Title
Initial Production Capacity Investments for Commercializing Pharmaceutical Products

Author
Yuen, Ming Kwan

Publication Date
2012

Peer reviewed|Thesis/dissertation

Initial Production Capacity Investments for Commercializing Pharmaceutical Products

by

Ming Kwan Yuen

A dissertation submitted in partial satisfaction of the

requirements for the degree of

Doctor of Philosophy

in

Engineering – Industrial Engineering and Operations Research

in the

Graduate Division

of the

University of California, Berkeley

Committee in charge:
Professor Philip M. Kaminsky, Co-chair
Professor Xin Guo, Co-chair
Professor Robert M. Anderson

Spring 2012

Initial Production Capacity Investments for Commercializing Pharmaceutical Products

Abstract

Initial Production Capacity Investments for Commercializing Pharmaceutical Products

by

Ming Kwan Yuen

Doctor of Philosophy in Engineering – Industrial Engineering and Operations Research

University of California, Berkeley

Professor Philip M. Kaminsky, Co-chair
Professor Xin Guo, Co-chair

This thesis is motivated by the investment problems pharmaceutical manufacturing firms face when introducing new drug products. We consider two different types of resources that play a role in determining the initial production capacity in pharmaceutical manufacturing: generic production resources, and specialized facilities and equipment. Due to the differences in availability of each resources type and the dynamics of information the firm receives with respect to the underlying uncertainties, the investment problems posed by these two resource types are very different. We build two separate dynamic optimization models to analyze the respective investment strategies.

For procurement of generic resources, we consider the firm facing random demand while the drug approval arrives at a random time. The firm can either increase or decrease inventory of the resources by buying or selling on a spot market where price fluctuates randomly over time. The firm's goal on this operation is to maximize expected discounted profit over the procurement process. We first show that this optimization problem is equivalent to a two-dimensional singular control problem. We then show that the optimal policy is completely characterized by a simple price-dependent two-threshold policy.

For specialized equipment, we consider a model where the firm must balance two conflicting objectives: on one hand, the delay in scaling-up production once the product is approved must be minimized, and on the other hand, the risk of investing in ultimately unused capacity must be minimized. We develop a stylized model of this type of capacity investment problem, where the firm re-evaluates its capacity investment strategy as information about the potential success of the product is continually updated (for example, via clinical trial results). We identify settings in which by continually reviewing the building strategy, the firm can substantially reduce both the delay of the commercial launch of the new product, and the risk of lost investment.

Although, our focus here is on the investment decisions in introducing a new drug product in the pharmaceutical industry, the models described in the following subsections can be

applied to the introduction of products which require specialized equipment to manufacture and have a long research and development phase.

For my parents, Lai Sheung Yuen-Wong and Hing Tong Yuen.

Contents

List of Figures

List of Tables

Acknowledgments

I would like to acknowledge the advice and guidance of Dr. Philip M. Kaminsky and Dr. Xin Guo, chairpersons of my thesis committee. I would also like to thank Dr. Robert M. Anderson, committee member, for his support. I thank the Berkeley Chapter of CELDi and its industrial research partners; Bayer, Genentech, and BioMarin; for invaluable industry experience and financial support. I thank Dr. Lee W. Schruben for sharing his experience and knowledge. I also acknowledge Dr. Pascal Tomecek, whose research was an important building block of this thesis. I thank the faulty and students of the Berkeley Industrial Engineering and Operations Research Department for providing a rich and stimulating research environment. Special thanks to staff member Mr. Michael Campbell for his professionalism and enthusiasm.

I thank my family. Their support and encouragement made my graduate career at Berkeley possible. Finally, I thank my girlfriend Stephanie Sapp for her cheerful spirit; especially, during the stressful final stages of writing this thesis.

Chapter 1

Introduction

In the pharmaceutical industry, introducing a new drug is a complex and time consuming process. Research and development alone is risky, expensive, and time consuming. The discovery of a new drug is only the beginning. The firm has to demonstrate safety and efficacy of the new drug for a particular disease or set of symptoms to regulatory agencies via a series of clinical trials in order to gain the approval to put the new drug on the market. Specifically, the drug is tested in three phases of trials: the first involves testing for safety on a relatively small group (typically 10 to 100 healthy individuals) patients; the second involves testing for effectiveness (typically on groups of 20-300 patients); and Phase III involves comparisons of existing treatments with the new drug (typically on 300 to 3000 patients).[1] On average, it takes eight years for all of the clinical trials to be completed and a drug to be approved.

The construction or acquisition and licensing of physical production capacity for a new drug can also be an expensive and time consuming project. Although firms acquire some production capacity to support clinical testing, this capacity is typically not sufficient to meet demand once a drug is approved. The firms have to secure sufficient production capacity to manufacture promptly to meet market demand upon the approval of the new drug, recover from investments in research and development, and meet regulatory requirements if any. It is of importance for firms to optimize the investments in acquiring production capacities on time while keeping costs as low as possible. There are multiple sources of randomness at each of the decision points in acquiring production capacity such as the outcome of the clinical trial (hence, the new drug application), the approval date, demand quantity at product launch, and the fluctuation of the prices of raw materials.

Traditionally, firms commit to building or acquiring production capacity early enough in the trial process that they can ensure that capacity will be ready at the time that the drug receives final approval. While firms do consult with their R&D group to assess the likelihood of the drug receiving regulatory approval, these estimates are usually made when the data from the on-going trial and the laboratory findings is limited, and are most often not updated to improve their accuracy as more information is obtained.

[1] For a detailed and comprehensive reference on clinical trials see (Friedman et al. 2010).

Of course, there is no guarantee that a new drug will pass all of the required phases of clinical trials, and in fact most drugs fail to show effectiveness at Phase II and Phase III of clinical studies. By committing to building or investing in a facility without the appropriate level of evidence that a drug will pass the required clinical trials, and by not updating this decision as information becomes available, the firm may be taking on excessive capital investment risk.[2] On the other hand, a firm has a limited window of exclusive sales rights for any drug before generic drug-makers enter the market and drive down prices, and this is typically the firm's primary opportunity to recover enormous initial investment costs. Under US patent law, for example, a firm has 20 years of exclusive rights to market a new drug in the United States, where the 20 years starts at the beginning of Phase I of the clinical studies. Thus, if the firm underestimates the likelihood of approval, it may act too conservatively when making production capacity investment decisions, and thus be unable to take full advantage of its limited window of exclusivity. Perhaps more importantly, patients will suffer if the availability of safe and approved drugs is delayed. Indeed, for the extreme case of life-saving orphan drugs, which are developed to treat diseases affecting less than 200,000 people in the United States, the FDA can choose to prematurely allow competition if the manufacturer is unable to provide a sufficient reliable supply of the drug (see `www.fda.gov` for more details).

1.1 Investing in Production Capacity

Production capacity is the measure of the abilities and limitations of a network of production resources. This include building manufacturing sites and equipment, and procuring the necessary raw materials and supplies in order to produce the first batch of commercial drug. We consider two different types of resources in determining the production capacity in pharmaceutical manufacturing: generic production resources, and specialized facilities and equipment. Due to the differences in availabilities of each resources types and the dynamics of information the firm receives with respect to the underlying uncertainties, the investment problems posted by these two resources type are very different from one another. We build two separate dynamic optimization models to analyze the respective investment strategies. Although, our focus here is on the investment decisions in introducing a new drug product in the pharmaceutical industry, the models described in the following subsections can be applied to the introduction of products which require specialized equipment to manufacture and have a long research and development phase.

[2] In practice, firms usually have multiple pipeline products on clinical trials to hedge against this risk, such that the new capacity built can be used by the next product in the pipeline in case their first drug fails. However, in that case, the excess capacity may stay idle for a few years, and the manufacturing facilities may need to be abandoned or require expensive retrofitting if the primary product for which they are being built is not approved.

1.1.1 Generic Production Resources

In biopharmaceutical manufacturing for example, due to the economies of scale in the fermentation process, it is a common practice to use production campaigns when manufacturing active pharmaceutical ingredients (API). Firms produce in a single campaign to satisfy the "yet-to-be-certain" annual or semi-annual demand of the drug. Therefore, the procurement process for raw materials and supplies such as fetal bovine serum (widely used as the growth medium in mammalian cell cultures), filters for purification (to harvest of raw active drug ingredient), buffer solutions for filtration columns, etc. is of importance for the launching of the new drug product to the market. Since they are widely available, the leadtimes for acquiring these generic materials are generally short. Thus, drug makers can wait until later in the on going clinical trial to achieve a higher confidence level in a successful commercial launch of the new drug before committing to the purchase. However, the spot prices of such materials fluctuate according to the market dynamics, so the firm faces supply uncertainties on material costs. Together with the perishable nature and high storage costs of the generic production resources, the procurement of such items resembles the classical "newsvendor" problem, except that the firm can buy and sell the raw materials at the respective spot markets before the launch date.

Spot market supply purchases are increasingly considered an important operational tool for the firm facing the risk of higher than anticipated demand for goods (see, e.g., (Simchi-Levi and Simchi-Levi 2008) and the references therein). For example, Hewlett-Packard manages the risks associated with electronic component procurement by utilizing a portfolio of long term and option contracts and the spot market (Billington 2002). Indeed, there has been a recent stream of research focusing on determining an optimal mix of long term fixed commitment and options/procurement contracts. In these models, the spot market is typically employed if supply requirements exceed the contracted amount of the fixed commitment contract, or if the spot price happens to be lower than the exercise price of the procurement options.

We argue that if effectively utilized, the spot market can be used to hedge against much more than just excess demand which is typically suggested in the literature. In many cases, the spot market can be a powerful tool for hedging against both supply cost uncertainty and demand price uncertainty in the supply chain, especially for pharmaceutical production where raw materials are often expensive. To explore this concept, we develop a stylized model of a firm that has a random period of random time[3] to increase or decrease inventory by purchasing or selling on the spot market before facing a single annual demand [4] of random magnitude, the revenue of which is a function of the spot market price when the demand is realized. We demonstrate that in many cases, the firm can use purchases and sales on the supply spot market to increase expected profits(savings), and thus to guard against both low prices for its products and high prices for its product components.

[3]The approval time of the new drug is a function of patients enrollment of the clinical trial and the performance of the new drug at the clinical stage which is exogenous to the firm.

[4]We assume the manufacturer use the common batch production method in fermentation campaigns in pharmaceutical manufacturing to take advantage of the economies of scale in terms of material yield.

1.1.2 Specialized Facilities and Equipment

Production facilities and equipment are often tailored to a particular production technology; for instance, the fermentation tanks, filtration columns, storage bins, refrigerators, pipelines and infrastructure of biopharmaceutical production facilities. Building and licensing traditional commercial-scale production facilities and equipment can take 4-5 years and cost up to $US 1 Billion. Due to the long lead times in acquiring the facility and equipment, the firm has to make investment commitments early enough to ensure that capacity will be ready at the time of commercial launch.

Recently, novel approaches to capacity expansion have become viable options – these options typically take less time to ramp up, but are more expensive to build and/or to operate. For example, there is growing interest in one-time-use disposable tanks and mixers for pharmaceutical production. These are typically pre-validated, pre-sterilized, and stocked by vendors, which significantly reduces the lead time associated with acquiring usable capacity. On the other hand, much of this equipment requires an expensive infrastructure, and has to, by definition, be replaced each time it is used, significantly increasing operating costs.

We propose a more sophisticated approach to production capacity investments on facilities and equipment, in which regular information updates are explicitly considered, might be worthwhile. To do this, we consider a stylized model of a firm that potentially takes such an atypically sophisticated approach. This firm frequently reviews the progress of ongoing trials to update an estimate of the ultimate likelihood of the product being approved. As additional results are observed, the quality of the estimate increases, and the firm can then adjust investment decisions to minimize penalties and investment costs. In addition to decisions related to whether or not to invest at specific times, the firm may consider alternative production technologies for a given drug, some of which are more expensive but require less time to build (thus improving information quality at higher cost), and some which are less expensive but require more time to build (thus sacrificing information for cost).

Given this model, we explore a variety questions associated with the appropriate timing and course of action when building or acquiring production capacity (for a new drug or something similar). For example, when is the approach we model (with regular reevaluation of capacity investment decisions as information is updated) worthwhile? If a firm has the opportunity to start building a production facility at any time as clinical data is gathered, when is the best time to build? If the firm can pause or abandon the project as additional data is received, under what conditions should it do so? What are the circumstances under which it makes sense to build more expensive capacity with a shorter lead time? Should a firm ever start building one type of capacity, and then switch to building an alternative type of capacity?

We characterize the optimal strategy under a variety of conditions, and then complete a computational study to develop additional insights. We see, for example, that continual adjustment of investment strategies is generally valuable, but particularly so when the quality of preliminary data is relatively poor or when the planning horizon is relatively long. Given the parameters of our computational study, for example, we see that regular reevaluation of investment decisions as information is updates the reevaluation approach can save the firm

up to 40% in expected construction cost and penalty costs. We also develop insights into the impact of increasing construction cost, and into the value of the types of alternative production technologies mentioned above.

Chapter 2

Literature Review

2.1 Inventory Control with Market Price Risk

There is a long history of research focusing on inventory strategies when the cost of the inventory is random, typically with the objective of minimizing inventory cost. (Kalymon 1971; Sethi and Cheng 1997; Cheng and Sethi 1999; Chen and Song 2001), for example, considered versions of periodic review models where component costs (and sometimes other problem parameters) are Markov-modulated, usually demonstrating the optimality of state dependent base-stock or (s, S) policies. Under these policies, there exist inventory thresholds such that it is optimal for manager to order up-to a preset inventory level S at every review periods (base-stock policy), or when the current inventory level falls below a lower inventory threshold s ((s, S) policy). Another stream of literature modeled deterministic demand and characterized optimal policies, either in a periodic model with random raw material prices (Golabi 1985; Berling 2008), increasing (Gavirneni and Morton 1999), or decreasing (Wang 2001) prices, or with one or two different levels of constant continuous time demand and occasional supply price discounts (Moinzadeh 1997; Goh and Sharafali 2002; Chaouc 2006).

There are two major streams of research that consider the impact of a spot market on supply chain operations; one of these focuses on contract valuation (see references in (Haksoz and Seshadri 2007)), and another considers optimal procurement from the spot market. This latter category of work can be further subdivided into single and multi-period models. (Seifert and Hausman 2004) considered a single period model where a supply contract is signed at the start of the period, demand is realized, and then the buyer can either make purchases on the spot market to meet demand, or salvage excess inventory. Optimal purchasing quantities were determined for this setting. (Akella and Kleinknecht 2001) considered a similar setting, and focused on reserving an appropriate capacity level to meet demand. (Wu D. J. and Zhang 2002) studied a single period optimal procurement contract model in the presence of a spot market, and (Golovachkina and Bradley 2003) considered a single period model designed to determine how the spot market impacts supply chain coordination. (Fu and Teo 2006) investigated a single period model in which the firm must first select from a set of supply contracts, and then random demand is realized and the firm must meet that

demand by utilizing either supply contracts or the spot market. Optimal contract selection and utilization policies were characterized. In (Li and Kouvelis 1999), deterministic demand must be met after a deterministic time period, but the firm has a contract to procure supply on the spot market at some point before demand is realized, where spot market price is a continuous random process. The optimal purchase time was derived numerically. (Yi and Scheller-Wolf 2003) considered a discrete time multi-period model in which the firm can either buy at a fixed price from a long-term supplier or buy from a spot market incurring variable purchasing cost and a fixed cost of using the spot market, and characterized the optimal purchasing policy. In (Cohen and Agrawal 1999), an optimal long term contract was compared to utilizing the spot market for a series of periods, where each period was essentially an independent news-vendor problem. (Martinez-de Albeniz and Simchi-Levi 2005) proposed a multi-period model in which a portfolio of supply contracts must be selected at the start of the horizon, and then in each period after demand is realized and the spot market price is observed, the decision to utilized contracts or buy on the spot market to meet demand must be made.

Other than the few exceptions noted above, these papers modeled the spot market with a single spot market price or a discretely realized series of spot market prices, and typically allowed one opportunity to buy or sell on the spot market following each demand realization.

2.2 Capacity Expansion

Models intended to determine when, how much and what kind of production capacity to build can be found in the literature as early as the 1960's. (Luss 1982; Van Mieghem 2003) provide comprehensive surveys of the capacity expansion literature. By assuming a deterministic growth rate of product demand, early work by (Sinden 1960; Erlenkotter 1977; Smith 1979), models production capacity as a type of inventory and approach the problem with Economic Order Quantity-like techniques in an infinite horizon setting.

Stochastic demand was introduced to the capacity expansion literature by (Manne 1961). By assuming steady state demand growth and no lead time for building capacity, this paper (as well as those by subsequent authors such as (Bean et al. 1992)) develops deterministic equivalents to the stochastic models. Explicit consideration of lead times significantly increases complexity of these models. (Elsanosi et al. 2000) and (Ryan 2004) use optimal stopping time models with appropriate adjustments to model the setting in which there is a lead time between when a capacity expansion decision is made and when the capacity is available in a stochastic setting. (Davis et al. 1987) utilize a stochastic control model, where the firm continuously controls the amount to invest in building capacity to meet stochastic demand with random lead time.

In contrast to most of the traditional stochastic capacity expansion literature, our goal is to model the likelihood of a successful clinical trial, rather than demand, as the primary source of randomness, as we discuss in Chapter 4. In addition, since the quality of the estimate of the likelihood of approval increases over time, a stationary model does not capture the salient features of our problem. Indeed, this is also the key difference between our model

and traditional real options analysis, a class of models that analyze and evaluate capital investment projects. Real options analysis often uses stochastic control techniques developed in the financial options literature, which also assume stationary stochastic processes. (Dixit and Pindyck 2004; Trigeorgis 1996) offer extensive outlines of this topic.

The multi-armed bandit problem originally introduced by (Robbins 1952) is often used to model the choice among the bets with unknown odds to maximize profits or among R&D projects with known success rates to maximize potential returns. In this setting, the firm must sequentially select among several experiment types, each with an unknown probability of success, with the long run goal of maximizing the number of successes. After each experiment, the firm updates its estimate of the success rate of that experiment type. Although the Bayesian update of the unknown parameters of the multi-armed bandit models captures the notion of updating key estimates used to select investments, these bandit models generally do not capture the complexities inherent in ensuring that capacity is ready for commercial launch of the product.

The inventory models described in (Scarf 1959; Azoury 1985; Lariviere and Porteus 1999) also use Bayesian learning to model the progressive improvement over time of the decision maker's knowledge of the underlying stochastic demand. In these discrete time models, the demand in each period is modeled by an independent and identically distributed random variable with unknown parameters. By assuming a class of continuous distributions with known conjugate prior, such as the "newsvendor distributions" described in (Braden and Freimer 1991), the decision maker in these models progressively updates the posterior distribution of the unknown parameters using the observed demand in each period, and then uses this information to optimize the inventory strategy to maximize profit over the planning horizon. Although we are also concerned with updating an estimate over time in order to ensure that demand can be met, the nature of the investment risk described in our capacity investment problem is very different than that of the inventory models. In contrast to the setting in those papers, our goal is to build a model which explicitly considers the multiple investments over time necessary to acquire production capacity, and "all-or-nothing" demand depending on the success of the clinical trial.

In Chapter 4, we develop a stochastic process which models change in the firm's estimate of the likelihood of a successful commercial launch of the new drug as the clinical trial progresses. We also model the decision process of building the initial production capacity.

Chapter 3

Investing in Generic Production Resources – Optimal Spot Market Inventory Strategies in the Presence of Cost and Price Risk

3.1 Introduction

We consider a continuous time model of spot market price evolution, and determine how the firm can buy and sell in the spot market *repeatedly* in order to guard against *both* supply cost uncertainty of raw materials. Specifically, we model the inventory level of the firm at time t, Y_t, with a pair of controls (ξ_t^+, ξ_t^-) so that $Y_t = Y_0 + \xi_t^+ - \xi_t^-$. Here ξ_t^+ and ξ_t^- are non-decreasing processes and they represent the cumulative buying and selling quantities at the spot market up to time t respectively. In case of a shortage of any of the required generic production resources at the start of production, time τ, the firm will make additional purchases at the spot market price. We further assume the firm uses an internal cost accounting method, such that it is equivalent to consider the revenue maximization problem where the procurement department is maximizing the profit from selling the raw materials to the manufacturing division of the firm at market price at time τ. We assume that the price of each unit of inventory is stochastic and modeled by a Brownian motion. We assume that the approval time of the new drug, and the demand starting at that time, are random. The revenue associated with the demand is assumed to be a function of the amount of that demand and the spot market price at the time when the demand arrives. In addition to the running holding cost, there are costs whenever inventory level is increased or decreased by selling or buying at the spot market: this adjustment cost is a function of the spot price and the amount of the adjustment, plus a proportional transaction cost. Note that this cost can be negative when selling inventory. Given this cost structure, the goal is to maximize expected discounted profit over an infinite time horizon. To facilitate our analysis, we assume no fixed cost and focus on explicitly characterizing the optimal policy.

In particular, we show that the optimal inventory policy depends on both the spot price and inventory level, and that it is in principle a simple and *not necessarily* continuous (F, G) policy. Given a spot price p and inventory level z, if (p, z) falls between $(F(z), G(z))$, no action is taken; if (p, z) falls above $F(z)$ (below $G(z)$), the inventory level is reduced to $F(z)$ (raised to $G(z)$).

Our model is closely related to a stream of research (Bather 1966; Archibald 1981; Constantinides and Richard 1978; Harrison and Taksar 1983; Harrison et al. 1983; Taksar 1985; Sulem 1986; Ormeci and Vate 2008) focusing on continuous time inventory models via impulse controls (i.e. with a fixed cost) or singular controls (i.e. without a fixed cost) formulation. Most of these papers (with the exception of (Archibald 1981) where the demand process is Poisson) considered a one product inventory model where the inventory level is a controlled Brownian motion. That is, the inventory level without intervention was modeled by a Brownian motion, and the continuous adjustment of the inventory level was assumed additive to the Brownian motion and with a linear cost plus a possibly fixed cost. Subject to an additional holding cost and shortage penalty, the objective in these papers was to minimize either the expected discounted cost or the average cost (Bather 1966; Ormeci and Vate 2008) over an infinite time horizon. Except for (Taksar 1985; Ormeci and Vate 2008), most of the models assumed no constraints on the inventory level besides restricting it to the positive real line. Assuming a fixed cost, (Constantinides and Richard 1978) proved the existence of an optimal (d, D, U, u) policy for this system: do nothing when inventory is in the region of (d, u), and adjust the inventory level to D (or U) whenever the inventory level falls to d (or rises to u). This optimal policy and the solution structure were more explicitly characterized under various scenarios in (Harrison and Taksar 1983; Harrison et al. 1983; Taksar 1985; Sulem 1986; Ormeci and Vate 2008).

The main contribution of our model is best discussed in light of several crucial elements underlying all previous control-theoretic inventory analysis. First, the price of the inventory was typically assumed to be constant so that the cost of the inventory control was linear. Secondly, the inventory control was usually additive to a Brownian motion, and as a result the inventory level was either unconstrained on the positive real line, or an infinite penalty cost was needed to ensure an upper bound on the inventory level (Taksar 1985; Ormeci and Vate 2008). These two characteristics ensured the control problem to be one-dimensional, to facilitate the analysis of the value function. The solution approach was to apply the Dynamic Programming Principle and to solve some form of Hamilton-Jacobi-Bellman equations or Quasi-Variational-Inequalities, with a priori assumptions on the regularity conditions.

In contrast, in our model, the adjustment cost is no longer linear and depends on the spot price, the transaction cost, and the amount of adjustment, and the inventory control variable is modeled directly, and is no longer necessarily additive to the underlying Brownian motion process. Thus, lower and upper bounds on the inventory level (that is, capacity constraints) are modeled directly. This approach to modeling capacity constraints has an additional advantage – it can be easily extended to more complex constraints on inventory levels without further technical difficulty. In essence, the introduction of price dynamics leads to a higher dimensional singular control problem for which previous analysis cannot be directly generalized. The derivation in this model is thus based on a new solution approach,

which allows us to bypass the possible non-regularity of the value functions. The key idea is to break down the two-dimensional control problem by "slicing" it into pieces of one-dimensional problem, which is an explicitly solvable two-state switching problem, and to show that this re-parametrization is valid by the notion of "consistency" established in (Guo and Tomecek 2008). (For more details, see the discussion section 3.5).

In the next section, we formally introduce our model. In Section 3.4, we translate the problem into an equivalent singular control problem and develop explicit analytical expressions for the optimal policy for this model. In Section 3.6, we computationally explore some of the implications of our results.

