discussion

In this chapter, we have studied a manufacturer's problem of dynamically optimizing the market entry timing and process improvement decisions for a new product. In contrast to the existing studies on new product introduction decisions, we have considered a dynamic strategy under which the manufacturer of the new product makes the market entry decision depending on the realization of two major uncertainties: competitors' movements and the readiness of the production process. We have developed a two-stage stochastic decision process for the manufacturer's problem. The first-stage is an optimal stopping problem that determines the timing for introducing a new product to the market, and process improvement decisions for the new product. The second-stage is a production and pricing decision problem for the new product. To solve this problem, we have formulated a two-stage dynamic program from which we establish the optimality of several threshold-type market entry policies and characterize optimal production and pricing decisions. Via numerical studies, we have shown that the dynamic strategy can increase the manufacturer's profit while substantially reducing the variability of profit. By examining the value of the dynamic strategy under various environments, we have shown that the dynamic strategy is effective when (i) the outcome of learning activities involve a great amount of uncertainties, (ii) the manufacturer can significantly reduce the production cost by improving the production process, (iii) flexible management of learning activities is possible, (iv) changes in the market potential are difficult to anticipate, (v) demand is highly uncertain, and (vi) the size of salvage market is small.

We have studied the market entry decision from a single firm's perspective. When multiple firms employ the same dynamic market entry strategies, the equilibrium outcome would be different from the conventional equilibrium outcome. Although the dynamic market entry strategy may trigger a more severe competition among competing firms, the impact of the dynamic strategy is not always negative for manufacturing firms. For example, when a firm is aware of the competitors' abilities to respond to its market entry decision, the firm may not want to enter the market with an ill-understood production process to simply beat the competition. Investigating the equilibrium outcome when multiple manufacturing firms adopt dynamic market entry strategies would be an interesting research problem.

Our study also highlights the importance of managing risks in new product introduction. Although balancing the trade-off between the time-to-market and the completeness of the production process has been the main subject of many studies, the risk involved in new product introduction decisions has received little attention. We have shown that the dynamic market entry strategy significantly reduces the variability of profit, i.e., the risk involved in new product introduction decisions, while increasing profit. We hope that our research paves a new line of research that investigates new product introduction decisions that take risk into consideration.

Appendix A

Chapter 2 Appendices

A.1. Stochastic Monotonicities of the State Transition

In this section, we provide some results that can be used to prove stochastic monotonicities of state transition models. We first provide a theorem.

Theorem A.1. *Suppose that $\tilde{x}(\theta) = \phi(\theta, \tilde{\xi})$, where ϕ is a deterministic function and $\tilde{\xi}$ is a random vector.*

1. *If ϕ is linear in θ for every realization of $\tilde{\xi}$, then $\tilde{x}(\theta)$ is stochastically convex in θ.*
2. *If ϕ is increasing in θ for every realization of $\tilde{\xi}$, then $\tilde{x}(\theta)$ is stochastically increasing in θ.*

Proof. For the first part, suppose that for any realization ξ of $\tilde{\xi}$, $\phi(\theta, \xi)$ is a deterministic linear function of θ. Then, for any convex function u, $u(\phi(\theta, \xi))$ is convex in θ, because the composition of a convex function and a linear function is convex. Then, $E[u(\phi(\theta, \tilde{\xi}))]$ is convex in θ, because taking expectation preserves the convexity. Therefore, $\tilde{x}(\theta)$ is stochastically convex. Part 2 can be proved in a similar way. Suppose for any realization ξ of $\tilde{\xi}$, $\phi(\theta, \xi)$ is a deterministic increasing function of

θ. Then, for any increasing function of u, $u(\phi(\theta, \xi))$ is increasing in θ for each ξ. Therefore, $E[u(\phi(\theta, \tilde{\xi}))]$ is increasing in θ, which concludes the theorem. □

This theorem can be used to prove stochastic monotonicities of a commonly used set of state transition models listed below. With the exception of the first three, the other parameterized random variables are from Shaked and Shanthikumar (2007). Additional results can be found in Müller and Stoyan (2002).

1. $\tilde{x}_{t+1}(x_t) = \tilde{\xi} + x_t$ is stochastically convex and stochastically increasing.

2. $\tilde{x}_{t+1}(x_t) = \tilde{\xi} x_t$ for a positive random variable $\tilde{\xi}$ is stochastically convex and stochastically increasing.

3. $\tilde{x}_{t+1}(x_t) = x_t + (1 - x_t)\tilde{\xi}$ for $x_t \in [0, 1]$ and $\tilde{\xi} \in [0, 1]$ is stochastically convex and stochastically increasing.

4. Suppose $\tilde{x}_{t+1,1}(x_{t,1}, x_{t,2})$ is a normal random variable with mean $x_{t,1}$ and standard deviation $x_{t,2}$. Then it is stochastically increasing in $x_{t,1}$ and stochastically convex in $x_{t,2}$.

5. Suppose $\tilde{x}_{t+1}(x_t)$ is a Poisson random variable with mean x_t. Then it is stochastically increasing in x_t.

6. Suppose $\tilde{x}_{t+1,1}(x_{t,1}, x_{t,2})$ is a binomial random variable with mean $x_{t,1}x_{t,2}$ and variance $x_{t,1}x_{t,2}(1 - x_{t,2})$. Then it is stochastically increasing in x_t.

7. Suppose $\tilde{x}_{t+1}(x_t)$ is uniformly distributed on $[0, x_t]$. Then it is stochastically increasing in $x_t > 0$.

8. Suppose $\tilde{x}_{t+1}(x_t)$ is uniformly distributed on $\{0, 1, \ldots, x_t - 1\}$. Then it is stochastically increasing in $x_t > 0$.

9. Let $Y(m)$, $m = 1, 2, \ldots$, be a sequence of non-negative i.i.d. random variables. Then $\tilde{x}_{t+1}(x_t) = \sum_{m=1}^{x_t} Y(m)$ is stochastically increasing in x_t.

10. Let $Y_i(j)$ for $i = 1, 2, \ldots, d$ and $j = 1, 2, \ldots$, be a sequence of non-negative i.i.d. random variables. Then $\tilde{x}_{t+1}(x_t) = (\sum_{j=1}^{x_{t,1}} Y_1(j), \ldots, \sum_{j=1}^{x_{t,d}} Y_d(j))$ is stochastically increasing in x_t.

A.2. Proofs

Proof of Proposition 2.2. The proof is based on an induction argument. At period $t = T-1$, we have $B_{T-1}(x_{T-1}) = M_{T-1}(x_{T-1})$, which is convex in x_{T-1}. Next assume for the induction argument that $B_{t+1}(x)$ is convex in x. The composition $\max\{0, B_{t+1}(x)\}$ is convex in x because function $\max\{0, x\}$ is convex increasing. The stochastic convexity of the state transition implies that $E[\max\{0, B_{t+1}(\tilde{x}_{t+1}(x_t))\}]$ is also convex in x_t. Because convexity is preserved under addition, the benefit function $B_t(x_t) = M_t(x_t) + \alpha E[\max\{0, B_{t+1}(\tilde{x}_{t+1}(x_t))\}]$ is convex in x_t, which concludes the induction argument and hence the proof of Part 1. Part 2 can be proved in a similar way to the proof of Proposition 2.1 Part 2. □

Proof of Proposition 2.3. The proof is based on an induction argument. We consider the increasing case. At period $t = T-1$, we have $B_{T-1}(x_{T-1}) = M_{T-1}(x_{T-1})$, which is increasing in $x_{T-1,i}$. Next assume for the induction argument that $B_{t+1}(x_{t+1})$ is increasing in $x_{t+1,i}$. The composition $\max\{0, B_{t+1}(x_{t+1})\}$ is also increasing in $x_{t+1,i}$. For a given $x_{t,-i}$, define $f(x_{t+1,i}) \equiv E[\max\{0, B_{t+1}(x_{t+1,i}, \tilde{x}_{t+1,-i}(x_{t,-i}))\}]$. Because taking expectation preserves the increasing property, $f(x_{t+1,i})$ is increasing in $x_{t+1,i}$. Then the stochastic increasing property of $\tilde{x}_{t+1,i}(x_t)$ in $x_{t,i}$ implies that $E[\max\{0, B_{t+1}(\tilde{x}_{t+1,i}(x_t), \tilde{x}_{t+1,-i}(x_{t,-i}))\}] = E[f(\tilde{x}_{t+1,i}(x_t)]$ is increasing in $x_{t,i}$. Because increasing property is preserved under addition,

$$B_t(x_t) = M_t(x_t) + \alpha E[\max\{0, B_{t+1}(\tilde{x}_{t+1}(x_t))\}]$$

is also increasing in $x_{t,i}$, which concludes the induction argument, hence the proof of the first part. □

Proof of Proposition 2.4. The proof is based on an induction argument. At period

$t = T-1$, we have $B_{T-1}(x_{T-1}) = M_{T-1}(x_{T-1})$, which is convex in $x_{T-1,i}$. Next assume for the induction argument that $B_{t+1}(x_{t+1})$ is convex in $x_{t+1,i}$. Then the composition $\max\{0, B_{t+1}(x_{t+1})\}$ is also convex in $x_{t+1,i}$. For a given $x_{t,-i}$, define $f(x_{t+1,i}) \equiv E[\max\{0, B_{t+1}(x_{t+1,i}, \tilde{x}_{t+1,-i}(x_{t,-i}))\}]$. Because taking expectation preserves convexity, $f(x_{t+1,i})$ is convex in $x_{t+1,i}$. Then the stochastic convex property of $\tilde{x}_{t+1,i}(x_t)$ in $x_{t,i}$ implies that $E[\max\{0, B_{t+1}(\tilde{x}_{t+1}(x_t))\}] = E[\max\{0, B_{t+1}(\tilde{x}_{t+1,i}(x_t), \tilde{x}_{t+1,-i}(x_{t,-i}))\}] = E[f(\tilde{x}_{t+1,i}(x_t)]$ is convex in $x_{t,i}$. Then, $B_t(x_t) = M_t(x_t) + \alpha E[\max\{0, B_{t+1}(\tilde{x}_{t+1}(x_t))\}]$ is convex in $x_{t,i}$, which concludes the induction argument, hence the proof of the first part. □

Proof of Proposition 2.5. The proof is based on an induction argument. At period $t = T-1$, we have $B_{T-1}(x_{T-1}) = M_{T-1}(x_{T-1})$, which is increasing in x_{T-1}. Next assume for the induction argument that $B_{t+1}(x_{t+1})$ is increasing in each element of x_{t+1}. The composition of an increasing function with $\max\{0, x\}$ is also increasing, hence, $\max\{0, B_{t+1}(x_{t+1})\}$ is increasing in x_{t+1}. Because the state transition $\tilde{x}_{t+1}(x_t)$ is stochastically increasing in x_t, $E[\max\{0, B_{t+1}(\tilde{x}_{t+1}(x_t))\}]$ is increasing in x_t. Because the multi-dimensional increasing property is preserved under addition, the benefit function $B_t(x_t) = M_t(x_t) + \alpha E[\max\{0, B_{t+1}(\tilde{x}_{t+1}(x_t))\}]$ is also an increasing function, which concludes the induction argument and the proof of the first part. □

Proof of Proposition 2.6. We define a set $A_t(x_{t,-i}) \equiv \{x_{t,i} : B_t(x_{t,i}, x_{t,-i}) \leq 0, x_t \in X\}$. To prove Part 1, first consider the case when $B_t(x_t)$ is increasing in both $x_{t,i}$ and $x_{t,j}$ for $i \neq j$. Then for given y_t and z_t such that $z_{t,-(i,j)} = y_{t,-(i,j)}$ and $y_{t,j} \geq z_{t,j}$, the following inequality holds:

$$\begin{aligned} B_t(\overline{x}_{t,i}(y_{t,-i}), z_{t,j}, z_{t,-(i,j)}) &= B_t(\overline{x}_{t,i}(y_{t,-i}), z_{t,j}, y_{t,-(i,j)}) \\ &\leq B_t(\overline{x}_{t,i}(y_{t,-i}), y_{t,j}, y_{t,-(i,j)}) = 0. \end{aligned}$$

The inequality stems from the fact that $B_t(x_t)$ is increasing in $x_{t,j}$, and the last equality stems from the definition of $\overline{x}_{t,i}(y_{t,-i})$. Therefore, $\overline{x}_{t,i}(y_{t,-i}) \in A_t(z_{t,-i})$, which in turn implies that $\overline{x}_{t,i}(z_{t,-i}) \geq \overline{x}_{t,i}(y_{t,-i})$. This concludes the proof of Part 1 for the increasing case. The decreasing case can be proved in a similar way.

To prove Part 2, note that for a given y_t and z_t such that $y_{t,-(i,j)} = z_{t,-(i,j)}$ and $y_{t,j} \leq z_{t,j}$, the following inequality holds:

$$\begin{aligned} B_t(\overline{x}_{t,i}(y_{t,-i}), z_{t,j}, z_{t,-(i,j)}) &= B_t(\overline{x}_{t,i}(y_{t,-i}), z_{t,j}, y_{t,-(i,j)}) \\ &\leq B_t(\overline{x}_{t,i}(y_{t,-i}), y_{t,j}, y_{t,-(i,j)}) = 0. \end{aligned}$$

The inequality stems from the fact that $B_t(x_t)$ is decreasing in $x_{t,j}$, and the equality stems from the definition of $\overline{x}_{t,i}(y_{t,-i})$. Therefore, $\overline{x}_{t,i}(y_{t,-i}) \in A_t(z_{t,-i})$, hence, $\overline{x}_{t,i}(y_{t,-i}) \leq \overline{x}_{t,i}(z_{t,-i}) \equiv \sup A_t(z_{t,-i})$. The argument for the increasing property of $\underline{x}_{t,j}(x_{t,-j})$ in $x_{t,i}$ can be proved in a similar way.

