Solvability in Discrete, Nonstationary, Infinite Horizon Optimization

by

Timothy David Lortz

A dissertation submitted in partial fulfillment
of the requirements for the degree of
Doctor of Philosophy
(Industrial and Operations Engineering)
in the University of Michigan
2008

Doctoral Committee:

Professor Robert L. Smith, Co-Chair
Assistant Professor Archis V. Ghate, University of Washington, Co-Chair
Assistant Professor Hyun-Soo Ahn
Assistant Professor Goker Aydin

Acknowledgements

Although I wrote this dissertation myself, it is a cumulative result of the tapestry of life I've woven over the past five years. Thus, I feel compelled to highlight some individuals who have had key contributions to the successful completion of the dissertation.

Chris and Tina: department support staff who make IOE a great place and who make sure you're never out of luck.

Blake, Esra, and Irina: officemates supreme and friends off the court as well.

Archis Ghate: mentor and confidant; having access to him saved me wasted effort and helped me develop better writing skills.

Robert Smith: my advisor from day one, without whom I would have had no clue in approaching this topic. The notions of "well-posed" and "coalescent" are his creations, and he always asked questions that pushed me to better understand the topic.

Mom and Dad: cheerleaders and friends; growing up in their loving and supportive home allowed me to get to graduate school in the first place.

Maya and Delaney: amazing and beautiful little girls; sources of hope and motivation; for better or for worse, they have forced me to make good use of the time I have available.

Kiersten: my friend, my love, my rock; *we* did this!

Heavenly Father and Jesus Christ: source of all truth, object of worship, and hidden strength for overcoming all obstacles along this path.

Contents

List of Figures

List of Tables

List of Appendices

Abstract

For several time-staged operations management problems, the optimal immediate decision is dependent on the choice of problem horizon. When that horizon is very long or indefinite, an appropriate modeling technique is infinite horizon optimization. For problems that have stationary data over time, optimizing system performance over an infinite horizon is generally no more difficult than optimizing over a finite horizon. However, restricting problem data to be stationary can render the models unrealistic, failing to include nonstationary aspects of the real world.

The primary difficulty in nonstationary, infinite horizon optimization is that the problem to solve can never be known in its entirety. Thus, solution techniques must rely upon increasingly longer finite horizon problems. Ideally, the optimal immediate decisions to these finite horizon problems converge to an infinite horizon optimum. When finite detection of that optimal decision is possible, we call the underlying infinite horizon problem well-posed. The literature on nonstationary, infinite horizon optimization has generally relied upon either uniqueness of the optimal immediate decision or monotonicity of that decision as a function of horizon length. In this thesis, we require neither of these, instead developing a more general structural condition called coalescence that is equivalent to well-posedness.

Chapters 2-4 study infinite horizon variants of three deterministic optimization applications: concave cost production planning, single machine replacement, and capacitated inventory planning. For each problem, we show that coalescence is equivalent to well-posedness. We also give a solution procedure for each application that will uncover an infinite horizon optimal immediate decision for any well-posed problem.

In Chapter 5, we generalize the results of these applications to a generic classes of optimization problems expressible as dynamic programs. Under two different sets of assumptions concerning the finiteness of and reachability between states, we show that coalescence and well-posedness are equivalent. We also give solution procedures that solve any well-posed problem under each set of assumptions. Finally, in Chapter 6, we introduce a stochastic application: the infinite horizon asset selling problem, and again show that coalescence and well-posedness are equivalent and give a solution procedure to solve any such well-posed problem.

Chapter 1

Introduction to Infinite Horizon, Nonstationary Optimization

1.1 Motivation for Infinite Horizon, Nonstationary Optimization

Some of the earliest models in Operations Research involve making optimal decisions at fixed time intervals over an infinite time horizon. For example, the Economic Order Quantity characterizes an optimal solution to the problem of determining the minimum cost order size when demand, as well as fixed and variable production costs, are constant over time (see, for example, [28], p. 50). Turning our attention to a stochastic problem, the literature on Markov Decision Processes generally requires that the sets of actions and their associated rewards and transition probabilities remain constant over time. This assumption facilitates the restriction of the search for optimal policies to those that are also stationary in time (see, for example, Puterman [35], p. 119).

Indeed, for a time-staged optimization problem with infinitely many parameters, unless those parameters are stationary in time or at least follow a predictable pattern, *the entire problem can never be known at once*! That is the key challenge motivating this thesis. See Figure 1.1 for a depiction of the impossibility of capturing an entire

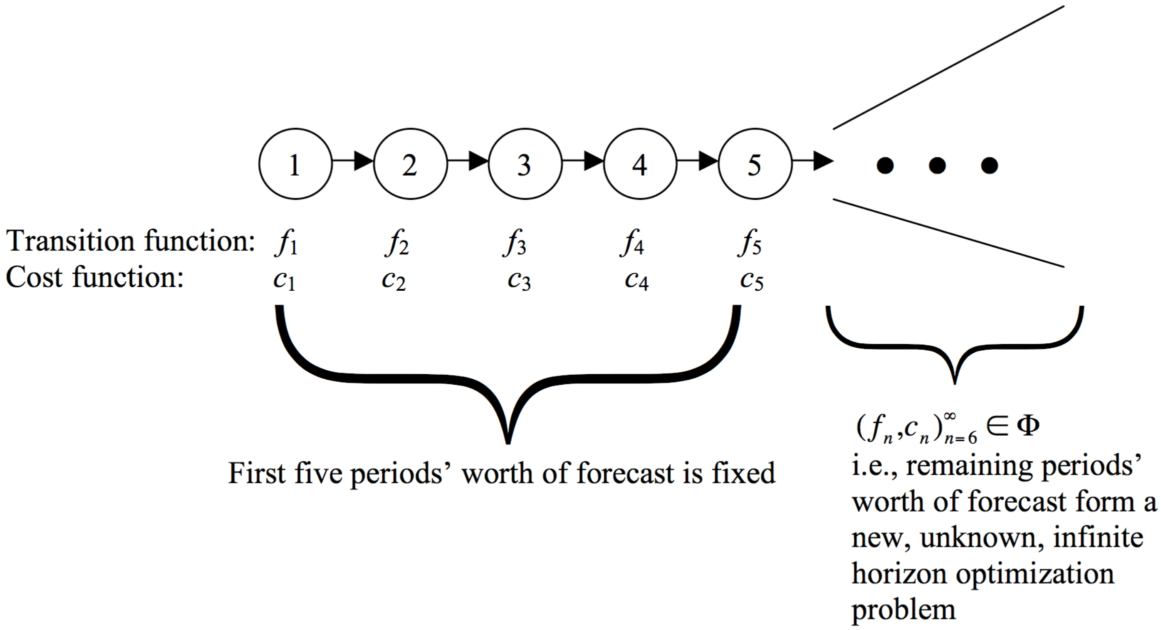

Figure 1.1: Impossibility of Capturing an Entire Nonstationary Infinite Horizon Forecast

nonstationary forecast. In that figure, each period n has a state transition function f_n and a cost function c_n, so that the problem could be solved using dynamic programming with a minimization objective. The class of all permissible forecasts is denoted by Φ. Although the figure shows only the first five periods' worth of forecast being fixed, that number could be arbitrarily large and the difficulty would persist.