3.2 The Model

We consider a firm that purchases supply from a spot market in which the price of the supply component fluctuates over time. At a random time τ, the firm faces a random customer demand D. The firm meets demand if possible (we assume one supply component meets one unit of demand), charging an exogenously determined price that is a function of the spot market component price, and then salvages any excess inventory. At any time $t \in [0, \tau)$, the firm can instantaneously increase inventory of the component up to some upper bound on capacity $b \leq \infty$ or instantaneously decrease inventory down to some lower bound on inventory $a \geq 0$. However, the firm cannot buy inventory of the component to satisfy demand at time τ, and the firm can only buy inventory a finite number of times in a finite interval. Net gain at time τ is from selling to arriving customers and liquidating excess inventory minus penalty associated with not meeting demand, and thus is a function of the selling price and the inventory level at time τ, and the demand distribution. Moreover, at any time $t \in [0, \tau)$, inventory increase is associated with the purchase price of per unit at the supply spot market price P_t, plus possibly additional proportional transaction cost K^+. Similarly, inventory reduction is associated with the spot market price P_t, minus possibly additional proportional transaction cost K^-. Finally, there is a running holding cost for each unit of inventory C_h.

To capture this scenario in mathematical terms, we start with a complete and filtered probability space $(\Omega, \mathcal{F}, \mathbb{P})$, and assume that the arrival time of the request, τ, is exponentially distribution with rate λ (so that the average arrival time is $1/\lambda$). D, the random variable representing the demand at time τ is described by distribution function F_D. Meanwhile, the component spot market price $(P_t)_{t\geq 0}$ is stochastic and its dynamics are governed by a geometric Brownian motion such that [1]

$$dP_t = P_t(\mu dt + \sqrt{2}\sigma dW_t), \quad P_0 = p. \tag{3.1}$$

Here W_t is the standard Brownian motion on the probability space $(\Omega, \mathcal{F}, \mathbb{P})$, and μ and σ represent respectively the expected spot market price appreciation and the potential price risk. We express the net gain at request time τ by $H(Y_\tau, D)P_\tau$, where $H(Y_\tau, D)$ represents

[1] The extra term $\sqrt{2}$ is for notational convenience in the main text.

the revenue multiplier associated with selling each unit of the inventory, as well as a penalty associated with each unit of unmet demand and salvage associated with each unit of excess inventory. Specifically,

$$H(y, D) = \alpha \min(D, y) + \alpha_o (y - D)^+ - \alpha_u (D - y)^+, \tag{3.2}$$

where $\alpha \geq 1$ is the mark-up multiplier for each unit of met demand, $\alpha_u \geq 0$ is the penalty price multiplier for each unit the firm is short, and $0 \leq \alpha_o \leq 1$ is the fraction of price the firm is able to get by salvaging excess inventory.

To define admissible inventory policies, we specify the filtration $\mathbb{F}$ representing the information on which inventory decisions are based. Given λ and the distribution of D, it is clear that $\mathbb{F} = (\mathcal{F}_t)_{t \geq 0}$ is the filtration generated by P_t. Given $\mathbb{F}$, we define a pair of left-continuous with right limit, adapted, and non-decreasing processes ξ_t^+ and ξ_t^- to be the cumulative increases and decreases in supply inventory (purchases and sales, respectively) up to time t. Therefore, Y_t, the inventory level at time $t \in [0, \tau)$, is given by

$$Y_t = y + \xi_t^+ - \xi_t^-, \tag{3.3}$$

where y is the initial inventory amount.

To be consistent with the restriction that the firm can only purchase supply inventory on the spot market a finite number of times in a finite interval, Y_t is assumed to be a finite variation process. Meanwhile, for uniqueness of expression (3.3), (ξ^+, ξ^-) are supported on disjoint sets. Furthermore, ξ^+ and ξ^- are adapted to $\mathbb{F}$ implying that the firm is not clairvoyant. Y is left-continuous, capturing the restriction that the commodity cannot be purchased at time τ to satisfy demand. Also, note that given the upper and lower bounds on capacity discussed above, there exists $0 \leq a < b \leq \infty$ [2] such that an admissible control policy must satisfy $Y_t \in [a, b]$ for all $t \leq \tau$. Finally, for well-posedness of the problem, we assume $\mathbb{E}\left[\int_0^\infty e^{-\rho t} d\xi_t^+ + \int_0^\infty e^{-\rho t} d\xi_t^-\right] < \infty$.

To account for the time value between $[0, \tau]$, we define $r \geq 0$ to be a discount rate. Thus, at time $t \in [0, \tau)$, increases in the inventory incur a cost $-e^{-rt}(P_t + K^+) d\xi_t^+$ per unit, and decreases in the inventory generate revenue $e^{-rt}(P_t - K^-) d\xi_t^-$ per unit. In addition, assuming a running holding cost C_h for each unit of inventory, the holding cost between $(t, t + dt) \subset [0, \tau)$ is $e^{-rt} C_h Y_t dt$.

Given this setting and any admissible control policy (ξ^+, ξ^-), the expected return to the firm is:

$$\begin{aligned} J(p, y; \xi^+, \xi^-) = &\text{ payoff at transaction time } \tau - \text{running holding cost between } [0, \tau] \\ &- \text{cost of inventory control (via buying and selling) between } [0, \tau] \\ = &\ \mathbb{E}\left[e^{-r\tau} H(Y_\tau, D) P_\tau - \int_0^\tau e^{-rt} C_h Y_t dt \right. \\ &\left. - \int_0^\tau e^{-rt}(P_t + K^+) d\xi_t^+ + \int_0^\tau e^{-rt}(P_t - K^-) d\xi_t^- \right]. \end{aligned}$$

[2] Although the general result applies to the, here, we can indeed consider $a = 0$ and $b < \infty$, where b is the physical constraint on the storage space of the production facility which we will consider in the next chapter.

Assuming that τ is independent of $\mathbb{F}$ and D is independent of both τ and $\mathbb{F}$, a simple and standard conditioning argument gives an equivalent form of this expected return:

$$\begin{aligned} J(p,y;\xi^+,\xi^-) =& \mathbb{E}\left[e^{-r\tau}H(Y_\tau,D)P_\tau - \int_0^\tau e^{-rt}C_hY_t dt \right. \\ & \left. - \int_0^\tau e^{-rt}(P_t+K^+)d\xi_t^+ + \int_0^\tau e^{-rt}(P_t-K^-)d\xi_t^-\right] \\ =& \mathbb{E}\left[\int_0^\infty \lambda e^{-(r+\lambda)t}H(Y_t,D)P_t dt - \int_0^\infty e^{-(r+\lambda)t}C_hY_t dt \right. \\ & \left. - \int_0^\infty e^{-(r+\lambda)t}(P_t+K^+)d\xi_t^+ + \int_0^\tau e^{-(r+\lambda)t}(P_t-K^-)d\xi_t^-\right]. \end{aligned} \tag{3.4}$$

3.3 The Optimization Problem

The firm's goal is to manage inventory in order to maximize the expected discounted value over all possible admissible control policies (ξ^+,ξ^-). In other words, the firm must solve the following optimization problem:

$$W(p,y) = \sup_{(\xi^+,\xi^-)\in\mathcal{A}'_y} J(p,y;\xi^+,\xi^-), \tag{3.5}$$

subject to

$$\begin{aligned} & Y_t := y+\xi_t^+-\xi_t^- \quad \in [a,b], \quad y\in[a,b], \\ & dP_t := \mu P_t dt + \sqrt{2}\sigma P_t dW_t, \quad P_0 := p>0, \\ & C_h\in\mathbb{R}, \quad K^++K^->0; \end{aligned} \tag{3.6}$$

$$\begin{aligned} & \mathcal{A}'_y := \big\{(\xi^+,\xi^-) : \xi^\pm \text{ are left continuous, non-decreasing processes,} \\ & \quad y+\xi_t^+-\xi_t^-\in[a,b], \quad \xi_0^\pm=0; \\ & \quad \mathbb{E}\left[\int_0^\infty e^{-\rho t}d\xi_t^+ + \int_0^\infty e^{-\rho t}d\xi_t^-\right]<\infty; \\ & \mathbb{E}\left[\int_0^\infty e^{-\rho t}P_t d\xi_t^+ + \int_0^\infty e^{-\rho t}P_t d\xi_t^-\right]<\infty.\big\} \end{aligned} \tag{3.7}$$

A few remarks about this formulation: First, we assume that $r+\lambda>\mu$ to ensure the finiteness of the value function. Secondly, to avoid an arbitrage opportunity in the market, we assume that $K^++K^->0$. In addition, we assume without loss of generality that $K^+>0$, and consider only a bounded inventory level. Finally, for ease of exposition, we fix without loss of generality $a=0$, thus $H(0,D)$=0.

Summarizing, we impose the following standing assumptions:

Assumption 3.3.1 $\rho := r+\lambda>\mu$.

Assumption 3.3.2 $K^+ > 0$.

Assumption 3.3.3 $a = 0$ *and* $[0, b]$ *is bounded.*

To simplify subsequent notation, we define $m < 0 < 1 < n$ to be the roots of $\sigma^2 x^2 + (\mu - \sigma^2)x - \rho = 0$, so that

$$m, n = \frac{-(\mu - \sigma^2) \pm \sqrt{(\mu - \sigma^2)^2 + 4\sigma^2\rho}}{2\sigma^2}. \tag{3.8}$$

Finally, from the identity $\rho = -\sigma^2 mn$, we define a useful quantity η such that

$$\eta := \frac{1}{\rho - \mu} = \frac{-mn}{(n-1)(1-m)\rho} = \frac{1}{\sigma^2(n-1)(1-m)} > 0. \tag{3.9}$$

3.4 Solutions

3.4.1 An Equivalent Problem $V(p, y)$

Assuming that τ is independent of $\mathbb{F}$ and that D is independent of both τ and $\mathbb{F}$, the one period optimization problem in the previous section is in fact equivalent to the following singular control problem over an infinite time horizon.

Theorem 3.4.1 *Assume that* τ *is independent of* $(W_t)_{t\geq 0}$*, and* D *is independent of both* τ *and* $(W_t)_{t\geq 0}$*. The value function* $W(\cdot, \cdot)$ *in Eq. (3.5) satisfies*

$$W(p, y) = (-\frac{C_h}{\rho} + p)y + V(p, y). \tag{3.10}$$

where $V(\cdot, \cdot)$ *is the value function for the following optimization problem*

$$\begin{aligned} V(p, y) = \sup_{(\xi^+, \xi^-) \in \mathcal{A}_y} \quad & \mathbb{E}\left[\int_0^\infty e^{-\rho t}\tilde{H}(Y_t)P_t dt - (K^+ + \frac{C_h}{\rho})\int_0^\infty e^{-\rho t} d\xi_t^+ \right. \\ & \left. -(K^- - \frac{C_h}{\rho})\int_0^\infty e^{-\rho t} d\xi_t^- \right], \end{aligned} \tag{3.11}$$

with

$$\begin{aligned} \tilde{H}(y) = \quad & \lambda\mathbb{E}[H(y, D)] - (\rho - \mu)y \\ = \quad & \lambda(\alpha + \alpha_u - \alpha_o)\left[y(1 - F_D(y)) - \int_y^\infty z f_D(z) dz\right] \\ & +\lambda(\alpha - \alpha_o)\mathbb{E}[D] + (\lambda\alpha_0 + \mu - r - \lambda)y; \qquad (3.12) \\ Y_t \quad & = y + \xi_t^+ - \xi_t^- \quad \in [0, b], \quad y \in [0, b]; \qquad (3.13) \end{aligned}$$

$$dP_t = \mu P_t dt + \sqrt{2}\sigma P_t dW_t, \quad P_0 = p > 0; \tag{3.14}$$

$$C_h \in \mathbb{R}, \quad K^+ + K^- > 0; \tag{3.15}$$

$$\begin{aligned}\mathcal{A}_y = \quad & \big\{(\xi^+, \xi^-) : \xi^{\pm} \textit{ are left continuous, non-decreasing processes,} \\ & y + \xi_t^+ - \xi_t^- \in [0, b], \quad \xi_0^{\pm} = 0; \\ & \mathbb{E}\left[\int_0^\infty e^{-\rho t} d\xi_t^+ + \int_0^\infty e^{-\rho t} d\xi_t^-\right] < \infty.\big\}\end{aligned} \tag{3.16}$$

This equivalence statement suggests that the incorporation of the dynamics of the price process leads to a two-dimension singular control problem formulation with a state space (p, y), and generalizes the one-dimensional singular control problem studied extensively in (Harrison and Taksar 1983).

3.4.2 Preliminary Analysis

Proposition 3.4.2 *(Finiteness of Value Function)* $V(p, y) \leq \eta M p + \frac{C_h}{\rho} b$, *where* $M = \sup_{y \in [0,b]} |\tilde{H}(y)| < \infty$.

Lemma 3.4.3 $\tilde{H}(y)$ *is concave* y *for ANY distribution of* F_D *with finite expectation. In particular,*

$$\tilde{H}(y_2) - \tilde{H}(y_1) = \int_{y_1}^{y_2} \tilde{h}(z) dz$$

with $\tilde{h}(y)$ *decreasing in* y, *and*

$$\tilde{h}(y) := \lambda[(\alpha + \alpha_u - \alpha_o)[1 - F_D(y)] + \lambda \alpha_o + \mu - \rho. \tag{3.17}$$

Furthermore,

$$\mathbb{E}\left[\int_0^\infty |e^{-\rho t} \tilde{H}(Y_t) P_t| dt\right] < \infty, \quad \mathbb{E}\left[\int_0^\infty |e^{-\rho t} \tilde{h}(Y_t)| dt\right] < \infty.$$

It is worth mentioning that intuitively, function $\tilde{H}(\cdot)$ captures the ultimate potential benefit of carrying inventory over time and its derivative $\tilde{h}(y)$ represents the impact on business of increasing or decreasing inventory levels. We shall see that this $\tilde{h}(\cdot)$ is a key quantity for characterizing optimal policies.

3.4.3 Solving $V(p, y)$

Now, we solve $V(p, y)$ explicitly. Our solution approach relies on the following critical lemma connecting the value function of the singular control problem and that of a switching control problem ((Guo and Tomecek 2008) [Theorem 3.7]). (For related background, see the Appendix).

Step 1: Singular control → switching control

Lemma 3.4.4 *The value function in problem (3.11) is given by*

$$V(p,y) = \int_0^y v_1(p,z)dz + \int_y^b v_0(p,z)dz, \tag{3.18}$$

where v_0 and v_1 are solutions to the following optimal switching problems

$$v_k(p,z) := \sup_{\substack{\alpha\in\mathcal{B}\\ \kappa_0=k}} \mathbb{E}\left[\int_0^\infty e^{-\rho t}\left[\tilde{h}(z)P_t\right] I_t dt - \sum_{n=1}^\infty e^{-\rho\tau_n} K_{\kappa_n}\right], \tag{3.19}$$

provided that the corresponding optimal switching controls are consistent (per definition in the Appendix) and the resulting singular control is integrable. Here, $\alpha = (\tau_n, \kappa_n)_{n\geq 0}$ is an admissible two-state switching control, $\mathcal{B}$ is the subset of admissible switching controls $\alpha = (\tau_n, \kappa_n)_{n\geq 0}$ such that $\mathbb{E}\left[\sum_{n=1}^\infty e^{-\rho\tau_n}\right] < \infty\}$, with $\kappa_0 = K^- - \frac{C_h}{\rho}, \kappa_1 = K^+ + \frac{C_h}{\rho}$, and I_t the regime indicator function for any given $\alpha \in \mathcal{B}$. A singular control (ξ^+, ξ^-) is integrable if

$$\mathbb{E}\left[\int_0^\infty e^{-\rho t}|\tilde{H}(Y_t)P_t|dt + \int_{[0,\infty)} e^{-\rho t}|K^+|d\xi_t^+ + \int_{[0,\infty)} e^{-\rho t}|K^-|d\xi_t^-\right] < \infty. \tag{3.20}$$

The detailed proof of this lemma can be found in (Guo and Tomecek 2008). Meanwhile, this lemma enables us to translate our original control problem to a two-state switching control problem between two regimes 0 and 1: for a given inventory level z, switching from state 0 to 1 corresponds to inventory increase and switching from state 1 to 0 corresponds to inventory decrease. The cost for inventory increase and decrease is given by $K^+ + \frac{C_h}{\rho}$ and $-K^- + \frac{C_h}{\rho}$ respectively, and the benefit of being at state 1 is accumulated at rate $\tilde{h}(y)$. Furthermore, if there exists a consistent collection of switching controls so that the resulting singular control is integrable, then we have

$$V(p,y) = \int_y^b v_0(p,z)dz + \int_0^y v_1(p,z)dz.$$

where v_0 and v_1 are the corresponding value functions for switching controls. Moreover, explcit solution to v_0 and v_1 can be described analytically according to (Ly Vath and Pham 2007).

Step 2: Solving switching controls and v_0, v_1

Proposition 3.4.5 *v_0 and v_1 are the unique C^1 viscosity solutions with linear growth condition to the following system of variational inequalities:*

$$\min\left\{-\mathcal{L}v_0(p,z),\ v_0(p,z) - v_1(p,z) + K^+ + \frac{C_h}{\rho}\right\} = 0, \tag{3.21}$$

$$\min\left\{-\mathcal{L}v_1(p,z) - \tilde{h}(z)p,\ v_1(p,z) - v_0(p,z) + K^- - \frac{C_h}{\rho}\right\} = 0, \tag{3.22}$$

with boundary conditions $v_0(0^+,z) = 0$ *and* $v_1(0^+,z) = \max\{-K^- + \frac{C_h}{\rho}, 0\}$. *Here* $\mathcal{L}u(p,z) = \sigma^2 u_{pp}(p,z) + \mu u_p(p,z) - \rho u(p,z)$.

To solve for v_0, v_1, we see by modifying the argument in (Ly Vath and Pham 2007, Theorem 3.1) that for any given $z \in [0,b]$ and $k \in \{0,1\}$, an optimal switching control exists and can be described in terms of switching regions: there exist $0 < F(z) < G(z) < \infty$ such that it is optimal to switch from regime 0 to regime 1 (to increase the inventory at level z) when $P_t \in [G(z), \infty)$, and to switch from regime 1 to regime 0 (decrease the inventory at level z) when $P_t \in [0, F(z)]$. Furthermore, based on (Ly Vath and Pham 2007, Theorem 4.2), we see that for each $z \in [0,b]$, the switching regions are described in terms of $F(z)$ and $G(z)$, which take values in $(0,\infty]$ and can be explicitly derived as follows.

Case I: $K^- - \frac{C_h}{\rho} \geq 0$**.** First, for each $z \in [0,b]$ such that $\tilde{h}(z) = 0$, it is never optimal to do anything, so we take $F(z) = \infty = G(z)$, and $v_0(p,z) = 0 = v_1(p,z)$.

Secondly, for z such that $\tilde{h}(z) > 0$, $G(z) < \infty$ and it is optimal to switch from regime 0 to regime 1 (to increase the inventory at level z) when $P_t \in [G(z), \infty)$. Since $K^- - \frac{C_h}{\rho} \geq 0$, it is never optimal to switch from regime 1 to regime 0 (i.e. $F(z) = \infty$). Furthermore, we have

$$\begin{aligned} v_0(p,z) &= \begin{cases} A(z)p^n, & p < G(z), \\ \eta\tilde{h}(z)p - (K^+ + \frac{C_h}{\rho}), & p \geq G(z), \end{cases} \\ v_1(p,z) &= \eta\tilde{h}(z)p. \end{aligned}$$

Since v_0 is C^1 at $G(z)$, we get

$$\begin{cases} A(z)G(z)^n &= \eta\tilde{h}(z)G(z) - (K^+ + \frac{C_h}{\rho}), \\ nA(z)G(z)^{n-1} &= \eta\tilde{h}(z). \end{cases}$$

That is,

$$\begin{cases} G(z) &= \frac{\nu}{\tilde{h}(z)}, \\ A(z) &= \frac{K^+ + \frac{C_h}{\rho}}{(n-1)} G(z)^{-n} = \frac{K^+ + \frac{C_h}{\rho}}{(n-1)} \nu^{-n} \tilde{h}(z)^n, \end{cases}$$

where $\nu = (K^+ + \frac{C_h}{\rho})\sigma^2 n(1-m)$.

Finally, when $\tilde{h}(z) < 0$, it is optimal to switch from regime 1 to regime 0 (reduce inventory at level z) when $P_t \in [F(z), \infty)$. Since $K^+ + \frac{C_h}{\rho} > 0$, it is never optimal to switch from regime 0 to regime 1 (i.e. $G(z) = \infty$). The derivation of the value function proceeds analogously to the derivation for the case of $\tilde{h}(z) > 0$.

Case II: $K^- - \frac{C_h}{\rho} < 0$**.**

First of all, for each $z \in [0,b]$ such that $\tilde{h}(z) \le 0$, it is always optimal to reduce the inventory because $K^- - \frac{C_h}{\rho} < 0$. That is, $F(z) = \infty = G(z)$. In this case, clearly $v_0(p,z) = 0$ and $v_1(p,z) = -K^- + \frac{C_h}{\rho}$.

Next, for each $z \in [0,b]$ such that $\tilde{h}(z) > 0$, it is optimal to switch from regime 0 to regime 1 (i.e. to increase in the inventory at level z) when $P_t \in [G(z), \infty)$, and to switch from regime 1 to regime 0 (i.e. to decrease in the inventory at level z) when $P_t \in (0, F(z)]$, where $0 < F(z) < G(z) < \infty$.

Moreover, v_0 and v_1 are given by

$$v_0(p,z) = \begin{cases} A(z)p^n, & p < G(z), \\ B(z)p^m + \eta x\tilde{h}(z) - (K^+ + \frac{C_h}{\rho}), & p \ge G(z), \end{cases}$$

$$v_1(p,z) = \begin{cases} A(z)p^n - (K^- - \frac{C_h}{\rho}), & p \le F(z), \\ B(z)p^m + \eta p\tilde{h}(z), & p > F(z). \end{cases}$$

Smoothness of $V(p,z)$ at $p = G(z)$ and $p = F(z)$ leads to

$$\begin{cases} A(z)G(z)^n & = B(z)G(z)^m + \eta G(z)\tilde{h}(z) - (K^+ + \frac{C_h}{\rho}), \\ nA(z)G(z)^{n-1} & = mB(z)G(z)^{m-1} + \eta\tilde{h}(z), \\ A(z)F(z)^n & = B(z)F(z)^m + \eta F(z)\tilde{h}(z) + (K^- - \frac{C_h}{\rho}), \\ nA(z)F(z)^{n-1} & = mB(z)F(z)^{m-1} + \eta\tilde{h}(z). \end{cases} \tag{3.23}$$

Eliminating $A(z)$ and $B(z)$ from (3.23) yields

$$\begin{cases} (K^+ + \frac{C_h}{\rho})G(z)^{-m} + (K^- - \frac{C_h}{\rho})F(z)^{-m} & = \frac{-m}{(1-m)\rho}\tilde{h}(z)(G(z)^{1-m} - F(z)^{1-m}), \\ (K^+ + \frac{C_h}{\rho})G(z)^{-n} + (K^- - \frac{C_h}{\rho})F(z)^{-n} & = \frac{n}{(n-1)\rho}\tilde{h}(z)(G(z)^{1-n} - F(z)^{1-n}). \end{cases} \tag{3.24}$$

Since the viscosity solutions to the variational inequalities are unique and C^1, for every z there is a unique solution $F(z) < G(z)$ to (3.24). Let $\kappa(z) = F(z)\tilde{h}(z)$, $\nu(z) = G(z)\tilde{h}(z)$, then the following system of equations for $\kappa(z)$ and $\nu(z)$ is guaranteed to have a unique solution for each z:

$$\begin{cases} (K^+ + \frac{C_h}{\rho})\nu(z)^{-m} + (K^- - \frac{C_h}{\rho})\kappa(z)^{-m} & = \frac{-m}{(1-m)\rho}(\nu(z)^{1-m} - \kappa(z)^{1-m}), \\ (K^+ + \frac{C_h}{\rho})\nu(z)^{-n} + (K^- - \frac{C_h}{\rho})\kappa(z)^{-n} & = \frac{n}{(n-1)\rho}(\nu(z)^{1-n} - \kappa(z)^{1-n}). \end{cases}$$

Moreover, these equations depend on z only through $\nu(z)$ and $\kappa(z)$, implying that there exist unique constants κ, ν such that $\kappa(z) \equiv \kappa$ and $\nu(z) \equiv \nu$ for all z. Hence $F(z) = \kappa\tilde{h}(z)^{-1}$, $G(z) = \nu\tilde{h}(z)^{-1}$, with $\kappa < \nu$ being the unique solutions to

$$\begin{cases} \frac{1}{1-m}\left[\nu^{1-m} - \kappa^{1-m}\right] & = -\frac{\rho}{m}\left[(K^+ + \frac{C_h}{\rho})\nu^{-m} + (K^- - \frac{C_h}{\rho})\kappa^{-m}\right], \\ \frac{1}{n-1}\left[\nu^{1-n} - \kappa^{1-n}\right] & = \frac{\rho}{n}\left[(K^+ + \frac{C_h}{\rho})\nu^{-n} + (K^- - \frac{C_h}{\rho})\kappa^{-n}\right]. \end{cases}$$

Given $F(z)$ and $G(z)$, $A(z)$ and $B(z)$ are solved from Eq. (3.23),

$$\begin{cases} B(z) & = -\frac{G(z)^{-m}}{n-m}\left(\frac{G(z)\tilde{h}(z)}{\sigma^2(1-m)} - n(K^+ + \frac{C_h}{\rho})\right) = -\frac{F(z)^{-m}}{n-m}\left(\frac{F(z)\tilde{h}(z)}{\sigma^2(1-m)} + n(K^- - \frac{C_h}{\rho})\right), \\ A(z) & = \frac{G(z)^{-n}}{n-m}\left(\frac{G(z)\tilde{h}(z)}{\sigma^2(n-1)} + m(K^+ + \frac{C_h}{\rho})\right) = \frac{F(z)^{-n}}{n-m}\left(\frac{F(z)\tilde{h}(z)}{\sigma^2(n-1)} - m(K^- - \frac{C_h}{\rho})\right). \end{cases}$$

Step 3: Establishing the optimal control

According to Lemma 3.4.4, it suffices to check the consistency of the switching control and the integrability of the corresponding singular control.