To prove part 3, first consider the case when $B_t(x_t)$ is increasing in $x_{t,i}$ and convex in $x_{t,j}$. Then for a given y_t and z_t such that $y_{t,-(i,j)} = z_{t,-(i,j)}$ and $y_{t,i} \leq z_{t,i}$, the following inequality holds: $B_t(y_{t,i}, \overline{x}_{t,j}(z_{t,-i}), y_{t,-(i,j)}) = B_t(y_{t,i}, \overline{x}_{t,j}(z_{t,-i}), z_{t,-(i,j)}) \leq B_t(z_{t,i}, \overline{x}_{t,j}(z_{t,-i}), z_{t,-(i,j)}) = 0$. The inequality stems from the fact that $B_t(x_t)$ is increasing in $x_{t,i}$, and the last equality stems from the definition of $\overline{x}_{t,i}(z_{t,-i})$. Therefore, $\overline{x}_{t,i}(z_{t,-i}) \in A_t(y_{t,-i})$, hence, $\overline{x}_{t,i}(z_{t,-i}) \leq \overline{x}_{t,i}(y_{t,-i})$. Similarly, $B_t(y_{t,i}, \underline{x}_{t,j}(z_{t,-i}), y_{t,-(i,j)}) = B_t(y_{t,i}, \underline{x}_{t,j}(z_{t,-i}), z_{t,-(i,j)}) \leq B_t(z_{t,i}, \underline{x}_{t,j}(z_{t,-i}), z_{t,-(i,j)}) = 0$. Therefore, $\underline{x}_{t,i}(z_{t,-i}) \in A_t(y_{t,-i})$, hence, $\underline{x}_{t,i}(z_{t,-i}) \geq \underline{x}_{t,i}(y_{t,-i})$, which concludes Part 3. □

Proof of Proposition 2.7. For the first part, we define $A_t \equiv \{x \in X : B_t(x) \leq 0\}$. The decreasing property of $B_t(x)$ in t implies that $A_t \subset A_{t+1}$ for every t. Therefore, $\overline{x}_t = \sup A_t \leq \sup A_{t+1} = \overline{x}_{t+1}$, and $\underline{x}_t = \inf A_t \geq \inf A_{t+1} = \underline{x}_{t+1}$. For the second part, we define $A_t(x_{t,-i}) \equiv \{x_{t,i} : B_t(x_{t,i}, x_{t,-i}) \leq 0, x_t \in X\}$. The decreasing property of $B_t(x)$ in t implies that $A_t(x_{t,-i}) \subset A_{t+1}(x_{t,-i})$ for every t and every $x_{t,-i}$. Therefore, $\overline{x}_t(x_{t,-i}) = \sup A_t(x_{t,-i}) \leq \sup A_{t+1}(x_{t,-i}) = \overline{x}_{t+1}(x_{t,-i})$, and $\underline{x}_t(x_{t,-i}) = \inf A_t(x_{t,-i}) \geq \inf A_{t+1}(x_{t,-i}) = \underline{x}_{t+1}(x_{t,-i})$. □

Proof of Proposition 2.8. For the first part, we define two sets $A_t \equiv \{x \in X : B_t(x) \leq 0\}$ and $O_t \equiv \{x \in X : M_t(x) \leq 0\}$. By definition, $B_t(x) \geq M_t(x)$ for every x, which implies that $A_t \subset O_t$. Therefore, $\overline{x}_t = \sup A_t \leq \sup O_t$, and $\underline{x}_t = \inf A_t \geq \inf O_t$. For the second part, we define two sets $A_t(x_{t,-i}) \equiv \{x_{t,i} : B_t(x_{t,i}, x_{t,-i}) \leq 0, x_t \in X\}$ and $O_t(x_{t,-i}) \equiv \{x_{t,i} : M_t(x_{t,i}, x_{t,-i}) \leq 0, x_t \in X\}$. By definition, $B_t(x) \geq M_t(x)$ for every

x, which implies that $A_t(x_{t,-i}) \subset O_t(x_{t,-i})$. Therefore, $\overline{x}_t(x_{t,-i}) = \sup A_t(x_{t,-i}) \leq \sup O_t(x_{t,-i})$, and $\underline{x}_t(x_{t,-i}) = \inf A_t(x_{t,-i}) \geq \inf O_t(x_{t,-i})$. □

Proof of Proposition 2.10. The proof is based on an induction argument. At period $t = T-1$, we have $B_{T-1}(a_t, x) = M_{T-1}(a_t, x)$. Let $a_t^*(x) = \arg\max_{a_t \in A_t} B_t(a_t, x)$. For any x^1, x^2, and $\beta \in (0,1)$, we have $\overline{B}_{T-1}(\beta x^1 + (1-\beta)x^2) = B_{T-1}(a_{T-1}^*(\beta x^1 + (1-\beta)x^2), \beta x^1 + (1-\beta)x^2) \leq \beta B_{T-1}(a_{T-1}^*(\beta x^1 + (1-\beta)x^2), x^1) + (1-\beta) B_{T-1}(a_{T-1}^*(\beta x^1 + (1-\beta)x^2), x^2) \leq \beta B_{T-1}(a_{T-1}^*(x^1), x^1) + (1-\beta) B_{T-1}(a_{T-1}^*(x^2), x^2)$, where the first inequality is from the convexity of $B_{T-1}(a, x)$, and the second inequality is by the definition of $a_{T-1}^*(x)$. Next assume for the induction argument that the benefit function $B_{t+1}(x_{t+1})$ is convex in x_{t+1}. The composition of a convex function and $\max\{0, x\}$ is also convex, hence, $\max\{0, \overline{B}_{t+1}(x)\}$ is a convex function. Because the state transition $\tilde{x}_{t+1}(a_t, x_t)$ is stochastically convex in x_t, $E[\max\{0, \overline{B}_{t+1}(\tilde{x}_{t+1}(a_t, x_t))\}]$ is convex in x_t for each fixed a_t. Because the convexity is preserved under summation, the benefit function $B_t(a_t, x_t) = M_t(a_t, x_t) + \alpha E[\max\{0, \overline{B}_{t+1}(a_t, \tilde{x}_{t+1}(x_t))\}]$ is convex in x_t for each a_t. By applying the same argument that we applied on $\overline{B}_{T-1}(x)$, we conclude that $\overline{B}_t(x_t) = \sup_{a_t \in A_t} B_t(a_t, x_t)$ is convex in x_t, which concludes the induction hypothesis and the proof of the Proposition. □

Proof of Proposition 2.11. For Part 1, we first prove that $V_t(x|T)$ is increasing in T for every t by an induction argument. Note that $V_t(x|t+1) = \max\{S(x), T(x) + \alpha E[V_{t+1}(\tilde{x}(x)|t+1)]\} \geq S(x) = V_t(x|t)$. Assume for an induction argument $V_t(x|T) \geq V_t(x|T-1)$. Because $\tilde{x}(x)$ is time-homogeneous and the reward functions are time-invariant, $V_t(x|T) = V_{t+1}(x|T+1)$, which implies

$$\begin{aligned} V_t(x|T+1) &= \max\{S(x), T(x) + \alpha E[V_{t+1}(\tilde{x}(x)|T+1)]\} \\ &= \max\{S(x), T(x) + \alpha E[V_t(\tilde{x}(x)|T)]\}. \end{aligned} \tag{A.1}$$

Therefore, we have $V_t(x|T+1) = \max\{S(x), T(x) + \alpha E[V_t(\tilde{x}(x)|T)]\} \geq \max\{S(x), T(x) + \alpha E[V_t(\tilde{x}(x)|T-1)]\} = V_t(x|T)$ for every t, which concludes the induction argument.

If $V_2(x|T)$ is increasing in T, $B_1(x|T) = \alpha E[V_2(\tilde{x}(x)|T)] + C(x) - S(x)$ is also increasing in T for every x. Therefore, $\lim_{T\to\infty} B_1(x|T)$ is well-defined if we allow it to take an infinite value. To conclude Part 1, we need to prove that the limit of $B_1(x|T)$ is indeed $B^*(x)$. To do so, we first prove that

$$\lim_{T\to\infty} E[V_2(\tilde{x}(x)|T)] = E[\lim_{T\to\infty} V_2(\tilde{x}(x)|T)] = E[V^*(\tilde{x}(x))], \tag{A.2}$$

under Assumption 2.1. For notational convenience, we denote $V_2(\tilde{x}(x)|T)$ by Y_T and $V^*(\tilde{x}(x))$ by Y^*. The increasing property of $V_t(x|T)$ in T and Assumption 2.1 imply that $Y_T \uparrow Y^*$ almost surely. However, we cannot directly apply the monotone convergence theorem (Durrett 1996 p. 464) to (A.2) because the monotone convergence theorem is applicable to non-negative random variables. Hence, we instead consider $Y_T = (Y_T - Y_2) + Y_2$. Because $Y_T \geq Y_2$ for every $T \geq 2$, $Y_T - Y_2 \geq 0$ almost surely. In addition, $Y_2 = V_2(\tilde{x}(x)|2) = S(\tilde{x}(x))$ is integrable from the assumption that $E|S(x_t)| < \infty$. When $(Y_T - Y_2) \geq 0$ and Y_2 is integrable, $E[Y_T] = E[(Y_T - Y_2) + Y_2] = E[Y_T - Y_2] + E[Y_2]$ (Durrett 1996 p. 455). From the same reason, $E[Y^*] = E[Y^* - Y_2] + E[Y_2]$. Therefore, (A.2) can be derived as $\lim_{T\to\infty} E[Y_T] = \lim_{T\to\infty} E[(Y_T - Y_2) + Y_2] = \lim_{T\to\infty} E[Y_T - Y_2] + E[Y_2] = E[\lim_{T\to\infty}(Y_T - Y_2)] + E[Y_2] = E[Y^* - Y_2] + E[Y_2] = E[Y^*]$, where the third equality is by the monotone convergence theorem. Finally, we have $\lim_{T\to\infty} B_1(x|T) = \lim_{T\to\infty} \alpha E[V_2(\tilde{x}(x)|T)] + C(x) - S(x) = \alpha E[V^*(\tilde{x}(x))] + C(x) - S(x) = B^*(x)$, which concludes the proof of Part 1.

Next we consider Part 2. We prove $\overline{x}_{1|T} \downarrow \overline{x}$ as $T \to \infty$, then the other cases can be proved in a similar way. From Part 1, $B_1(\overline{x}_{1|T+1}|T) \leq B_1(\overline{x}_{1|T+1}|T+1) = 0$, which implies $\overline{x}_{1|T} \geq \overline{x}_{1|T+1}$ for every T. Similarly, $B_1(\overline{x}|T) \leq B^*(\overline{x}) = 0$, which implies $\overline{x}_{1|T} \geq \overline{x}$ for every T. Therefore, $\lim_{T\to\infty} \overline{x}_{1|T} \geq \overline{x}$. Now suppose $y \equiv \lim_{T\to\infty} \overline{x}_{1|T} > \overline{x}$. This inequality implies that $B^*(y) > 0$. Because $B_1(y|T)$ converges to $B^*(y)$ as $T \to \infty$, there exits $W < \infty$ such that $B_1(y|W) > 0$, which implies that $y = \lim_{T\to\infty} \overline{x}_{1|T} > \overline{x}_{1|W}$. This inequality contradicts the decreasing property of $\overline{x}_{1|T}$ in T. Therefore, $\lim_{T\to\infty} \overline{x}_{1|T} = \overline{x}$, which concludes the proof. □

Appendix B

Chapter 3 Appendices

B.1. Notation

We denote the supplier by (s), the manufacturer by (m), and the centralized decision maker by (c).

Demand and Forecast
X_{N+1} : demand during the sales period
e_j : random variable representing the impact of event j on demand
E_{n_s,n_m} : set of events whose information is obtained by (s) at period n_s and by (m) at period n_m
$\delta_{n_s,n_m} = \prod_{j \in E_{n_s,n_m}} e_j$: random variable representing the total information obtained by (s) at period n_s and by (m) at period n_m
X_n^s : (s)'s demand forecast
X_n^m : (m)'s demand forecast
$\Delta_n^s = X_{n+1}^s - X_n^s$: difference between (s)'s subsequent forecasts
$\Delta_n^m = X_{n+1}^m - X_n^m$: difference between (m)'s subsequent forecasts
$A_n = X_n^m - X_n^s$: difference between the (s) and (m)'s forecasts
$\delta_n^s = X_{n+1}^s / X_n^s$: random variable representing the ratio of (s)'s successive forecasts
$\delta_n^m = X_{n+1}^m / X_n^m$: random variable representing the ratio of (m)'s successive forecasts
$\epsilon_n = \prod_{k=n}^N \delta_k^m$: random variable representing demand uncertainty at period n
$\xi_n = \prod_{k=n}^N \delta_k^s / \epsilon_n$: random variable representing information asymmetry at period n
$X_{N+1} = X_n^m \epsilon_n = X_n^s \xi_n \epsilon_n$
σ_Z : standard deviation of $\log(Z)$ of the log-normal random variable Z
$G_n(.)$, $g_n(.)$: c.d.f. and p.d.f. of ϵ_n
$F_n(.)$, $f_n(.)$: c.d.f. and p.d.f. of ξ_n

Cost Parameters
r : per unit retail price
w : per unit wholesale price
c : per unit production cost
c_n : per unit capacity cost at period n
C_n : fixed capacity cost at period n
$\underline{\pi}^m$: (m)'s reservation profit
(Optimal) Decision variables
u_n : (s)'s stopping decision
$\{K(.), P(.)\}$: menu of contracts
(K_n^{dc}, P_n^{dc}) : optimal contract
$[L_n, U_n]$: (s)'s optimal control-band
$n^* = \arg\min_n \hat{\pi}_n$
K_n^{cs} : (c)'s optimal capacity level
$[L_n^{cs}, U_n^{cs}]$: (c)'s optimal control-band
$n^{cs} \equiv \arg\min_n \hat{\pi}_n^{cs}$
Profit Functions
$\Pi_n^s(K(\xi), P(\xi), \xi, X_n^s)$: (s)'s expected profit
$\Pi_n^m(K(\xi), P(\xi), \xi, X_n^s)$: (m)'s expected profit
$\Pi_n^{tot}(K(\xi), P(\xi), \xi, X_n^s)$: total expected profit
$\Pi_n^{cs}(K, X_n^m)$: (c)'s expected profit
Optimal Profit Functions
$\pi_n(X_n^s)$: (s)'s optimal expected profit
$\hat{\pi}_n$: (s)'s normalized expected profit
$V_n(X_n^s)$: (s)'s optimal value-to-go function
$\pi_n^{cs}(X_n^m)$:(c)'s optimal expected profit
$\hat{\pi}_n^{cs}$: (c)'s normalized expected profit
$V_n^{cs}(X_n^m)$: (c)'s optimal value-to-go function

B.2. Additive Case

In Chapter 3, we have used the multiplicative MMFE for the capacity planning problem. Although the multiplicative MMFE fits actual data better than the additive MMFE (Hausman 1969, Heath and Jackson 1994), the additive model has also been used in the literature. In this subsection, we first extend the additive MMFE to the cases of multiple decision makers, and then investigate how the capacity planning problem changes when we employ the additive MMFE.