Real-world problems typically are not stationary and are subject to noise from outside. When the time-varying influences can be forecast, then the infinite horizon problem can have at least some of its parameters captured. For example, a job shop may have some work regularly scheduled and have other custom jobs that arrive over time. The sales team can give estimates of the custom jobs until some point in the future, but the reliability of those estimates deteriorates as the planning horizon increases. However, the scheduling team needs to make a production schedule now. Thus, they must rely upon a finite horizon truncation of the infinite horizon forecast, where the forecast consists of the jobs with associated machining requirements, deadlines, and revenues.

1.2 Background

If the decision maker wishes to determine an infinite horizon optimal immediate decision, then he or she must choose a horizon sufficiently long that an optimal immediate decision, given the parameters revealed through that horizon, is also optimal for *any* future parameters. When this condition holds, the given horizon is called a *forecast horizon*, and in this thesis, we will call such a problem *well-posed*. However, there is no guarantee that a forecast horizon exists. Examples of nonstationary, infinite horizon optimization problems for which no forecast horizon exists can be found in the context of general infinite horizon optimization [1], production planning [8], and asset selling [16].

On the other hand, there is a substantial literature both in establishing general conditions under which a forecast horizon exists, and in developing solution procedures that will yield an infinite horizon optimal initial decision for an infinite horizon nonstationary optimization problem. An excellent classified bibliography of research endeavors in both of these directions over a broad spectrum of applications and theory was recently made in [10]. In light of that article, we will only mention other research as it pertains to this thesis.

The largest portion of the forecast horizon literature deals with algorithm development and implementation. Generally speaking, most of these algorithms rely upon possessing a dynamic programming formulation of the problem under study so that the state of the system can be characterized given all previous decisions and the problem parameters, carefully choosing a set of states through which all optimal state sequences must pass, then checking for agreement in the initial decision for the finite horizon problem restricted to pass through each of these states. If there is agreement, then one stops. Otherwise, one moves forward in time, reconstructs the state set, and repeats.

The primary difficulty in constructing an effective forecast horizon detection procedure is balancing the conflict between being too conservative in constructing the state sets, and thereby omitting some optimal states that could indicate that the problem is not well-posed, or in being too liberal in the construction by including non-optimal states and thereby failing to notice agreement over the truly optimal states. A set of states in a time-staged problem chosen so that every optimal state sequence must pass through that set is called a *regeneration set*, and a regeneration set that contains only states optimal for some infinite horizon problem identical to the one under consideration through at least the time of that set is called a *minimal regeneration set*. When agreement in the optimal initial decision occurs over all the finite horizon problems with terminal states in a minimal regeneration set, then an infinite horizon optimal initial decision has been discovered. As Lundin and Morton [32] explain, "the quest for [forecast] horizons and other related procedures may be reduced to the quest for regeneration sets." While regeneration sets can often be rather easy to construct, minimal regeneration sets can become exceedingly difficult. The problems studied in this work will have straightforward characterizations of their minimal regeneration sets, facilitating the development of effective solution procedures.

1.3 Overview of Approach

The problems chosen in this thesis all satisfy the condition that there exists a feasible solution with finite total cost. Therefore, the optimality criterion for each problem will be that of total cost, whether the objective is to minimize (cost) or to maximize (rewards). Under the total cost optimality criterion and the existence of feasible infinite horizon solutions with finite total cost, many researchers have pointed out that one can bound the maximum deviation from optimal total cost by solving to optimality a finite horizon problem. That is, one can obtain optimal solutions over a finite horizon that ensure total cost within ϵ of the optimal infinite horizon cost for any $\epsilon > 0$. See, for example, [30, 17] in the context of production planning and [4] in the context of equipment replacement. This thesis makes the distinction that the decision maker desires convergence by looking at incrementally longer finite horizon problems not just in terms of *cost*, but also in terms of *policy*. Thus, even if solving a sufficiently long finite horizon problem can yield a policy that, if followed and optimally appended over the horizon, can bring the total cost within an acceptable amount of the true optimal cost, this thesis adopts the convention that failure to converge in policy as well is unacceptable.

Fortunately, it is generally the case that there exists a sequence of finite horizon optimal initial decisions that converges, in finite time, to an infinite horizon optimal initial decision. Bean and Smith [1, 2] explore this topic extensively, as does Lasserre [29]. We will follow a course similar to those works, except that we explicitly drop the assumption that the optimal initial decision (or any decision, for that matter) is unique. The assumption of uniqueness of the optimal initial decision acts as a surrogate (in fact, it is a sufficient condition) for the *solvability* of a particular problem instance. However, since there exist examples of problem instances for which the optimal initial decision is not unique, this thesis will permit the existence of multiple optima. *In this thesis, solvability implies the existence of a finite time horizon sufficiently long that knowledge of problem parameters beyond that horizon is not necessary for determining an infinite horizon optimal initial decision.* Solvability is independent of the computational burden required by a solution procedure for detecting an infinite horizon optimal initial decision, and is even independent of the existence of such an algorithm. This notion of solvable will henceforth be referred to as *well-posed*. For a treatment of well-posedness in the context of non-homogeneous Markov Decision Processes, see Cheevaprawatdomrong et al [16].

Under some circumstances, the existence of multiple optima poses no difficulty to convergence to an infinite horizon optimal initial decision. Yet under other circumstances, when multiple optima exist, convergence (in the horizon length) of the optimal initial policy may never occur. The difference between these two events can be explained by a general condition called *coalescence* that will frequently appear in this thesis. In short, this work will show that for a variety of applications, a problem instance is well-posed if and only if coalescence is satisfied. Although in Chapter 5 coalescence is defined for a large class of optimization problems, explicitly defining coalescence for specific applications can lend some insight as to the solvability of those applications and into the construction of effective solution procedures. Thus, coalescence will be recast for each application treated in this thesis, even if the results concerning solvability are similar.