First, given the solution to the switching problems, clearly the optimal switching control for any level $z \in (0, b)$ is given by the following:

Case I: For $z \in (0, b)$ and $p > 0$, let F and G be as given in for Case I. The switching control $\hat{\alpha}_k(p, z) = (\hat{\tau}_n(p, z), \hat{\kappa}_n(z))_{n\geq 0}$, starting from $\hat{\tau}_0(p, z) = 0$ and $\hat{\kappa}_0(z) = k$ is given by, for $n \geq 1$

- If $k = 0$, $\hat{\tau}_1(p, z) = \inf\{t > 0 : P_t \in [G(z), \infty)\}$ and for $n \geq 2$, $\hat{\tau}_n(z) = \infty$,
- If $k = 1$, $\hat{\tau}_1(p, z) = \inf\{t > 0 : P_t \in [F(z), \infty)\}$ and for $n \geq 2$, $\hat{\tau}_n(z) = \infty$.

Case II: For $z \in (0, b)$ and $p > 0$, F and G as given for case II. The switching control $\hat{\alpha}_k(p, z) = (\hat{\tau}_n(p, z), \hat{\kappa}_n(z))_{n\geq 0}$, starting from $\hat{\tau}_0(p, z) = 0$ and $\hat{\kappa}_0(z) = k$ is given by, for $n \geq 1$

- If $\hat{\kappa}_{n-1}(z) = 0$, $\hat{\tau}_n(p, z) = \inf\{t > \tau_{n-1} : X_t^x \in [G(z), \infty)\}$, $\hat{\kappa}_n(z) = 1$.
- If $\hat{\kappa}_{n-1}(z) = 1$, $\hat{\tau}_n(p, z) = \inf\{t > \tau_{n-1} : X_t^x \in (0, F(z)]\}$, $\hat{\kappa}_n(z) = 0$,

Now, define the collection of admissible switching controls $(\hat{\alpha}(p, z))_{z\in(0,b)}$ so that $\hat{\alpha}(p, z) = \hat{\alpha}_0(p, z)$ for $z > y$ and $\hat{\alpha}(p, z) = \hat{\alpha}_1(p, z)$ for $z \leq y$. Then,

Proposition 3.4.6 *The collection of switching controls $(\hat{\alpha}(p, z))_{z\in(0,b)}$ is consistent.*

Next, this consistent collection of optimal switching control corresponds to an admissible singular control $(\hat{\xi}^+, \hat{\xi}^-) \in \mathcal{A}_y$ in the following way according to (Guo and Tomecek 2008).

Lemma 3.4.7 (From Switching Controls to Singular Controls) *Given $y \in (0, b)$ and a consistent collection of switching controls $(\hat{\alpha}(z))_{z\in\mathcal{I}}$, define two processes $\hat{\xi}^+$ and $\hat{\xi}^-$ by setting $\hat{\xi}_0^+ = 0$, $\hat{\xi}_0^- = 0$, and for $t > 0$: $\hat{\xi}_t^+ := \int_{\mathcal{I}} I_t^+(z)dz$, $\hat{\xi}_t^- := \int_{\mathcal{I}} I_t^-(z)dz$. Then*

1. *The pair $(\hat{\xi}^+, \hat{\xi}^-) \in \mathcal{A}_y$ is an admissible singular control,*
2. *For all t, we almost surely have*

$$\hat{Y}_t = \operatorname{ess\,sup}\{z \in \mathcal{I} : I_t(z) = 1\} = \operatorname{ess\,inf}\{z \in \mathcal{I} : I_t(z) = 0\},$$

where $\operatorname{ess\,sup}\emptyset := \inf\mathcal{I}$ and $\operatorname{ess\,inf}\emptyset := \sup\mathcal{I}$.

Moreover,

Proposition 3.4.8 *The corresponding admissible singular control $(\hat{\xi}^+, \hat{\xi}^-) \in \mathcal{A}_y$ is integrable.*

Step 4. Solution

Combining these results, we see that the ordering region is given by $\{(p,z) : p \geq G(z)\}$ and the downsizing region by $\{(p,z) : p \leq F(z)\}$. It is optimal to take no action in the continuation region, given by $\{(p,z) : F(z) < x < G(z)\}$. If (p,y) is in the ordering (or downsizing) region, then a jump is exerted at time zero to make $\hat{Y}_{0+} = G^{-1}(p)$ (or $\hat{Y}_{0+} = F^{-1}(p)$).

Finally, by (Guo and Tomecek 2008, Theorem 3.10), we have

$$V(p,y) = \int_y^b v_0(p,z)dz + \int_0^y v_1(p,z)dz,$$

with

$$v_0(p,z) = \begin{cases} A(z)p^n, & p < G(z), \\ B(z)p^m + \eta\tilde{h}(z)p - K^+ - \frac{C_h}{\rho}, & p \geq G(z), \end{cases}$$

$$v_1(p,z) = \begin{cases} A(z)p^n - K^- + \frac{C_h}{\rho}, & p \leq F(z), \\ B(z)p^m + \eta\tilde{h}(z)p, & p > F(z). \end{cases}$$

3.4.4 Main Result

In summary, we see that the optimal value function is characterized below for two distinct cases. In the first, $K^- - \frac{C_h}{\rho} \geq 0$, implying that the proportional loss incurred upon selling inventory is greater than the gain from reduced future holding cost. In the second, $K^- - \frac{C_h}{\rho} < 0$, implying that reducing holding cost dominates the transaction cost.

Theorem 3.4.9 [Optimal value function for $K^- - \frac{C_h}{\rho} \geq 0$]

$$V(p,y) = \int_0^y v_1(p,z)dz + \int_y^b v_0(p,z)dz, \tag{3.25}$$

where v_0 and v_1 are given by

1. *For each $z \in [0,b]$ such that $\tilde{h}(z) = 0$: $v_0(p,z) = v_1(p,z) = 0$.*
2. *For each $z \in [0,b]$ such that $\tilde{h}(z) > 0$:*

$$\begin{cases} v_0(p,z) & = \begin{cases} A(z)p^n, & p < G(z), \\ \eta\tilde{h}(z)p - K^+ - \frac{C_h}{\rho}, & p \geq G(z), \end{cases} \\ v_1(p,z) & = \eta\tilde{h}(z)p, \end{cases}$$

where $G(z) = \frac{\nu}{\tilde{h}(z)}$, and $A(z) = \frac{K^+ + \frac{C_h}{\rho}}{(n-1)}G^{-n}(z)$, with $\nu = (K^+ + \frac{C_h}{\rho})\sigma^2 n(1-m)$.

3. For each $z \in [0,b]$ such that $\tilde{h}(z) < 0$:

$$\begin{cases} v_0(p,z) & = 0, \\ v_1(p,z) & = \begin{cases} B(z)p^n + \eta\tilde{h}(z)p, & p < F(z), \\ -K^- + \frac{C_h}{\rho}, & p \geq F(z), \end{cases} \end{cases}$$

where $F(z) = -\frac{\kappa}{\tilde{h}(z)}$, and $B(z) = \frac{K^- - \frac{C_h}{\rho}}{(n-1)}\kappa^{-n}F^{-n}(z)$, with $\kappa = (K^- - \frac{C_h}{\rho})\sigma^2 n(1-m)$.

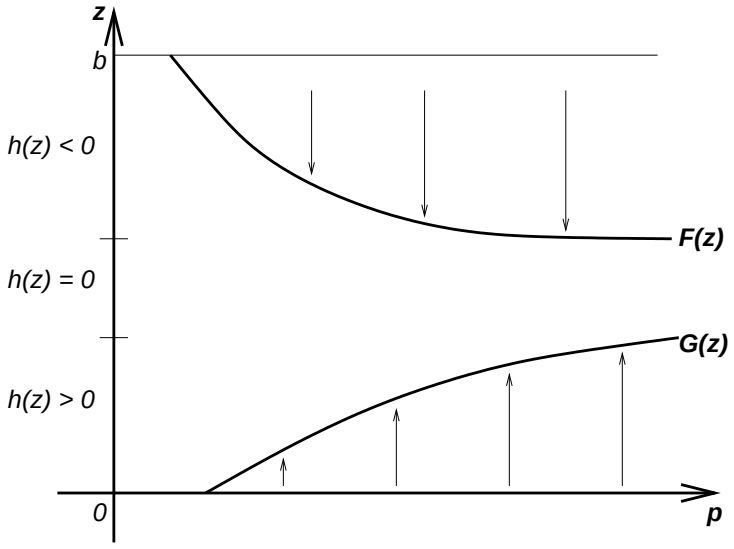

Figure 3.1 Policy when $K^- - \frac{C_h}{\rho} \geq 0$, with $F(z) = -\frac{\kappa}{\tilde{h}(z)}$ and $G(z) = \frac{\nu}{\tilde{h}(z)}$.

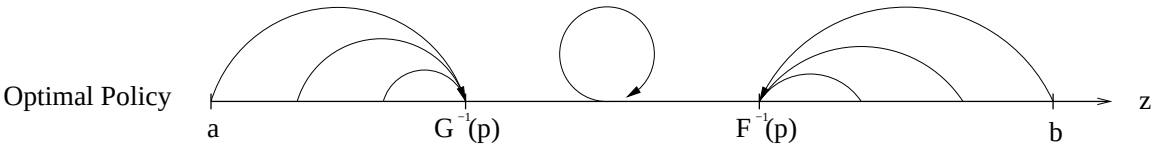

Figure 3.2 Illustration of the two-threshold order policy for fixed price p when $K^- - \frac{C_h}{\rho} \geq 0$.

Theorem 3.4.10 [Optimal value function for $K^- - \frac{C_h}{\rho} < 0$]

$$V(p,y) = \int_0^y v_1(p,z)dz + \int_y^b v_0(p,z)dz, \tag{3.26}$$

where v_0 and v_1 are given by

1. For each $z \in [0,b]$ such that $\tilde{h}(z) \leq 0$: $v_0(p,z) = 0, v_1(p,z) = -K^- + \frac{C_h}{\rho}$.

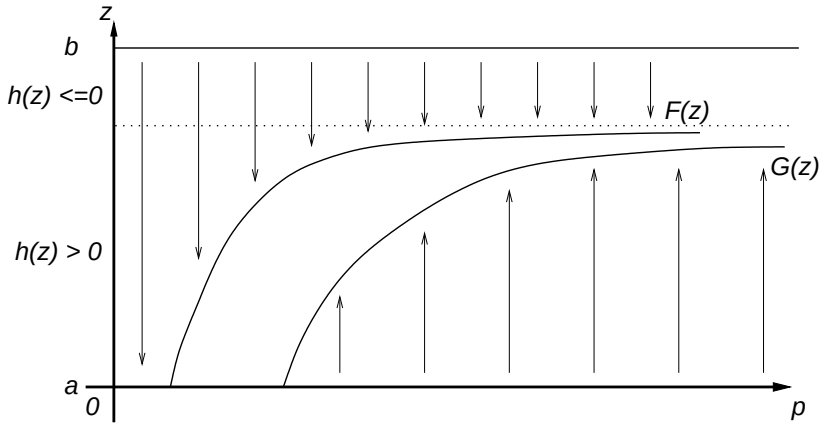

Figure 3.3 Policy when $K^- - \frac{C_h}{\rho} < 0$.

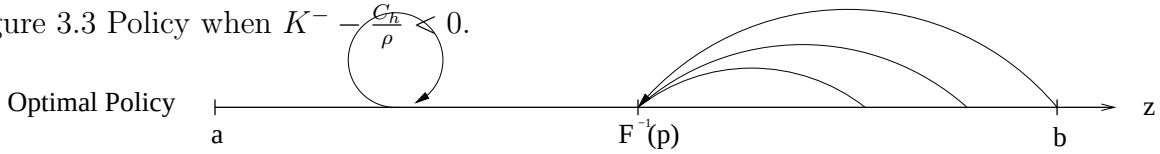

Figure 3.4 Illustration of the one-threshold order policy for fixed (low) price p and when $K^- - \frac{C_h}{\rho} < 0$.

2. For each $z \in [0, b]$ such that $\tilde{h}(z) > 0$:

$$v_0(p,z) = \begin{cases} A(z)p^n, & p < G(z), \\ B(z)p^m + \eta\tilde{h}(z)p - K^+ - \frac{C_h}{\rho}, & p \geq G(z), \end{cases} \tag{3.27}$$

$$v_1(p,z) = \begin{cases} A(z)p^n - K^- + \frac{C_h}{\rho}, & p \leq F(z), \\ B(z)p^m + \eta\tilde{h}(z)p, & p > F(z). \end{cases} \tag{3.28}$$

Here

$$A(z) = \frac{\tilde{h}(z)^n}{(n-m)\nu^n}\left(\frac{\nu}{\sigma^2(n-1)} + m(K^+ + \frac{C_h}{\rho})\right) \tag{3.29}$$

$$B(z) = \frac{-\tilde{h}(z)^m}{(n-m)\nu^m}\left(\frac{\nu}{\sigma^2(1-m)} - n(K^+ + \frac{C_h}{\rho})\right). \tag{3.30}$$

The functions F and G are non-decreasing with

$$F(z) = \frac{\kappa}{\tilde{h}(z)} \quad \textit{and} \quad G(z) = \frac{\nu}{\tilde{h}(z)}, \tag{3.31}$$

where $\kappa < \nu$ are the unique solutions to

$$\frac{1}{1-m}\left[\nu^{1-m} - \kappa^{1-m}\right] = -\frac{\rho}{m}\left[(K^+ + \frac{C_h}{\rho})\nu^{-m} + (K^- - \frac{C_h}{\rho})\kappa^{-m}\right], \tag{3.32}$$

$$\frac{1}{n-1}\left[\nu^{1-n} - \kappa^{1-n}\right] = \frac{\rho}{n}\left[(K^+ + \frac{C_h}{\rho})\nu^{-n} + (K^- - \frac{C_h}{\rho})\kappa^{-n}\right]. \tag{3.33}$$

Theorem 3.4.11 [Optimal control for $K^- - \frac{C_h}{\rho} \geq 0$] *For each $z \in [0, b]$, the optimal control is described in terms of $F(z)$ and $G(z)$ from Theorem 3.4.9 such that*

- *For z such that $\tilde{h}(z) > 0$, it is optimal to increase inventory past level z when $P_t \in [G(z), \infty)$, and never decreases.*
- *When $\tilde{h}(z) < 0$, it is optimal to decrease below inventory level z when $P_t \in [F(z), \infty)$, and it is never optimal to increase. When $\tilde{h}(z) = 0$, it is optimal to do nothing (i.e. $F(z) = \infty = G(z)$).*

Theorem 3.4.12 [Optimal control for $K^- - \frac{C_h}{\rho} < 0$] *For each $z \in [0, b]$, the optimal control is described in terms of $F(z)$ and $G(z)$ from Theorem 3.4.10 such that*

- *For z such that $\tilde{h}(z) > 0$, it is optimal to increase inventory past level z when $P_t \in [G(z), \infty)$, and to decrease inventory below level z when $P_t \in (0, F(z)]$.*
- *For z such that $\tilde{h}(z) \leq 0$, it is always optimal to decrease inventory level.*

The optimal policy is illustrated in Figure 3.1 and Figure 3.2 for the $K^- \geq \frac{C_h}{\rho}$ case, and Figure 3.3 and Figure 3.4 for the case of $K^- < \frac{C_h}{\rho}$. When $K^- < \frac{C_h}{\rho}$, implying a relatively high holding cost, for low prices inventory is decreased regardless of the value of $\tilde{h}(z)$; for high prices inventory is increased or decreased as necessary; for intermediate prices, no action is taken in general if inventory is low enough, but otherwise it is decreased (where Figure 3.4 illustrates this last case). In contrast, when $K^- \geq \frac{C_h}{\rho}$, the relatively low holding cost introduces a different policy: above a certain threshold price, except around the $\tilde{h}(z) = 0$ region, inventory is typically decreased for negative $h(z)$ values, and increased for positive $h(z)$ values, as illustrated in Figure 3.2. In general, depending on the holding cost C_h and the cost of selling K^-, reducing holding cost is a key driver, so conditions must be more favorable before inventory is increased, and it is more likely that inventory will be decreased.

Remark 3.4.13 *We emphasize that these results are quite general, and indeed hold for any $\tilde{H}(\cdot)$ function that is concave. Thus, $F(z)$ and $G(z)$ are not necessarily continuous. Nevertheless, when $\tilde{H}$ is continuously differentiable, strictly increasing, and strictly concave, we will have the regularity condition for F and G and for the value function, as postulated in the current literature.*

i

3.5 Discussion

As mentioned earlier, our solution approach is different from traditional approaches in which variational inequalities are solved directly.

To see this more closely, we first review the traditional dynamic programming approach and related variational inequalities. Clearly the optimization problem (3.11) for V has the state space $\{p, z\}$. Given any price p and the inventory level z at time 0, there are three options: do nothing, increase the inventory by purchasing on the spot market, or reduce the inventory by selling on the spot market.

If a quantity is purchased on the spot market, the inventory level jumps from z to $z + \Delta z$, thus the value function is at least as good as choosing over all possible jumps of size Δz with proportional cost $(K^+ + \frac{C_h}{\rho})\Delta z$. That is,

$$V(p, z) \geq \sup_{\Delta z}(-(K^+ + \frac{C_h}{\rho})\Delta z + V(p, z + \Delta z)),$$

which, by simple Ito's calculus, leads to $V_y(p, y) \leq K^+ + \frac{C_h}{\rho}$, with$V_y(\cdot, \cdot)$ the derivative of the value function V with respect to y. Similarly we see $V_y(p, y) \geq -K^- + \frac{C_h}{\rho}$ if choosing the option of reducing the inventory. Meanwhile, if no action is taken between time 0 and an infinitesimal amount of time dt, then expressing the value function at time 0 in terms of the value function at time dt though dynamic programming and Ito's calculus (as in (Constantinides and Richard 1978)) yields $\sigma^2 p^2 V_{pp}(p, y) + \mu p V_p(p, y) - rV(p, y) + \tilde{H}(y) \leq 0$. Combining these observations, we get the following (quasi)-Variational Inequalities

$$\begin{aligned} \max\{&\sigma^2 p^2 V_{pp}(p, y) + \mu p V_p(p, y) - rV(p, y) + \tilde{H}(y)p, \\ &V_y(p, y) - K^+ - \frac{C_h}{\rho}, -V_y(p, y) - K^- + \frac{C_h}{\rho}\} = 0. \end{aligned} \tag{3.34}$$

Moreover, the optimal policy (if it exists) can be characterized by explicitly finding the action and continuation regions where

$$\begin{cases} \text{(Inventory increase region)} = & \{(p, y) : V_y(p, y) = -K^- + \frac{C_h}{\rho}\}, \\ \text{(Inventory decresse region)} = & \{(p, y) : V_y(p, y) = K^+ + \frac{C_h}{\rho}\}, \\ \text{(No action region)} = & \{(p, y) : V_{y^-}(p, y) > -K^- + \frac{C_h}{\rho}, V_{y^+}(p, y) < K^+ + \frac{C_h}{\rho}, \\ & \quad \sigma^2 p^2 V_{pp}(p, y) + \mu p V_p(p, y) - rV(p, y) + \tilde{H}(y)p = 0\}. \end{cases}$$

A typical explicitly solvable optimal policy is a two-threshold bang-bang type policy or some degenerate form. (For more detailed derivation and background, interested readers are referred to e.g., (Constantinides and Richard 1978) or (Harrison and Taksar 1983)).

Taking this analysis one step further, one would expect, due to the stochastic nature of the P_t, a state-dependent threshold policy, where inventory is lowered if it is above the upper threshold, and increased if it is below the lower threshold (see, e.g, (Constantinides and Richard 1978)). That is, we would expect a downsizing region for inventory: $\{(p, z) : p \geq G(z)\}$, an ordering region: $\{(p, z) : p \leq F(z)\}$, and a (continuation) no-action region: $\{(p, z) : F(z) < p < G(z)\}$.

However, there is a serious issue in this straightforward extension. In order to derive a complete characterization of the optimal policy and the value function, one in general would assume *a priori* enough smoothness for the value function and the boundary to solve the

QVI. Unfortunately, the regularity conditions for this two-dimensional control problem do not hold in general. (See counter-examples in (Guo and Tomecek 2008)). Indeed, the value function may not be C^1 in p (although it is C^1 in y) and F, G may not be continuous. This possible irregularity especially in F and G makes explicit solution or "guessing" of F and G particularly difficult.

This is where we depart from the traditional approach: instead of solving variational inequalities directly, we translate the singular control problem (3.11) into a switching control problem, following (Guo and Tomecek 2008). The key idea is that by fixing each level z_0, we effectively will be solving for a one-dimensional two-state switching control problem, where switching from 0 to 1 corresponds to inventory increases and switching from 1 to 0 represents inventory reduction. In order for this approach to work for all z, meaning we can break down the two-dimensional control problem by slicing it into pieces of one-dimensional problem, we need to make sure the resulting control policies at different levels of z are "consistent". Intuitively, this consistency requires that for a given price p at level z_0, if it is optimal to reduce the inventory level, then it is also optimal to reduce the inventory level given the same p and a higher level $z(> z_0)$. This is the essence of Lemma 3.4.4 and our solution approach.

3.6 Computational Experiments and Observations

We have characterized the optimal solution for the material procurement problem of a pharmaceutical production firm facing random demand at the end of a single period of random length which models the new drug approval time, when the firm has the opportunity to trade on the spot market while waiting for demand. Of course, the opportunity to actively utilize the spot market to guard against cost and price risks must be balanced against the increasing complexities of actively trading inventory prior to experiencing demand. Below, we describe computational experiments that explore this trade-off.

3.6.1 Modified Newsvendor

To that end, we compare the expected results from the optimal policy as described here to those based on the use of a modified newsvendor solution to this model. We elected to use this modified newsvendor approach as a reasonable proxy for how a procurement manager who is not interested in repeatedly buying and selling on the spot market might manage the system.

In this modified newsvendor model, the manager purchases the inventory at the beginning of the investment, using a version of the well-known newsvendor solution adapted for the specifics of this setting. In particular, the newsvendor decision is based on the expected time until demand arrival, $E[\tau]$, the expected discounted sales price, $\alpha E[e^{-r\tau}p_\tau]$, and the expected discounted penalty cost and salvage value, $\alpha_u E[e^{-r\tau}p_\tau]$ and $\alpha_o E[e^{-r\tau}p_\tau]$. The

modified newsvendor inventory level, y^*, is given as follows:

$$y^* = \min\left\{b, \max\left\{0, F_D^{-1}\left(\frac{(\alpha+\alpha_u)E[e^{-r\tau}p_\tau] - (p_0 + K^+ + C_h E[\tau])}{(\alpha - \alpha_0 + \alpha_u)E[e^{-r\tau}p_\tau]}\right)\right\}\right\},$$

where 0 and b are the lower and upper bounds on the inventory level, $p_0 + K^+ + C_h E[\tau]$ is the expected cost per unit of acquiring (and holding) inventory, and F_D is the cumulative distribution of the random demand, D.

Recall τ is an exponential random variable with parameter λ, and p_t is a geometric Brownian motion with drift μ and volatility σ. Thus, $E[\tau] = \frac{1}{\lambda}$ and $E[e^{-r\tau}p_\tau] = \frac{\lambda p_0}{\lambda + r - \mu}$. Substituting, we get

$$y^* = \min\left\{b, \max\left\{0, F_D^{-1}\left(\frac{(\alpha+\alpha_u)\lambda P_0 - (p_0 + K^+ + C_h/\lambda)(\lambda + r - \mu)}{(\alpha - \alpha_0 + \alpha_u)\lambda P_0}\right)\right\}\right\}.$$

Finally, the expected value of implementing this modified newsvendor policy is

$$E[\alpha e^{-r\tau}P_\tau \min(y^*, D) - (p_0 + K^+ + C_h\tau)y^* + \alpha_o e^{-r\tau}P_\tau \max(y^* - D, 0) - \alpha_u e^{-r\tau}P_\tau \max(D - y^*, 0)].$$

3.6.2 Scenarios

Recall that we have two different cases when characterizing the optimal value function: $K^- \geq C_h/\rho$, where the proportional loss incurred on selling inventory is greater than the gain from reduced future holding cost, and $K^- < C_h/\rho$, where reducing holding cost dominates transaction costs. Based on solution structures in Section 3.4.4, we consider three scenarios in these experiments:

Scenario 1 $K^- \geq C_h/\rho$, and $y_0 = 0$ ($h(y_0) > 0$), so that in the optimal solution the manager may buy more inventory before time τ.