B.2.1 The Additive MMFE

We develop the additive MMFE for multiple decision makers. We define the difference between successive forecasts as $\delta_n^i \equiv X_{n+1}^i - X_n^i$, for $n < N$ and $\delta_N^i \equiv X_{N+1} - X_N^i$. As before, we assume that there are in total K events that affect demand, and let e_j be the random variable that models the impact of event j. Unlike in the multiplicative case, now we assume that the change in the forecast due to each event is independent of the current forecast. In other words, after obtaining the information of event j, decision maker i updates his forecast from X_n^i to $X_n^i + e_j$. Following this explanation, we first express demand as $X_{N+1} = \sum_{j=1}^{K} e_j$. Next, we define E_{n_s,n_m} as the set of events whose information is obtained by the supplier during period n_s and by the manufacturer during period n_m. Using this set, we define $\delta_{n_s,n_m} \equiv \sum_{j \in E_{n_s,n_m}} e_j$, which indicates the total demand information obtained by the supplier at period n_s and by the manufacturer at period n_m. We assume that each δ_{n_s,n_m} is normally distributed has a mean value of 0 except $\delta_{0,0}$[1]. When E_{n_s,n_m} is an empty set, $\delta_{n_s,n_m} = 0$.

Given this construction, we can express demand as $X_{N+1} = \sum_{n_s=0}^{N} \sum_{n_m=0}^{N} \delta_{n_s,n_m}$. The supplier's information set at the beginning of period n is

$$\mathcal{F}_n^s \equiv \sigma([\delta_{0,0}, \ldots, \delta_{0,N}], \ldots, [\delta_{n-1,0}, \ldots, \delta_{n-1,N}]).$$

Then, the supplier's demand forecast is $X_n^s = E[X_{N+1}|\mathcal{F}_n^s] = \sum_{n_s=0}^{n-1} \sum_{n_m=0}^{N} \delta_{n_s,n_m}$, and the difference between successive forecasts is $\delta_n^s = \sum_{n_m=0}^{N} \delta_{n,n_m}$. Because the sum of normal random variables is also a normal random variable, δ_n^s is also normally distributed. The manufacturer's demand forecast can be expressed in a similar way. Figure B.1 illustrates the information structure of the additive MMFE for two decision makers. From this construction, we can fully characterize the evolution of X_n^s and X_n^m by determining the value of $\delta_{0,0}$ and the variances of δ_{n_s,n_m}. The two variants of the MMFE are identical except that multiplication operators are replaced by addition operators and log-normal random variables are replaced by normal random variables

[1]Both decision makers have the information $\delta_{0,0}$ before the beginning of the forecast horizon. Hence, $\delta_{0,0}$ is a deterministic value. Note also that when $E[\delta_{n_s,n_m}] \neq 0$ for some (n_s, n_m), we can push this information to $\delta_{0,0}$ and normalize δ_{n_s,n_m} by $\delta_{n_s,n_m} - E[\delta_{n_s,n_m}]$. Hence, the assumption $E[\delta_{n_s,n_m}] = 0$ is without loss of generality.

Figure B.1: Information Structure of the additive MMFE

n	0		1		$\cdots$		N	(s)
0	$\delta_{0,0}$	+	$\delta_{0,1}$	+	$\cdots$	+	$\delta_{0,N}$	X_1^s
	+		+		+		+	+
1	$\delta_{1,0}$	+	$\delta_{1,1}$	+	$\cdots$	+	$\delta_{1,N}$	δ_1^s
	+		+		+		+	+
$\vdots$	$\vdots$	+	$\vdots$	+	$\ddots$	+	$\vdots$	$\vdots$
	+		+		+		+	+
N	$\delta_{N,0}$	+	$\delta_{N,1}$	+	$\cdots$	+	$\delta_{N,N}$	δ_N^s
(m)	X_1^m	+	δ_1^m	+	$\ldots$	+	δ_N^m	X_{N+1}

in the additive case.

We can define the *additive Martingale Model of Asymmetric Forecast Evolution* (a-MMAFE) in a similar way to the m-MMAFE. Because the supplier obtains no information strictly earlier than the manufacturer, we have $\delta_{n_s,n_n} = 0$ for every $n_m > n_s$. In this case, the manufacturer's demand uncertainty at the beginning of period n is given as $\epsilon_n \equiv \sum_{k=n}^{N} \delta_k^m$, and the manufacturer's private information is given as $\xi_n \equiv \sum_{k=n}^{N} \delta_k^s - \epsilon_n = \sum_{k=n}^{N} \sum_{n_m=0}^{n-1} \delta_{k,n_m}$. Finally, demand and forecasts have the following relation: $X_{N+1} = X_n^m + \epsilon_n = X_n^s + \xi_n + \epsilon_n$.

B.2.2 Determining the Optimal Time to Offer an Optimal Mechanism

We next revisit the capacity planning problem with the a-MMAFE. The problem setting is the same as before except that now demand forecasts follow an a-MMAFE. We first consider the second-stage mechanism design problem. Suppose that the supplier has decided to offer a screening contract at period n. Then, the supplier has to determine the optimal screening contract given demand forecast X_n^s, demand uncertainty ϵ_n, and the manufacturer's private information ξ_n. When demand forecasts follow an a-MMAFE, the random variables ϵ_n and ξ_n are normally distributed.

As before, to determine the optimal screening contract, the supplier solves (3.3). However, because the demand and the forecast have the relationship; $X_{N+1} = X_n^s +$

$\xi_n + \epsilon_n$, now the supplier's expected profit is

$$\Pi_n^s(K(\check{\xi}), P(\check{\xi}), \xi_n, X_n^s) \equiv (w-c)E_{\epsilon_n}[\min(X_n^s+\xi_n+\epsilon_n, K(\check{\xi}))]+P(\check{\xi})-(c_nK(\check{\xi})+C_n),$$

and the manufacturer's expected profit is

$$\Pi_n^m(K(\check{\xi}), P(\check{\xi}), \xi_n, X_n^s) \equiv (r-w)E_{\epsilon_n}[\min(X_n^s+\xi_n+\epsilon_n, K(\check{\xi}))] - P(\check{\xi}).$$

Özer and Wei (2006) solve this problem in a static setting, but they provide structural properties of the optimal contract when ξ_n has an increasing probability density function. Because ξ_n does not have an increasing probability density in our model, we extend their result to the class of random variables ξ_n that have increasing failure rates (IFR).[2]

To solve the supplier's problem, we first introduce an equivalent formulation of (3.3), and then define a normalized problem. From Lemma 1 of Özer and Wei (2006), the optimization problem (3.3) has the following equivalent formulation:

$$\begin{aligned} \pi_n(X_n^s) \quad &\equiv \quad \max_{K(.)} E_{\xi_n}\Big[(r-c)E[\min(X_n^s+\xi_n+\epsilon_n, K(\xi_n))] - (c_nK(\xi_n)+C_n) \\ &\quad -\frac{1-F_n(\xi_n)}{f_n(\xi_n)}(r-w)G_n(K(\xi_n)-\xi_n-X_n^s)\Big] - \underline{\pi}^m \qquad \text{(B.1)} \\ s.t. \quad &\quad K(\xi) \text{ is increasing.} \end{aligned}$$

After determining the optimal capacity reservation function $K_n^{dc}(.)$ from (B.1), we derive the corresponding payment function as

$$\begin{aligned} P_n^{dc}(\xi) \quad &= \quad (r-w)E_{\epsilon_n}[\min(X_n^s+\xi+\epsilon_n, K_n^{dc}(\xi))] \\ &\quad - \int_{\underline{\xi}_n}^{\xi}(r-w)G_n(K_n^{dc}(x)-x-X_n^s)dx - \underline{\pi}^m. \qquad \text{(B.2)} \end{aligned}$$

[2]The failure rate of a random variable is defined as $\frac{f(x)}{1-F(x)}$, where $f(x)$ and $F(x)$ are the p.d.f. and the c.d.f. of the random variable. Normal random variables have IFRs.

Next, we define the normalized version of (B.1) as follows:

$$\begin{aligned} \hat{\pi}_n^s \quad \equiv \quad & \max_{K(.)} E_{\xi_n}\Big[(r-c)E_{\epsilon_n}[\min(\xi_n+\epsilon_n, K(\xi_n))] - c_n K(\xi_n) \\ & -\frac{1-F_n(\xi_n)}{f_n(\xi_n)}(r-w)G_n(K(\xi_n)-\xi_n)\Big] \\ s.t. \quad & K(\xi) \text{ is increasing.} \end{aligned} \tag{B.3}$$

We denote the optimal solution of (B.3) by $\hat{K}_n^{dc}(.)$, and then we can derive the normalized payment function as

$$\hat{P}_n^{dc}(\xi) = (r-w)E_{\epsilon_n}[\min(\xi_n+\epsilon_n, \hat{K}_n^{dc}(\xi))] - \int_{\underline{\xi}_n}^{\xi}(r-w)G_n(\hat{K}_n^{dc}(x)-x)dx - \underline{\pi}^m.$$

Based on these functions, we derive the following structural properties of the optimal contract and the expected profit:

Theorem B.1. *The following statements are true for all n:*

(a) $K_n^{dc}(\xi) = X_n^s + \hat{K}_n^{dc}(\xi)$.

(b) $P_n^{dc}(\xi) = X_n^s(r-w) + \hat{P}_n^{dc}(\xi) - \underline{\pi}^m$.

(c) $\pi_n(X_n^s) = X_n^s(r-c-c_n) + \hat{\pi}_n^s - C_n - \underline{\pi}^m$.

Theorem 3.4 and Theorem B.1 show the major difference between the multiplicative and the additive models. When demand uncertainty is additive to demand forecast, the impact of ϵ_n and ξ_n are independent of the forecast level, X_n^s. Hence, the normalized capacity reservation and prices, $(\hat{K}_n^{dc}, \hat{P}_n^{dc})$, are additive to X_n^s. Similarly, the normalized expected profit $\hat{\pi}_n^s$ is also additive to X_n^s.

Next, we discuss how to solve the optimization problem (B.3). The solution approach is the same as before. We first relax the constraint that $K(.)$ is increasing, and then verify that the optimal solution of the unconstraint problem is increasing in ξ.

Theorem B.2. *If ϵ_n and ξ_n have IFRs, then the following properties hold:*

(a) $\hat{K}_n^{dc}(\xi)$ *is the unique solution of the first-order condition*

$$(r-c)(1-G_n(K-\xi)) - c_n - \frac{1-F_n(\xi)}{f_n(\xi)}(r-w)g_n(K-\xi) = 0.$$

(b) *Both* $K_n^{dc}(\xi)$ *and* $P_n^{dc}(\xi)$ *are increasing in* ξ.

(c) *Both* $\Pi_n^s(K_n^{dc}(\xi), P_n^{dc}(\xi), \xi, X_n^s)$ *and* $\Pi_n^m(K_n^{dc}(\xi), P_n^{dc}(\xi), \xi, X_n^s)$ *are increasing in* ξ.

(d) $P_n^{dc}(K_n^{dc})$ *is an increasing concave function of* K_n^{dc}.

The first-order condition in Part (a) is slightly different from that of Theorem 3.5, but other structural properties of the optimal contract remain the same in the additive case.

We next consider the first-stage optimal stopping problem. The formulation of the first-stage problem is also the same as before except that now the state variable is updated $X_{n+1}^s = \delta_n^s + X_n^s$, and the reward of stopping at period n is $\pi_n(X_n^s) = X_n^s(r-c-c_n) + \hat{\pi}_n^s - C_n - \underline{\pi}^m$. The following theorem characterizes the optimal stopping policy:

Theorem B.3. *The following statements are true for all* n*:*

(a) *A control band policy that offers a capacity reservation contract at period* n *if* $X_n^s \in [L_n, U_n]$*, is optimal.*

(b) *When* $c_{n+1} > c_n$ *for all* n*, the upper threshold,* U_n*, is* ∞ *for all* n*. Hence, a lower threshold policy that offers the capacity reservation contract at period* n *if* $X_n^s \geq L_n$*, is optimal.*

(c) *Let* $n^* \equiv \arg\max_n \hat{\pi}_n - C_n$*. When* $c_{n+1} = c_n$ *for all* n*, the lower and upper thresholds satisfy that* $L_n = U_n = 0$ *for* $n \neq n^*$*, and that* $L_n = 0$ *and* $U_n = \infty$ *for* $n = n^*$*. Hence, a state-independent policy that offers the capacity reservation contract at period* n^* *is optimal.*

For both the additive and multiplicative cases, a control band policy is optimal. However, when both c_n and C_n are strictly increasing in n, the optimal thresholds of the two models have different limits: the multiplicative model satisfies $L_n = 0$ and the additive model satisfies $U_n = \infty$. When c_n is increasing in n, the supplier incurs more unit capacity costs by delaying the capacity decision, and the additional costs are proportionally increasing in the forecast level. In contrast, the impact of ϵ_n and ξ_n on profit is independent of the forecast level. Hence, delaying the capacity decision becomes less beneficial as the demand forecast increases, which implies the optimality of the lower threshold policy. Although both variants of the MMFE have been widely used in the literature, the multiplicative model is superior over the additive model in terms of consistency with empirical data (Hausman 1969, Heath and Jackson 1994). Hence, the multiplicative MMFE would be more appropriate to use for the capacity planning problem.

B.3. Proofs

Proof of Theorem 3.1. For Part (a), we first prove that X_n is integrable. This property holds directly from the square-integrability of X_{N+1}. Then, we have $E[X^i_{n+1}|\mathcal{F}^i_n] = E[E[X_{N+1}|\mathcal{F}^i_{n+1}]|\mathcal{F}^i_n] = E[X_{N+1}|\mathcal{F}^i_n] = X^i_n$, where the second equality is from the tower property of conditional expectation. Therefore, Part (a) is true. For Part (b), first note that $\sigma(X^i_n) \subseteq \mathcal{F}^i_n$. Then, again from the tower property, we have $E[X_{N+1}|X^i_n] = E[E[X_{N+1}|F^i_n]|X^i_n] = E[X^i_n|X^i_n] = X^i_n$. For Part (c), note that $\sigma(X^i_n) \subseteq \mathcal{F}^i_{n+l}$ for every $l \geq 0$. Then, from the tower property, we have $E[X^i_{n+l}|X^i_n] = E[E[X_{N+1}|\mathcal{F}^i_{n+l}]|X^i_n] = E[X_{N+1}|X^i_n] = X^i_n$. Finally, for Part (d), we have $E[\Delta^i_n] = E[X^i_{n+1} - X^i_n] = E[X^i_{n+1}] - E[X^i_n] = E[E[X_{N+1}|\mathcal{F}^i_{n+1}]] - E[E[X_{N+1}|\mathcal{F}^i_n]] = E[X_{N+1}] - E[X_{N+1}] = 0$. In addition, for any random variable Y that is measurable in $\mathcal{F}^i_n$, we have $E[\Delta^i_l Y] = E[E[\Delta^i_l Y|\mathcal{F}^i_n]] = E[Y E[\Delta^i_l|\mathcal{F}^i_n]] = E[Y E[X^i_{l+1} - X^i_l|\mathcal{F}^i_n]] = E[Y(X^i_n - X^i_n)] = 0$. Therefore, Δ^i_l is uncorrelated with $\mathcal{F}^i_n$ for every $l \geq n$, which concludes the proof. □

Proof of Theorem 3.2. By definition, we have $F^i_n \subseteq F^{cf}_n$ for every $i \in \{s, m\}$ and n.