For the applications in this work, for any finite horizon truncation of an infinite horizon forecast, the problem class contains an infinite number of problem extensions. Thus, in attempting to solve the problem, a decision maker must account for *all* potential future extensions. An appropriate paradigm for the forecasting function is an oracle that reveals a period's worth of parameters each time it is called. The oracle knows all the parameters over the infinite horizon and must reveal the parameters for any specified period when called to do so. The decision maker, querying the oracle, must make a decision based upon only the parameters revealed by the oracle and any bounds on the future behavior of the oracle (such as application-specific assumptions). However, it may be possible for future oracle calls to reveal that finite horizon optimal decisions based upon previously revealed parameters are no longer optimal. A problem is well-posed if it is possible to make an infinite horizon optimal initial decision after only finitely many oracle calls.

Chapters 2-4 will present infinite horizon nonstationary versions of the following discrete-time deterministic optimization problems: concave cost production planning, single machine replacement economy, and capacitated inventory planning. Each chapter will characterize the optimal strategies, define coalescence, show that it is necessary and sufficient for the solvability of a problem instance, and describe solution procedures guaranteed to solve any well-posed problem instance for the respective problem. Chapter 5 describes a more general class of deterministic optimization problems, and for each of three different sets of structural assumptions of varying strength, exposes the critical role of coalescence in the solvability of problem instances and describes a solution procedure guaranteed to solve any well-posed problem instance for each set of assumptions.

Lastly, in Chapter 6, it is shown that coalescence is also synonymous with the solvability of a classic stochastic problem - the asset selling problem. Interestingly, the most concise structural meaning of coalescence, among all the applications in this work, arises from this stochastic problem. A new solution procedure that will solve any well-posed asset selling problem appears in that chapter along with a detailed analysis of its performance.

Chapter 2

An Infinite Horizon Concave Cost Production Planning Problem

2.1 Introduction

The goal in a deterministic production planning problem is to meet demand for a single product at minimum production and inventory holding costs. Even though several efficient solution procedures exist to solve a variety of production planning problems over a *given* finite horizon, it is often not clear whether a specific finite horizon is sufficiently long for planning purposes when the data are non-stationary owing to technological and other economic changes. Initial production decisions for a problem with a short horizon are likely to prove non-optimal when future data are revealed, whereas planning for a very long finite horizon is challenging owing to the difficulty in forecasting data. Moreover, the actual horizon for which a manufacturing firm is likely to be in business is often very long and indefinite. A model that effectively addresses all of these issues is the infinite horizon, nonstationary production planning model.

Unfortunately, non-stationary, infinite horizon models introduce another complication. They are characterized by infinite data that clearly cannot be known at once. When a problem in which only n periods' worth of data has been revealed possesses an n-horizon optimal initial decision that is also optimal for any potential realiza-

Problem Horizon	x_1^*	x_2^*	x_3^*	x_4^*	x_5^*	x_6^*	x_7^*
5	9	0	14	0	0	-	-
6	14	0	0	14	0	0	-
7	18	0	0	0	15	0	0

Table 2.1: Finite Horizon Optimal Strategies for an Unsolvable Production Planning Problem

tion of future parameters, n called a *forecast horizon*. As a result, the notion of a forecast horizon is paramount in this context. In this chapter, we give a novel condition called coalescence which is both necessary and sufficient for finite detection of a forecast horizon, along with a solution procedure that will find a forecast horizon in finite time whenever one exists. In general, a forecast horizon may not exist, rendering it impossible to detect an optimal initial decision, as the following example shows.

Example of an unsolvable problem

Bhaskaran and Sethi [8] present a pair of undiscounted, deterministic, infinite horizon production planning problems in which there is never agreement in the optimal initial decisions when solving the corresponding incrementally longer finite horizon truncations. We will focus on the first of their two examples, modifying it only through applying a discount factor. Let the production and holding cost forecasts, respectively, be

$$\begin{aligned} c_n(x) &= \begin{cases} 1, & x > 0 \\ 0, & \text{otherwise} \end{cases}, \\ h_n(x) &= .05x \end{aligned}$$

for all periods n, with demand forecast $d = (5, 4, 5, 4, 5, 5, 5, \cdots)$. For any single-period discount factor $\alpha \geq .87$, Table 2.1 characterizes the optimal production decisions in each period of the 5-, 6-, and 7-period problems. Let $x_n^*, n = 1, \ldots, 7$ denote the optimal production quantity in period n. Moreover, the stationary problem beginning in period 5 has a unique infinite horizon optimal production plan x^*

Strategy	x_1^*	x_2^*	x_3^*	x_4^*	x_5^*	x_6^*	x_7^*	x_8^*	$\cdots$
1	9	0	14	0	0	15	0	0	$\cdots$
2	14	0	0	14	0	0	15	0	$\cdots$
3	18	0	0	0	15	0	0	15	$\cdots$

Table 2.2: Infinite Horizon Optimal Strategies for an Unsolvable Production Planning Problem

in which

$$x_n^* = \begin{cases} 15, & n \bmod 3 = 1 \\ 0, & \text{otherwise} \end{cases} .$$

Consider the three infinite horizon strategies resulting from acting optimally in the 5-, 6- and 7-period problems followed by acting optimally in the stationary problem by producing 15 units every three periods, as shown in the table below. The costs of these three strategies converge to the same infinite horizon total cost, and thus, all three initial decisions are infinite horizon optimal.

Thus, there are three infinite horizon optimal initial decisions: produce 9, 14, or 18 units. However, when solving a finite horizon problem whose horizon length is a multiple of 3, the unique optimal initial production quantity is 9 units. Likewise, when that horizon length divided by 3 has remainder 1, the unique optimal initial production quantity is 14, and when the remainder is 2, the optimal quantity is 18. These optimal solutions alternate as the horizon length increases. However, raising the production costs to

$$c_n(x) = \begin{cases} 3, & x > 0 \\ 0, & \text{otherwise} \end{cases}$$

for example, in all periods $n \geq n'$ for some $n' > 8$, while keeping all other parameters fixed, can render the optimal initial decision resulting in a setup in period n' non-optimal. Recall that only finitely many periods' worth of parameters can be known at once, so although the problem may appear stationary indefinitely long, there is no assurance that it will always remain so. Thus, finite horizon solution methods can never be guaranteed to return an infinite horizon optimal initial decision, so no forecast horizon exists for this problem and it is not well-posed.

2.1.1 Literature Review

In this chapter, we will focus on nonstationary, infinite horizon concave cost production planning models, including as a special case the infinite horizon version of the well-known dynamic lot size model (DLSM). The problem of detecting a forecast horizon for the DLSM has long attracted researchers. Over the years, both the robustness and the computational effort of algorithms devoted to this pursuit have improved. What follows is a brief discussion of solution procedures in the concave cost production planning literature, with an emphasis on those designed for the DLSM.