Scenario 2 $K^- \geq C_h/\rho$, and $y_0 = b$ ($h(y_0) < 0$), so that in the optimal solution the manager may sell off excessive inventory before time τ

Scenario 3 $K^- < C_h/\rho$ (and $y_0 = 0$), so that in the optimal solution may buy more inventory or sell off excessive inventory before time τ.

3.6.3 Parameters

For simplicity, we set the discount rate $r = 0$, the demand arrival rate $\lambda = 1$, the salvage multiplier $\alpha_0 = 0$, and the penalty cost multiplier $\alpha_u = 0$. To create the three scenarios described above, we set the mark-up multiplier $\alpha = 1.3$, the holding cost $C_h = 1$, the transaction cost for buying $K^+ = 1$, for Scenarios 1 and 2 the transaction cost of selling $K^- = 1.5$, and for Scenario 3 the transaction cost of selling $K^- = 0.5$. We set the inventory upper bound $b = 200$. To model the random demand, D, we assume $\log(D)$ follows a normal

distribution with parameters $\mu_d = 5$, and $\sigma_d = 0.7$. Hence, the mean and the variance of D are:

$$E[D] = e^{\mu_d + \frac{\sigma_d^2}{2}} = 189.61,$$

$$var[D] = (e^{\sigma_d^2} - 1)e^{2\mu_d + \sigma_d^2} = 22734.40.$$

For the price process p_t, we set the initial price $p_0 = 5$, and the drift rate $\mu = 0.3$. We vary the volatility σ from $0.05/\sqrt{2}$ to $140/\sqrt{2}$ to assess the impact of volatility.

3.6.4 Results and Observations

We use Theorems 3.4.1, 3.4.9 and 3.4.10 and the parameters listed above to calculate the optimal expected value for Scenarios 1, 2, and 3, and in figures 3.5, 3.6, and 3.7, we graph the relative gain from using the optimal policy rather than the modified newsvendor policy (that is, the difference between the optimal expected value and the newsvendor value divided by the optimal expected value) for different price volatilities.

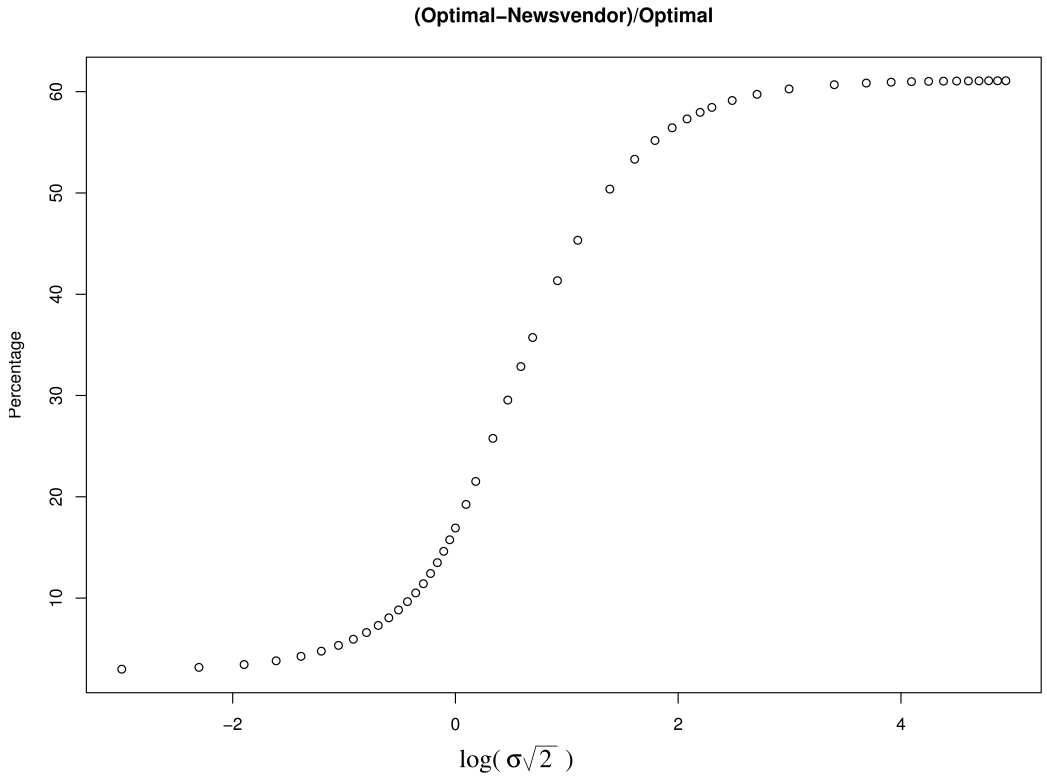

Figure 3.5 Scenario 1, percentage difference versus change in price volatility.

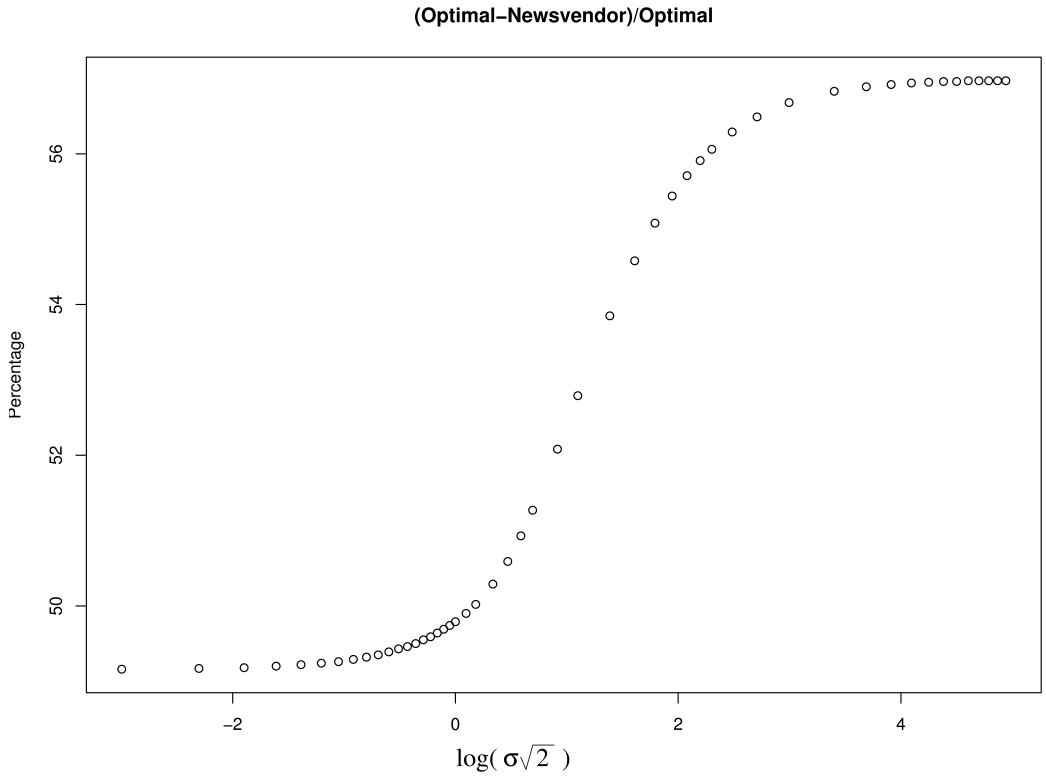

Figure 3.6 Scenario 2, percentage difference versus change in price volatility.

For all three scenarios, the advantage of the optimal policy over the modified newsvendor policy increases rapidly then slowly saturates at around 55-60% as σ increases. To better understand the policy leading to these results, we redraw figures 3.1 and 3.3 for different price volatilities.

In figure 3.8, we observe that both F^{-1} and G^{-1} shift to the right as σ increases. This suggests that as price volatility increases, the spot market price at which the manager should take action increases, enabling the manager to take advantage of this increased volatility. In addition, as volatility increases, F^{-1} shifts upwards and G^{-1} shifts downwards so that the "no action" region increases in size, suggesting more conservative behavior on the part of the manager. Similarly, in figure 3.9, G^{-1} shifts in a similar way as the volatility increases, again suggesting more "conservative" purchasing as volatility increases (that is, smaller amounts bought for a given price). However, F^{-1} shifts up and to the left as volatility increases, implying that as volatility increases, the manager in this case should wait until the price process drops further down and make fewer adjustments as volatility increases.

3.7 Summary

We have completely characterized the optimal material procurement policy for a pharmaceutical firm facing a random demand after the new drug approval, the manager is able to buy and sell on the spot market. In computational tests, we observed that this policy

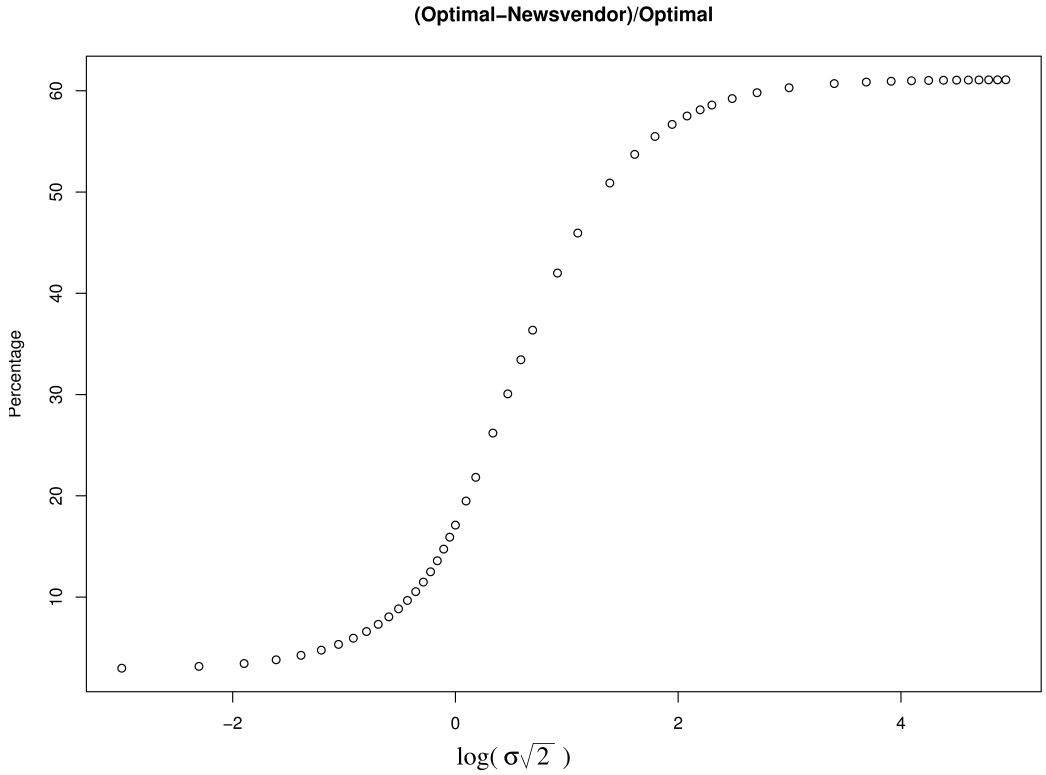

Figure 3.7 Scenario 3, percentage difference versus change in price volatility.

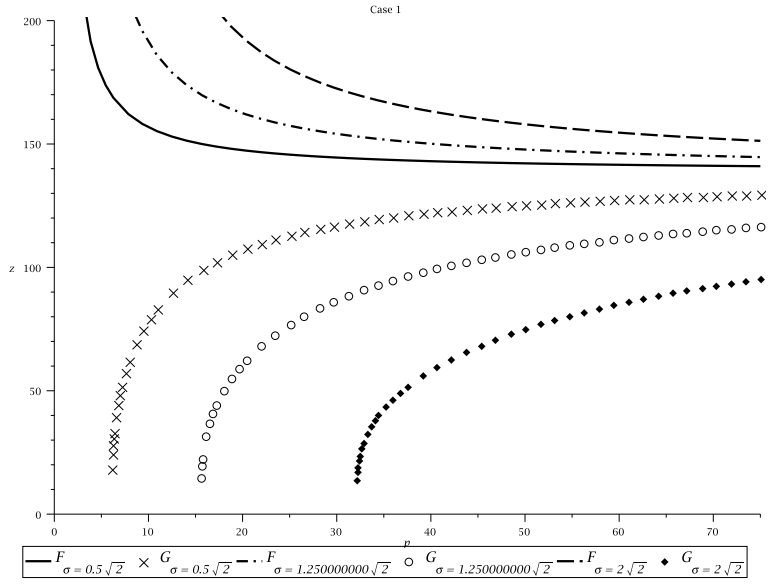

Figure 3.8 Figure 3.1 redrawn for different price volatilities.

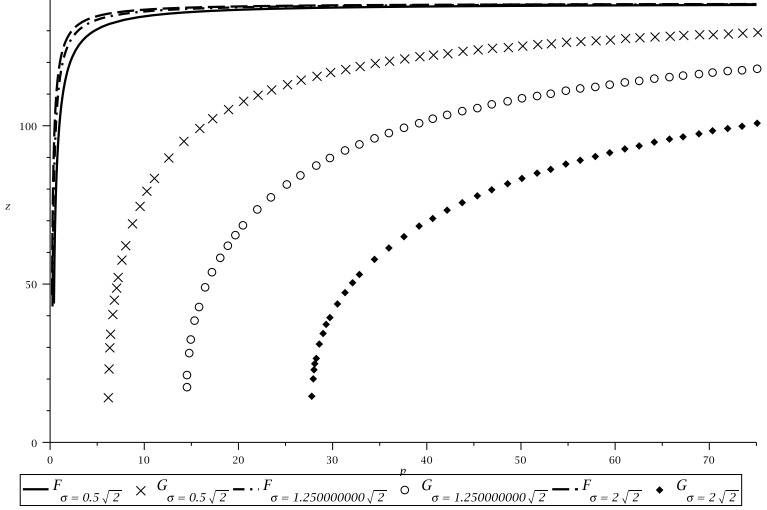

Figure 3.9 Figure 3.3 redrawn for different price volatilities.

performs significantly better than a version of the traditional newsboy policy (utilized as a proxy for a reasonable inventory management policy for firms not interested in trading on the spot market), most notably when price volatility is relatively high.

Although the results presented provide insight into the value to a firm of effectively utilizing the spot market, and contribute to the state of the art in continuous time inventory control, there are significant extensions possible to this work from both technical and modeling perspectives.

For example, the addition of a fixed inventory ordering cost changes the singular control to a more difficult impulse control problem. It will be interesting to see if the analogous state dependent version of (d, D, U, u) policy still holds for this two-dimensional problem, and whether the regularity property holds as well. Additionally, the price process could be modeled by stochastic processes other than a Brownian motion. For instance, it would be interesting to explore whether or not the two-threshold policy holds for the case of a mean-reverting process. More sophisticated constraints on the inventory, such as a requirement that the inventory either be 0 or above some minimum level, are in theory not more difficult by our solution approach, but it would be interesting to complete this analysis. Finally, multi-period models, with multiple demand opportunities and inventory carried between periods will be significantly more difficult to analyze, but may yield interesting insights on effective inventory management in the presence of a spot market.

Chapter 4

Investing in Specialized Equipment – Production Facility Investment with Trial Result Updates

4.1 Introduction

We consider a stylized model of a firm that receives a stream of information updates while making a series of capacity investment decisions. The model applies to any firm that completes a series of experiments to determine if a new product should be introduced, but in light of our motivating problem we refer to these tests as experiments or test cases that make up clinical trials, and the product as a drug. Although the actual clinical trial processes in this industry are quite complex –with multiple phases, multiple endpoints as performance measures, etc.– for analytical tractability, and to develop insight, we specifically focus on a simplified version of a firm receiving clinical trial information and making pharmaceutical production capacity investment decisions.

We assume that there is essentially one (or one remaining) type or phase of trial, made up of a series of independent tests or experiments with the same unknown success rate. A firm must complete this series of tests over a time horizon with T discrete periods in order to determine the effectiveness of a product. At the start of period $T + 1$, if the clinical trial sufficiently demonstrates the efficacy of the drug, the firm has to produce to satisfy demand, which we assume to be a known constant rate D per unit time since our goal here is to focus on the pre-approval capacity acquisition process. The firm's goal is to build sufficient production capacity to support this demand rate D if the drug is proven effective. If the clinical trial is not successful, the product will not be produced and the production capacity is not needed. We assume building sufficient capacity takes an exogenously determined number of periods, that the firm can revisit its capacity building decision at the start of each period, and that the result of a single case or experiment is revealed at the end of each of the first T periods.

We model the outcome of clinical cases or experiments as independent identical Bernoulli random variables. We motivate this by observing that before a new drug can be marketed, its manufacturer must demonstrate that its safety and efficacy via clinical trials or tests on individual patients. We let p be the probability that the drug successfully treats an individual patient in the trial, so that each result represents the outcome of testing on a specific patient. We assume that the Bernoulli rate p is unknown to the firm. This differs from typical assumptions in the real options and capacity investment literature, but better captures the nature of uncertainty in clinical trials. In our model, instead of using a fixed estimate of the unknown p throughout the capacity planning process, the firm updates the estimate of p as the trial progresses.

Regulatory agencies typically require a conventional *frequentist* approach to assessing the safety and efficacy of new drugs. Nevertheless, the firm typically has access to a large amount of preliminary data (from lab results, early trial phases, expert opinions, etc.), which it has the freedom to incorporate into its decision-making process when it makes other decisions that depend on clinical trial results. In particular, to make capacity-related decisions, the firm is free to use a Bayesian approach that utilizes this preliminary data to estimate a prior distribution of the success rate, and then in each period updates this estimate in the form of a posterior distribution that incorporates new results. Specifically, in our model we make use of the fact that the Beta distribution is the conjugate prior of the Bernoulli distribution (that is, if the prior distribution of the success rate p is a Beta, then the posterior distribution of p is also Beta), and specify the underlying stochastic process of our model as follows:

4.2 The Model

?? A pair of parameters (γ^0, ζ^0) model any preliminary data about the potential success of the trial (or more specifically, about the success rate p). The prior distribution of the success rate p is thus Beta(γ^0, ζ^0), where $\gamma^0 = \zeta^0 = 1$ means the firm does not use preliminary data, so that the prior distribution of p is Uniform$(0, 1)$, a special case of the Beta distribution with $\gamma^0 = \zeta^0 = 1$. The pair (γ^0, ζ^0) determines the shape of density function Beta(γ^0, ζ^0) over the $(0, 1)$ interval, where the density function has a negative (positive) skew if $\gamma^0 > (<)\zeta^0$ and the variance of the density function decreases in $\gamma^0 + \zeta^0$.

At the end of each of the first T periods (after the investment decision is already made for that period), the firm observes an experimental result, represented by the binary random variable δ^t, where 1 means a successful experiment, and 0 an unsuccessful one. Thus, the stochastic process

$$\gamma^{t-1} := \gamma^0 + \sum_{i=1}^{t-1} \delta^i,$$

records the number of successful cases observed up to period $t \in \{1, 2, \ldots, T\}$ (but not including the result, δ^t) adjusted for preliminary data by γ^0. The observation γ^{t-1} is then the most up-to-date information the firm has when the decision is made for period t. We

also define ζ^t:

$$\zeta^{t-1} := \zeta^0 + \sum_{i=1}^{t-1}(1-\delta^i),$$

to count the total number of failed experiments observed (adjusted for preliminary data by adding ζ_0), such that

$$\zeta^{t-1} := \gamma_0 + \zeta_0 + (t-1) - \gamma^{t-1}.$$

Therefore, the posterior distribution of the success rate p at the beginning of period t is Beta$(\gamma^{t-1}, \zeta^{t-1})$. Since the posterior estimate of p is given by:

$$\frac{\gamma^0 + \sum_{i=1}^{t-1}\delta^i}{\gamma^0 + \zeta^0 + t - 1}$$

the larger the magnitude of the pair (γ^0, ζ^0), the stronger the influence of the preliminary data on the estimate of the success rate, p, in each future period t. In our model, we specifically assume that a large (γ^0, ζ^0) implies more, and thus higher quality, initial data.

After T periods, the product is considered sufficiently effective if the total number of successful experiments, γ^T, meets the exogenously specified minimum requirement on the total number of successful experiments, γ'^c (so that $\{\sum_{i=1}^{T} \geq \gamma'^c\}$). To simplify the exposition, since γ^0 is known, we define $\gamma^c := \gamma'^c + \gamma^0$ so that the event

$$\{\gamma^T \geq \gamma^c\}$$

implies the drug has successfully demonstrated safety and efficacy.

When building its capacity, the firm (in our most general model) has I possible types of facilities to select among in each period, and each facility type $i \in I$ requires s_i^0 (not necessarily consecutive) time periods to complete construction. The vector $\mathbf{s^{t-1}}, t = 1, 2, \ldots, T+1$ is an I component vector, where the i^{th} component represents the remaining number of periods required to complete facility type i at the start of period t. Note that $\mathbf{s^T}$ contains information about how much if any additional time is required to complete capacity at the time when this capacity might be needed. [1] Therefore, s_i^0 represents the total required construction time for building facility type i. Note that we assume each project can be built within the planning horizon T so that $s_i^0 \leq T$ for each $i \in I$.

The cost for one period of construction of a facility type $i \in I$ in period t is c_i for each $i \in I$. For some of what follows, we consider the impact of a setup cost – an additional cost paid when construction of a facility takes place during period t, but didn't take place during period $t-1$. If this is the case, the setup cost for facility type i is $K_i \geq 0, i \in I$.

At the start of each period, a capacity investment decision (that is, the type of capacity to invest in for the next period, if any) must be made for that period. Once an investment decision is made, the corresponding investment cost is incurred, and one period's worth of construction is completed. At the end of the horizon, in period $T+1$, if the total number

[1] We use boldface to denote vectors, and subscript to denote the components of vectors, so that, for example, s_i^t is the i^{th} component of vector $\mathbf{s^t}$.

of successes γ^T is less than γ^c, the product is not produced and sold. Otherwise, if the production capacity is not yet ready, the firm pays to complete construction, and in addition pays a positive penalty cost $\Pi(s)$ (intended to model the delay in receiving revenue from the product, loss of patent life, loss of revenue, loss of goodwill, etc.) which is a function of the number of periods s required to complete construction. The objective is to develop a capacity investment strategy that minimizes the total sum of discounted expected cost of building the necessary capacity and expected penalty cost.

To model this as a stochastic dynamic program, we number periods forward from 1 to $T+1$. At the start of each period t, $t = 1, 2, ..., T$, the firm observes the remaining construction time for each of the potential projects $\mathbf{s^{t-1}}$ (where if a project has not been started, the remaining construction time is equal to the full construction time), the number of successful experiments up to and including the result that arrived at the end of the previous period γ_{t-1}, and the construction completed in the previous period. To represent the firm's investment decision in period $t-1$, we define an I component vector $\mathbf{a^{t-1}}$ where if the firm elects to build facility type $i, i \in I$ at time $t-1$, the i^{th} component of $\mathbf{a^{t-1}}$ will equal 1; all other vector components will be 0. If the firm is potentially selecting between multiple project types, we assume that the firm can invest in only one project type in any period, so that $\mathbf{a^t} \in \{0,1\}^I$ and $\sum_i a_i^t \leq 1$ for $t = 1, 2, \ldots, T$. We assume that the firm has no ongoing capacity investments before period $t = 1$, so that $\mathbf{a^0}$ is a zero vector. We let $\Omega_t, t = 1, 2, \ldots, T$ denote all admissible investment strategies at time t. At the end of the period, after completing construction, the firm observes the next experimental result δ^t.