Then, we have the following inequality:

$$\begin{aligned} E[(X_{N+1}-X_n^i)^2|\mathcal{F}_n^{cf}] &= E[(X_{N+1})^2|\mathcal{F}_n^{cf}] - 2E[X_{N+1}X_n^i|\mathcal{F}_n^{cf}] + E[(X_n^i)^2|\mathcal{F}_n^{cf}] \\ &= E[(X_{N+1})^2|\mathcal{F}_n^{cf}] - 2X_n^iE[X_{N+1}|\mathcal{F}_n^{cf}] + (X_n^i)^2 \\ &= E[(X_{N+1})^2|\mathcal{F}_n^{cf}] - 2X_n^iX_n^{cf} + (X_n^i)^2 \\ &\geq E[(X_{N+1})^2|\mathcal{F}_n^{cf}] - (X_n^{cf})^2 = E[(X_{N+1}-X_n^{cf})^2|\mathcal{F}_n^{cf}], \end{aligned}$$

where the second equality is from the fact that X_n^i is measurable on $\mathcal{F}_n^{cf}$, the inequality is from $(X_n^{cf})^2 - 2X_n^iX_n^{cf} + (X_n^i)^2 \geq 0$, and the last equality is from $E[X_{N+1}X_n^{cf}|\mathcal{F}_n^{cf}] = X_n^{cf}E[X_{N+1}|\mathcal{F}_n^{cf}] = (X_n^{cf})^2$. By taking expectation on both sides, we have $E[(X_{N+1} - X_n^i)^2] \geq E[(X_{N+1} - X_n^{cf})^2]$, which concludes the proof. □

Proof of Theorem 3.3. For Part (a), we first note that $\sigma(X_n^s) \subseteq \mathcal{F}_n^s \subseteq \mathcal{F}_{n+l}^m$ for every $l \geq 0$. Then, from the tower property, we have $E[X_{n+l}^m|X_n^s] = E[E[X_{N+1}|\mathcal{F}_{n+l}^m]|X_n^s] = E[X_{N+1}|X_n^s] = X_n^s$. For Part (b), first note that $X_n^m = X_n^s + A_n \in \sigma(X_n^s, A_n)$. Because both X_n^m and X_n^s are measurable in $\mathcal{F}_n^m$, and we also have $\sigma(X_n^s, A_n) \subseteq \mathcal{F}_n^m$. Then, we have $E[X_{N+1}|X_n^s, A_n] = E[E[X_{N+1}|F_n^m]|X_n^s, A_n] = E[X_n^m|X_n^s, A_n] = X_n^m$. For Part (c), we need to note that $(F_1^s, ..., F_n^s, F_n^m)$ forms a filtration. Then, Part (c) follows Part (d) of Theorem 3.1. □

Proof of Lemma 3.1. To prove Part (a), we first show that we can replace PC with

$$\text{PC}': \Pi_n^m(K(\underline{\xi}_n), P(\underline{\xi}_n), \underline{\xi}_n, X_n^s) = \underline{\pi}^m$$

under IC. For any $\xi^1 < \xi^2$, we have

$$\begin{aligned} \Pi_n^m(K(\xi^1), P(\xi^1), \xi^1, X_n^s) &\leq \Pi_n^m(K(\xi^1), P(\xi^1), \xi^2, X_n^s) \\ &\leq \Pi_n^m(K(\xi^2), P(\xi^2), \xi^2, X_n^s), \end{aligned}$$

where the first inequality is from the fact that the profit increases with ξ_n for a fixed (K, P) and the second inequality is from IC. Therefore, PC for $\xi > \underline{\xi}_n$ is redundant once it is satisfied for $\underline{\xi}_n$. In addition, the supplier can increase $P(.)$ uniformly without breaking IC until $\Pi_n^m(K(\underline{\xi}_n), P(\underline{\xi}_n), \underline{\xi}_n, X_n^s)$ becomes $\underline{\pi}^m$. Therefore, (IC,PC′) is equivalent to (IC,PC) for this optimization problem.

Next, we show that (IC,PC′) implies the two conditions of Part (a) for (3.3). IC implies that $\max_{\check{\xi}}\Pi_n^m(K(\check{\xi}),P(\check{\xi}),\xi,X_n^s)=\Pi_n^m(K(\xi),P(\xi),\xi,X_n^s)$. Hence, by applying the envelope theorem on $\Pi_n^m(K(\xi),P(\xi),\xi,X_n^s)$, we can derive

$$\frac{d\Pi_n^m(K(\xi),P(\xi),\xi,X_n^s)}{d\xi}=\left.\frac{\partial\Pi_n^m(K(\check{\xi}),P(\check{\xi}),\xi,X_n^s)}{\partial\xi}\right|_{\check{\xi}=\xi}=(r-w)X_n^s\int_{\underline{\epsilon}_n}^{\frac{K(\xi)}{\xi X_n^s}}yg_n(y)dy.$$

By integrating both sides from $\underline{\xi}_n$ to ξ with the boundary condition of PC′, we can derive Condition (i).

Next we prove that (IC,PC′) implies Condition (ii). First, note that we have the following equation for every $\check{\xi}$ and ξ:

$$\begin{aligned}
\Pi_n^m(K(\check{\xi}),P(\check{\xi}),\xi,X_n^s) &= \int_{\underline{\xi}_n}^{\xi}\frac{\partial\Pi_n^m(K(\check{\xi}),P(\check{\xi}),x,X_n^s)}{\partial x}dx+\Pi_n^m(K(\check{\xi}),P(\check{\xi}),\underline{\xi}_n,X_n^s)\\
&= \Pi_n^m(K(\check{\xi}),P(\check{\xi}),\check{\xi},X_n^s)+\int_{\check{\xi}}^{\xi}\left[(r-w)X_n^s\int_{\underline{\epsilon}_n}^{\frac{K(\check{\xi})}{xX_n^s}}yg_n(y)dy\right]dx\\
&= \Pi_n^m(K(\xi),P(\xi),\xi,X_n^s)-\int_{\check{\xi}}^{\xi}\left[(r-w)X_n^s\int_{\underline{\epsilon}_n}^{\frac{K(x)}{xX_n^s}}yg_n(y)dy\right]dx\\
&\quad+\int_{\check{\xi}}^{\xi}\left[(r-w)X_n^s\int_{\underline{\epsilon}_n}^{\frac{K(\check{\xi})}{xX_n^s}}yg_n(y)dy\right]dx\\
&= \Pi_n^m(K(\xi),P(\xi),\xi,X_n^s)\\
&\quad+\int_{\check{\xi}}^{\xi}\left[(r-w)X_n^s\left(\int_{\underline{\epsilon}_n}^{\frac{K(\check{\xi})}{xX_n^s}}yg_n(y)dy-\int_{\underline{\epsilon}_n}^{\frac{K(x)}{xX_n^s}}yg_n(y)dy\right)\right]dx\\
&= \Pi_n^m(K(\xi),P(\xi),\xi,X_n^s)+\int_{\check{\xi}}^{\xi}\left[(r-w)X_n^s\left(-\int_{\frac{K(\check{\xi})}{xX_n^s}}^{\frac{K(x)}{xX_n^s}}yg_n(y)dy\right)\right]dx.
\end{aligned}$$

Then, IC implies that $\int_{\check{\xi}}^{\xi}\left[(r-w)X_n^s\left(-\int_{\frac{K(\check{\xi})}{xX_n^s}}^{\frac{K(x)}{xX_n^s}}yg_n(y)dy\right)\right]dx\le 0$ for every $\check{\xi}<\xi$, which in turn implies that $\int_{\frac{K(\check{\xi})}{\xi X_n^s}}^{\frac{K(\xi)}{\xi X_n^s}}yg_n(y)dy\ge 0$ for every $\check{\xi}<\xi$. Because $yg_n(y)\ge 0$, this inequality holds only if $K(.)$ is increasing. Therefore, (IC,PC′) implies the two conditions of Part (a).

For the converse, we prove that the two conditions of Part (a) imply (IC,PC′). First note that Condition (i) with $\xi = \underline{\xi}_n$ directly implies PC′. For IC, we recall the equation

$$\begin{aligned} & \Pi_n^m(K(\check{\xi}), P(\check{\xi}), \xi, X_n^s) - \Pi_n^m(K(\xi), P(\xi), \xi, X_n^s) \\ = & \int_{\check{\xi}}^{\xi} \left[(r-w)X_n^s \left(\int_{\underline{\epsilon}_n}^{\frac{K(\check{\xi})}{xX_n^s}} y g_n(y)dy - \int_{\underline{\epsilon}_n}^{\frac{K(x)}{xX_n^s}} y g_n(y)dy \right) \right] dx \qquad \text{(B.4)} \\ = & -\int_{\xi}^{\check{\xi}} \left[(r-w)X_n^s \left(\int_{\underline{\epsilon}_n}^{\frac{K(\check{\xi})}{xX_n^s}} y g_n(y)dy - \int_{\underline{\epsilon}_n}^{\frac{K(x)}{xX_n^s}} y g_n(y)dy \right) \right] dx. \qquad \text{(B.5)} \end{aligned}$$

If $\xi > \check{\xi}$, the integrand of (B.4) is non-positive, because Condition (ii) implies that $K(.)$ is increasing. Similarly, if $\xi < \check{\xi}$, the integrand of (B.5) is non-negative, which implies that $\Pi_n^m(K(\check{\xi}), P(\check{\xi}), \xi, X_n^s) \leq \Pi_n^m(K(\xi), P(\xi), \xi, X_n^s)$ for every $\check{\xi}$ and ξ. Therefore, (IC,PC) and the two conditions of Part (a) are equivalent for the optimization problem (3.3).

To prove Part (b), we first derive $E_{\xi_n}[\Pi_n^s(K(\xi_n), P(\xi_n), \xi_n, X_n^s)]$ as

$$E_{\xi_n}\left[\Pi_n^{tot}(K(\xi_n), P(\xi_n), \xi_n, X_n^s) - \int_{\underline{\xi}_n}^{\xi_n} \left[(r-w)X_n^s \int_{\underline{\epsilon}_n}^{\frac{K(x)}{xX_n^s}} y g_n(y)dy\right] dx\right] - \underline{\pi}^m.$$

Finally, we can derive (3.4) by applying integration by parts and Condition (ii) to this equation, which concludes the proof. □

Proof of Theorem 3.4. We define $\hat{K}(.) = K(.)/X_n^s$. By construction, $K(.)$ is increasing if and only if $\hat{K}(.)$ is increasing. Using $\hat{K}(.)$, we can derive the objective function of (3.4) as

$$X_n^s \left[(r-c)E[\min(\xi_n\epsilon_n, \hat{K}(\xi_n))] - c_n\hat{K}(\xi_n) - \frac{1-F_n(\xi_n)}{f_n(\xi_n)}(r-w)\int_{\underline{\epsilon}_n}^{\frac{\hat{K}(\xi_n)}{\xi_n}} y g_n(y)dy\right]$$
$$-C_n - \underline{\pi}^m.$$

Because the term inside of [.] is the objective function of (3.6), Part (a) and Part (c) hold. Finally, by applying Part (a) to (3.5), we can prove Part (b). □

Proof of Theorem 3.5. For Part (a), we first define

$$H(K,\xi) \equiv (r-c)E[\min(\epsilon_n \xi, K)] - c_n K - \frac{1-F_n(\xi)}{f_n(\xi)}(r-w)\int_{\underline{\epsilon}_n}^{\frac{K(\xi)}{\xi}} y g_n(y) dy.$$

Without the constraint that $K(.)$ is increasing, the objective function (3.6) can be maximized with the values of K that maximizes $H(K,\xi)$ for each ξ. If the maximizer of the unconstraint problem is increasing, it is indeed the optimal solution for (3.6). We prove that this approach works when ϵ_n and ξ_n have IGFRs. We first prove that $H(K,\xi)$ is quasi-concave in K and have a finite maximizer. The first-order derivative is

$$\begin{aligned}\frac{\partial H(K,\xi)}{\partial K} &= (r-c)(1-G_n(\frac{K}{\xi})) - c_n - \frac{1-F_n(\xi)}{\xi f_n(\xi)}(r-w)\frac{K}{\xi}g_n(\frac{K}{\xi}) \\ &= (1-G_n(\frac{K}{\xi}))\left(r-c-\frac{1-F_n(\xi)}{\xi f_n(\xi)}(r-w)\frac{\frac{K}{\xi}g_n(\frac{K}{\xi})}{1-G_n(\frac{K}{\xi})}\right) - c_n.\end{aligned}$$

Then, the second order derivative at the point in which $\frac{\partial H(K,\xi)}{\partial K}=0$ can be derived as

$$\begin{aligned}\left.\frac{\partial^2 H(K,\xi)}{\partial K^2}\right|_{\frac{\partial H(K,\xi)}{\partial K}=0} &= -\frac{1}{\xi}g_n(\frac{K}{\xi})\left(\frac{c_n}{1-G_n(\frac{K}{\xi})}\right) \\ &+ (1-G_n(\frac{K}{\xi}))\left(-\frac{1-F_n(\xi)}{\xi f_n(\xi)}(r-w)\frac{d}{dK}\left(\frac{\frac{K}{\xi}g_n(\frac{K}{\xi})}{1-G_n(\frac{K}{\xi})}\right)\right) < 0.\end{aligned}$$

The inequality is from the fact that $\frac{\frac{K}{\xi}g_n(\frac{K}{\xi})}{1-G_n(\frac{K}{\xi})}$ is increasing in K due to the IGFR assumption. The inequality is strict because $g_n(\frac{K}{\xi})\left(\frac{c_n}{1-G_n(\frac{K}{\xi})}\right)$ is strictly positive. In other words, $H(K,\xi)$ is quasi-concave and the slope of the fuction is strictly negative at the point it crosses 0. Finally, we have $\frac{\partial H(K,\xi)}{\partial K}|_{K<\underline{\xi}_n} = r-c-c_n > 0$, and $\lim_{K\to\infty}\frac{\partial H(K,\xi)}{\partial K} = -c_n < 0$. Therefore, there exists a finite solution of $\frac{\partial H(K,\xi)}{\partial K}=0$. In addition, because $\frac{\partial H(K,\xi)}{\partial K}$ is strictly decreasing at $\frac{\partial H(K,\xi)}{\partial K}=0$, there is only one solution that satisfies the first-order condition.