Early work in this area, including Wagner and Whitin [42] (hereafter W-W) and Zabel [43], demonstrate how to solve the DLSM with stationary unit production costs. Their stopping rule specifies that when $l(n)$, the last optimal setup period in an n-period problem, is equal to n, then n is a forecast horizon. Moreover, they show that $l(n)$ is monotone nondecreasing in n. Eppen et al [21] allow the variable production costs to vary with time instead of holding them stationary as their predecessors did, and they find that the last finite horizon optimal production point might not be monotonic. In order to obtain the W-W planning horizon theorem for the general model, they then require that the "marginal cost of production does not rise as rapidly over time as the sum of holding costs." They show that the case of nondecreasing setup costs, for example, results in this condition. Where this condition on the marginal production costs is not met, they try to decrease the number of potential last production points by deriving a "violator set" of periods whose marginal production costs rise slowly enough that they could potentially supplant $l(n)$ (they implicitly assume $L(n) = \{l(n)\}$ for all n, i.e., the optimal last production period is always unique) as an optimal last production point prior to period n for some longer horizon problems. When the violator set is empty, a planning horizon has been found.

However, like W-W, this stopping rule may never be satisfied even when a planning horizon exists.

For the same problem, Lundin and Morton [32] were the first to explicitly identify a *regeneration set*, a set of periods Z (with upper bound z^u) such that any finite horizon problem in agreement with the problem through period z^u has at least one optimal solution with a regeneration point (zero-inventory point) in Z. Further procedural improvements were offered by Chand and Morton [11] and Federgruen and Tzur [22].

Bensoussan et al [6] allow for general piecewise linear, concave production and holding cost functions in their work. Given a finite forecast (the terminal period of which they call a forecast horizon, though it differs from the notion of a forecast horizon predominant in the literature), they provide an algorithm which is designed to detect any existing planning horizons. Recently, Dawande et al [18] developed integer programming formulations to detect minimal forecast horizons, when they exist, for a subclass of the DLSM when future demand is restricted to integer multiples of of a given positive real number.

On the other hand, few researchers have discussed problem classes for which an algorithm designed to detect forecast horizons for the DLSM will actually stop. For the restricted class of problems in which demand is stationary for some initial number of periods, Chand et al [13] give existence conditions for forecast horizons in the undiscounted DLSM. Chand et al [14] give existence conditions for forecast horizons in the discounted DLSM. Using continuous discounting and a finite set of policy alternatives at each decision epoch, Bean and Smith [1] show that many sequential decision problems (including concave cost production planning) have unique infinite horizon optimal decisions, and that a forecast horizon thus exists for most such problems. Their results are applied in the paper by Bean et al [5] which considers the general

DLSM.

Federgruen and Tzur [22] base their claim that the event of failing to detect a forecast horizon for the DLSM is "presumably rare" on empirical evidence gleaned from experiments that they ran. While their observation is quite likely to be true, it would be advantageous to know specific conditions for existence of forecast horizons for the DLSM, and for concave cost production planning problems in general.

2.1.2 Overview of Approach

The remainder of this chapter will proceed as follows. In §2.2, we rigorously define a class of infinite horizon, deterministic, concave cost production planning problem without backlogging. This class of problems includes, but is not limited to, the DLSM. In §2.3, we introduce the coalescence property, which for this problem class, arises when some infinite horizon optimal initial decision may be continued by other infinite horizon optimal decisions to share at least one zero-inventory point with any other infinite horizon optimal production plan. We show that satisfying the coalescence condition is necessary and sufficient for finite detection of an infinite horizon optimal initial decision, and we employ coalescence to identify solvable problem classes.

In §2.4, we give a new solution algorithm that will solve any solvable problem in our class, meaning it will detect an infinite horizon optimal initial decision as long as one can be detected by *some* exact algorithm. This solution algorithm runs with effort that is constant in a bound on the time between optimal zero-inventory points, no matter the length of the forecast horizon. Whereas solution procedures in the literature designed specifically for the DLSM can fail for more general concave cost problems, our algorithm successfully detects a forecast horizon any time one exists for a general class of concave cost production planning problems. We offer concluding remarks in §2.5.

2.2 Model and Assumptions

We first present the infinite horizon, deterministic, discounted, concave cost production planning problem without backlogging. Its variables and parameters appear below.

d_n	$=$	demand forecast for period n (we assume that $d_n \in \mathbb{Z}_+$)
d	$=$	$d = (d_1, d_2, \ldots)$, the infinite horizon demand forecast
$c_n(\cdot)$	$=$	undiscounted production cost function in period n
c	$=$	$(c_1(\cdot), c_2(\cdot), \ldots)$, the infinite horizon production cost forecast
$h_n(\cdot)$	$=$	undiscounted inventory holding cost function in period n
h	$=$	$(h_1(\cdot), h_2(\cdot), \ldots)$, the infinite horizon inventory holding cost forecast
p	$=$	(d, c, h), a triple referred to as the *problem* or *forecast*
x_n	$=$	decision variable representing the quantity produced in period n
i_n	$=$	inventory level beginning period n
α	$=$	one-period discount factor ($0 < \alpha < 1$)

The optimization problem is then

$$\begin{array}{lll} \min & \sum_{n=1}^{\infty} \alpha^{n-1}[c_n(x_n) + h_n(i_n)] & \\ \text{subject to} & i_n + x_n - i_{n+1} = d_n, & n = 1, 2, \ldots \\ & x_n, i_n \in \mathbb{Z}_+, & n = 1, 2, \ldots \end{array} \tag{2.1}$$

The initial inventory $i_1 = 0$. Note that i_n is determined by the demand forecast and production decisions through period $n - 1$.

Definition 2.1. *We say that $x = (x_1, x_2, \ldots)$ is a* feasible production plan *(with resulting* inventory plan *i) for forecast p if x and i satisfy program (1).*

At this point, we introduce the three main assumptions used in this chapter.