We formulate the dynamic program for periods $t = 1, 2, \ldots, T$ below:

$$V_t(\gamma^{t-1}, \mathbf{s^{t-1}}, \mathbf{a^{t-1}})$$
$$= \min_{\mathbf{a^t} \in \Omega_t} \left\{ \sum_{i=1}^{I} \left[a_i^t c_i + K_i \cdot (a_i^t - a_i^{t-1})^+ \right] + \alpha E \left[V_{t+1}(\gamma^{t-1} + \delta^t, \mathbf{s^{t-1}} - \mathbf{a^t}, \mathbf{a^t}) \middle| \gamma^{t-1} \right] \right\}, \quad (4.1)$$

where $\alpha \in [0,1]$ is the discount factor. In period $T+1$, the firm observes the final experimental result, and thus one of two possible outcomes: either the product is approved and goes into production so the firm pays a penalty if production capacity is not yet ready, or the product is not approved and is not produced. This is captured in the following terminal value function:

$$V_{T+1}(\gamma^T, \mathbf{s^T}, \mathbf{a^T}) = \begin{cases} F(\mathbf{s^T}, \mathbf{a^T}) & \text{if } \gamma^T \geq \gamma^c \\ 0 & \text{otherwise} \end{cases} \quad (4.2)$$

where

$$F(\mathbf{s^T}, \mathbf{a^T}) = \min_{1 \leq i \leq I} \left\{ s_i^T c_i + \Pi(s_i^T) + K_i \cdot (1 - a_i^T)^+ \right\}. \quad (4.3)$$

Unless specified, we assume a constant per period penalty so that $\Pi(s) := \pi s$ where π is a positive constant.

We explore the implications of the firm's capacity investment problem. We begin in the next section with a preliminary analysis of the evolution of the stochastic process γ^{t-1}. In Section 4.4, we characterize the optimal investment strategy in a variety of settings of increasing complexity. We initially consider a single capacity type with constant per period

penalty, and then settings with setup costs, settings where penalty cost is an increasing convex function of s^T, and settings with alternative technology options. In Section 4.6, we utilize a computational study to develop additional insights into appropriate capacity investment strategies for firms in these setting.

4.3 Preliminaries

The stochastic process γ^{t-1} is the key to developing the investment strategy for building the appropriate capacity in this model. It yields the posterior distribution of the unknown success rate of the Bernoulli experiments, and its terminal value γ^T indicates whether or not the firm experiences demand at rate D at the conclusion of the clinical trial. We consider the transition probability given in the following lemma.

Lemma 4.3.1 *γ^{t-1} is a Markov process and its transition probability is given as follow:*

$$Pr\left(\delta^1 = 1 \middle| \gamma^0\right) = \frac{\gamma^0}{\gamma^0 + \zeta^0}, \tag{4.4}$$

for $t \in 2, 3, \ldots, T$:

$$Pr\left(\sum_{j=1}^{m} \delta^{t-1+j} = k \middle| \gamma^{t-1}\right) = \frac{\binom{m}{k} \prod_{j=0}^{k-1}(\gamma^{t-1}+j) \prod_{j=0}^{m-k-1}(\zeta^{t-1}+j)}{\prod_{j=0}^{m-1}(\gamma^{t-1}+\zeta^{t-1}+j)}, \tag{4.5}$$

for $k \in \{0, 1, \ldots, m\}$.

The Markov property results from the assumptions that the clinical experiments are independent and identical. Thus, the order in which the experiments are revealed does not matter and each result has equal weight in the posterior distribution of p, so γ^{t-1} is the sufficient statistic for the upcoming transitions. Observe in equation (4.5) that the impact of marginal results on the transition probability decreases over time. This is not surprising – as the end of the horizon approaches, it is natural for the firm to have a better sense of whether or not the experiment will be a success. Thus, in contrast to the typical assumptions of dynamic programming models, our process γ^t is time non-homogeneous, with a state space increasing in time. However, the planning horizon in our model is finite, and the overall success of the clinical trial, which in turn give raise to the demand D, is revealed at a predetermined period T. The firm's strategy in building the capacity depends heavily on its estimate of the likelihood of the event $\{\gamma^T \geq \gamma^c\}$. The following lemma is the key to our subsequent analysis:

Lemma 4.3.2 *$Pr(\gamma^T \geq \gamma^c | \gamma^{t-1} = \gamma)$ is an increasing function of γ.*

This is intuitive for two reasons. First, more successes observed up to time t implies that the firm's estimate of p is higher, which implies a higher expected likelihood of more successes in the remaining test cases. Second, a large γ^t means fewer successful observations are required in the remaining $T - t$ cases in order to meet the exogenous requirement γ^c. The following lemma is an immediate consequence of Lemma 4.3.2.

Lemma 4.3.3 $Pr(\gamma^T \geq \gamma^c | \gamma^t = \gamma + 1) \geq Pr(\gamma^T \geq \gamma^c | \gamma^{t-1} = \gamma)$.

In other words, a successful result improves the chance of meeting γ^c. Similarly, we have:

Lemma 4.3.4 $Pr(\gamma^T \geq \gamma^c | \gamma^t = \gamma) \leq Pr(\gamma^T \geq \gamma^c | \gamma^{t-1} = \gamma)$.

Thus, an unsuccessful trial decreases the chance of meeting γ^c.

4.4 Investment Strategies

Next, we characterize the optimal investment strategies in various settings. Recall that we assume it takes longer to complete the clinical trial than it does to build each type of production facility, so that $T \geq s_i^0$ for all $i \in I$. Nevertheless, given the increasing quality of information, the firm should intuitively delay production as long as there is no risk of delaying production. Indeed, by employing the appropriate coupling technique, this intuition can be proven:

Lemma 4.4.1 *Define* $s^0_{\max} = \max\{s_i^0 | i \in I\}$. *Suppose* $T > s^0_{\max}$, *then*

$$V_1(\gamma^0, \mathbf{s^0}, \mathbf{a^0}) = \alpha^{T - s^0_{\max}} E\left[V_{T - s^0_{\max} + 1}(\gamma^{T - s^0_{\max}}, \mathbf{s^0}, \mathbf{a^0}) \middle| \gamma^0 \right].$$

Hence, it is optimal for the firm to delay production until period $t = T - s^0_{\max}$.

Thus, for any starting period $t - s^0_{max}$, the firm can write an equivalent problem with a planning horizon of $s^0_{\max}$ periods by setting $(\gamma'^0, \zeta'^0) = (\gamma^{T - s^0_{\max}}, \zeta^{T - s^0_{\max}})$. Hence, without loss of generality, we assume $s^0_{\max} = T$ in our subsequence analysis.

In the following subsections, we characterize the firm's optimal investment strategies under various settings of the dynamic program (that is, of equations (4.1) and (4.2)). We begin with a relatively straightforward setting, and then progressively consider more complicated settings.

4.4.1 One Capacity Type/Stationary Cost/No Fixed Cost

We first model the case where the firm is considering one production facility type, with per period construction cost, $c > 0$. The firm can start, pause and restart construction in each of the first T periods without incurring additional setup costs, so that $K = 0$. If at the end of the horizon the new drug is approved but the production facility is not yet ready, the firm is then charged a fixed per-period-delayed penalty cost, $\Pi(s) = \pi s$ for each of the periods needed to complete construction. Note that without any setup cost, the firm's decision at period t is independent of its investment history at $t-1$. In this setting, we do not have to consider the decision in the previous period, so:

$$V_t(\gamma^{t-1}, s^{t-1}, a^{t-1}) = V_t(\gamma^{t-1}, s^{t-1}).$$

Thus:

$$V_t(\gamma^{t-1}, s^{t-1}) = \min\{c + \alpha E[V_{t+1}(\gamma^t, s^{t-1} - 1)|\gamma^{t-1}], \alpha E[V_{t+1}(\gamma^t, s^{t-1})|\gamma^{t-1}]\},$$

with

$$V_t(\gamma^{t-1}, 0) = 0,$$
$$V_{T+1}(\gamma^T, s^T) = \begin{cases} s^T \cdot (c+\pi) & \text{if } \gamma^T \geq \gamma^c, \\ 0 & \text{otherwise} \end{cases} \tag{4.6}$$

for $t = 1, 2, \ldots, T$. Since the penalty cost incurred at $T+1$ is linear to the number of periods that production will be delayed s^T (or in other words, the number of periods required to complete production):

$$\begin{aligned} V_t(\gamma^{t-1}, s^{t-1}) &= \min\{c + \alpha E[V_{t+1}(\gamma^t, s^{t-1} - 1)|\gamma^{t-1}], \alpha E[V_{t+1}(\gamma^t, s^{t-1})|\gamma^{t-1}]\} \\ &= \min\{c, \alpha^{T-t+1}(c+\pi)Pr(\gamma^T \geq \gamma^c|\gamma^{t-1})\} + \alpha E[V_{t+1}(\gamma^t, s^{t-1} - 1)|\gamma^{t-1}]. \end{aligned}$$

Hence, it is optimal for firm to invest in construction in period t if

$$c \leq \alpha^{T-t+1}(c+\pi)Pr(\gamma^T \geq \gamma^c|\gamma^{t-1}). \tag{4.7}$$

In other words, the firm's decision involves choosing between the cost of investing at the current period t and the expected marginal penalty incurred by delaying the investment.

Observation 1 *Observe in (4.7) that the time to-build-parameter, s^t, is not a factor in the investment decision. By Lemma 4.4.1, we can assume without loss of generality that $s^0 = T$, so $s^t \geq T - t$. Therefore, the decision is monotonic in s^t; that is, if it is optimal for the firm to invest when the time-to-build $s^t = s$, it is also optimal for the firm to invest at any time-to-build $s^t \geq s$.*

Observation 2 *We are able to compute the optimal investment strategy without explicitly solving the dynamic program. This is a consequence of the fact that the penalty term $s^T(c+\pi)$ is linear in the time-to-build term s^T.*

We use inequality (4.7) to characterize the optimal policy in this case.

Theorem 4.4.2 *Given a facility to build with constant per unit penalty costs and no setup cost, at each period* $t \in \{1, 2, \ldots, T\}$*, there exists a* $\gamma_*^{t-1} \in \{\gamma^0, \gamma^0 + 1, \ldots, \gamma^0 + t - 1\} \cup \{\infty\}$ *such that the optimal construction strategy is given by:*

$$a_*^t = \begin{cases} 1, & \text{if } \gamma^{t-1} \geq \gamma_*^{t-1} \\ 0, & \text{otherwise} \end{cases}$$

where

$$\gamma_*^{t-1} := \inf \left\{ \gamma \in \{\gamma^0, \ldots, \gamma^0 + t - 1\} \,\middle|\, Pr(\gamma^T \geq \gamma^c | \gamma^{t-1} = \gamma) \geq \frac{c}{\alpha^{T-t+1}(c + \pi)} \right\}.$$

[2]This result is a direct consequence of Lemma 4.3.2. This threshold policy indicates that in this setting and given the problem parameters, the optimal investment strategy depends only on t and γ^{t-1} , is monotonic in γ^{t-1}, and is independent of the current progress s^{t-1}. Because of the linearity of the penalty function, the firm can compare the construction cost and the discounted expected marginal penalty directly. Moreover, Lemma 4.3.2 implies that a higher γ^{t-1} implies a higher likelihood that the new drug will be sold in the market, so that the firm is more likely to be penalized if the production capacity is not ready by time $T + 1$. Thus, if is optimal for the firm to invest if $\gamma^{t-1} = \gamma$ for some γ, it is also optimal to invest if $\gamma^{t-1} = \gamma + 1$. Thus,

Corollary 4.4.3 *Suppose* $\gamma^{t-1} \geq \gamma_*^{t-1}$ *and* $\gamma^t = \gamma^{t-1} + 1$*, then* $\gamma^t \geq \gamma_*^t$ *for each* $t \in \{1, 2, \ldots, T\}$*.*

In other words, if it is optimal for firm to invest in construction in period t and the firm sees a success from latest test result at the end of period t, then it is also optimal for the firm to invest in period $t + 1$. Moreover:

Corollary 4.4.4 *For* $\alpha = 1$ *and* $t \in \{2, \ldots, T\}$*, if* $\gamma_*^{t-1} < \infty$*, then* $\gamma_*^t \geq \gamma_*^{t-1}$

However, this may not be true when the discount rate α is less than 1. Consider the following counterexample:

Example 1 Consider the setting in which the firm has preliminary data $\gamma^0 = \zeta^0 = 5$, the adjusted threshold, $\gamma^c = 24$, there are 48 periods in the planning horizon, the per period building cost $c = \$10$ Million, the per period penalty cost $\pi = \$100$ Million, and the per period discount rate $\alpha = 0.9$. In this case, the threshold is *decreasing* in periods 26, 27 and 28 (see Figure 4.1).

There is no monotonicity in this case because the firm is comparing the investment cost and the discounted expected marginal penalty, and for a long enough planning horizon, the discount factor, which is exponential in time, makes the investment in the current period less attractive than in subsequent ones.

[2]Here, we use $t - 1$ as the time index to match the realization of the stochastic process, γ^{t-1} at period t.

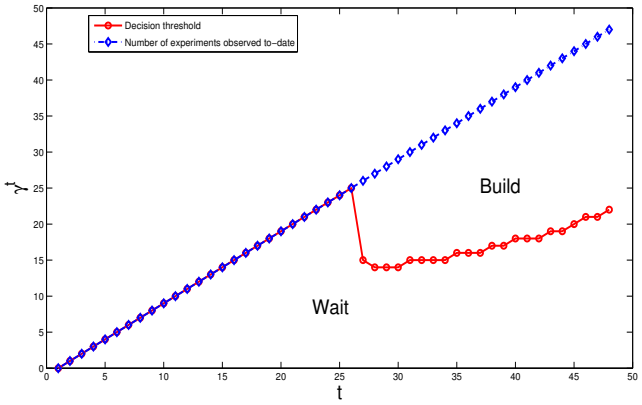

Figure 4.1 Non-monotonic investment threshold with $\alpha < 0$.

4.4.2 Positive Setup Cost, $K > 0$

It can often be expensive to start or restart a construction project due to hiring, resource procurement, etc. To explore the implications of these additional costs, we extend the model from the previous subsection by adding a positive setup cost $K > 0$ that is incurred when a project is started or restarted after construction has not taken place in the previous period.

In this setting, the optimal investment strategy in period t is characterized by a pair of thresholds depending on whether or not the construction project was underway in the previous period:

Theorem 4.4.5 *In the setting described above, with linear terminal penalty function and a single capacity type with positive setup cost $K > 0$ charged when construction is started or restarted after an idle period, for each $t \in \{1, \ldots, T\}$ there exists a pair of threshold values $\gamma_{*0}^{t-1}, \gamma_{*1}^{t-1} \in \{\gamma^0, \gamma^0 + 1, \ldots, \gamma^0 + t - 1\} \cup \{\infty\}$:*

$$\gamma_{*0}^{t-1} := \inf \left\{ \gamma \in \{\gamma^0, \ldots, \gamma^0 + t - 1\} \,\middle|\, Pr(\gamma^T \geq \gamma^c | \gamma^{t-1} = \gamma) \geq \frac{c + K + \alpha E[W_{t+1}(\gamma^t, s^{t-1} - 1) | \gamma^{t-1}]}{\alpha^{T-t+1}(c + \pi)} \right\},$$

$$\gamma_{*1}^{t-1} := \inf \left\{ \gamma \in \{\gamma^0, \ldots, \gamma^0 + t - 1\} \,\middle|\, Pr(\gamma^T \geq \gamma^c | \gamma^{t-1} = \gamma) \geq \frac{c + \alpha E[W_{t+1}(\gamma^t, s^{t-1} - 1) | \gamma^{t-1}]}{\alpha^{T-t+1}(c + \pi)} \right\},$$

where

$$W_{t+1}(\gamma, s) := V_{t+1}(\gamma, s, 1) - V_{t+1}(\gamma, s, 0). \tag{4.8}$$

such that

- $\gamma_{*0}^{t-1} \geq \gamma_{*1}^{t-1}$,

- *If* $a^{t-1} = 0$, $a_*^t = 1$ *if* $\gamma^{t-1} \geq \gamma_{*_0}^{t-1}$ *otherwise* $a_*^t = 0$, *and*
- *If* $a^{t-1} = 1$, $a_*^t = 1$ *if* $\gamma^{t-1} \geq \gamma_{*_1}^{t-1}$ *otherwise* $a_*^t = 0$.

In other words, there are two thresholds $\gamma_{*_0}^{t-1}$ and $\gamma_{*_1}^{t-1}$ governing the firm's investment strategy, depending on action in the previous period. If the firm invested in construction project in the previous period so that $a^{t-1} = 1$, then the firm should invest in period t if and only if $\gamma^{t-1} \geq \gamma_{*_1}^{t-1}$. Likewise, if the firm did not invest in the previous period, $a^{t-1} = 0$, then the firm should invest in period t if and only if $\gamma^{t-1} \geq \gamma_{*_0}^{t-1}$.

The relationship between the two thresholds $\gamma_{*_0}^{t-1} \geq \gamma_{*_1}^{t-1}$ reflects that fact that the setup cost introduces an entrance fee to the investment, and thus the firm needs to have more confidence in the event $\gamma^T \geq \gamma^c$ before investing if it didn't in the previous period. To compare the solution to the no setup cost case from Subsection 4.4.1 , we observe the following:

Corollary 4.4.6 *Given the same set of parameters as the model characterized in Theorem 4.4.2 with an optimal threshold in period* t *of* γ_*^{t-1}, *then*

$$\gamma_{*_1}^{t-1} \leq \gamma_*^{t-1} \leq \gamma_{0_*}^{t-1}.$$

In other words, the decision threshold level from the no setup cost model is sandwiched between the pair of threshold levels for the model with setup cost.

4.4.3 Increasing Marginal Penalty Cost

Up to now, we have implicitly assumed the penalty cost function models lost sales during the time periods that the commercial production of the new drug is delayed due to lack of production capacity. We therefore modeled it as a constant per unit cost, $\Pi(s) = \pi s$. In some cases, however, we need to account for loss of goodwill, patent life, patients' health, etc. In general, we can consider a penalty cost function $F(s^T = s, a^T = a)$, where the amount penalized is a function of the number of periods s production is delayed. Indeed, in many cases, the longer the production delay, the more damage the firm suffers, so that the marginal penalty cost is an increasing function of the number of periods s required to complete construction after time T, and thus:

$$F(s+1, a) - F(s, a) \geq F(s, a) - F(s-1, a). \tag{4.9}$$

Note that we are considering a discrete time model so $s \geq 0$ takes integer values. Inequality (4.9) implies that $F(s, a)$ is a convex function on the discrete set s.

Theorem 4.4.7 *In the setting described above with a non-negative, increasing, convex penalty function satisfying condition (4.9), zero setup cost, and a single type of construction project,*

then for each $s \geq T - t + 1$ *and* $t \in \{1, 2, \ldots, T\}$*, there exists a threshold value* $\gamma_*^{t-1}(s) \in \{\gamma^0, \gamma^0 + 1, \ldots, \gamma^0 + t - 1\} \cup \{\infty\}$*, such that:*

$$a_*^t(s) = \begin{cases} 1, & \text{if } \gamma^{t-1} \geq \gamma_*^{t-1}(s) \\ 0, & \text{otherwise.} \end{cases}$$

Moreover,

- $\gamma_*^{t-1}(s) \geq \gamma_*^{t-1}(s+1)$,
- $V_t(\gamma^{t-1}, s) - V_t(\gamma^{t-1}, s-1) \geq 0$ *and increases in* γ^{t-1}*, and*
- $V_t(\gamma^{t-1}, s) \geq 0$ *is convex in* s *for all* $t \in \{1, 2, \ldots, T\}$.

Observe that $\gamma_*^{t-1}(s)$ is decreasing in s, which is due to the increasing marginal penalty cost. This suggests that when there is still much to be built and the end of the horizon is approaching, the firm employs an increasingly aggressive investment strategy in order to reduce the duration of the potential delay in production. The expected total cost in an increasing function of s and the differences increases in γ^{t-1} because the expected penalty cost increases in both the remaining time-to-build and the probability of a successful trial.

We prove Theorem 4.4.7 using the fact that our dynamic program (4.1) preserves convexity and positivity of $V_t(\cdot, \cdot)$, which ensure the existence of the decision boundary $\gamma_*^{t-1}(s)$ at each period t and time-to-build s, and the monotonicity of $\gamma_*^{t-1}(s)$ in s. However, the monotonicity of $\gamma_*^{t-1}(s)$ in s does not hold if the penalty function is not convex. For example, consider the following counterexample where $F(s, a)$ is concave in s.

Example 2 Consider an example with the same parameters setting as in Figure 4.1 but with a penalty cost function

$$F(s, a) = \$100 \text{ Million} \times s^{0.5},$$

which is increasing but concave in the discrete set s. At period $t = T$, the firm should only build if $\gamma^{47} \geq 35$ when $1 \leq s \leq 10$, only build if $\gamma^{47} \geq 36$ for $11 \leq s \leq 20$, and not build at all if $s > 20$.

4.4.4 Two Investment Projects

Up to this point, our models have reflected one key trade-off faced by the firm – the trade-off between investing early on to ensure that sufficient production capacity will be available at product launch, and waiting until additional information is collected to avoid unnecessary investment. Increasingly, however, firms are considering alternative types of production capacity that are faster to build, and thus allow firms to delay capacity projects without impacting the availability of the product if it is approved, such as the use of disposable (one-time-use) technologies, or the outsourcing of part of all of the production process.

For example, in bio-pharmaceutical production, disposable tanks and mixers can be utilized that are pre-sterilized, pre-validated and self-contained, significantly reducing capacity acquisition lead time. Clearly, if this capacity is also less expensive than traditional capacity, it dominates the traditional capacity. Typically, however, it costs more than its long-lead time counterpart.

To capture this dynamic, we consider a setting in which our stylized firm has two capacity investment projects to select from, both of which produce the same amount production capacity, but at different costs and with different lead times. Project 1 takes longer to build (s_1^0 periods), but at a lower total cost, while Project 2 can be built in fewer periods ($s_2^0 < s_1^0$ periods) but at a higher total cost. Here, we assume no setup cost. As in Section 4.4.1, the firm faces a fixed per period penalty cost $\Pi(s) = \pi > 0$, and is penalized for the number of periods until production can start after the approval of the new drug:

$$F(\mathbf{s^T}) = \min\{s_1^T \cdot (\pi + c_1), s_2^T \cdot (\pi + c_2)\}. \tag{4.10}$$

From Lemma 4.4.1, we know that the firm will not consider investing in Project 2 until period $t = T - s_2^0$. The problem can be reduced to a two-stage problem, where at each period in Stage 1, $t \in \{1, \ldots, T - s_2^0 - 1\}$, it is feasible to either investment in Project 1, or to not invest, while for each period in Stage 2, $t \in \{T - s_2^0, \ldots, T\}$, it is feasible to invest in either project, or neither project. Although the firm will not invest in Project 2 until the start of Stage 2, the availability of this more expensive alternative may impact the firm's decision making in Stage 1. To characterize the optimal strategy, we will first approach the Stage 2 problem.

The Stage 2 Problem

For each period in the Stage 2 problem, we partially characterize the optimal policy as follows:

Lemma 4.4.8 *For each period t in Stage 2, there exists two boundaries*

$$\phi_1^t(x) = \frac{\pi + c_1}{\pi + c_2}x + k_1^t,$$

$$\phi_2^t(x) = \frac{\pi + c_1}{\pi + c_2}x + k_2^t,$$

$k_1^t \geq k_t^2$, k_1^t increase in t, and k_2^t decreasing in t, such that.

$$V_t(\gamma^{t-1}, s_1^{t-1}, s_2^{t-1}) = V_t(\gamma^{t-1}, s_1^{t-1})$$

if $s_2^{t-1} \geq \phi_1^t(s_1^{t-1})$, and

$$V_t(\gamma^{t-1}, s_1^{t-1}, s_2^{t-1}) = V_t(\gamma^{t-1}, s_2^{t-1})$$

if $s_2^{t-1} < \phi_2^t(s_1^{t-1})$.

In other words, for each period t in Stage 2, we can identify regions in the (s_1^{t-1}, s_2^{t-1})–space where we can reduce the 2 Project model to an equivalent 1 Project model. At each period t in Stage 2, these regions are defined by two parallel lines with slope $(\pi + c_1)/(\pi + c_2)$ as illustrated in Figure 4.2. Above the upper boundary, the firm will not consider building Project 2; below the lower boundary, the firm will not consider building Project 1. In the intermediate region, both projects must be considered. Figure 4.2 illustrates these regions for $t = \{T-4, T-3, \ldots, T+1\}$.

Clearly, the firm's decision depends on the remaining time to complete each of the projects, and the larger the difference, the clearer the decision. If the firm has been progressing on Project 1, it may be able to drop Project 2 from consideration relatively early. Conversely, if the firm has been delaying construction of Project 1, it may reach the point relatively early where there is not value to considering Project 1. In addition, although the "gray area" (where both of the projects are still under consideration) expands as we move further from the terminal period, in our computational experiments (in the next section) the firm rarely switches project in Stage 2. Intuitively, this "gray area" is the region that comes into play if the current data suggests the trial is on the borderline between passing and not passing, but in most of our experiments it was extremely rare for the problem to remain within this region for an extended period of time.