Next we prove that the function $K(.)$ that satisfies the first-order conditions is increasing in ξ. Note that

$$\begin{aligned}\left.\frac{\partial^2 H(K,\xi)}{\partial K \partial \xi}\right|_{\frac{\partial H(K,\xi)}{\partial K}=0} &= \frac{1}{\xi^2} g_n(\frac{K}{\xi}) \left(\frac{c_n}{1-G_n(\frac{K}{\xi})} \right) \\ &+ (1-G_n(\frac{K}{\xi})) \left(-\frac{d}{d\xi} \left[\frac{1-F_n(\xi)}{\xi f_n(\xi)} \right] (r-w) \left(\frac{\frac{K}{\xi} g_n(\frac{K}{\xi})}{1-G_n(\frac{K}{\xi})} \right) \right) \\ &+ (1-G_n(\frac{K}{\xi})) \left(-\frac{1-F_n(\xi)}{\xi f_n(\xi)} (r-w) \frac{d}{d\xi} \left(\frac{\frac{K}{\xi} g_n(\frac{K}{\xi})}{1-G_n(\frac{K}{\xi})} \right) \right) > 0,\end{aligned}$$

where the inequality is from the fact that both $\frac{1-F_n(\xi)}{\xi f_n(\xi)}$ and $\frac{\frac{K}{\xi} g_n(\frac{K}{\xi})}{1-G_n(\frac{K}{\xi})}$ are decreasing in ξ due to the IGFR assumption. Therefore, $K_n^{dc}(\xi)$ is increasing in ξ, which concludes Part (a).

We already proved the increasing property of $K_n^{dc}(.)$ of Part (b) in the proof of Part (a). The increasing property of $P_n^{dc}(.)$ is from

$$\frac{dP_n^{dc}(\xi)}{d\xi} = (r-w)(1-G_n(\frac{K_n^{dc}(\xi)}{\xi}))\frac{dK_n^{dc}(\xi)}{d\xi} > 0,$$

which concludes Part (b).

The increasing property of $\Pi_n^m(K_n^{dc}(\xi), P_n^{dc}(\xi), \xi, X_n^s)$ in Part (c) stems directly from Part (a) of Lemma 3.1. Note that the supplier's expected profit is

$$\begin{aligned}\Pi_n^s(K_n^{dc}(\xi), P_n^{dc}(\xi), \xi, 1) &= (r-c)E[\min(\epsilon_n \xi, K_n^{dc}(\xi))] - c_n K_n^{dc}(\xi) \\ &- \int_{\underline{\xi}_n}^{\xi} \left[(r-w) \int_{\underline{\epsilon}_n}^{\frac{K_n^{dc}(x)}{x}} y g_n(y) dy \right] dx - \underline{\pi}^m,\end{aligned}$$

from which we can derive

$$\begin{aligned}
&\frac{d\Pi_n^s(K_n^{dc}(\xi), P_n^{dc}(\xi), \xi, 1)}{d\xi} \\
= &\ (w-c)\int_{\epsilon_n}^{\frac{K_n^{dc}(\xi)}{\xi}} y g_n(y)dy + \frac{dK_n^{dc}(\xi)}{d\xi}\left((r-c)(1-G_n(\frac{K_n^{dc}(\xi)}{\xi})) - c_n\right) \\
= &\ (w-c)\int_{\epsilon_n}^{\frac{K_n^{dc}(\xi)}{\xi}} y g_n(y)dy + \frac{dK_n^{dc}(\xi)}{d\xi}\frac{1-F_n(\xi)}{\xi f_n(\xi)}(r-w)\frac{K_n^{dc}(\xi)}{\xi}g_n(\frac{K_n^{dc}(\xi)}{\xi}) > 0,
\end{aligned}$$

where the second equality is from Part (a). This inequality concludes Part (c).

Finally, for Part (d), we first derive

$$\frac{dP_n^{dc}}{dK_n^{dc}} = \frac{dP_n^{dc}/d\xi}{dK_n^{dc}/d\xi} = (r-w)(1-G_n(\frac{K_n^{dc}(\xi)}{\xi})).$$

Next, we prove that $\frac{d}{d\xi}\left(\frac{K_n^{dc}(\xi)}{\xi}\right) \geq 0$. For notational simplicity, we define $A(\xi) = \frac{K_n^{dc}(\xi)}{\xi}$. From Part (a), we have

$$(1-G_n(A(\xi)))\left(r - c - \frac{1-F_n(\xi)}{\xi f_n(\xi)}(r-w)\frac{A(\xi)g_n(A(\xi))}{1-G_n(A(\xi))}\right) - c_n = 0.$$

By taking derivative on this equation, we can derive

$$\begin{aligned}
&\frac{dA(\xi)}{d\xi}\Big\{ - g_n(A(\xi))\left(\frac{c_n}{1-G_n(A(\xi))}\right) \\
+ &\ (1-G_n(A(\xi)))\left(-\frac{1-F_n(\xi)}{\xi f_n(\xi)}(r-w)\left[\frac{A(\xi)g_n(A(\xi))}{1-G_n(A(\xi))}\right]'\right)\Big\} \\
= &\ (1-G_n(A(\xi)))\frac{d}{d\xi}\left(\frac{1-F_n(\xi)}{\xi f_n(\xi)}\right)(r-w)\frac{A(\xi)g_n(A(\xi))}{1-G_n(A(\xi))}.
\end{aligned}$$

Because the term inside $\{.\}$ and the R.H.S. are negative, we have $\frac{dA(\xi)}{d\xi} > 0$. Therefore, we have $\frac{d}{d\xi}\left(\frac{dP_n^{dc}}{dK_n^{dc}}\right) = -(r-w)g_n(\frac{K_n^{dc}(\xi)}{\xi})\frac{d}{d\xi}(A(\xi)) \leq 0$. Finally, we have

$$\frac{d^2P_n^{dc}}{(dK_n^{dc})^2} = \frac{d(\frac{dP_n^{dc}}{dK_n^{dc}})/d\xi}{dK_n^{dc}/d\xi} \leq 0,$$

which concludes the proof of theorem. □

Proof of Theorem 3.6. To determine the structure of the optimal stopping policy, we use the two-step method that we have proposed in Chapter 2. As before, we define

$$M_n(X_n^s) \equiv E[\pi_{n+1}(X_{n+1}^s)|X_n^s] - \pi_n(X_n^s),$$

and and

$$B_n(X_n^s) \equiv E[V_{n+1}(X_{n+1}^s)|X_n^s] - \pi_n(X_n^s).$$

Then, Part (c) of Theorem 3.4 implies that

$$\begin{aligned} M_n(X_n^s) &= E[X_{n+1}^s\hat{\pi}_{n+1} - C_{n+1} - \underline{\pi}^m|X_n^s] - (X_n^s\hat{\pi}_n - C_n - \underline{\pi}^m) \\ &= X_n^s(\hat{\pi}_{n+1} - \hat{\pi}_n) - (C_{n+1} - C_n), \end{aligned} \tag{B.6}$$

which is a linear function of X_n^s. Because every linear function is convex, and the state transition is stochastically convex, a control-band policy is optimal from Proposition 2.2. □

Proof of Theorem 3.7. We first prove Part (a). From the proof of Theorem 3.6, $B_n(X_n^s)$ is convex in X_n^s. We will prove that $B_n(0) \leq 0$ if $C_n < C_{n+1}$ for every n. When $X_n^s = 0$, $X_l^s = 0$ almost surely for every $l \geq n$. Therefore, $V_n(0) = -C_n - \underline{\pi}^m$, which implies that $B_n(0) = -(C_{n+1} - C_n) \leq 0$. If a convex function satisfies $B_n(0) \leq 0$, then $B_n(X_n^s)$ can cross 0 at most once from below to above in $(0, \infty)$. Therefore, the lower threshold, L_n, is 0, and the upper threshold policy is optimal.

For Part (b), we first define $\eta_n \equiv \max_{m>n} \hat{\pi}_m$. We prove by induction that $B_n(X_n^s) = (\eta_n - \hat{\pi}_n)X_n^s$ for all n. For period $n = N - 1$, we have

$$B_n(X_n^s) = M_n(X_n^s) = X_n^s(\hat{\pi}_{n+1} - \hat{\pi}_n) - (C_{n+1} - C_n) = (\hat{\pi}_{n+1} - \hat{\pi}_n)X_n^s,$$

where $\eta_n = \hat{\pi}_{n+1}$ by definition. Next assume for an induction argument that $B_{n+1}(X) =$

$(\eta_{n+1} - \hat{\pi}_{n+1})X_n^s$. If $\eta_{n+1} \geq \hat{\pi}_{n+1}$, then $\eta_n = \eta_{n+1}$ and

$$\begin{aligned} B_n(X_n^s) &= E[\max\{0, B_{n+1}(X_{n+1}^s)\}|X_n^s] + M_n(X_n^s) \\ &= (\eta_{n+1} - \hat{\pi}_{n+1})X_n^s + M_n(X_n^s) = (\eta_{n+1} - \hat{\pi}_{n+1})X_n^s + (\hat{\pi}_{n+1} - \hat{\pi}_n)X_n^s \\ &= (\eta_{n+1} - \hat{\pi}_n)X_n^s = (\eta_n - \hat{\pi}_n)X_n^s. \end{aligned}$$

In contrast, if $\eta_{n+1} < \hat{\pi}_{n+1}$, then $\eta_n = \hat{\pi}_{n+1}$ and

$$\begin{aligned} B_n(X_n^s) &= E[\max\{0, B_{n+1}(X_{n+1}^s)\}|X_n^s] + M_n(X_n^s) \\ &= 0 + M_n(X_n^s) = (\hat{\pi}_{n+1} - \hat{\pi}_n)X_n^s = (\eta_n - \hat{\pi}_n)X_n^s, \end{aligned}$$

which concludes the induction argument.

We next prove that the optimal policy always stops at period n^*. For $n < n^*$, $\eta_n = \hat{\pi}_{n^*}$, hence $\eta_n > \hat{\pi}_n$ by the definition of n^*. In this case, $B_n(X_n^s) \geq 0$ for all X_n^s, and it is always optimal to continue the process. For $n = n^*$, we have $\eta_{n^*} \leq \hat{\pi}_n$, which implies that $B_n(X_n^s) \leq 0$. Hence, the optimal policy always stops at period n^*. □

Proof of Theorem 3.8. For Parts (a) to (c), we can apply exactly the same method as in the proofs of Theorems 3.6 and 3.7. In the problem formulation, the only two changes are the replacement of $\hat{\pi}_n$ by $\hat{\pi}_n^{cs}$ and the replacement of X_n^s by X_n^m. Because the proofs of Theorems 3.6 and 3.7 do not depend on the values of $\hat{\pi}_n$, the replacement of $\hat{\pi}_n$ do not change the proof. In addition, the only required property of X_n^s for the proofs is the Martingale property, which X_n^m also satisfies. Therefore, Parts (a) to (c) stem directly from the two theorems.

For Part (d), recall from the proof of Theorem 3.5 that $K_n^{dc}(\xi)$ maximizes $H(K,\xi) \equiv (r-c)E[\min(\epsilon_n\xi, K)] - c_nK - \frac{1-F_n(\xi)}{f_n(\xi)}(r-w)\int_{\underline{\epsilon}_n}^{\frac{K(\xi)}{\xi}} yg_n(y)dy$. The optimal K_n^{cs} satisfies the following first-order condition: $(r-c)(1-G_n(\frac{K}{\xi})) - c_n = 0$. Then, we have $\frac{\partial H(K,\xi)}{\partial K}|_{K=K_n^{cs}(\xi)} = -\frac{1-F_n(\xi)}{\xi f_n(\xi)}(r-w)\frac{K}{\xi}g_n(\frac{K}{\xi}) < 0$, which implies that $K_n^{dc}(\xi) \leq K_n^{cs}(\xi)$. □

Proof of Theorem 3.9. We first prove the optimality of the contract $w_n^{dc}(\xi) = r$, $K_n^{dc}(\xi) = X_n^s\xi G_n^{-1}(\frac{r-c-c_n}{r-c})$, and $P_n^{dc}(\xi) = -\underline{\pi}^m$ by showing that the supplier cannot attain a higher expected profit than the expected profit from the proposed contract. The expected profits of the supplier, manufacturer and the total supply chain under the contract $\{K(\xi), P(\xi), w(\xi)\}$ are

$$\begin{aligned}
\Pi_n^s(K(\check{\xi}), P(\check{\xi}), w(\check{\xi}), \xi_n, X_n^s) &\equiv (w(\check{\xi}) - c)E_{\epsilon_n}[\min(X_n^s \xi_n \epsilon_n, K(\check{\xi}))] + P(\check{\xi}) \\
&\quad -(c_n K(\check{\xi}) + C_n), \\
\Pi_n^m(K(\check{\xi}), P(\check{\xi}), w(\check{\xi}), \xi_n, X_n^s) &\equiv (r - w(\check{\xi}))E_{\epsilon_n}[\min(X_n^s \xi_n \epsilon_n, K(\check{\xi}))] - P(\check{\xi}), \\
\Pi_n^{tot}(K(\check{\xi}), \xi_n, X_n^s) &\equiv (r - c)E_{\epsilon_n}[\min(X_n^s \xi_n \epsilon_n, K(\check{\xi}))] - (c_n K(\check{\xi}) + C_n),
\end{aligned}$$

where $(K(\check{\xi}), P(\check{\xi}), w(\check{\xi}))$ is the manufacturer's choice of the contract.

Note that $\Pi_n^m(K(\check{\xi}), P(\check{\xi}), w(\check{\xi}), \xi_n, X_n^s) \geq \underline{\pi}^m$ from the participation constraint. We also have that $\Pi_n^{tot}(K(\check{\xi}), \xi_n, X_n^s) \leq \pi_n^{cs}(X_n^s \xi_n)$, because $\pi_n^{cs}(X_n^s \xi_n)$ is the optimal expected profit of the centralized supply chain. Hence, the supplier's expected profit under any contract is bounded above as

$$\begin{aligned}
\Pi_n^s(K(\check{\xi}), P(\check{\xi}), w(\check{\xi}), \xi_n, X_n^s) &= \Pi_n^{tot}(K(\check{\xi}), \xi_n, X_n^s) - \Pi_n^m(K(\check{\xi}), P(\check{\xi}), w(\check{\xi}), \xi_n, X_n^s) \\
&\leq \pi_n^{cs}(X_n^s \xi_n) - \underline{\pi}^m.
\end{aligned}$$

We show that the proposed contract achieves this upper bound for every ξ_n.