Assumption 2.1. *For all n, $c_n \geq 0$ and $h_n \geq 0$ are non-negative concave functions, with $c_n(0) = h_n(0) = 0$ for all n.*

Assumption 2.2. *There exists an upper bound $L < \infty$ on the maximum number of periods over which it could be optimal to carry a unit of inventory, independent of the problem parameters.*

Assumption 2.3. *For any forecast p, there exists a feasible production plan with finite total discounted cost.*

We are now in a position to set forth some additional notation.

$$
\begin{array}{rcl}
\mathcal{D} & = & \{d : \ d_n \in \mathbb{Z}_+, d_n \leq \bar{d} < \infty, \text{ for all } n\} \\
\mathcal{C} & = & \{c : \ c \text{ satisfies Assumption 2.1}\} \\
\mathcal{H} & = & \{h : \ h \text{ satisfies Assumption 2.1}\} \\
\mathcal{P} & = & \{p = (d, c, h) : d \in \mathcal{D}, c \in \mathcal{C}, h \in \mathcal{H}\}, \text{ the set of all possible forecasts} \\
\mathcal{P}^n(p) & = & \text{the set of all } p' \in \mathcal{P} \text{ that are in agreement with forecast } p \text{ through} \\
 & & \text{period } n \\
p^n & = & (d^n, c^n, h^n), \text{ the } n\text{-period truncation of } p \\
\underline{p}^n & = & (\underline{d}^n, \underline{c}^n, \underline{h}^n), \text{ the } n\text{-period truncation of } p \text{ with uniformly zero demand} \\
 & & \text{as well as cost function forecasts beginning period } n+1 \\
\mathcal{X}^*(p) & = & \text{the set of all infinite horizon optimal production plans for forecast } p \\
\mathcal{X}_1^*(p) & = & \text{the set of all infinite horizon optimal first period production decisions} \\
 & & \text{for forecast } p \in \mathcal{P} \\
\mathcal{X}_1^{*n}(p) & = & \cap_{p' \in \mathcal{P}^n(p)} X_1^*(p'), \text{ the set of first period decisions that are infinite} \\
 & & \text{horizon optimal for all forecasts in agreement with forecast } p \text{ through} \\
 & & \text{period } n \\
V(p, x) & = & \sum_{n=1}^{\infty} \alpha^{n-1} c_n(x_n) + h_n(i_n), \text{ the total cost of production plan } x \text{ and} \\
 & & \text{resulting inventory plan } i \text{ for problem } p
\end{array}
$$

Throughout this chapter, whenever we mention an optimal production plan x^*, we mean that x^* is an optimal production plan for the infinite horizon problem p, unless noted otherwise. Also, the first n optimal production decisions for a problem $\underline{p}^n$ are optimal for the n-horizon forecast p^n.

We note here that we will limit consideration of optimal production plans to those in which

$$x_n^* \cdot i_n^* = 0, \ \forall n, \tag{2.2}$$

where i^* is the inventory plan resulting from x^* and demand forecast d. That there is no loss of optimality in doing this is described in A.1.

We note that if the zero-inventory production policy restriction (2.2) arose through natural constraints on the problem, then we could relax the concave cost assumption Assumption 2.1. The concave cost assumption is included only to obtain the zero-inventory production policy restriction. This restriction is useful in that production decisions concern only the *number* of periods' worth of demand to satisfy at each production point, rather than considering all possible production levels.

2.3 Forecast Horizons for the Concave Cost Production Planning Problem

In this section, we will give necessary and sufficient conditions for the existence of a forecast horizon, or in other words, for the ability to find infinite horizon optimal production decisions for some initial set of periods.

2.3.1 Solvability Definitions

We define here both *well-posed* and *forecast horizon* - two definitions that formalize the meaning of solvability - as they pertain to concave cost production planning problems. The second definition has been used extensively in the literature, while the first is less common. We feel the first definition captures more concisely the essence of solvability.

Definition 2.2. *Forecast p is* well-posed with respect to $\mathcal{P}$ *if there exists some period $\bar{N}$ such that $\mathcal{X}_1^{*\bar{N}}(p) \neq \emptyset$.*

A well-posed problem requires only a finite number of periods' worth of forecast in order to detect an optimal initial decision that is also optimal for any problem in agreement with it through its forecast horizon, which we now define.

Definition 2.3. *Period N^* is a* forecast horizon *for problem p with respect to class $\mathcal{P}$ if there exists an infinite horizon optimal initial decision x_1^* such that $x_1^* \in \mathcal{X}_1^{*N}(p)$ for all $N \geq N^*$.*

In other words, existence of a forecast horizon N^* implies that one can find an initial optimal solution to any problem in agreement with p through the first N^* periods by solving the finite horizon problem p^{N^*}. Well-posedness or the existence of a forecast horizon is highly dependent upon the problem class. For example, if $\mathcal{P} = \{p\}$, a singleton, then p is certainly well-posed. It is straightforward to see that a problem is well-posed with respect to its problem class if and only if a forecast horizon exists for that problem with respect to its problem class. Moreover, if either condition is satisfied, the problem is solvable.

2.3.2 Necessary and Sufficient Conditions for Well-Posed Problems

As we saw in the example at the beginning of this chapter, difficulty in detecting an infinite horizon optimal initial decision can arise when ties in minimum total cost exist among multiple production plans with distinct initial decisions. Although such cases might be rare in practice, it is important that they be considered so that a solution procedure does not fail in their presence. We provide in this section a rigorous analysis of problems with multiple optima and distinguish solvable problems from unsolvable problems by the behavior of their optima.

Definition 2.4. *Let x^* and x^{**} be infinite horizon optimal production plans for problem p, with corresponding inventory plans i^* and i^{**}, respectively. We say that x^{**} is* optimally reachable *from x^* if $\exists y^* \in \mathcal{X}^*(p)$ with inventory plan j^* such that $j^*_m = i^*_m = 0$ and $j^*_n = i^{**}_n = 0$ for some m and n, $n \geq m$.*

In other words, one optimal production plan is optimally reachable from another if some of the optimal decisions of the first plan can be appended by some optimal continuation to share a common zero inventory position with the second plan in some period. Clearly, if i^*_n and $i^{**}_n = 0$ in period n, then i^* optimally reaches i^{**}, though this need not be the case. In the case that all infinite horizon optimal production plans are optimally reachable from one or more plans sharing some optimal initial

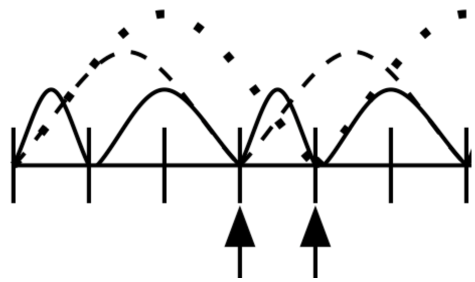

Figure 2.1: Sketch of Coalescence in a Production Planning Problem

decision(s), the problem has the following characteristic.

Definition 2.5. *We say that p satisfies the* coalescence condition *(or is* coalescent*) if there exists an optimal plan $x^* \in \mathcal{X}^*(p)$ called a* source plan *such that any other optimal production plan is optimally reachable from x^*.*

For a graphical illustration of coalescence, see Figure 2.1, which shows a sketch of the infinite horizon optimal decisions in the first six periods of some problem. The arcs represent carrying inventory from one production point to the next. The continuous, dashed, and dotted lines represent the three different optimal production plans; call them 1, 2 and 3, respectively. Plans 1 and 2 join at period three, marked by the first arrow, and plans 1 and 3 join at period 4, marked by the second arrow. Thus, plan 1 is a source plan and this problem is coalescent.