The Stage 1 Problem

In Stage 1, the firm decides whether or not to build Project 1 in each period. The corresponding dynamic program for periods $t \in \{1, \ldots, T - s_2^0 - 1\}$ can be written as:

$$V_t(\gamma^{t-1}, s_1^{t-1}, s_2^0) = \min_{a_1^t \in \{0,1\}} \left\{ a_1^t c_1 + \alpha E\left[V_{t+1}(\gamma^{t-1} + \delta^t, s_1^{t-1} - a_1^t, s_2^0) \middle| \gamma^{t-1} \right] \right\}, \tag{4.11}$$

with terminal value function,

$$V_{T-s_2^0}(\gamma^{T-s_2^0-1}, s_1^{T-s_2^0-1}, s_2^0). \tag{4.12}$$

Although this dynamic programming formulation resembles the problem in Section 4.4.1, it turns out that in general the optimal building decision of the dynamic program (4.11,4.12) is dependent on the time-to-build s_1^{t-1} and is not monotonic in s_1^{t-1}.

Example 3 Consider an example with preliminary data $\gamma^0 = \zeta^0 = 5$, adjusted threshold $\gamma^c = 24$, 48 periods in the planning horizon, and per period building costs for Project 1 and Project 2 of $c_1 = \$10$ Million and $c_2 = \$20$ Million respectively. The per period penalty cost is $\pi = \$100$ Million, and the per period discount rate is $\alpha = 0.99$. The lead time of Project 2 is 36 periods. At period $t = 11$, with $s_1^{10} = 41$ the firm should build at period t if $\gamma^{10} \geq 3$ and with $s_1^{10} = 42$ the firm should only build when $3 \leq \gamma^{10} \leq 7$.

In this case, if the firm does not build for more than 4 periods by $t = 11$ and $\gamma^{10} \geq 7$ implying a good chance of passing the trial, the potential penalty of delaying for more than

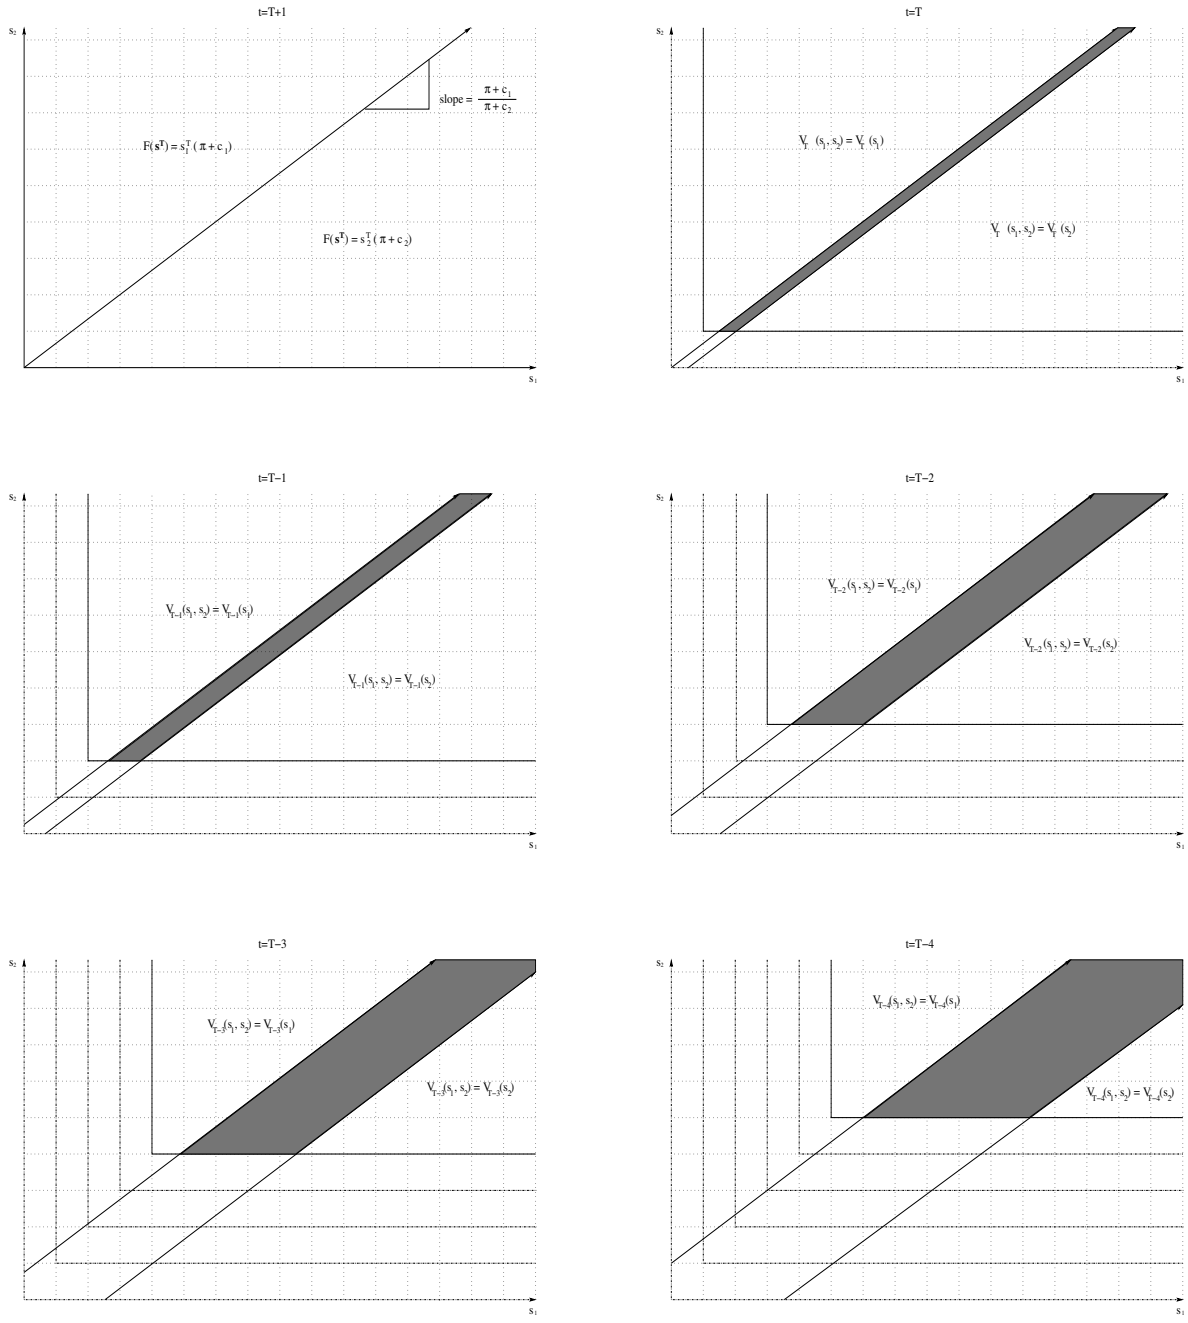

Figure 4.2 The pair of decision boundaries where the 2 project problems can be reduced to a single project problem.

4 periods is too high, so the firm will abandon Project 1, and focus on Project 2 in Stage 2 instead. However, if the firm misses only 3 periods by period $t = 11$, the savings from building the more economical option, Project 1, may compensate for the lost in delaying production, and the firm will build Project 1 for period $t = 11$.

4.5 Multiple Trial Results per Review Period

In previous subsections 4.4.1, 4.4.2 ,4.4.3, and 4.4.4, we assume that the firm makes one observation from the clinical trial in each review period; however, in practice, patients arrive at the treatment centers at random. In this section, we consider the single product, no setup cost and linear penalty cost case, as in section 4.4.1 with one difference: the firm receives multiple individual trial results (modeled by discrete random variables) in each period. The firm has to see at least γ_{crit} number of successes in the first N clinical cases to claim overall success of the clinical trial and file the new drug application.

We let n_t be the number clinical cases revealed at the end of period t, which is unknown to the firm prior to the review at period t. We assume n_t are i.i.d. with discrete distribution $\phi(\cdot)$, and we also assume $\gamma_{t-1} \perp n_j$ for all $j \geq t$. We also let $m_t = \sum_{i=0}^{t} n_i$ such that:

$$\tau := \inf\{t \geq 0 | m_t \geq N\}$$

is the first time there are enough clinical cases seen to ensure the significant level of the clinical trial. For simplicity, we assume the firm only sees the first N results of the trial.

With τ defined, we can rewrite our dynamic program as, for $t = 1, 2, \ldots, \tau$:

$$\begin{aligned} V_t(\gamma^{t-1}, m_{t-1}, s^{t-1}) = \min_{a^t \in \{0,1\}} \big\{ & ca^t + \alpha E\left[V_{t+1}(\gamma^t, m_t, s^{t-1} - a^t) | \gamma^{t-1}, m_t\right], \\ & \alpha E\left[V_{t+1}(\gamma^t, m_t, s^{t-1}) | \gamma^{t-1}, m_t\right] \big\} \end{aligned} \tag{4.13}$$

$$V_{\tau+1}(\gamma^\tau, s^\tau) = \pi(s^\tau)^+ 1_{\{\gamma^\tau \geq \gamma_{crit}\}}. \tag{4.14}$$

Since for each sample path of the trial result updates, the penalty cost is linear, we can rewrite (4.13) as before:

$$\begin{aligned} V_t(\gamma^{t-1}, m_{t-1}, s^{t-1}) = & \min\{c, (\pi + c) E\left[\alpha^{\tau - t + 1} 1_{\{\gamma^\tau \geq \gamma_{crit}, s^{t-1} - 1 > \tau - t\}} | \gamma^{t-1}, m_{t-1}\right]\} \\ & + \alpha E\left[V_{t+1}(\gamma^t, m_t, s^{t-1} - 1) | \gamma^{t-1}, m_{t-1}\right]. \end{aligned}$$

Note:

$$E\left[\alpha^{\tau - t + 1} 1_{\{\gamma^\tau \geq \gamma_{crit}, s^{t-1} - 1 > \tau - t\}} | \gamma^{t-1}, m_{t-1}\right] \tag{4.15}$$

$$= \sum_{j=t}^{s^{t-1}+t-1} \alpha^{j-t+1} Pr(\gamma^j \geq \gamma_{crit} | \tau = j, \gamma^{t-1}, m_{t-1}) Pr(\tau = j | \gamma^{t-1}, m_{t-1}) \tag{4.16}$$

$$= \sum_{k=(\gamma_{crit}-\gamma^{t-1})}^{N-m_{t-1}} \frac{\binom{N-m_{t-1}}{k} \prod_{i=0}^{k-1}(\gamma^{t-1}+i) \prod_{i=0}^{N-m_{t-1}-k-1}(\zeta^{t-1}+i)}{\prod_{i=0}^{N-m_{t-1}-1}(m_{t-1}+i)} \cdot \sum_{j=0}^{s^{t-1}-1} \alpha^{j+1} Pr(\tau - t = j | m_{t-1}), \tag{4.17}$$

where by Lemma 4.3.2 for each fixed m_{t-1}, $E\left[\alpha^{\tau-t+1} 1_{\{\gamma^\tau \geq \gamma_{crit}, s^{t-1}-1>\tau-t\}} | \gamma^{t-1}, m_{t-1}\right]$ is increasing with γ^{t-1}. Therefore, we have the following lemma.

Lemma 4.5.1 *For each m_{t-1} and s^{t-1} there exists a threshold $\gamma^*(m_{t-1}, s^{t-1}) \in \{0, \ldots, m_{t-1}\}$ such that when $\gamma^{t-1} \geq \gamma^*(m_{t-1}, s^{t-1})$ the firm should invest for period t.*

4.6 Computational Study

In the previous sections, we developed a stylized model to characterize decision-making about capacity building decisions as information is updated during a clinical trial. Building on our characterization of optimal facility construction strategies, we present the results of a series of computational experiments to give insight into the following questions:

- How does an approach in which building decisions are regularly re-evaluated as available information changes compare to more traditional approaches? Under what conditions is re-evaluation necessary and beneficial?
- Given that it decides to do so, how frequently should a firm pause and restart building projects?
- How does the quality of preliminary data affect the performance of each approach? What if the firm ignores the preliminary data altogether?
- How do the construction and penalty costs impact these observations?
- Under what condition should the firm choose to build expensive alternative production capacity (such as facilities utilizing disposable technologies)?
- How does the random arrivals of patients at the the treatment centers affect the firm's building decisions and investment performance?

4.6.1 Experimental Design

In our experiments, we compare a variety of strategies to make capacity building decisions. In addition to the **Optimal** approach described, we model several other strategies that a firm might use.

Traditionally, firms decide to build in time to ensure that new production capacity will be ready when the product is approved, estimating the probability of success using initially available data. For our experiments, we model this **Traditional** approach as follows: preliminary data, prior lab results, and expert opinions are collectively modeled by the preliminary data (γ^0, ζ^0) analogous to the preliminary data in our model, and as in our model, we assume that the firm uses Beta(γ^0, ζ^0) to describe the distribution of the true success rate p given the preliminary data. In contrast to the Optimal approach, in our model of the Traditional approach, the firm simply uses this distribution of p, Beta(γ^0, ζ^0), to calculate the probability of a successful trial $\gamma^T \geq \gamma_c$ at the terminal period $T+1$, in order to decide whether or not to commit to building the production capacity. Once the decision is made, the firm does not change its strategy throughout the planning horizon.

In addition to the Optimal and Traditional approaches, we model several other approaches that a firm might use to make capacity investment decisions. For example, a firm might track successful and failed experiments and use this information without updating its assessment of the underlying probability of a successful experiment p. We call this approach the **Intermediate I** approach. In this case, the firm uses $\gamma^0/(\gamma^0+\zeta^0)$ as the point estimate of p at each review period, so that the change in probability of passing the clinical trial is based solely on the change in the number of successes necessary for a clinical trial as the firm sees successful experiments over the horizon. Similarly, a firm may ignore preliminary data (as is typically required by regulatory agencies to assess drug efficacy), and base decisions solely on the results observed in the clinical trial. We call this the **Intermediate II** approach.

For each of our experiments, we consider three performance measures: the expected value of total cost over the horizon, which includes building cost (regardless of whether or not the product is approved) and penalty cost if capacity is not ready in time; conditional expected penalty cost if the product is approved; and conditional expected building cost if the product is not approved. These last two measures in some sense capture the cost of wrong decisions.

For single patient arrival per period, in each set of experimental parameters, we compute the optimal decisions in all relevant states (as well as the decisions using each of our other policies), simulate $100,000$ sample paths of γ^t, and calculate our performance measures for each of our policies. In each of the subsequent subsections, we assess the impact of changing a subset of the problem parameters (such as costs, data quality, etc.) on the performance of each of our approaches. To facilitate our analysis, we employ a reference set of parameters as a starting point, and then vary them as appropriate. These reference problem parameters include a problem horizon (that is, both trial length and construction project length) of 48 months (periods) with a discount rate of $\alpha = 0.99$ per month. The new drug is approved if the number of successful clinical cases exceeds $\gamma^c - \gamma^0 = 40$. It costs the firm \$10 million per month for 48 months to build the production capacity, and upon approval of the new drug, if the production capacity is not yet complete, the firm faces a penalty of $\pi =$ \$100 million for each month of delay.

For random patient arrival, we compute the optimal decision in all relevant states (m_t, γ_t and s), and simulate $50,000$ sample paths of the firms observation of the clinical trial

and calculate the expected costs of the optimal policy. In order to observe the impact of the randomness of the random clinical trial conclusion time toward the firm's investment performance, we use i.i.d. binomial random variables to model the patient arrivals. Such that at each review period, there are a maximum n number of patients the treatment center can receive and each patient will enter the trial and give a result independently with a common probability $\rho \in (0.1]$. Note that the deterministic arrival is a special case of this setting with $\rho = 1$. To ensure the firm will have enough time build the production capacity, we select $N = 48$ and $s = 15$, such that $N > n\rho * s$ with $\gamma_0 = \zeta_0 = 1$. We select four arrival distributions: Binomial($n = 2, \rho = 1$), Binomial($n = 4, \rho = 1/2$), Binomial($n = 6, \rho = 1/3$), and Binomial($n = 8, \rho = 1/4$), to observe the impact the volatility of the number of results revealed at each period on the firms decision. The building cost $c = \$10$ million, the penalty cost $\pi = \$100$ million, and the discount rate $\alpha = 0.99$ are used for the case of random arrival.

4.6.2 Single Project Type Experiments

We have completed a variety of experiments to help answer the questions raised above. For a single project type, we start with the base case described above, and then vary the amount (and thus implicitly, the quality) of preliminary data (Table 4.1), construction costs (Table 4.2), and penalty costs (Table 4.3). Below, we discuss the insights we've developed based on our experiments.

$n_0 = 4$	Traditional	Intermediate I	Intermediate II	Optimal
E[Total Cost]	192.02	177.69	125.47	113.73
E[Penalty$\|\gamma^T \geq \gamma^c$]	217.74	570.06	74.39	71.07
E[Bldg. Cost$\|\gamma^T < \gamma^c$]	106.82	20.87	53.86	39.12
$n_0 = 12$	Traditional	Intermediate I	Intermediate II	Optimal
E[Total Cost]	154.97	126.53	125.47	107.16
E[Penalty$\|\gamma^T \geq \gamma^c$]	210.38	235.40	74.39	61.62
E[Bldg. Cost$\|\gamma^T < \gamma^c$]	63.25	21.24	53.86	33.23
$n_0 = 20$	Traditional	Intermediate I	Intermediate II	Optimal
E[Total Cost]	140.11	115.76	125.47	104.54
E[Penalty$\|\gamma^T \geq \gamma^c$]	164.09	164.45	74.39	59.41
E[Bldg. Cost$\|\gamma^T < \gamma^c$]	53.26	21.80	53.86	30.81

Table 4.1 The change in expected cost with respect to the change in preliminary data quality (in million of dollars).

4.6.3 How does an approach in which building decisions are regularly re-evaluated as available information changes compare to more traditional approaches? Under what conditions is re-evaluation necessary and beneficial?

We see that the firm benefits significantly from re-evaluating capacity construction decisions as additional experimental results are obtained. Depending on problem parameters, we see that on average the Optimal approach costs the firm from 25% to 40% less than the Traditional approach. Clearly, effectively utilizing up-to-date information to make decisions about delaying or expediting construction can have significant benefits. The firm is also penalized significantly less on average when using the Optimal approach. Depending on parameter settings, the firm pays between 34% and 71% less in penalty costs when there is production delay (given a successful trial), and the firm also spends 42% to 63% less in construction cost given that the new drug fails the trial and the facility is not used. In other words, the Optimal approach significantly reduce the risk that the firm faces.

While the Traditional approach is significantly inferior to the Optimal approach, it is also dominated by approaches that use partial information, the Intermediate I and Intermediate II approaches. In general, the Intermediate II approach dominates the Intermediate I approach for most parameter settings, and the firm pays substantially more in penalty cost on average when using the Intermediate I approach. One possible explanation for this behavior is that in the Intermediate I approach the estimate of p is fixed over the planning horizon, and therefore in particular doesn't change in response to an upward trend in the number of successful individual experiments. As a result, the firm is prone to longer delays in completing the production capacity. Although the firm also only partially utilizes information in the Intermediate II approach, it does adapt its estimates based on incoming data.

4.6.4 Given that it decides to do so, how frequently should a firm pause and restart building projects?

One key feature of our proposed Optimal approach is that the firm may decide to pause and restart construction as information is updated. Overall, with our selected set of parameters the firm pauses and restarts construction in about 16% of the samples. When the firm does pause and restart construction, it does so on average 1.8 times over the planning horizon. This tendency to pause and restart is impacted by the quality of preliminary data as we detail in the next subsection (and indeed, the firm tends to change its building decisions more frequently when the quality of preliminary data is low). However, the impact of construction and penalty costs on the gains from stopping and restarting projects is unclear.

4.6.5 How does the quality of preliminary data affect the performance of each approach? What if the firm ignores the preliminary data altogether?

In our model, the magnitude of $n_0 := \gamma^0 + \zeta^0$ represents the quality of the preliminary data, where as n_0 increases, the firm is able to use preliminary data (such as prior lab data, experts opinions, etc) to better estimating the success rate p. This allows us to explore the impact of the quality of this preliminary data on decision-making under various strategies. As the key to our Optimal approach is the updating of estimates of success as additional information is acquired, we would expect to particularly see the advantages of updating these estimates if we have limited preliminary data.

We consider three levels of preliminary data quality, $n_0 = 4, 12, 20$, and summarize performance measures in Table 4.1. As expected, the Optimal approach significantly outperform the other approaches, and the performance differences between our approaches are most prominent when the quality of preliminary data is low.

For instance, when n_0 changes from 20 to 12, the total costs incurred in the Optimal approach are 25% and 31% less than that of the Traditional approach respectively. For the case when the quality of the preliminary data is low, $n_0 = 4$, the Optimal approach saves the firm 40% in total cost over the Traditional approach. The Optimal approach also greatly reduces the risk of delaying the launch of the new drug and the risk of wasting construction cost in case the drug fails. Here, when $n_0 = 4$, the firm pays 67% less in penalty and 63% less in wasted building cost on average relative to the Traditional approach.

Observe that even with limited preliminary data, the Optimal approach still significantly outperforms the Intermediate II approach, in which the firm ignores preliminary data. This suggests that even when preliminary data is limited or of relatively poor quality, it is still valuable for the firm to utilize it in the decision-making process. In the subsequent subsections, we discuss cost and information settings in which re-evaluation is particularly beneficial.

As mentioned above, the firm is more likely to pause and restart construction when the quality of the preliminary data is low. In this case, the firm has only a rough estimate of the success rate p early in the planning horizon, and thus its investment decision is more sensitive to changes in the overall prospect of the drug as marginal individual result are revealed. Specifically, the percentage of runs in which projects are paused and restarted decreases from 21.8% to 15.6% to 13.4% as n_0 increases from 4 to 12 to 20. When the quality of the preliminary data is low, we expect to see the firm change its building decision more frequently. However, given that the firm pauses and restarts construction at least once, the expected frequency of such switches increases from 1.71 times to 1.80 times to 1.83, as n_0 increases from 4 to 12 to 20. This counterintuitive behavior is likely due to the significantly smaller number of times that the firm pauses projects as n_0 increases, as can be seen from the overall unconditioned frequency of switching, which decreases from 0.376, 0.281 and 0.246 as n_0 increases from 4 to 12 to 20. We suspect that as the quality of preliminary information increases, switches only occur in scenarios where switching is more valuable, and thus more

likely.

4.6.6 How do the construction and penalty costs impact these observations?

Ultimately, in this setting the firm is trading off building cost on one hand, and penalty due to production delay on the other. Thus, we explore how the relationship between these costs impacts our policy performance. In our experiments, we consider per period building costs of $c = \$5$ million, \$10 million, and \$15 million; and per period penalty costs of $\pi = \$50$ million, \$100 million, and \$150 million.

The expected total cost increases sub-linearly in the incremental building cost c for each of the approaches. In each of the 4 approaches, the firm balances the expected building cost over the horizon and the expected penalty in case if the production capacity is not ready for the approved drug. As the incremental building cost c increases, the firm becomes more conservative in building capacity. This can also be observed from the change in performance of the firm in the Optimal approach, where the expected construction cost in case of a failed trial is 57% and 97% more as c increases from \$5 million to \$10 million to \$15 million.

The expected total cost also increases sub-linearly in π. As the incremental penalty cost π increases, the firm becomes more liberal in building the capacity across the 4 approaches in general, although the net effect of increasing π on the change in expected penalty cost is unclear. For example, when the Traditional approach is employed, the total expected penalty first increases and then decreases. In the Optimal approach, however, it is clear that the total expected penalty cost indeed decreases in π, and the expected "wasted" building cost increases in π as expected.

Note that while increasing either the building cost or the penalty cost when holding the other constant, the firm adjusts the decision to balance the costs and penalties. For example, when building cost increases and penalty cost is held constant, the firm will be more conservative in its building decision, and thus more willing to risk shortage penalties, and vice-versa, thus explaining the sub-linear increase in expected total cost as these costs increase.

The impact of changing construction and penalty costs on the likelihood of pausing and restarting construction is unclear. When the per period construction cost c decreases from \$15 million to \$10 million to \$5 million, the firm pauses and restarts the project 14.4%, 15.6% and 15.4% of the time, respectively, and the expected switching frequencies are 0.26, 0.28 and 0.27. Similarly, as penalty costs change from \$50 million to \$100 million to \$150 million, firm pause and restarts 16.8%, 15.6% and 17.2% of the time, with expected switching frequencies 0.32, 0.28 and 0.31, so no significant difference or trends are apparent.

$c = 5$	Traditional	Intermediate I	Intermediate II	Optimal
E[Total Cost]	84.42	74.61	67.09	57.22
E[Penalty$\|\gamma^T \geq \gamma^c$]	57.65	170.50	37.12	28.52
E[Bldg. Cost$\|\gamma^T < \gamma^c$]	48.24	13.01	31.93	21.13
$c = 10$	Traditional	Intermediate I	Intermediate II	Optimal
E[Total Cost]	154.97	126.53	125.47	107.16
E[Penalty$\|\gamma^T \geq \gamma^c$]	210.38	235.40	74.39	61.62
E[Bldg. Cost$\|\gamma^T < \gamma^c$]	63.25	21.24	53.86	33.23
$c = 15$	Traditional	Intermediate I	Intermediate II	Optimal
E[Total Cost]	213.86	173.39	177.33	155.05
E[Penalty$\|\gamma^T \geq \gamma^c$]	210.14	302.57	191.36	100.16
E[Bldg. Cost$\|\gamma^T < \gamma^c$]	94.75	27.52	54.21	41.71

Table 4.2 The change in expected costs with respect to the change in construction cost (in million of dollars).