We first fix $w(\xi) = r$ and $P(\xi) = -\underline{\pi}^m$. Then, the supplier's profit is

$$\Pi_n^s(K(\check{\xi}), P(\check{\xi}), w(\check{\xi}), \xi_n, X_n^s) = \Pi_n^{tot}(K(\check{\xi}), \xi_n, X_n^s) - \underline{\pi}^m,$$

and the manufacturer's profit is $\Pi_n^m(K(\check{\xi}), P(\check{\xi}), w(\check{\xi}), \xi_n, X_n^s) = \underline{\pi}^m$. In this case, the manufacturer always attains her reservation profit $\underline{\pi}^m$ regardless of her report $\check{\xi}$, hence she has no incentive to inflate her demand forecast, i.e., the supplier would report her forecast truthfully. Then, the supplier determines the capacity level to maximize $\Pi_n^{tot}(K(\xi_n), \xi_n, X_n^s) - \underline{\pi}^m$, which is the same as the objective function of the centralized supply chain minus $\underline{\pi}^m$. Hence, the optimal $K(\xi)$ when we fix $w(\xi) = r$ and $P(\xi) = -\underline{\pi}^m$ is $K(\xi) = K_n^{cs}(X_n^s \xi) = X_n^s \xi G_n^{-1}(\frac{r-c-c_n}{r-c})$, which implies that $\Pi_n^s(K(\xi), P(\xi), w(\xi), \xi, X_n^s) = \pi_n^{cs}(X_n^s \xi) - \underline{\pi}^m$, which is the same as the upper bound.

Therefore, the proposed contract is the optimal contract, and the supplier's optimal expected profit at period n is give as

$$\pi_n^s(X_n^s) = E_{\xi_n}[\pi_n^{cs}(X_n^s \xi_n) - \underline{\pi}^m] = X_n^s \hat{\pi}_n^{cs} - C_n - \underline{\pi}^m.$$

□

Proof of Theorem B.1. We define $\hat{K}(.) = K(.) - X_n^s$. By construction, $\hat{K}(.)$ is increasing if and only if $K(.)$ is increasing.Using $\hat{K}(.)$, we can derive (B.1) as

$$\begin{aligned}
& (r-c)E[\min(X_n^s + \xi_n + \epsilon_n, K(\xi_n))] - (c_n K(\xi_n) + C_n) \\
& -\frac{1-F_n(\xi_n)}{f_n(\xi_n)}(r-w)G_n(K(\xi_n) - \xi_n - X_n^s) \\
= \; & (r-c)E[\min(\xi_n + \epsilon_n, \hat{K}(\xi_n))] - c_n\hat{K}(\xi_n) - \frac{1-F_n(\xi_n)}{f_n(\xi_n)}(r-w)G_n(\hat{K}(\xi_n) - \xi_n) \\
& +(r-c-c_n)X_n^s - C_n \\
= \; & \left[(r-c)E[\min(\xi_n + \epsilon_n, \hat{K}(\xi_n))] - c_n\hat{K}(\xi_n) - \frac{1-F_n(\xi_n)}{f_n(\xi_n)}(r-w)G_n(\hat{K}(\xi_n) - \xi_n)\right] \\
& +(r-c-c_n)X_n^s - C_n, \qquad \text{(B.7)}
\end{aligned}$$

where the term inside of [.] is identical to the objective function of (B.3), which instantly implies Part (a) and (c). Finally, by applying $\hat{K}(.)$ to (B.2), we can prove that Part (b) also holds. □

Proof of Theorem B.2. We first define a function $H(K,\xi) \equiv (r-c)E[\min(\epsilon_n+\xi_n), K)] - c_n K - \frac{1-F_n(\xi)}{f_n(\xi)}(r-w)G_n(K-\xi)$. For Part (a), we prove that $H(K,\xi)$ is quasi-concave in K and have a finite maximizer. First, we have

$$\begin{aligned}
\frac{\partial H(K,\xi)}{\partial K} &= (r-c)(1-G_n(K-\xi)) - c_n - \frac{1-F_n(\xi)}{f_n(\xi)}(r-w)g_n(K-\xi) \\
&= (1-G_n(K-\xi))\left(r - c - \frac{1-F_n(\xi)}{f_n(\xi)}(r-w)\frac{g_n(K-\xi)}{1-G_n(K-\xi)}\right) - c_n,
\end{aligned}$$

from which we can derive

$$\left.\frac{\partial^2 H(K,\xi)}{\partial K^2}\right|_{\frac{\partial H(K,\xi)}{\partial K}=0} = -g_n(K-\xi)\left(\frac{c_n}{1-G_n(K-\xi)}\right)$$
$$+ \quad (1-G_n(K-\xi))\left(-\frac{1-F_n(\xi)}{f_n(\xi)}(r-w)\frac{d}{dK}\left(\frac{g_n(K-\xi)}{1-G_n(K-\xi)}\right)\right) < 0.$$

The inequality is from the fact that $\frac{g_n(K-\xi)}{1-G_n(K-\xi)}$ is increasing in K from the IFR assumption. The inequality is strict because $g_n(K-\xi)\left(\frac{c_n}{1-G_n(K-\xi)}\right)$ is strictly positive. In other words, $H(K,\xi)$ is quasi-concave where the slope when it crosses zero is strictly negative. Finally, we have $\frac{\partial H(K,\xi)}{\partial K}|_{K<\underline{\xi}_n} = r-c-c_n > 0$, and $\lim_{K\to\infty}\frac{\partial H(K,\xi)}{\partial K} = -c_n < 0$. Therefore, there exists a finite solution of $\frac{\partial H(K,\xi)}{\partial K} = 0$. In addition, because $\frac{\partial H(K,\xi)}{\partial K}$ is strictly decreasing at $\frac{\partial H(K,\xi)}{\partial K} = 0$, there is only one solution that satisfies the first-order condition.

Next we prove that the function $K(.)$ that satisfies the first-order conditions is increasing in ξ. We take second-order derivative on H, which gives

$$\left.\frac{\partial^2 H(K,\xi)}{\partial K\partial\xi}\right|_{\frac{\partial H(K,\xi)}{\partial K}=0} = g_n(K-\xi)\left(\frac{c_n}{1-G_n(K-\xi)}\right)$$
$$+ \quad (1-G_n(K-\xi))\left(-\frac{d}{d\xi}\frac{1-F_n(\xi)}{f_n(\xi)}(r-w)\left(\frac{g_n(K-\xi)}{1-G_n(K-\xi)}\right)\right)$$
$$+ \quad (1-G_n(K-\xi))\left(-\frac{1-F_n(\xi)}{f_n(\xi)}(r-w)\frac{d}{d\xi}\left(\frac{g_n(K-\xi)}{1-G_n(K-\xi)}\right)\right) > 0.$$

The inequality is from the fact that both $\frac{1-F_n(\xi)}{f_n(\xi)}$ and $\frac{g_n(K-\xi)}{1-G_n(K-\xi)}$ are decreasing in ξ from the IFR assumption. Therefore, the optimal $K(.)$ of the unconstraint problem is increasing in ξ, hence it is the optimal $K_n^{dc}(.)$. Therefore, Part (a) and the increasing property of K_n^{dc} in Part (c) are true. For the increasing property of $P_n^{dc}(\xi)$ in Part (c), we take the derivative on the equation (B.2), which gives

$$\frac{dP_n^{dc}(\xi)}{d\xi} = (r-w)(1-G_n(K_n^{dc}(\xi)-\xi))\frac{dK_n^{dc}(\xi)}{d\xi} > 0.$$

The increasing property of $\Pi_n^m(K_n^{dc}(\xi), P_n^{dc}(\xi), \xi, X_n^s)$ in Part (d) is directly from Part

(a) of Lemma 1 of Özer and Wei (2006). The rest of the proof is the same as the proof of Theorem 1 of Özer and Wei (2006). □

Proof of Theorem B.3. In the additive case, we have $M_n(X_n^s) = X_n^s(c_n - c_{n+1}) - (C_{n+1} - C_n) + (\hat{\pi}_{n+1} - \hat{\pi}_n)$, which is also a linear function of X_n^s. Because every linear function is convex, Part (a) holds from Proposition 2.2.

For Part (b), first note that $M_n(X_n^s)$ is decreasing in X_n^s if $c_n < c_{n+1}$ for every n. Then, Proposition 2.1 implies Part (b).

For Part (c), we first define $\eta_n \equiv \max_{m>n} \hat{\pi}_m - C_m$. We prove by induction that $B_n(X_n^s) = \eta_n - \hat{\pi}_n + C_n$ for all n. For period $n = N-1$, we have

$$B_n(X_n^s) = M_n(X_n^s) = (\hat{\pi}_{n+1} - C_{n+1}) - \hat{\pi}_n + C_n,$$

where $\eta_n = \hat{\pi}_{n+1} - C_{n+1}$ by definition. Next assume for an induction argument that $B_{n+1}(X) = \eta_{n+1} - \hat{\pi}_{n+1} + C_{n+1}$. If $\eta_{n+1} \geq \hat{\pi}_{n+1} - C_{n+1}$, then $\eta_n = \eta_{n+1}$ and

$$\begin{aligned} B_n(X_n^s) &= E[\max\{0, B_{n+1}(X_{n+1}^s)\}|X_n^s] + M_n(X_n^s) \\ &= \eta_{n+1} - \hat{\pi}_{n+1} + C_{n+1} + M_n(X_n^s) \\ &= \eta_{n+1} - \hat{\pi}_{n+1} + C_{n+1} + \hat{\pi}_{n+1} - C_{n+1} - \hat{\pi}_n + C_n \\ &= \eta_{n+1} - \hat{\pi}_n + C_n = \eta_n - \hat{\pi}_n + C_n. \end{aligned}$$

In contrast, if $\eta_{n+1} < \hat{\pi}_{n+1} - C_{n+1}$, then $\eta_n = \hat{\pi}_{n+1} - C_{n+1}$ and

$$\begin{aligned} B_n(X_n^s) &= E[\max\{0, B_{n+1}(X_{n+1}^s)\}|X_n^s] + M_n(X_n^s) \\ &= 0 + M_n(X_n^s) = \hat{\pi}_{n+1} - C_{n+1} - \hat{\pi}_n + C_n = \eta_n - \hat{\pi}_n + C_n, \end{aligned}$$

which concludes the induction argument.

We next prove that the optimal policy always stops at period n^*. For $n < n^*$, $\eta_n = \hat{\pi}_{n^*} - C_{n^*}$, hence $\eta_n > \hat{\pi}_n - C_n$ by the definition of n^*. In this case, $B_n(X_n^s) > 0$ for all X_n^s, and it is always optimal to continue the process. For $n = n^*$, we have $\eta_{n^*} \leq \hat{\pi}_n - C_n$, which implies that $B_n(X_n^s) \leq 0$. Hence, the optimal policy always stops at period n^*. □

Appendix C

Chapter 4 Appendices

C.1. Notation

Decision Variables	**State Variables and Updates**
u_t : stopping decision	s_t : market potential
i_t : investment decision	$w_t(s_t)$: reduction in market potential
Q : production quantity	$s_{t+1} = s_t - w_t(s_t)$
z : stocking factor	x_t : knowledge level
p_r : regular sales price	$k_t(i_t)$: knowledge improvement
p_s : salvage price	$x_{t+1} = x_t + k_t(i_t)$
Cost and Demand Parameters	
$c_i(i_t)$: cost of investment option i_t	
$c_p(x_t) = \delta_0 + \delta_1 e^{-\gamma x_t}$: unit production cost	
b : price elasticity of demand	
$D_n(s_t, p_n) = s_t A_n p_n^{-b}$: demand during the sales period $n \in \{r, s\}$	
$[\underline{A}_n, \overline{A}_n]$: support of A_n	
Profit Functions	
$J_r(s_t, Q)$: revenue-to-go function of the regular sales period	
$J_s(s_t, Q_s)$: revenue-to-go function of the salvage period	
$\Pi_t(s_t, x_t)$: expected profit of stopping at period t	
$V_t(s_t, x_t)$: value-to-go function of period t	
$M_t(s_t, x_t, i_t)$: one-step benefit function	
$B_t(s_t, x_t, i_t)$: benefit function	
$\overline{B}_t(s_t, x_t)$: maximal benefit function	

C.2. Proofs

Proof of Theorem 4.1. The first part is from Proposition 2 of Monahan et al. (2004). Next, because z_s is defined as $\frac{Q_s}{s_t p_s^{-b}}$, we have $p_s^*(s_t, Q_s) = \left(\frac{s_t z_s^*}{Q_s}\right)^{1/b}$. Finally, by the definition of J_s^* and Equation (4.7), we have $J_s(s_t, Q_s) = s_t \left(\frac{Q_s}{s_t}\right)^{1-1/b} J_s^*$. □

Proof of Theorem 4.2. We first prove that $z_r^* \geq \underline{A}_r$. When $z \leq \underline{A}_r$, the objective function is given as $f(z) = z^{1/b}$, because $z \leq \underline{A}_r \leq A_r$. In this case, $f(z)$ is strictly increasing in z, which implies that $z_r^* \geq \underline{A}_r$.

We next prove that $f(z)$ is quasi-concave in z for $z > \overline{A}_r$. When $z > \overline{A}_r$, the objective function $f(z)$ is given as

$$f(z) = z^{-1+1/b}\left(E[A_r] + \beta J_s^* E[(z - A_r)^{1-1/b}]\right).$$

We can derive the first-order derivative of $f(z)$ for $z > \overline{A}_r$ as

$$f'(z) = \left(1 - \frac{1}{b}\right) z^{-2+1/b}\left(-E[A_r] + \beta J_s^* E[A_r(z - A_r)^{-1/b}]\right),$$

and the second-order derivative of $f(z)$ at the point that satisfies that $f'(z) = 0$ as

$$f''(z)|_{f'(z)=0} = -\frac{1}{b}\left(1 - \frac{1}{b}\right) z^{-2+1/b} \beta J_s^* E[A_r(z - A_r)^{-1-1/b}],$$

which is strictly smaller than 0. Therefore, $f(z)$ is quasi-concave in z for $z > \overline{A}_r$, and the curvature of $f(z)$ at the mode is strictly concave. Hence, the optimal z_r^* is either the unique solution that satisfies $\beta J_s^* E[A_r(z - A_r)^{-1/b}] = E[A_r]$ or the maximizer of $f(z)$ in $[\underline{A}_r, \overline{A}_r]$.