We are now prepared to present the main results of this section. The proof of the main theorem to follow will flow more easily by first presenting a technical lemma. See A.3 for a proof.

Lemma 2.1. *Let $p \in \mathcal{P}$ and choose any subsequence of integers $\{n_j\}$ with associated forecasts $\{p(n_j)\}, p(n_j) \in \mathcal{P}^{n_j}(p)\ \forall j$. Then there exist a further subsequence $\{n_{j_k}\} \subseteq \{n_j\}$ with associated optimal production plans $\{\tilde{x}^*(n_{j_k})\}, \tilde{x}^*(n_{j_k}) \in \mathcal{X}^*\left(p(n_{j_k})\right)\ \forall k$, and some $\tilde{x}^* \in \mathcal{X}^*(p)$ such that $\tilde{x}^{*(n_{j_k})} \to \tilde{x}^*$, where convergence is componentwise.*

Theorem 2.2. *Problem p is well-posed with respect to $\mathcal{P}$ if and only if it satisfies the coalescence condition.*

Proof. Suppose that p is well-posed with respect to $\mathcal{P}$. Choose $x^* \in \mathcal{X}^*(p)$ and N such that x_1^* is an infinite horizon optimal initial decision for all $p' \in \mathcal{P}^N(p)$. Then in particular, x_1^* is an optimal initial decision for $\underline{p}^n$ (and hence p^n) for all $n \geq N$. By (2.2), every optimal plan for p must have a zero-inventory point at some period in $\{N, N+1, \ldots, N+L\}$. Since x_1^* is an optimal initial decision for each of the problems $\underline{p}^N, \underline{p}^{N+1}, \ldots, \underline{p}^{N+L}$, it is straightforward how to construct an optimal production plan beginning with x_1^* that shares a zero-inventory point with any other optimal plan for p. Thus, well-posedness with respect to $\mathcal{P}$ implies coalescence.

Now suppose that p is coalescent but is not well-posed with respect to $\mathcal{P}$. Let x^{**} be a source plan. Then there exists a subsequence $\{n_j\}_{j=1}^{\infty}$ and $p(n_j) \in \mathcal{P}^{n_j}(p)$ such that $x_1^{**} \notin \mathcal{X}_1^*(p(n_j))$ for all j. But Lemma 2.1 states that, resorting to a further subsequence $\{n_{j_k}\}_{k=1}^{\infty}$, there exists some sequence of infinite horizon optimal production plans $x^*(n_{j_k}) \in \mathcal{X}^*(p(n_{j_k}))$ for all k such that $x^*(n_{j_k}) \to \tilde{x}^* \in \mathcal{X}^*(p)$.

By assuming that p is coalescent, $\tilde{x}^*$ is reached at some period M by some infinite horizon optimal plan with initial decision x_1^{**}. Let $K(M)$ be large enough that $x^*(n_{j_k})$ is in agreement with $\tilde{x}^*$ up through and including period M for all $k \geq K(M)$. Note that such a $K(M)$ exists in light of Lemma 2.1. By the Principle of Optimality, there now exists a $p(n_{j_k})$-optimal plan $\forall k \geq K(M)$ with initial decision x_1^{**}, and which is in agreement with $x^*(n_{j_k})$ beginning period M. Hence, $x_1^{**} \in \mathcal{X}_1^*(p(n_{j_k}))\ \forall k \geq K(M)$. We've reached a contradiction and thus conclude that satisfying the coalescence condition is sufficient for well-posedness. □

For the problem class $\mathcal{P}$, well-posedness and coalescence of a fixed infinite horizon forecast are equivalent. The choice of $\mathcal{P}$ is important, for we will see in the example in section 2.3.3 that the equivalence between well-posedness and coalescence may not hold when the problem class differs from $\mathcal{P}$.

2.3.3 Examples of Well-Posed Problems

In this section, we will expose several examples of problem classes whose solvability is easily determined by applying the coalescence condition. Recalling the example given in §2.1, it is easy to see that that problem is not coalescent since the three optimal production plans have no common zero-inventory points. Therefore, it is not well-posed.

Unique infinite horizon optimal solution

Suppose that $p \in \mathcal{P}$ has a unique infinite horizon optimal solution. Then, in particular, there is a unique optimal initial decision, and a unique first (after the first period) zero-inventory position. That the coalescence condition is satisfied follows trivially, and therefore, any such problem is well-posed. Compare the ease of this approach with the proof that forecast horizons exist for a general class of infinite horizon discounted problems involving sequential decisions having unique infinite horizon optimal solutions in Bean and Smith [1].

Stationary forecast and cost parameters

Let the stationary demand be $d_1 = d_2 = \cdots = a \in \mathbb{Z}_+$ with stationary cost forecasts $c_1(x) = c_2(x) = \cdots = c(x)$ and $h_1(i) = h_2(i) = \cdots = h(i)$, for all $x, i \in \mathbb{Z}_+$. At each zero-inventory point, the decision maker is faced with a constant set of alternatives. The set consists of the number of periods' worth of demand to satisfy, and the undiscounted costs of each alternative in the set do not change from production point to production point. Specifically, there are $L+1$ alternatives at each zero-inventory point, with lengths $f_1 = 1, f_2 = 2, f_3 = 3, \ldots, f_{L+1} = L+1$ and costs $p_1 = c(a), p_2 = c(2a) + \alpha h(a), \ldots, p_{L+1} = c((L+1)a) + \sum_{j=1}^{L} \alpha^j h\left((L-j+1)a\right)$. This problem is an instance of the infinite horizon discounted knapsack problem of Shapiro and Wagner [40].

Following the procedure in [40], one can obtain all turnpike policies, that is, all alternatives resulting in the lowest annualized discounted cost. Let there be $1 \leq M \leq L+1$ turnpike policies. If $M = 1$, then there is a unique infinite horizon optimal production plan and the coalescence condition is trivially satisfied as noted under the last heading. If $M \geq 2$, then let $j_1, j_2, \ldots, j_M$ be the turnpike policies, each in terms of the number of succeeding periods whose demand is satisfied by current production. Note that the set of infinite horizon optimal production plans is the infinite Cartesian product of the j_i's, $i = 1, \ldots, M$, so that each optimal plan is a string of turnpike policies initiated at zero-inventory points. Without loss of generality, let the optimal source plan be to choose j_1 at each zero-inventory point.

We can construct an optimal plan from the source plan reaching any other optimal plan y^* as follows. Because any infinite horizon optimal plan consists of an infinite number of decisions, and there are a finite number of turnpike policies, y^* has an infinite number of decisions of some length $j_{i'}$.