$\pi = 50$	Traditional	Intermediate I	Intermediate II	Optimal
E[Total Cost]	138.55	109.67	111.87	99.84
E[Penalty$\|\gamma^T \geq \gamma^c$]	113.93	183.96	128.65	72.11
E[Bldg. Cost$\|\gamma^T < \gamma^c$]	63.50	16.35	30.27	23.44
$\pi = 100$	Traditional	Intermediate I	Intermediate II	Optimal
E[Total Cost]	154.97	126.53	125.47	107.16
E[Penalty$\|\gamma^T \geq \gamma^c$]	210.38	235.40	74.39	61.62
E[Bldg. Cost$\|\gamma^T < \gamma^c$]	63.25	21.24	53.86	33.23
$\pi = 150$	Traditional	Intermediate I	Intermediate II	Optimal
E[Total Cost]	164.26	137.28	131.00	112.41
E[Penalty$\|\gamma^T \geq \gamma^c$]	92.16	290.76	79.74	60.85
E[Bldg. Cost$\|\gamma^T < \gamma^c$]	96.29	24.04	58.27	39.07

Table 4.3 The change in expected costs with respect to the change in penalty cost (in million of dollars).

4.6.7 Alternative Technologies/Capacity Types

To explore the impact of using novel technologies that are more expensive but have reduced capacity availability lead times, we completed a series of experiments considering a variety of low lead time, expensive alternatives to the base case described above:

- 36 month lead time, \$15 million or \$20 million per month to build,
- 24 month lead time, \$30 million or \$40 million per month to build,
- 12 month lead time, \$60 million, \$80 million or \$100 million per month to build.

(s_2^0, c_2)	$(36, 15)$	$(36, 20)$	$(24, 30)$	$(24, 40)$	$(12, 60)$	$(12, 80)$	$(12, 100)$
E[Total Cost]	102.47	124.5	110.42	121.47	96.94	113.45	121.64
E[Bldg. Cost]	92.92	111.5	102.25	109.78	91.14	98.45	111.62
E[Penalty$\|\gamma^T \geq \gamma^c$]	51.82	71.13	45.13	64.23	31.32	82.6	54.51
% 1 only & Success	0.00%	16.16%	0.00%	16.86%	0.00%	13.36%	16.65%
% 1 only & Failure	0.00%	78.35%	0.00%	80.35%	0.00%	32.21%	80.91%
% 2 only & Success	18.43%	0.00%	18.11%	0.00%	18.52%	1.46%	0.00%
% 2 only & Failure	13.23%	0.00%	6.26%	0.00%	2.63%	0.41%	0.00%
% $1 \rightarrow 2$ & Success	0.00%	2.13%	0.00%	1.36%	0.00%	3.35%	1.73%
% $1 \rightarrow 2$ & Failure	0.00%	3.36%	0.00%	1.44%	0.00%	0.91%	0.71%
% Not Building & Failure	68.34%	0.00%	75.63%	0.00%	78.84%	48.31%	0.00%

Table 4.4 The change in expected costs and building strategy with respected to the change in construction cost and build time of the expensive option ($\gamma_0 = \zeta_0 = 1$, in million of dollars).

In addition to performance measures we considered for the single capacity type experiments, we are also interested in understanding when the firm uses the less expensive longer lead time base case project, when the firm uses the expensive alternative, and when the firm starts with the base case project, and then later switches to the more expensive alternative. To isolate the effect of preliminary data, we set $(\gamma_0, \zeta_0) = (1, 1)$. The results are summarized in Table 4.4.

For all three construction lead times, 12, 24, and 36 months, the expensive alternative's role changes as a function of its building cost. When the building cost of this capacity type is low, the firm tends to wait for more information and chooses to build the more expensive capacity rather than the base case project. For example, for the case when $(s_2^0, c_2) = (36 \text{ months}, \$15 \text{ million})$, 32% of the time firm builds capacity during the planning horizon, and when it does, it chooses the expensive option over the base case capacity type. However, when this alternative costs more, \$20 million, the firm chooses the base case most of the time, and the expensive option is only used as a back up. In this case, each time the expensive technology is used, the firm has already made some investment in the basic capacity, and switches to the expensive option later to make up for "lost time" after it pauses construction of the cheaper capacity.

In other words, when the expensive alternative is relatively cheap, the firm is better off paying the premium to build this capacity in order to "buy" more time for the prospect of a final approval to become clearer. Although this option is more expensive than the base case, the firm on average pays less for "mistakes" made in building either too early or too late, and hence, the firm is better off by paying this premium in exchange for better information. Comparing the two settings in which the novel technology takes 36 months to build, we see that the firm is able to make the "right" decision not to build 68% of the time when $c_2 = \$15$ million. In contrast, when $c_2 = \$20$ million the firm has to build the basic capacity earlier on in the horizon (when the firm confidence level on the estimate of p is low) in order to hedge against the possibility of production delay or the need to use the expensive option; as a result, the firm invests in the basic capacity type when there ultimately isn't

any demand 78% of the time.

We see similar results when the expensive alternative takes 24 months and 12 months to build – the availability of an expensive alternative generally improves the solution to this capacity investment problem. However, the benefit of having this alternative capacity type decreases as the construction cost becomes relatively expensive. For example, when $s_2^0 = 12$ months, the expensive alternative is used less as a backup as c_2 increases from \$80 million to \$100 million.

4.6.8 How does the random arrivals of patients at the the treatment centers affect the firm's building decisions and investment performance?

In additional to the likelihood of a successful trial, the estimation of the arrival time of the clinical trial conclusion is important to firm's capacity building decision. In practice, patients arrival at the treatment centers at random, hence, the number of marginal results at each review periods and the drug approval time are not known to the firm ahead of time. At each review period, the firm's decision to build at each period should be a function of both the progress of the trial and the construction in order to finish the production capacity on time, as we have seen in section 4.5. The firm is expected have a greater risk exposure given this additional uncertainty, in this section, we examined the firm's capacity investment performance at the presence of random drug approval time. We assumed the prior information $(\gamma_0, \zeta_0) = (1, 1)$, the building cost $c =$ \$10 million, the penalty cost $\pi =$ \$100 million, the discount rate $\alpha = 0.99$, the construction time $s = 15$ and the trial size $N = 48$. We assume the number of marginal results revealed at each period are i.i.d. Binomial(n, ρ). We look at four settings of (n, ρ): $(2, 1)$, $(4, 1/2)$, $(6, 1/3)$, and $(8, 1/4)$, such that, the expected number of marginal results are identical in each setting, but with different variances.

n	ρ	$E[TotalCost]$	$E[Penalty\|Success]$	$E[Building\|Failure]$
2	1	\$56.51	\$4.92	\$21.28
4	1/2	\$59.76	\$55.72	\$18.68
6	1/3	\$61.36	\$65.81	\$18.15
8	1/4	\$60.38	\$63.39	\$18.71

Table 4.5 The change in expected costs respected to the change in the distribution of patients arrival ($\gamma_0 = \zeta_0 = 1$, in million of dollars).

From Table(4.5), we observed that the production capacity is more expensive when the trial conclusion time τ is random, and the decrease in the expected building cost spend given a failed trial suggests that the firm build more conservatively when the number of results observed at each period is random. Furthermore, the firm is expected to pay more in penalty cost from a delayed production in the cases where the drug is approved when the approval

time is random, but it is uncertain if the increase in the variability of patient arrival at each period has any effect on the firm's performance. This can be explained by the conclusion time τ, as shown in Table(4.6), there is a wide range of possible value τ can take when the patient arrival is random. When τ is small, the firm may not have enough time to react the early approval, and hence, a higher expected penalty cost in case of a successful trial.

n	ρ	$\max(\tau)$	$\min(\tau)$
2	1	24	24
4	1/2	36	16
6	1/3	38	16
8	1/4	38	14

Table 4.6 The change in clinical trial conclusion time with respected to the change in the distribution of patients arrival ($\gamma_0 = \zeta_0 = 1$, in million of dollars).

4.7 Summary & Discussion

We studied the capacity investment problem faced by a pharmaceutical firm that requires new production capacity for a product that is still undergoing testing. Under pressure to recover research and development costs as quickly as possible, firms traditionally make a one-time commitment to build or acquire production capacity so that the new drug can be launched on time. However, it is typically difficult to gauge the likelihood of passing required clinical trials early in the process. We developed a stylized model to study the value of periodically reviewing data from the on-going clinical trial and appropriately adjusting the ongoing capacity acquisition project, and characterized the optimal building strategy under various model settings. Our analysis suggests that in many cases, reevaluation dramatically reduces both the total expected investment cost, as well as the costs associated with the "mistakes" made in building too early or too late, i.e. the investment cost wasted in case of a failed clinical trial and the penalty cost of delayed production.

Of course, this stylized model significantly simplifies both capacity acquisition and clinical trial data collection. Nevertheless, we believe that the key take-away from this research will hold in more complex settings: regularly reviewing clinical data and reassessing ongoing capacity investment projects can significantly reduce the risk associated with these complex capital investment projects. Moreover, in many cases, but not always, it is often optimal for the firm to a premium to build an alternative novel type of production facility if it has significantly shorter construction time.

To better reflect reality, we intend to consider a set of natural extensions to our model. We will consider more realistic models of clinical trials, focusing on the combination of safety and efficacy that is often used in practice to compare alternative and standard treatments in phase three clinical trials; the survival time of trial participants that is often used in measuring the performance of drugs for cancer and other terminal deceases; and the random

arrival of trial participants to test centers and the resultant random treatment schedule. We also plan to consider more realistic and detailed models of capacity acquisition, modeling stages of construction where the specificity of the capacity increases as construction moves forward, so that early in the construction project, the capacity may be appropriate for a variety of possible products.

Chapter 5

Appendices

5.1 Appendix I – Preliminary

Definition 5.1.1 *A switching control $\alpha = (\tau_n, \kappa_n)_{n\geq 0}$ consists of an increasing sequence of stopping times $(\tau_n)_{n\geq 0}$ and a sequence of new regime values $(\kappa_n)_{n\geq 0}$ that are assumed immediately after each stopping time.*

Definition 5.1.2 *A two-state switching control $\alpha = (\tau_n, \kappa_n)_{n\geq 0}$ is called admissible if the following hold almost surely: $\tau_0 = 0$, $\tau_{n+1} > \tau_n$ for $n \geq 1$, $\tau_n \to \infty$, and for all $n \geq 0$, $\kappa_n \in \{0, 1\}$ is $\mathcal{F}_{\tau_n}$ measurable, with $\kappa_n = \kappa_0$ for even n and $\kappa_n = 1 - \kappa_0$ for odd n.*

Proposition 5.1.3 *There is a one-to-one correspondence between admissible switching controls and the regime indicator function $I_t(\omega)$, which is an $\mathbb{F}$-adapted càglàd process of finite variation, so that $I_t(\omega) : \Omega \times [0, \infty) \to \{0, 1\}$, with*

$$I_t := \sum_{n=0}^{\infty} \kappa_n 1_{\{\tau_n < t \leq \tau_{n+1}\}}, \quad I_0 = \kappa_0. \tag{5.1}$$

Definition 5.1.4 *Let $\mathcal{I} = (0, b)$, $y \in \bar{\mathcal{I}}$ be given, and for each $z \in \mathcal{I}$, let $\alpha(z) = (\tau_n(z), \kappa_n(z))_{n\geq 0}$ be a switching control. The collection $(\alpha(z))_{z\in\mathcal{I}}$ is consistent if*

$$\alpha(z) \text{ is admissible for Lebesgue-almost every } z \in \mathcal{I}, \tag{5.2}$$

$$I_0(z) := \kappa_0(z) = 1_{\{z\leq y\}}, \text{ for Lebesgue-almost every } z \in \mathcal{I}, \tag{5.3}$$

and for all $t < \infty$,

$$\int_{\mathcal{I}} (I_t^+(z) + I_t^-(z))dz < \infty, \text{ almost surely, and} \tag{5.4}$$

$$I_t(z) \text{ is decreasing in } z \text{ for } \mathbb{P} \otimes dz\text{-almost every } (\omega, z). \tag{5.5}$$

Here $I_t^+(z)$ and $I_t^-(z)$ are defined by

$$I_t^+ := \sum_{n>0,\kappa_n=1}^{\infty} 1_{\{\tau_n<t\}}, \quad I_0^+ = 0 \qquad \textit{and} \qquad I_t^- := \sum_{n>0,\kappa_n=0}^{\infty} 1_{\{\tau_n<t\}}, \quad I_0^- = 0.$$

5.2 Appendix II – Proofs

Proof of Theorem 3.4.1:

First, given any $(\xi^+, \xi^-) \in \mathcal{A}'_y$, from the integration by parts formula,

$$\mathbb{E}\left[\int_0^\infty e^{-\rho t} C_h Y_t dt\right] = \frac{C_h}{\rho} y + \frac{C_h}{\rho} \mathbb{E}\left[\int_0^\infty e^{-\rho t} d\xi_t^+ - \int_0^\infty e^{-\rho t} d\xi_t^-\right].$$

Next, let $Z_t = e^{-\rho t} P_t$. by Ito's formula, we have

$$dZ_t = -(\rho - \mu) Z_t dt + \sqrt{2}\sigma Z_t dW_t, \quad Z_0 = p > 0.$$

Note that $\rho - \mu > 0$ implies $\mathbb{E}[\int_0^\infty Z_s ds] < \infty$. Moreover $\lim_{t\to 0} Z_t = 0$ almost surely and in L^1. Now, from the integration by parts formula in (Protter 2004, p. 68) and noting that Z and Y are left continuous,

$$\begin{aligned}\int_0^t Z_{s-} dY_s &= Z_t Y_t - \int_0^t Y_{s-} dZ_s - [Z, Y]_t \\ &= Z_t Y_t + (\rho - \mu) \int_0^t Y_s Z_s ds - \sqrt{2}\sigma \int_0^t Y_s Z_s dW_s - py,\end{aligned}$$

because $[Z, Y]_t = Z_0 Y_0 = py$ for all t from the finite variation of Y.

In addition, from (Protter 2004, p. 63), the process $\int_0^t Y_s Z_s dW_s$ is a martingale, so for all $t > 0$,

$$\mathbb{E}\left[\int_0^t Z_s dY_s\right] = \mathbb{E}\left[Z_t Y_t + (\rho - \mu) \int_0^t Y_s Z_s ds\right] - py.$$

And since $Z_t \to 0$ in L^1 and $Y_t \in [0, b]$ is bounded, $\lim_{t\to\infty} \mathbb{E}[Z_t Y_t] = 0$.

Furthermore, since $(\xi^+, \xi^-) \in \mathcal{A}'_y$ and Z_t is integrable, the dominated convergence theorem gives

$$\begin{aligned}\mathbb{E}\left[\int_0^\infty Z_s dY_s\right] &= \lim_{t\to\infty} \mathbb{E}\left[\int_0^t e^{-\rho s} P_s dY_s\right] = \lim_{t\to\infty} \mathbb{E}\left[Z_t Y_t + (\rho - \mu) \int_0^t Y_s Z_s ds\right] - py \\ &= \mathbb{E}\left[(\rho - \mu) \int_0^\infty Y_s Z_s ds\right] - py.\end{aligned}$$

Hence,

$$\mathbb{E}\left[-\int_0^\infty e^{-\rho t} P_t d\xi_t^+ + \int_0^\infty e^{-\rho t} P_t d\xi_t^-\right] = py - \mathbb{E}\left[(\rho - \mu) \int_0^\infty e^{-\rho t} Y_t P_t dt\right].$$

Now, putting all terms together, an easy application of the Dominated Convergence Theorem reveals the statement. □

Proof of Proposition 3.4.2: Let $p > 0$ and $y \in [0, b]$ be given. Since $\rho > \mu$ we have

$$\mathbb{E}\left[\int_0^\infty e^{-\rho t}[\tilde{H}(Y_t)P_t]dt\right] \leq \mathbb{E}\left[\int_0^\infty e^{-\rho t}[MP_t]dt\right] \leq \eta M p.$$

Note that for any given $(\xi^+, \xi^-) \in \mathcal{A}_y$, $-y \leq \xi_t^+ - \xi_t^- \leq b - y$. From integration by parts, for any $T > 0$,

$$-\int_{[0,T)} e^{-\rho t} d\xi_t^+ \leq -\int_{[0,T)} e^{-\rho t} d\xi_t^- + y,$$

which, together with $K^+ + K^- > 0$ and $K^+ > 0$, implies

$$\begin{aligned}&\mathbb{E}\left[-(K^+ + \frac{C_h}{\rho})\int_0^\infty e^{-\rho t}d\xi_t^+ - (K^- - \frac{C_h}{\rho})\int_0^\infty e^{-\rho t}d\xi_t^-\right] \\ &\leq \frac{C_h}{\rho}y - (K^+ + K^-)\mathbb{E}\left[\int_0^\infty e^{-\rho t}d\xi_t^-\right] \leq \frac{C_h}{\rho}b.\end{aligned}$$

Since these bounds are independent of the control, we have

$$V(p, y) \leq \eta M p + \frac{C_h}{\rho} b < \infty.$$

□

Proof of Proposition 3.4.6: The consistency follows immediately from Definition 5.1.4 and the following monotonicity property of F and G: F is non-increasing and G is non-decreasing in Case I, and F is non-decreasing and G is non-increasing in Case II.

Proof of Proposition 3.4.8: Since $(0, b)$ is bounded, the integrability follows (Guo and Tomecek 2008, Theorem 3.10) as

$$\lim_{t\to\infty} \mathbb{E}\left[e^{-\rho t}G^{-1}(M_t)\right] = 0.$$

(See also Lemma 1 and Eqn. (23) in (Merhi and Zervos 2007)). □

Proof of Lemma 4.3.1: By direct calculation,

$$\begin{aligned}&Pr\left(\sum_{j=1}^m \delta^{t-1+j} = k \middle| \gamma^{t-1}, \gamma^{t-2}, \ldots, \gamma^0\right) \\ &= \binom{m}{k} \frac{\Gamma(\gamma^{t-1} + \zeta^{t-1})}{\Gamma(\gamma^{t-1})\Gamma(\zeta^{t-1})} \int_0^1 p^{\gamma^{t-1}+k-1}(1-p)^{\zeta^{t-1}+(m-k)-1}dp \\ &= \binom{m}{k} \frac{\prod_{j=0}^{k-1}(\gamma^{t-1}+j)\prod_{j=0}^{m-k-1}(\zeta^{t-1}+j)}{\prod_{j=0}^{m-1}(\gamma^{t-1}+\zeta^{t-1}+j)} \qquad (5.6)\end{aligned}$$

$$=Pr\left(\sum_{j=1}^{m}\delta^{t-1+j}=k\middle|\gamma^{t-1}\right). \tag{5.7}$$

□

Proof of Lemma 4.3.2:

$$Pr(\gamma^T\geq\gamma^c|\gamma^{t-1}=\gamma)=\frac{\sum_{k=(\gamma^c-\gamma)^+}^{T-t+1}\left(\begin{array}{c}T-t+1\\k\end{array}\right)\prod_{j=0}^{k-1}(\gamma^{t-1}+j)\prod_{j=0}^{T-t-k}(\zeta^{t-1}+j)}{\prod_{j=0}^{T-t}(\gamma^{t-1}+\zeta^{t-1}+j)}. \tag{5.8}$$

With some algebraic manipulation, we have:

$$\begin{aligned}&Pr(\gamma^T\geq\gamma^c|\gamma^{t-1}=\gamma)-Pr(\gamma^T\geq\gamma^c|\gamma^{t-1}=\gamma-1)\\&=\frac{\sum_{k=(\gamma^c-\gamma)^+}^{T-t}\left(\begin{array}{c}T-t+1\\k\end{array}\right)\prod_{j=0}^{k-1}(\gamma^{t-1}+j)\prod_{j=1}^{T-i-k}(\zeta^{t-1}+j)\left[\zeta^{t-1}-(\gamma^{t-1}-1)\frac{T-t-k+1}{k+1}\right]+\prod_{j=0}^{T-t}(\gamma^{t-1}+j)}{\prod_{j=0}^{T-t}(\gamma^{t-1}+\zeta^{t-1}+j)}.\end{aligned}$$

Note that the term

$$(\zeta^{t-1})(k+1)-(\gamma^{t-1}-1)(T-t-k+1)$$

is increasing in k, and changes sign at most once. Combined with the fact that all other terms remain positive, we observe that:

$$\begin{aligned}&Pr(\gamma^T\geq\gamma^c|\gamma^{t-1}=\gamma)-Pr(\gamma^T\geq\gamma^c|\gamma^{t-1}=\gamma-1)\\\geq&\min\{Pr(\gamma^T\geq\gamma^c|\gamma^{t-1}=\gamma^c)-Pr(\gamma^T\geq\gamma^c|\gamma^{t-1}=\gamma^c-1),\\&Pr(\gamma^T\geq\gamma^c|\gamma^{t-1}=\gamma^c-T+t-1)-Pr(\gamma^T\geq\gamma^c|\gamma^{t-1}=\gamma^c-T+t-2)\}\\\geq&0,\end{aligned}$$

and by transitivity, we have:

$$Pr(\gamma^T\geq\gamma^c|\gamma^{t-1}=\gamma_1)\geq Pr(\gamma^T\geq\gamma^c|\gamma^{t-1}=\gamma_2),$$

for all $\gamma_1>\gamma_2$. □

Proof of Lemma 4.3.3: From (4.3.1), we show that γ^{t-1} is a Markov process, and by conditioning on γ^t, we have:

$$\begin{aligned}&Pr(\gamma^T\geq\gamma^c|\gamma^t=\gamma+1)-Pr(\gamma^T\geq\gamma^c|\gamma^{t-1}=\gamma)\\=&Pr(\gamma^T\geq\gamma^c|\gamma^t=\gamma+1)-\sum_{i=0}^{1}Pr(\gamma^T\geq\gamma^c|\gamma^t=\gamma+i)]Pr(\delta^t=i|\gamma^{t-1}=\gamma)\geq 0.\end{aligned}$$

The inequality results from a direct application of Lemma 4.3.2. □

Proof of Lemma 4.3.4: Similar to the proof of Lemma 4.3.3, we have:

$$Pr(\gamma^T\geq\gamma^c|\gamma^t=\gamma)-Pr(\gamma^T\geq\gamma^c|\gamma^{t-1}=\gamma)$$

$$=Pr(\gamma^T \geq \gamma^c|\gamma^t = \gamma) - \sum_{i=0}^{1} Pr(\gamma^T \geq \gamma^c|\gamma^t = \gamma + i)Pr(\delta^t = i|\gamma^{t-1} = \gamma) \leq 0.$$

The inequality again results from a direct application of Lemma 4.3.2. □

Proof of Lemma 4.4.1: Suppose $T > s^0_{\max}$ and consider an admissible strategy $\mathbf{a^1}$ as a potential candidate optimal strategy at period $t = 1$ with $\mathbf{a^1} \neq 0$. We define index I' to identify the entries such that $a^1_i = 1$. We further define a stopping time for each $i \in I'$, $\tau_i := \min\{t|a^t_i = 0\}$, and note that $\tau \leq s^0_{\max} < T$ *a.s.*, since it is a waste of resources to invest in a project which is already completed. We consider an alternative strategy $\mathbf{a'^1}$ such that $\mathbf{a'^1} = 0$, and for $t = 2, \ldots, T$. We couple $\mathbf{a'^t}$ with $\mathbf{a^t}$ for each outcome γ_{t-1}, such that $a'^t_i = a^t_i$ for $i \in I \setminus I'$ and for each $t = 2, 3, \ldots, T$. For facility types $i \in I'$, we couple a'^t_i with a^t_i such that,

$$a'^t_i = \begin{cases} 1 & \text{for } t = \tau_i, \\ a^t_i & \text{otherwise.} \end{cases}$$

Let $J(\gamma_0, \mathbf{s^0}, \mathbf{a^0}, \mathbf{a^1})$ be the expected total cost if the firm uses investment strategy $\mathbf{a^1}$ at period $t = 1$, then

$$\begin{aligned} J(\gamma_0, \mathbf{s^0}, \mathbf{a^0}, \mathbf{a^1}) &\geq J(\gamma_0, \mathbf{s^0}, \mathbf{a^0}, \mathbf{a'^1}) + E\left[\sum_{i \in I'} \sum_{j=1}^{\tau_i} (\alpha^{j-1} - \alpha^j) C_i(s^0_i - j + 1)|\gamma_0\right] \\ &\geq J(\gamma_0, \mathbf{s^0}, \mathbf{a^0}, \mathbf{a'^1}) \\ &\geq V_1(\gamma_0, \mathbf{s^0}, \mathbf{a^0}). \end{aligned}$$