By embedding $J_r(s_t, Q) = s_t \left(\frac{Q}{s_t}\right)^{1-1/b} J_r^*$ in (4.6), we can derive

$$\Pi_t(s_t, x_t) = s_t \left(\frac{Q}{s_t}\right)^{1-1/b} J_r^* - c_p(x_t) Q, \tag{C.1}$$

from which we can determine the optimal production quantity as $Q^*(s_t, x_t) = s_t \left(\frac{(b-1)J_r^*}{b c_p(x_t)}\right)^b$.

Then, because $z_r = \frac{Q}{s_t p_r^{-b}}$, the optimal regular sales price is given as $p_r^*(s_t, x_t) = \frac{(b-1)J_r^*}{bc_p(x_t)(z_r^*)^{1/b}}$. Finally, by embedding $Q^*(s_t, x_t)$ in (C.1), we can derive $\Pi_t(s_t, x_t) = s_t c_p(x_t)^{1-b}\pi^*$. □

Proof of Theorem 4.3. We prove that $\overline{B}_t(s_t, x_t)$ is decreasing in x_t at periods $t \geq \hat{t}$. To do so, we first show that $M_t(s_t, x_t, i_t)$ is decreasing in x_t at periods $t \geq \hat{t}$. We can derive $M_t(s_t, x_t, i_t)$ as

$$\begin{aligned} M_t(s_t, x_t, i_t) &= -c_i(i_t) + \alpha E[\Pi_{t+1}(s_{t+1} - w_t(s_t), x_t + k_t(i_t))] - \Pi_t(s_t, x_t) \\ &= -c_i(i_t) + \alpha E[s_t - w_t(s_t)]E[(c_p(x_t + k_t(i_t)))^{1-b}]\pi^* - s_t c_p(x_t)^{1-b}\pi^* \\ &= -c_i(i_t) + \{\alpha E[(s_t - w_t(s_t))] - s_t\} E[(c_p(x_t + k_t(i_t)))^{1-b}]\pi^* \qquad \text{(C.2)} \\ &\quad + s_t \left\{E[c_p(x_t + k_t(i_t))^{1-b}] - c_p(x_t)^{1-b}\right\} \pi^*. \end{aligned}$$

Because $\alpha E[(s_t - w_t(s_t))] \leq s_t$ and $c_p(x_t + k_t(i_t))$ is increasing in x_t for every realization of $k_t(i_t)$, the second term of (C.2) is decreasing in x_t. At periods $t \geq \hat{t}$, the function $c_p(x_t)^{1-b}$ is concave in x_t, which implies that $E[c_p(x_t + k_t(i_t))^{1-b}] - c_p(x_t)^{1-b}$ is decreasing in x_t. Hence, $M_t(s_t, x_t, i_t)$ is decreasing in x_t for every (s_t, i_t).

Using the decreasing property of $M_t(s_t, x_t, i_t)$, we next show that $\overline{B}_t(s_t, x_t)$ is decreasing in x_t for periods $t \geq \hat{t}$. This proof is based on an induction argument. The benefit function and the one-step benefit function satisfy the following recursion:

$$\begin{aligned} & B_T(s_T, x_T, i_T) = M_T(s_T, x_T, i_T) \\ & B_t(s_t, x_t, i_t) = -c_i(i_t) + \alpha E[V_{t+1}(s_{t+1} - w_t(s_t), x_t + k_t(i_t))] - \Pi_t(s_t, x_t) \\ & = \alpha E\left[\max\left\{0, \overline{B}_{t+1}(s_{t+1} - w_t(s_t), x_t + k_t(i_t))\right\} + \Pi_{t+1}(s_{t+1} - w_t(s_t), x_t + k_t(i_t))\right] \\ & -\Pi_t(s_t, x_t) - c_i(i_t) \\ & = M_t(s_t, x_t, i_t) + \alpha E\left[\max\left\{0, \overline{B}_{t+1}(s_{t+1} - w_t(s_t), x_t + k_t(i_t))\right\}\right], \text{ for } t < T. \qquad \text{(C.3)} \end{aligned}$$

Hence, at period $t = T$, we have $B_T(s_T, x_T, i_T) = M_T(s_T, x_T, i_T)$, which is decreasing in x_T. Then, for any $x^1 \leq x^2$, we have $\overline{B}_T(s_T, x^2) = B_T(s_T, x^2, i_T^*(s_T, x^2)) \leq B_T(s_T, x^1, i_T^*(s_T, x^2)) \leq B_T(s_T, x^1, i_T^*(s_T, x^1)) = \overline{B}_T(s_T, x^1)$, where the first inequality is from the fact that $B_T(s_T, x_T, i_T)$ is decreasing in x_T, and the second inequality is by the definition of $i_T^*(s_T, x_T)$. This result implies that $\overline{B}_T(s_T, x_T)$ is decreasing in x_T.

Next we assume for the induction argument that $\overline{B}_{t+1}(s_{t+1}, x_{t+1})$ is decreasing in x_{t+1}. Because the composition of a decreasing function and $\max\{0, x\}$ is also decreasing, $\max\{0, \overline{B}_{t+1}(s_{t+1}, x)\}$ is a decreasing function of x. Then, the independence between $k_t(i_t)$ and x_t implies that $E[\max\{0, \overline{B}_{t+1}(s_t - w_t(s_t), x_t + k_t(i_t))\}]$ is decreasing in x_t. Finally, the benefit function $B_t(s_t, x_t, i_t)$ is also decreasing in x_t, because $M_t(s_t, x_t, i_t)$ is decreasing in x_t for periods $t \geq \hat{t}$. If $B_t(s_t, x_t, i_t)$ is decreasing in x_t, then so is $\overline{B}_t(s_t, x_t) = \sup_{i_t \in I_t} B_t(s_t, x_t, i_t)$ based on the same argument that we have applied to $\overline{B}_T(s_T, x_T)$. Hence, $\overline{B}_t(s_t, x_t)$ is decreasing in x_t for periods $t \geq \hat{t}$.

When $\overline{B}_t(s_t, x_t)$ is decreasing in x_t, $\overline{B}_t(s_t, x_t) \leq 0$ if and only if $x_t \geq \inf\{x : \overline{B}_t(s_t, x) \leq 0\} = \underline{x}_t(s_t)$ for each given s_t. Hence, at periods $t \geq \hat{t}$, it is optimal to stop process design if $x_t \geq \underline{x}_t(s_t)$, and it is optimal otherwise to continue process design. □

Proof of Theorem 4.4. We first prove that if

$$\frac{dE[w_t(s_t)]}{ds_t} \leq 1 - \min_{i_t \in I_t} \frac{c_p(x_t)^{1-b}}{\alpha E[c_p(x_t + k_t(i_t))^{1-b}]},$$

then $M_t(s_t, x_t, i_t)$ is increasing in s_t for every (x_t, i_t). Recall that $M_t(s_t, x_t, i_t)$ is given as

$$M_t(s_t, x_t, i_t) \quad = \quad -c_i(i_t) + \alpha E[(s_t - w_t(s_t))] E[(c_p(x_t + k_t(i_t)))^{1-b}]\pi^* - s_t c_p(x_t)^{1-b}\pi^*,$$

from which we can derive

$$\frac{\partial M_t(s_t, x_t, i_t)}{\partial s_t} \quad = \quad \alpha(1 - \frac{dE[w_t(s_t)]}{ds_t}) E[(c_p(x_t + k_t(i_t)))^{1-b}]\pi^* - c_p(x_t)^{1-b}\pi^*.$$

Hence, when the sufficient condition (4.11) holds, $M_t(s_t, x_t, i_t)$ is increasing in s_t for every (x_t, i_t).

Using the increasing property of $M_t(s_t, x_t, i_t)$ in s_t, we next show that $\overline{B}_t(s_t, x_t)$ is increasing in s_t for all periods. The proof is based on an induction argument.

At period $t = T$, we have $B_T(s_T, x_T, i_T) = M_T(s_T, x_T, i_T)$, which is increasing in s_T. Then, for any $s^1 \leq s^2$, we have $\overline{B}_T(s^1, x_T) = B_T(s^1, x_T, i^*_T(s^1, x_T)) \leq B_T(s^2, x_T, i^*_T(s^1, x_T)) \leq B_T(s^2, x_T, i^*_T(s^2, x_T)) = \overline{B}_T(s^2, x_T)$, where the first inequality is from the fact that $B_T(s_T, x_T, i_T)$ is increasing in s_T, and the second inequality is by the definition of $i^*_T(s_T, x_T)$. This result implies that $\overline{B}_T(s_T, x_T)$ is increasing in s_T.

Next assume for the induction argument that $\overline{B}_{t+1}(s_{t+1}, x_{t+1})$ is increasing in s_{t+1}. Because the composition of an increasing function and $\max\{0, x\}$ is also increasing, $\max\{0, \overline{B}_{t+1}(s, x_{t+1})\}$ is an increasing function of s. Then, the stochastic increasing property of $s_{t+1} = s_t - w_t(s_t)$ in s_t implies that $E[\max\{0, \overline{B}_{t+1}(s_t - w_t(s_t), x_t + k_t(i_t), i_t)\}]$ is increasing in s_t. Finally, the benefit function $B_t(s_t, x_t, i_t) = M_t(s_t, x_t, i_t) + \alpha E\left[\max\{0, \overline{B}_{t+1}(s_t - w_t(s_t), x_t + k_t(i_t), i_t)\}\right]$ is increasing in s_t, because $M_t(s_t, x_t, i_t)$ is also increasing in s_t. If $B_t(s_t, x_t, i_t)$ is increasing in s_t, then so is $\overline{B}_t(s_t, x_t)$ based on the same argument that we have applied to $\overline{B}_T(s_T, x_T)$. Hence, $\overline{B}_t(s_t, x_t)$ is increasing in s_t for all periods.

When $\overline{B}_t(s_t, x_t)$ is increasing in s_t, $\overline{B}_t(s_t, x_t) \leq 0$ if and only if $s_t \leq \sup\{s : \overline{B}_t(s, x_t) \leq 0\} = \overline{s}_t(x_t)$ for each given x_t. Hence, at each period t, it is optimal to stop process design if $s_t \leq \overline{s}_t(x_t)$, and it is optimal otherwise to continue process design. □

Proof of Theorem 4.5. The proof of this theorem is similar to that of Theorem 4.4. We first prove that if

$$\frac{dE[w_t(s_t)]}{ds_t} \geq 1 - \max_{i_t \in I_t} \frac{c_p(x_t)^{1-b}}{\alpha E[c_p(x_t + k_t(i_t))^{1-b}]},$$

then $M_t(s_t, x_t, i_t)$ is increasing in s_t for every (x_t, i_t). Recall from the proof of Theorem 4.4 that

$$\frac{\partial M_t(s_t, x_t, i_t)}{\partial s_t} = \alpha(1 - \frac{dE[w_t(s_t)]}{ds_t})E[(c_p(x_t + k_t(i_t)))^{1-b}]\pi^* - c_p(x_t)^{1-b}\pi^*.$$

Hence, when the sufficient condition (4.12) holds, $M_t(s_t, x_t, i_t)$ is increasing in s_t for every (x_t, i_t).

Using the decreasing property of $M_t(s_t, x_t, i_t)$ in s_t, we next show that $\overline{B}_t(s_t, x_t)$ is decreasing in s_t for all periods. The proof is based on an induction argument. At period $t = T$, we have $B_T(s_T, x_T, i_T) = M_T(s_T, x_T, i_T)$, which is decreasing in s_T. Then, for any $s^1 \leq s^2$, we have $\overline{B}_T(s^2, x_T) = B_T(s^2, x_T, i_T^*(s^2, x_T)) \leq B_T(s^1, x_T, i_T^*(s^2, x_T)) \leq B_T(s^1, x_T, i_T^*(s^1, x_T)) = \overline{B}_T(s^1, x_T)$, where the first inequality is from the fact that $B_T(s_T, x_T, i_T)$ is decreasing in s_T, and the second inequality is by the definition of $i_T^*(s_T, x_T)$. This result implies that $\overline{B}_T(s_T, x_T)$ is decreasing in s_T.

Next assume for the induction argument that $\overline{B}_{t+1}(s_{t+1}, x_{t+1})$ is decreasing in s_{t+1}. Because the composition of a decreasing function and $\max\{0, x\}$ is also decreasing, $\max\{0, \overline{B}_{t+1}(s, x_{t+1})\}$ is a decreasing function of s. Then, the stochastic increasing property of $s_{t+1} = s_t - w_t(s_t)$ in s_t implies that $E[\max\{0, \overline{B}_{t+1}(s_t - w_t(s_t), x_t + k_t(i_t), i_t)\}]$ is decreasing in s_t. Finally, the benefit function $B_t(s_t, x_t, i_t) = M_t(s_t, x_t, i_t) + \alpha E\left[\max\{0, \overline{B}_{t+1}(s_t - w_t(s_t), x_t + k_t(i_t), i_t)\}\right]$ is decreasing in s_t, because $M_t(s_t, x_t, i_t)$ is also decreasing in s_t. If $B_t(s_t, x_t, i_t)$ is decreasing in s_t, then so is $\overline{B}_t(s_t, x_t)$ based on the same argument that we have applied to $\overline{B}_T(s_T, x_T)$. Hence, $\overline{B}_t(s_t, x_t)$ is decreasing in s_t for all periods.

When $\overline{B}_t(s_t, x_t)$ is decreasing in s_t, $\overline{B}_t(s_t, x_t) \leq 0$ if and only if $s_t \geq \inf\{s : \overline{B}_t(s, x_t) \leq 0\} = \underline{s}_t(x_t)$ for each given x_t. Hence, at each period t, it is optimal to stop process design if $s_t \geq \underline{s}_t(x_t)$, and it is optimal otherwise to continue process design. □

Proof of Theorem 4.6. The proof is directly from Proposition 2.6 □

Bibliography

Aitchison, J., J. A. C. Brown. 1957. *The Lognormal Distribution*. Cambridge University Press, Cambridge, UK.

Akan, M., B. Ata, J. Dana. 2009. Revenue management by sequential screening. Working Paper. Northwestern University, Evanston, IL.

Alagoz, O., L. M. Maillart, A. J. Schaefer, M. S. Roberts. 2004. The optimal timing of living-donor liver transplantation. *Management Sci.* **50**(10) 1420–1430.

Alagoz, O., L. M. Maillart, A. J. Schaefer, M. S. Roberts. 2007a. Choosing among living-donor and cadaveric livers. *Management Sci.* **53**(11) 1702–1715.

Alagoz, O., L. M. Maillart, A. J. Schaefer, M. S. Roberts. 2007b. Determining the acceptance of cadaveric livers using an implicit model of the waiting list. *Oper. Res.* **55**(1) 24–36.

Altman, E., S. Stidham. 1995. Optimality of monotonic policies for two-action markovian decision processes. *Queueing Systems* **21** 267–291.