1. Reveal the decisions of y^* until j_1 decisions of length $j_{i'}$ (among m total decisions) have been observed.

2. Rearrange the decisions of y^* into another infinite horizon optimal plan $y^{*\prime}$ such that the first j_1 decisions are of length $j_{i'}$ and the other $m - j_1$ decisions are as in y^*.

3. $y^{*\prime}$ can be reached by an optimal plan beginning with the initial decision of the source plan after $j_{i'}$ decisions of length j_1, followed by appending the $m - j_1$ decisions among the first m decisions of y^*.

4. This composite plan then reaches y^* at its mth zero-inventory point.

Since y^* was chosen arbitrarily, it is seen that stationary problems satisfy the coalescence condition. Therefore, they are also well-posed. The strength of this result is that stationary problems are well-posed, even when they are contained in the class $\mathcal{P}$ of all possible forecasts. This result is of assistance in the following section.

Stationary cost, constant initial demand

The stationary cost DLSM was studied at length in Chand, Sethi and Sorger [14]; its undiscounted analogue was studied (with similar results) in Chand, Sethi and Proth [13]. In essence, those authors give an upper bound on the number of initial periods over which demand must also remain stationary in order to guarantee the existence of a forecast horizon. Using numerical studies, they find that the forecast horizons are generally quite short in practice, though the upper bound on the initial demand may be longer. While their proofs of the existence of forecast horizons for this class of problems are straightforward, it is even easier to see that the problems are well-posed by using the coalescence condition. In light of the previous example, stationary problems are well-posed, so that any forecast which is initially stationary long enough is also well-posed. In other words, forecast horizons exist for initially stationary forecasts in the class $\mathcal{P}$ of all possible forecasts, where both demands and costs are permitted to vary.

Cyclic forecast

The inclusion of this problem class is intended to illustrate that even some cyclic problems fail to be well-posed. Consider a cyclic production planning problem of cycle length 2, where the costs and demands for the first two periods are as follows.

Knowing the first two periods' worth of parameters describes the entire infinite horizon problem instance, which we denote p. Consider the two feasible production plans x' and x'', with production occurring in periods 1,3,5,7,... under x' and in periods 1,2,4,6,... under x''. The total costs $V(p, x')$ and $V(p, x'')$ attributable to x'

Period	**Demand**	**Production Cost**	**Holding Cost**
1	1	$K_1 + v_1 x$	$h_1 x$
2	2	$K_2 + v_2 x$	$h_2 x$

Table 2.3: General Problem Parameters for Cyclic Forecast under Fixed Plus Linear Costs

Period	**Demand**	**Production Cost**	**Holding Cost**
1	1	$K_1 + 2x$	$2.5x$
2	2	$2 + 3x$	$1.5x$

Table 2.4: Problem Parameters for Non-Well-Posed Cyclic Production Planning Forecast

and x'', respectively, can be calculated as

$$V(p, x') = \frac{K_1 + 3v_1 + 2\alpha h_2}{1 - \alpha^2} \tag{2.3}$$

$$V(p, x'') = K_1 + v_1 + \alpha \left(\frac{K_2 + 3v_2 + \alpha h_1}{1 - \alpha^2} \right) \tag{2.4}$$

Setting $V(p, x') = V(p, x'')$, we can obtain an expression for any one of the cost parameters in terms of the remaining parameters so that the two total costs are equal. Choosing K_1 to be the dependent variable, we can write it as

$$K_1 = \frac{-2v_1 - 2\alpha h_2 - \alpha^2 v_1 + \alpha K_2 + 3\alpha v_2 + \alpha^2 h_1}{\alpha^2}. \tag{2.5}$$

With K_1 thus chosen, if x' is an optimal solution for p, then so is x'', and vice versa. Of course, the other parameters must satisfy some feasibility constraints for K_1, namely, that (2.5) results in a nonnegative right hand side. For a variety of values for the cost parameters and discount factor, the production plans x' and x'' are indeed optimal. For example, consider the following instance of the above form, where K_1 can be determined by the other parameters (its value is 4.50 in this case). To see that the only two infinite horizon optimal initial decisions for the above problem are to produce for one and two periods, respectively, consider any production plan which produces for the first $n \geq 3$ periods initially. Then one must solve either the problem p again or the problem p' (although discounted by a factor of α^n) where the even-numbered periods of p' have parameters equal to those of the odd-numbered periods of p and vice versa in period $n+1$. However, it can be shown that total cost for both p and p' is minimized by production in every other period (as opposed to every one, three, four, five, etc. periods), and hence, there is no incentive to use any production plan other than x' and x''. Thus, since $V(p, x') = V(p, x'')$, we have that x' and x'' are both infinite horizon optimal for p. Clearly, x' and x'' do not coalesce, and this problem instance is not well-posed. We conclude that it is not true in general that all cyclic problems (or even all cyclic problems of cycle length 2) are well-posed with respect to the class $\mathcal{P}$ of forecasts.

On the other hand, consider the problem class $\mathcal{P}'$, which is equivalent to $\mathcal{P}$, but with the further restriction that all of its problems are cyclic with cycle length at most N (a fixed and known value imposed upon the oracle). Here, once the first $N+1$ periods' worth of p are known, the entire problem is known. Moreover, all future extensions are also known since $\mathcal{P}'^N(p) = \{p\}$, a singleton. Thus, any infinite horizon optimal initial decision for p is also optimal for all $p' \in \mathcal{P}^N(p)$, and p is well-posed by definition. Interestingly, by this argument, the problem p in the example above is well-posed *with respect to* $\mathcal{P}'$, even though we have shown that p is not coalescent. This apparent contradiction can be explained by the fact that $\mathcal{P}' \neq \mathcal{P}$.

2.4 Solution Procedures

In general, one cannot know a priori if a production planning problem has unique infinite horizon optimal initial decisions or if it has stationary parameters. However, there appears to be consensus in the DLSM forecast horizon literature that forecast horizons exist for virtually all randomly generated instances. Because of the impossibility of forecasting infinite nonstationary data, it is in general necessary to verify algorithmically (rather than by inspection) that a problem is well-posed.

2.4.1 A New Stopping Rule Algorithm

As discussed in §2.1, stopping rules developed to solve the DLSM may fail in the case of general concave costs. Others developed for general concave cost production planning problems may become computationally expensive as they fail to adequately trim the list of potential optimal last production points. It will be shown that the solution algorithm in this section will solve *any* solvable concave cost problem, and that it runs with effort that is constant, even as the horizon length grows arbitrarily long. Recall that L represents an (integer) upper bound on the number of periods over which it can be optimal to carry a unit of inventory. With L thus defined, any interval $[N, N+L]$ is a regeneration set, and a solution algorithm for a given problem p immediately follows.