The decision, $\mathbf{a^t}$, indeed gives a suboptimal solution at $t = 1$, therefore by contradiction,

$$V_1(\gamma^0, \mathbf{s^0}, \mathbf{a^0}) = \alpha E\left[V_2(\gamma^1, \mathbf{s^0}, \mathbf{a^0})\middle|\, \gamma^0\right].$$

□

Proof of Theorem 4.4.2: The result is a consequence of Lemma 4.3.2 and (4.7). □

Proof of Corollary 4.4.3: From Lemma 4.3.2, and the monotonicity stated in Theorem 4.4.2:

$$Pr(\gamma^T \geq \gamma^c|\gamma^t = \gamma + 1) \geq Pr(\gamma^T \geq \gamma^c|\gamma^{t-1} = \gamma) \geq Pr(\gamma^T \geq \gamma^c|\gamma^{t-1} = \gamma^{t-1}_*) \geq \frac{c}{\alpha^{T-t}(c + \pi)},$$

therefore, it is also optimal for the firm to invest in period t. □

Proof of Corollary 4.4.4: The corollary is a direct consequence of Theorem 4.4.2 and Lemma 4.3.4. □

Proof of Theorem 4.4.5: Here, we observe that the firm should build capacity for period t if:

$$Pr(\gamma^T \geq \gamma^c|\gamma^{t-1}) \geq \frac{c + K + \alpha E[W_{t+1}(\gamma^t, s^{t-1} - 1)|\gamma^{t-1}]}{\alpha^{T-t+1}(c + \pi)},$$

when $a^{t-1} = 0$; and

$$Pr(\gamma^T \geq \gamma^c|\gamma^{t-1}) \geq \frac{c + \alpha E[W_{t+1}(\gamma^t, s^{t-1} - 1)|\gamma^{t-1}]}{\alpha^{T-t+1}(c + \pi)},$$

when $a^{t-1} = 1$. Since $K > 0$, if we can show that for $W(\cdot, \cdot)$ as defined in equation (4.8), $E[W_{t+1}(\gamma^t, s^{t-1} - 1)|\gamma^{t-1}]$ is a non-increasing function in γ^{t-1} for all t, then for each t there exist a pair of investment thresholds $\gamma^{t-1}_{*_0} \geq \gamma^{t-1}_{*_1}$,

$$\gamma^{t-1}_{*_0} := \inf\left\{\gamma \in \{\gamma^0, \ldots, \gamma^0 + t - 1\}\,\middle|\, Pr(\gamma^T \geq \gamma^c|\gamma^{t-1} = \gamma) \geq \frac{c + K + \alpha E[W_{t+1}(\gamma^t, s^{t-1} - 1)|\gamma^t]}{\alpha^{T-t+1}(c + \pi)}\right\},$$

$$\gamma^{t-1}_{*_1} := \inf\left\{\gamma \in \{\gamma^0, \ldots, \gamma^0 + t - 1\}\,\middle|\, Pr(\gamma^T \geq \gamma^c|\gamma^{t-1} = \gamma) \geq \frac{c + \alpha E[W_{t+1}(\gamma^t, s^{t-1} - 1)|\gamma^t]}{\alpha^{T-t+1}(c + \pi)}\right\},$$

such that for $a^{t-1}=0$ the firm should built at period t if and only if $\gamma^{t-1}\geq\gamma_{*_0}^{t-1}$; and for $a^{t-1}=1$ if and only if $\gamma^{t-1}\geq\gamma_{*_1}^{t-1}$. To show $E[W_{t+1}(\gamma^t,s^{t-1})|\gamma^{t-1}]$ is non-increasing in γ^{t-1} for all $t\in\{1,2,\ldots,T+1\}$, we use induction. For $t=T+1$,

$$W_{T+1}(\gamma^T,s^T)=\begin{cases}0 & \text{for } \gamma^T<\gamma^c,\\ -K & \text{for } \gamma^T\geq\gamma^c.\end{cases}$$

Then by Lemma 4.3.2, we see $E[W_{T+1}(\gamma^T,s^T)|\gamma^{T-1}]$ is non-increasing in γ^{T-1}, and $\gamma_{*_0}^{T-1}\geq\gamma_{*_1}^{T-1}$. Now we suppose $E[W_{t+1}(\gamma^t,s^t)|\gamma^{t-1}]$ is non-increasing in γ^{t-1} at period t and need to show it is also true for period $t-1$. Observe:

$$W_t(\gamma^{t-1},s^{t-1})=\begin{cases}0 & \text{for } \gamma^{t-1}<\gamma_{*_1}^{t-1},\\ c+\alpha E[W_{t+1}(\gamma^t,s^t)|\gamma^{t-1}]-\alpha^{T-t+1}(c+\pi)Pr(\gamma^T\geq\gamma^c|\gamma^{t-1}) & \text{for } \gamma_{*_1}^{t-1}\leq\gamma^{t-1}\leq\gamma_{*_0}^{t-1},\\ -K & \text{for } \gamma^{t-1}\geq\gamma_{*_0}^{t-1}.\end{cases}$$

By Lemma 4.3.2, $E[W_t(\gamma^{t-1},s^{t-1})|\gamma^{t-2}]$ is non-increasing in γ^{t-2}. □

Proof of Theorem 4.4.7: Note that since the setup cost is zero, we rewrite the penalty function as follows:

$$F(s^T,a^T)=F(s^T)$$

The proof proceeds by backward induction:

i. Suppose period $t=T$,

$$\begin{aligned}&V_T(\gamma^{T-1},s^{T-1})\\ =&\min\{c,\alpha(F(s^{T-1})-F(s^{T-1}-1))Pr(\gamma^T\geq\gamma^c|\gamma^{T-1})\}+\alpha F(s^{T-1}-1)Pr(\gamma^T\geq\gamma^c|\gamma^{T-1}).\end{aligned}$$

By Lemma 4.3.2, there exists $\gamma_{*_s}^{T-1}$ for each $s^{T-1}=s$ such that the firm should invest if and only if $\gamma^{T-1}\geq\gamma_{*_s}^{T-1}$.

ii. Moreover, by (4.9), $\gamma_{*_s}^{T-1}$ is decreasing in s.

iii. After some algebraic manipulation, we have:

$$\begin{aligned}&V_T(\gamma^{T-1},s)-V_T(\gamma^{T-1},s-1)\\ =&\begin{cases}\alpha(F(s-1)-F(s-2))Pr(\gamma^T\geq\gamma^c|\gamma^{T-1}) & \text{for } \gamma^{T-1}\geq\gamma_{*_{s-1}}^{T-1},\\ c & \text{for } \gamma_{*_s}^{T-1}\leq\gamma^{T-1}<\gamma_{*_{s-1}}^{T-1},\\ \alpha(F(s)-F(s-1))Pr(\gamma^T\geq\gamma^c|\gamma^{T-1}) & \text{for } \gamma^{T-1}<\gamma_{*_s}^{T-1}.\end{cases}\quad\geq 0.\end{aligned}$$

By Lemma 4.3.2, we also see that $V_T(\gamma^{T-1},s)-V_T(\gamma^{T-1},s-1)$ is increasing in γ^{T-1}.

iv. For the convexity of $V_T(\gamma^{T-1},s)$ in s, we observe:

$$\begin{aligned}&V_T(\gamma^{T-1},s+1)-2V_T(\gamma^{T-1},s)+V_T(\gamma^{T-1},s-1)\\ =&\begin{cases}\alpha[F(s)-2F(s-1)+F(s-2)]Pr(\gamma^T\geq\gamma^c|\gamma^{T-1}) & \text{for } \gamma^{T-1}\geq\gamma_{*_{s-2}}^{T-1},\\ \alpha(F(s)-F(s-1))Pr(\gamma^T\geq\gamma^c|\gamma^{T-1})-c & \text{for } \gamma_{*_{s-1}}^{T-1}\leq\gamma^{T-1}<\gamma_{*_{s-2}}^{T-1},\\ c-\alpha(F(s)-F(s-1))Pr(\gamma^T\geq\gamma^c|\gamma^{T-1}) & \text{for } \gamma_{*_s}^{T-1}\leq\gamma^{T-1}<\gamma_{*_{s-1}}^{T-1},\\ \alpha[F(s+1)-2F(s)+F(s-1)]Pr(\gamma^T\geq\gamma^c|\gamma^{T-1}) & \text{for } \gamma^{T-1}<\gamma_{*_s}^{T-1}.\end{cases}\quad\geq 0.\end{aligned}$$

Therefore, $V_T(\gamma^{T-1},s)$ is convex in s.

Now we suppose that the above statements i, ii, iii and iv hold true for $t+1$ and show the they hold true for t. Observe:

$$V_t(\gamma^{t-1},s)=\min\{c,\alpha E[V_{t+1}(\gamma^t,s)-V_{t+1}(\gamma^t,s-1)|\gamma^{t-1}]\}+\alpha E[V_{t+1}(\gamma^t,s-1)|\gamma^{t-1}].$$

- By Lemma 4.3.2 and $V_{t+1}(\gamma^t, s) - V_{t+1}(\gamma^t, s-1)$ is increasing in γ^t. Therefore, we conclude that there exists a $\gamma_{*_s}^{t-1}$ for each $s^{t-1} = s$, such that the firm should invest if and only if $\gamma^{t-1} \geq \gamma_{*_s}^{t-1}$.
- In addition, by the convexity of $V_{t+1}(\gamma^t, s)$ in s, we see that $\gamma_{*_s}^{t-1}$ is decreasing in s.
- Again, after algebraic manipulation, we have:

$$\begin{aligned} & V_t(\gamma^{t-1}, s) - V_t(\gamma^{t-1}, s-1) \\ = & \begin{cases} E[V_{t+1}(\gamma^t, s-1) - V_{t+1}(\gamma^t, s-2)|\gamma^{t-1}] & \text{for } \gamma^{t-1} \geq \gamma_{*_{s-1}}^{t-1}, \\ c & \text{for } \gamma_{*_s}^{t-1} \leq \gamma^{t-1} < \gamma_{*_{s-1}}^{t-1}, \\ E[V_{t+1}(\gamma^t, s) - V_{t+1}(\gamma^t, s-1)|\gamma^{t-1}] & \text{for } \gamma^{t-1} < \gamma_{*_s}^{t-1}. \end{cases} \quad \geq 0. \end{aligned}$$

- Then by Lemma 4.3.2, we see that $V_t(\gamma^{t-1}, s) - V_t(\gamma^{t-1}, s-1)$ is increasing in γ^{t-1}.
- For the convexity of $V_t(\gamma^{t-1}, s)$ in s:

$$\begin{aligned} & V_t(\gamma^{t-1}, s+1) - 2V_t(\gamma^{t-1}, s) + V_t(\gamma^{t-1}, s-1) \\ = & \begin{cases} \alpha E[V_{t+1}(\gamma^t, s) - 2V_{t+1}(\gamma^t s - 1) + V_{t+1}(\gamma^t, s-2)|\gamma^{t-1}] & \text{for } \gamma^{t-1} \geq \gamma_{*_{s-2}}^{t-1}, \\ \alpha E[V_{t+1}(\gamma^t, s-1) - V_{t+1}(\gamma^t, s-2)|\gamma^{t-1}] - c & \text{for } \gamma_{*_{s-1}}^{t-1} \leq \gamma^{-1t} < \gamma_{*_{s-2}}^{t-1}, \\ c - \alpha E[V_{t+1}(\gamma^t, s-1) - V_{t+1}(\gamma^t, s-2)|\gamma^{t-1}] & \text{for } \gamma_{*_s}^{t-1} \leq \gamma^{t-1} < \gamma_{*_{s-1}}^{t-1}, \\ \alpha E[V_{t+1}(\gamma^t, s+1) - 2V_{t+1}(\gamma^t, s) + V_{t+1}(\gamma^t, s-1)|\gamma^{t-1}] & \text{for } \gamma^{t-1} < \gamma_{*_s}^{t-1}. \end{cases} \quad \geq 0. \end{aligned}$$

□

Prove of Lemma 4.4.8: Here, we calculate k_1^t and k_2^t directly. At $t = T+1$, we see that $k_1^{T+1} = k_2^{T+1} = 0$ from the penalty function (4.10).

At $t = T$, for $s_2 \geq s_1 \frac{\pi+c_1}{\pi+c_2} + 1$, since $E[V_{T+1}(\gamma^T, s_1, s_2 - 1)|\gamma^{T-1}] = E[V_{T+1}(\gamma^T, s_1, s_2)|\gamma^{T-1}] = E[V_{T+1}(\gamma^T, s_1)|\gamma^{T-1}]$, it is not optimal for the firm to build Project 2. Similarly, for $s_2 < (s_1 - 1)\frac{\pi+c_1}{\pi+c_2}$, it is not optimal for the firm to build Project 1. For $s_1 \frac{\pi+c_1}{\pi+c_2} \leq s_2 < s_1 \frac{\pi+c_1}{\pi+c_2} + 1$, the firm should choose Project 2 over Project 1 when

$$s_2 \geq s_1 \frac{\pi + c_1}{\pi + c_2} - \frac{c_2 - c_1}{\pi + c_2} \left(\frac{1}{\alpha Pr(\gamma^T \geq \gamma^c | \gamma^{T-1})} - 1 \right),$$

However, since it is always true that

$$s_2 \geq s_1 \frac{\pi + c_1}{\pi + c_2}$$

in this region, it is not optimal for the firm to build Project 1. For $s_1 \frac{\pi+c_1}{\pi+c_2} > s_2 \geq (s_1 - 1)\frac{\pi+c_1}{\pi+c_2}$, the firm's decision is based on a three way comparison:

$$\begin{aligned} \min\{\alpha s_2(\pi + c_2) Pr(\gamma^T \geq \gamma^c | \gamma^{T-1}), & c_1 + \alpha(s_1 - 1)(\pi + c_1) Pr(\gamma^T \geq \gamma^c | \gamma^{T-1}), \\ & c_2 + \alpha(s_2 - 1)(\pi + c_2) Pr(\gamma^T \geq \gamma^c | \gamma^{T-1})\}, \end{aligned}$$

Note that,

$$\frac{c_1}{\alpha[s_2(\pi + c_2) - (s_1 - 1)(\pi + c_1)]} \leq \frac{c_2 - c_1}{\alpha[(s_1 - 1)(\pi + c_1) - (s_2 - 1)(\pi + c_2)]},$$

if and only if

$$s_2 \geq (s_1 - 1)\frac{\pi + c_1}{\pi + c_2} + \frac{c_1}{c_2},$$

thus, it is not optimal for the firm to build Project 1 in the region $s_2 < (s_1 - 1)\frac{\pi+c_1}{\pi+c_2} + \frac{c_1}{c_2}$. Therefore, $k_1^T = 0$ and $k_2^T = \frac{c_1}{c_2} - \frac{\pi+c_1}{\pi-c_2}$.

At $t = T - 1$, for $s_2 \geq (s_1 - 1)\frac{\pi+c_1}{\pi+c_2} + \frac{c_1}{c_2} + 1$, we have

$$E[V_T(\gamma^{T-1}, s_1 - 1, s_2)|\gamma^{T-2}] = E[V_T(\gamma^{T-1}, s_1 - 1)|\gamma^{T-2}] = E[V_T(\gamma^{T-1}, s_1 - 1, s_2 - 1)|\gamma^{T-2}],$$

therefore, it is not optimal for the firm to build Project 2. Likewise, for $s_2 < (s_2 - 2)\frac{\pi+c_1}{\pi+c_2} + \frac{c_1}{c_2}$, we have

$$E[V_T(\gamma^{T-1}, s_1 - 1, s_2)|\gamma^{T-2}] = E[V_T(\gamma^{T-1}, s_2)|\gamma^{T-2}] = E[V_T(\gamma^{T-1}, s_1, s_2)|\gamma^{T-2}],$$

the cost to go of building Project 1 is always greater than the cost of waiting, therefore, it is not optimal for the firm to build Project 1. Thus, we have $k_1^{T-1} = \frac{c_1}{c_2} + \frac{c_2-c_1}{\pi+c_2}$ and $k_2^{T-1} = \frac{c_1}{c_2} - 2\frac{\pi+c_1}{\pi+c_2}$ With similar analysis, we see that for $0 \leq t < T_1$, $k_1^{T-t} = \frac{c_1}{c_2} + (T-t)\frac{c_2-c_1}{\pi+c_2}$ and $k_2^{T-t} = \frac{c_1}{c_2} - (T-t+1)\frac{\pi+c_1}{\pi+c_2}$. $\square$

Bibliography

Akella, R. V.F. Araman, J. Kleinknecht. 2001. B2b markets: Procurement and supplier risk management. P.M. Pardalos Geunes H.E.R.J, E. Romeijn, eds., *Supply Chain Management: Models, Applications, and Research Directions*. Kluwer: Dordrecht.

Archibald, A. 1981. Continuous review (s, s) policies with lost sales. *Management Science* **27** 1171–1177.

Azoury, Katy S. 1985. Bayes solution to dynamic inventory models under unknown demand distribution. *Management Science* **31**(9) pp. 1150–1160.

Bather, J. A. 1966. A continous time inventory model. *Journal of Applied Probability* **3** 538–549.

Bean, James C., Julia L. Higle, Robert L. Smith. 1992. Capacity expansion under stochastic demands. *Operations Research* **40** pp. S210–S216. URL `http://www.jstor.org/stable/3840810`.

Berling, P. 2008. The capital cost of holding inventory with stochastically mean-reverting purchase price. *European Journal of Operational Research* **186** 620–636.

Billington, C. 2002. Hp cuts risk with portfolio approach. *Purchasing Magazine Online* .

Braden, David J., Marshall Freimer. 1991. Informational dyanmics of censored obesrvations. *Management Science* **37**(11) pp. 1390–1404.

Chaouc, B. 2006. Inventory control and periodic price discounting campaigns. *Naval Research Logistics* **54**.

Chen, F., J-S. Song. 2001. Optimal polices for multiechelon inventory problems with markov-modulated demand. *Operations Resarch* **49**(2) 226–234.

Cheng, F., S.P. Sethi. 1999. Optimality of state-dependent (s,s) policies in inventory models with markov-modulated demand and lost sales. *Production and Operations Management* **8**(2) 183–192.

Cohen, M., N. Agrawal. 1999. An analytical comparison of long and short term contracts. *IIE Transactions* **31** 783–796.

Constantinides, George M., Scott F. Richard. 1978. Existence of optimal simple policies for discounted-cost inventoru and cash management in continuous time. *Operations Research* **26**(4).

Davis, M. H. A., M. A. H. Dempster, S. P. Sethi, D. Vermes. 1987. Optimal capacity expansion under uncertainty. *Advances in Applied Probability* **19**(1) pp. 156–176. URL `http://www.jstor.org/stable/1427378`.

Dixit, Avinash K., Robert S. Pindyck. 2004. *Investment Under Uncertainty*. Princeton University Press.

Elsanosi, Ismail, Bernt Oksendal, Agńes Sulem. 2000. Some solvable stochastic control problems with delay. *Stochastics and Stochastics Reports* **71** pp. 69–89.

Erlenkotter, Donald. 1977. Capacity expansion with imports and inventories. *Management Science* **23**(7) pp. 694–702. URL `http://www.jstor.org/stable/2630849`.

Friedman, Lawrence M., Curt D. Furberg, David L. DeMets. 2010. *Fundamentals of Clinical Trials*. 4th ed. Springer.

Fu, C-Y Lee, Q., C-P Teo. 2006. Procurement risk management using options: Random spot price and the portfolio effect. *IIE Transactions* **42** 793–811.

Gavirneni, S., T.E. Morton. 1999. Inventory control under speculation: Myopic heuristics and exact procedures. *European Journal of Operational Research* **117** 211–221.

Goh, M., M. Sharafali. 2002. Price-dependent inventory models with discount offers at random times. *Production and Operations Management* **11**(2) 139–156.

Golabi, K. 1985. Optimal inventory policies when ordering prices are random. *Operations Research* **33**(3) 575–588.

Golovachkina, N., J. Bradley. 2003. Supplier-manufacturer relationships under forced com-

pliance contracts. *Manufacturing & Service Operations Management* **5**(1) 67–69. URL http://ideas.repec.org/a/inm/ormsom/v5y2003i1p67-69.html.

Guo, Xin, Pascal Tomecek. 2008. Connections between singular control and optimal switching. *SIAM J. Control Optim.* **47**(1) 421–443.

Haksoz, C., S. Seshadri. 2007. Supply chain operations in the presence of spot market: A review with discussion. *Journal of the Operational Research Society* .

Harrison, J. M., T. M. Selke, A. J. Taylor. 1983. Impulse control of brownian motion. *Mathematics of Operations Research* **8**(3) 454–466.

Harrison, J. Michael, Michael I. Taksar. 1983. Instantaneous control of brownian motion. *Mathematics of Operations Research* **8**(3).

Kalymon, B. A. 1971. Stochastic prices in a single-item inventory purchasing model. *Operations Research* **19**(6) 1434–1458.

Lariviere, Martin A., Evan L. Porteus. 1999. Bayesian inventory management with unobserved lost sales. *Management Science* **45**(3) pp. 346–363.

Li, C-L, P. Kouvelis. 1999. Flexible and risk-sharing contracts under price uncertainty. *Management Science* **45**(10) 1378–1398.

Luss, Hanan. 1982. Operations research and capacity expansion problems: A survey. *Operations Research* **30**(5) pp. 907–947. URL http://www.jstor.org/stable/170351.

Ly Vath, V, H. Pham. 2007. Explicit solution to an optimal switching problem in the two-regime case. *SIAM Journal on Control and Optimization* **46**(2) 395–426.

Manne, Alan S. 1961. Capacity expansion and probabilistic growth. *Econometrica* **29**(4) pp. 632–649. URL http://www.jstor.org/stable/1911809.

Martinez-de Albeniz, V., D. Simchi-Levi. 2005. A portfolio approach to procurement contracts. *Production and Operations Management* (1) 90–114.

Merhi, A., M. Zervos. 2007. A model for reversible investment capacity expansion. *SIAM Journal on Control and Optimization* **46** 839–876.

Moinzadeh, K. 1997. Replenishment and stocking policies for inventory systems with random deal offerings. *Management Science* **43**(3) 334–342.

Ormeci, J. G. Dai, A., J. V. Vate. 2008. Impulse control of brownian motion: the constrained average cost case. *Operations Research* **56**(3) 618–629.

Protter, Philip E. 2004. *Stochastic Integration and Differential Equations*. Springer-Verlag, Berlin.

Robbins, H. 1952. Some aspects of the sequential design of experiments. *Bulletin of the American Mathematical Society* **55** pp. 527–535.

Ryan, Sarah M. 2004. Capacity expansion for random exponential demand growth with lead times. *Management Science* **50**(6) pp. 740–748. URL http://www.jstor.org/stable/30046114.

Scarf, Herbert. 1959. Bayes solutions of the statistical inventory problem. *The Annals of Mathematical Statistics* **30**(2) pp. 490–508.

Seifert, U. W. Thonemann, R. W., W. H. Hausman. 2004. Optimal procurement strategies for online spot markets. *European Journal of Operational Research* **152** 781–799.

Sethi, S. P., F. Cheng. 1997. Optimality of (s, s) policies in inventory models with markovian demand. *Operations Research* **45**(6) 931–939.

Simchi-Levi, P. Kaminsky, D., E. Simchi-Levi. 2008.

Sinden, Frank W. 1960. The replacement and expansion of durable equipment. *Journal of the Society for Industrial and Applied Mathematics* **8**(3) pp. 466–480. URL http://www.jstor.org/stable/2098983.

Smith, Robert L. 1979. Turnpike results for single location capacity expansion. *Management Science* **25**(5) pp. 474–484. URL http://www.jstor.org/stable/2630278.

Sulem, A. 1986. A solvable one-diemensional model of a diffusion inventory system. *Mathematics of Operations Research* **11**(1) 125–133.

Taksar, M. I. 1985. Average optimal singular sontrol and a related stopping problem. *Mathematics of Operations Research* **10** 63–81.

Trigeorgis, Lenos. 1996. *Real Options: Managerial Flexibility and Strategy in Resource Allocation*. MIT Press.

Van Mieghem, Jan A. 2003. Capacity management, investment, and hedging: Review and recent developments. *Manufacturing & Service Operations Management* **5**(4) pp. 267–302.

Wang, Y. 2001. The optimality of myopic stocking policies for systems with decreasing prices. *European Journal of Operational Research* **133** 153–159.

Wu D. J., P. R. Kleindorfer, J. E. Zhang. 2002. Optimal bidding and contracting strategies for capital-intensive goods. *European Journal of Operational Research* **137**(3) 657–676.

Yi, J., A. Scheller-Wolf. 2003. Dual sourcing from a regular supplier and a spot market. Carnegie Mellon Tepper School of Business.