Altug, M. S., A. Muharremoglu. 2009. Supply chain management with advance supply information. Working Paper. Columbia University, New York, NY.

Anily, S., A. Grosfeld-Nir. 2006. An optimal lot-sizing and offline inspection policy in the case of nonrigid demand. *Oper. Res.* **54**(2) 311–323.

Athey, S. 2000. Characterizing properties of stochastic objective functions. Working Paper. MIT, Boston, MA.

Aviv, Y. 2001. The effect of collaborative forecasting on supply chain performance. *Management Sci.* **47**(10) 1326–1343.

Aviv, Y. 2002. Gaining benefits from joint forecasting and replenishment processes: The case of auto-correlated demand. *Manufacturing Service Oper. Management* **4**(1) 55–74.

Aviv, Y. 2007. On the benefits of collaborative forecasting partnerships between retailers and manufacturers. *Management Sci.* **53**(5) 777–794.

Babich, V., M. J. Sobel. 2004. Pre-ipo operational and financial decisions. *Management Sci.* **50**(7) 935–948.

Bayus, B. L. 1997. Speed-to-market and new product performance trade-offs. *J. of Product Innovation Management* **14**(6) 485–497.

Bayus, B. L., S. Jain, A. G. Rao. 1997. Too little, too early: Introduction timing and new product performance in the personal digital assistant industry. *J. of Marketing Res.* **34**(1) 50–63.

Ben-Ameur, H., M. Breton, P. L'Ecuyer. 2002. A dynamic programming procedure for pricing american-style asian options. *Management Sci.* **48**(5) 625–643.

Bertsekas, D. P. 2005. *Dynamic Programming and Optimal Control, Vol. 1*. Athena Scientific, Belmont, MA.

Bertsekas, D. P. 2007. *Dynamic Programming and Optimal Control, Vol. 2*. Athena Scientific, Belmont, MA.

Black, F., M. Scholes. 1973. The pricing of options and corporate liabilities. *J. Political Econom.* **3** 637–654.

Bohn, R. E., C. Terwiesch. 1999. The economics of yield-driven processes. *J. Oper. Management* **18** 41–59.

Boyaci, T., Ö. Özer. 2009. Information acquisition via pricing and advance selling for capacity planning: When to stop and act? *Oper. Res.* Forthcoming.

Cachon, G. P., M. A. Lariviere. 2001. Contracting to assure supply: How to share demand forecasts in a supply chain. *Management Sci.* **47**(5) 629–646.

Carrillo, J. E., R. M. Franza. 2006. Investing in product development and production capabilities: The crucial linkage between time-to-market and ramp-up time. *Eur. J. Oper. Res.* **171** 536–556.

Carrillo, J. E., C. Gaimon. 2000. Improving manufacturing performance through process change and knowledge creation. *Management Sci.* **46**(2) 265–288.

Chen, A. H. Y. 1970. A model of warrant pricing in a dynamic market. *J. Finance* **25**(5) 1041–1059.

Chen, F. 1999. Decentralized supply chains subject to information delays. *Management Sci.* **45**(8) 1076–1090.

Chen, F. 2003. Information sharing and supply chain coordination. S. Graves, T. de Kok, eds., *Handbook of Operations Research and Management Science: Supply Chain Management*. Elsevier, Amsterdam, The Netherlands, 341–421.

Chen, J., D. D. Yao, S. Zheng. 1998. Quality control for products supplied with warranty. *Oper. Res.* **46**(1) 107–115.

Chen, L., H. L. Lee. 2009. Information sharing and order variability control under a generalized demand model. *Management Sci.* **55**(5) 781–797.

Chod, J., N. Rudi. 2006. Strategic investments, trading, and pricing under forecast updating. *Management Sci.* **52**(12) 1913–1929.

Chow, Y. S., S. Moriguti, H. Robbins, S. M. Samuels. 1964. Optimal selection based on relative rank. *Israel Journal of Mathematics* **2**(2) 81–90.

Chow, Y. S., H. Robbins, D. Siegmund. 1971. *Great Expectations: the Theory of Optimal Stopping*. Houghton Mifflin Company, Boston, MA.

Clark, N. 2007. U.P.S. walks away from A380 order. *The New York Times* (March 3).

Cohen, M. A., J. Eliashberg, T. Ho. 1996. New product development: the performance and time-to-market tradeoff. *Management Sci.* **42**(2) 173–186.

Corbett, C. J., X. de Groote. 2000. A supplier's optimal quantity discount policy under asymmetric information. *Management Sci.* **46**(3) 444–450.

Courty, P., H. Li. 2000. Sequential screening. *Review of Economic Studies* **67** 697–717.

Cox, J. C., S. A. Ross, M. Rubinstein. 1979. Option pricing: A simplified approach. *J. Financial Econom.* **7** 229–263.

Durrett, R. 1996. *Probability: Theory and Examples*. Duxbury Press, Belmont, CA.

Files, J. 2001. Economic downturn leaves Cisco with stacks of excess inventory. *San Jose Mercury News* (April 27).

Fine, C. H., E. L. Porteus. 1989. Dynamic process improvement. *Management Sci.* **37**(4) 580–591.

Freeman, P. R. 1983. The secretary problem and its extensions: A review. *Internat. Statistical Rev.* **51**(2) 189–206.

Gallego, G., Ö. Özer. 2001. Integrating replenishment decisions with advance demand information. *Management Sci.* **47**(10) 1344–1360.

Glasserman, P., D. D. Yao. 1994. *Monotone Structure in Discrete Event Systems*. Wiley, New York, NY.

Graves, S. C., D. B. Kletter, W. B. Hetzel. 1998. A dynamic model for requirements planning with application to supply chain optimization. *Oper. Res.* **46**(3) 35–49.

Graves, S. C., H. C. Meal, S. Dasu, Y. Qiu. 1986. Two-stage production planning in a dynamic environment. S. Axäter, C. Schneeweiss, E. Silver, eds., *Multi-Stage Production Planning and Control. Lecture Notes in Economics and Mathematical Systems, 26*. Springer-Verlag, Berlin, Germany, 9–43.

Ha, A. 2001. Supplier-buyer contracting: Asymmetric cost information and cutoff level policy for buyer participation. *Naval Res. Logist.* **48** 41–64.

Harrison, J. M., D. M. Kreps. 1979. Martingales and arbitrage in multiperiod securities markets. *J. Econom. Theory* **20** 381–408.

Hausman, W. H. 1969. Sequential decision problems: A model to exploit existing forecasters. *Management Sci.* **16**(2) B93–B111.

Hausman, W. H., R. Peterson. 1972. Multiproduct production scheduling for style goods with limited capacity, forecast revisions and terminal delivery. *Management Sci.* **18**(7) 370–383.

Heath, D. C., P. L. Jackson. 1994. Modeling the evolution of demand forecasts with application to safety stock analysis in production/distribution systems. *IIE Trans.* **26**(3) 17–30.

Hui, S. K., J. Eliashberg, E. I. George. 2008. Modeling dvd preorder and sales: An optimal stopping approach. *Oper. Res.* **27**(6) 1097–1110.

Iida, T., P. H. Zipkin. 2006. Approximate solutions of a dynamic forecast-inventory model. *Manufacturing Service Oper. Management* **8**(4) 407–425.

Iida, T., P. H. Zipkin. 2009. Competition and cooperation in a two-stage supply chain with demand forecasts. *Oper. Res.* Forthcoming.

Kalish, S, G. L. Lilien. 1986. A market entry timing model for new techniques. *Management Sci.* **32**(2) 194–205.

Kalyanaram, G., V. Krishnan. 1997. Deliberate product definition: Customizing the product definition process. *J. of Marketing Res.* **34**(2) 276–285.

Klastorin, T., W. Tsai. 2004. New product introduction: Timing, design, and pricing. *Manufacturing Service Oper. Management* **6**(4) 302–320.

Krishnan, V., K. T. Ulrich. 2001. Product development decisions: A review of the literature. *Management Sci.* **47**(1) 1–21.

Kumar, S., T. R. McCaffrey. 2003. Engineering economics at a hard disk drive manufacturer. *Technovation* **23** 749–755.

Laffont, J., J. Tirole. 1988. The dynamics of incentive contracts. *Econometrica* **56**(5) 1153–1175.

Laprè, M. A., A. S. Mukherjee, L. N. Van Wassenhove. 2000. Behind the learning curve: Linking learning activities to waste reduction. *Management Sci.* **46**(5) 597–611.

Lee, H. L., K. C. So, C. S. Tang. 2000. The value of information sharing in a two-level supply chain. *Management Sci.* **46**(5) 626–643.

Lovejoy, W. S. 2006. Optimal mechanisms with finite agent types. *Management Sci.* **52**(5) 788–803.

Lutze, H., Ö. Özer. 2008. Promised lead-time contracts under asymmetric information. *Oper. Res.* **56**(4) 898–915.

Martinez-de Albeniz, V., D. Simchi-Levi. 2003. Mean-variance trade-offs in supply contracts. *Naval Res. Logist.* **53**(7) 603–616.

McAfee, R. P., J. McMillan. 1987. Auctions and bidding. *J. of Economic Literature* **25**(2) 699–738.

Monahan, G. E., N. C. Petruzzi, W. Zhao. 2004. The dynamic pricing problem from a newsvendor's perspective. *Manufacturing Service Oper. Management* **6**(1) 73–91.

Müller, A., D. Stoyan. 2002. *Comparison Methods for Stochastic Models and Risks*. John Wiley and Sons, Chichester, UK.

Myerson, R.B. 1982. Optimal coordination mechanisms in generalized principal-agent problems. *Journal of Mathematical Economics* **10** 67–81.

Özer, Ö., O. Uncu. 2008. An integrated framework to optimize stochastic, dynamic qualification timing and production decisions. Working Paper.

Özer, Ö., O. Uncu, W. Wei. 2007. Selling to the newsvendor with a forecast update: Analysis of a dual purchase contract. *Eur. J. Oper. Res.* **182** 1150–1176.

Özer, Ö., W. Wei. 2006. Strategic commitments for an optimal capacity decision under asymmetric forecast information. *Management Sci.* **52**(8) 1239–1258.

Peskir, G., A. Shiryaev. 2006. *Optimal Stopping and Free-Boundary Problems*. Birkhäuser, Basel, Switzerland.

Petruzzi, N. C., M. Dada. 1999. Pricing and the news vendor problem: a review with extensions. *Oper. Res.* **47**(2) 183–194.

Pisano, G. P. 1996. Learning-before-doing in the development of new process technology. *Research Policy* **25**(7) 1097–1119.

Plambeck, E. L., S. A. Zenios. 2000. Performance-based incentives in a dynamic principal-agent model. *Manufacturing Service Oper. Management* **2**(3) 240–263.

Riley, J., R. Zeckhauser. 1983. Optimal selling strategies: When to haggle, when to hold firm. *Quarterly J. of Economics* **98**(2) 267–289.

Ross, S. 1983. *Introduction to Stochastic Dynamic Programming*. Academic Press Inc., New York, NY.

Sapra, A., P. L. Jackson. 2009. On the equilibrium behavior of a supply chain market for capacity. Working Paper. Indian Institute of Management Bangalore, Bangalore, India.

Savin, S., C. Terwiesch. 2005. Optimal product launch times in a duopoly: Balancing life-cycle revenues with product cost. *Oper. Res.* **53**(1) 26–47.

Schoenmeyr, T., S. C. Graves. 2009. Strategic safety stocks in supply chains with evolving forecasts. *Manufacturing Service Oper. Management* **11**(4) 657–673.

Shaked, M., J. G. Shanthikumar. 2007. *Stochastic Orders*. Springer, New York, NY.

Shane, S. A., K. T. Ulrich. 2004. Technological innovation, product development, and entrepreneurship in Management Science. *Management Sci.* **50**(2) 133–144.

Shiryaev, A. N. 1978. *Optimal Stopping Rules*. Springer-Verlag, New York, NY.

Skreta, V. 2006. Sequentially optimal mechanisms. *Review of Economic Stud.* **73**(4) 1085–1111.

Smith, J. E., K. F. McCardle. 2002. Structural properties of stochastic dynamic programs. *Oper. Res.* **50**(5) 796–809.

Stadje, W. 1991. Optimal stopping in a cumulative damage model. *OR Spectrum* **13**(1) 31–35.

Taub, E. 2007. Microsoft to spend $1.15 billion for xbox repairs. *The New York Times*, July 6.

Taylor, T. A., W. Xiao. 2009. Incentives for retailer forecasting: Rebates vs. returns. *Management Sci.* **55**(10) 1654–1669.

Terwiesch, C., R. E. Bohn. 2001. Learning and process improvement during production ramp-up. *Internat. J. Production Econom.* **70**(1) 1–19.

Terwiesch, C., Y. Xu. 2004. The copy-exactly ramp-up strategy: Trading-off learning with process change. *IEEE Trans. Engrg. Management* **51**(1).

Toktay, L. B., L. M. Wein. 2001. Analysis of a forecasting-production-inventory system with stationary demand. *Management Sci.* **47**(9) 1268–1281.

Topkis, D. M. 1978. Minimizing a submodular function on a lattice. *Oper. Res.* **26** 305–321.

Tsitsiklis, J. N., B. Van Roy. 1999. Optimal stopping of markov processes: Hilbert space theory, approximation algorithms, and an application to pricing high-dimensional financial derivatives. *IEEE Trans. Automatic Control* **44** 1840–1851.

Ülkü, S., L. B. Toktay, E. Yücesan. 2005. The impact of outsourced manufacturing on timing of entry in uncertain markets. *Production Oper. Management* **14**(3) 301–314.

Ulu, C., J. E. Smith. 2009. Uncertainty, information acquisition, and technology adoption. *Oper. Res.* **57**(3) 740–752.

Wang, Y., B. Tomlin. 2009. To wait or not to wait: Optimal ordering under lead time uncertainty and forecast updating. *Naval Res. Logist.* **56**(8) 766–779.

Wu, R., M. C. Fu. 2003. Optimal exercise policies and simulation-based valuation for american-asian options. *Oper. Res.* **51**(1) 52–66.

Yao, D. D., S. Zheng. 1999a. Coordinated quality control in a two-stage system. *IEEE Trans. Automatic Control* **44**(6) 1166–1179.

Yao, D. D., S. Zheng. 1999b. Sequential inspection under capacity constraints. *Oper. Res.* **47**(3) 410–421.

Zangwill, W. I., P. B. Kantor. 1998. Toward a theory of continuous improvement and the learning curve. *Management Sci.* **44**(7) 910–920.

Zhang, H., S. Zenios. 2008. A dynamic principal-agent model with hidden information: Sequential optimality through truthful state revelation. *Oper. Res.* **56**(3) 681–696.