Algorithm 2.1.

1. *Set* $M = 1$.
2. *If* $\bigcap_{n=M}^{M+L} \mathcal{X}_1^*(p^n) \neq \emptyset$, *stop. Return* $N^* = M$ *and any* $x^* \in \bigcap_{n=M}^{M+L} \mathcal{X}_1^*(p^n)$.
3. *Otherwise, set* $M \leftarrow M+1$. *Return to 2.*

Let

$$L(n) = \{j : j \text{ is an optimal last production point for the } n\text{-period problem}\},$$

and let $l(n) = \min\{j : j \in L(n)\}$, the minimum last optimal production point for the n-period problem. We rely upon Assumption 2.2 in that for any forecast in $\mathcal{P}$ and any n, $l(n)$ is a member of the interval $[n-L, n]$. Thus, one must simply store the set of optimal initial production decisions for each finite horizon problem and look at the intersection of those sets over $L+1$ consecutive finite horizon problems. The exact method of determining and storing $\mathcal{X}_1^*(p^n)$ in Step 2 may be determined by the implementor. In practice, it is generally the case that when multiple optimal solutions exist for some finite horizon problem p^n, they are relatively few in number.

Let $d_{i,j} = \sum_{k=i}^{j} d_k$. If $F(n)$ represents the cost of an optimal solution for p^n, then one would need to store the value of $F(n)$ and the set of optimal predecessors

$$S(n) = \{j < n | F(j) + \alpha^j c_{j+1}(d_{j+1,n}) + \sum_{k=j+2}^{n} \alpha^{k-1} h(d_{k,n}) = F(n)\}.$$

The following theorem shows that the algorithm will solve every solvable problem in $\mathcal{P}$.

Theorem 2.3. *Algorithm 2.1 will finitely terminate if and only if p is well-posed.*

Proof. Suppose that p is well-posed. Then there exists some period N^* such that $\mathcal{X}_1^{*N^*}(p) \neq \emptyset$. In particular, since the problems $\underline{p}^n, n = N^*, \ldots, N^* + L$ are in $\mathcal{P}^{N^*}(p)$, there is at least one x^* which is optimal for each of the $\underline{p}^n$ and which is also optimal for all $p' \in \mathcal{P}^{N^*}(p)$. Moreover, any optimal solution to the problem p^n also is an optimal policy for the first n periods of the problem $\underline{p}^n$ and vice versa. Thus, x^* is optimal for $p^n, n = N^*, \ldots, N^* + L$ as well, and it will be detected by Algorithm 2.1.

Now suppose that Algorithm 2.1 stops at some period N^* in solving problem p. Then there is an initial decision x_1^* which is optimal for each of the problems p^n, $n = N^* - L, \ldots, N^*$. By definition of L, any forecast of horizon length greater than N^* must have at least one zero-inventory point in the interval $[N^* - L, N^*]$. This includes the infinite horizon forecast p, as well as all forecasts $p' \in \mathcal{P}^{N^*}(p)$. By the Principle of Optimality, x_1^* is an infinite horizon optimal initial decision for all $p' \in \mathcal{P}^{N^*}(p)$, so $\mathcal{X}_1^{*N^*}(p) \neq \emptyset$ and p is well-posed. $\square$

If $p \in \mathcal{P}$ is not well-posed, then it should be highlighted that *no* exact solution algorithm (including this one) will detect an infinite horizon optimal initial decision. When the problem is not well-posed, no matter how many oracle calls are made, the algorithm will always find some period whose set of optimal initial decisions has none in common with those of the previous period. The algorithm will run interminably in this case.

If p is well-posed, it is possible to find an arbitrary number of periods' worth of infinite horizon optimal decisions by iteratively applying Algorithm 2.1. If x_1^*, the initial infinite horizon optimal decision as found by the algorithm, satisfies demand for the first n_1 periods, then one can roll forward to period $n_1 + 1$ and apply the algorithm anew. As a final clarifying remark, in the spirit of the discussion in 2.3.3, we note that the stopping rule algorithm will solve any well-posed problem with respect to the problem class $\mathcal{P}$, but we make no such guarantees for other problem classes. This statement should in no way diminish the power of Algorithm 2.1, for $\mathcal{P}$ is a fairly generic problem class. As we saw in 2.3.3, difficulties arose when we *restricted* $\mathcal{P}$. Thus, well-posedness may be a more *useful* (meaning easy to verify algorithmically) property when applied to a large problem class like $\mathcal{P}$, whereas a larger proportion of forecasts may be well-posed within a smaller problem class like $\mathcal{P}'$.

2.4.2 A Note Regarding Computational Complexity

The number of computations in an iteration of Algorithm 2.1 is highly dependent upon the bound L. The algorithm must find the optimal initial decisions for the first $L + 1$ finite horizon problems in the initial iteration. If it fails to stop at that time, then the optimal initial decisions must be found for the finite horizon problem of length $L + 2$ and compared to those of the previous L, and so on. Each finite horizon problem requires $2L$ additions and comparisons (and $2n$ when $n < L$) in order to find its optimal initial decision(s), provided that the costs to go are immediately available. However, the heaviest computational burden is incurred by updating the costs to go from each of periods $n - L, n - L + 1, \ldots, n - 1$ to period n at each iteration. This requires $O(L^2)$ work.

Admittedly, if both L and the forecast horizon N^* are large, Algorithm 2.1 can be computationally expensive. However, by choosing L to be as small as possible, one can minimize excess computations. Therefore, the algorithm is best suited to problems in which one can make accurate estimates of the bounds on the cost functions.

2.4.3 Solving A Problem Chosen at "Random"

The results of the authors mentioned earlier in this chapter agree in that nearly all DLSM instances, when randomly generated, can have their infinite horizon optimal initial decision found using a relatively small number of finite horizon problems. We wish to reinforce the notion that randomly generated problems are almost always well-posed by choosing a problem whose demand forecast all readers are familiar with: the digits of the constant π. We will show that for a variety of cost parameters, our stopping rule algorithm is satisfied and yields an infinite horizon optimal initial decision.

We let $\alpha = .99$, $c_n(x) = \begin{cases} K, & x > 0 \\ 0, & x = 0 \end{cases}$, and $h_n(x) = hx$ for all n. We choose this structure for simplicity: for fixed d, an optimal production plan remains optimal for any other values of K and h as long as the ratio K/h remains constant. Let $r = K/h$. By holding the cost parameters constant over time, it is also easy to derive the value of L. In this case, L is conservatively defined (since demand can be zero in some periods) as

$$L = \arg\min_n \left\{ K + \sum_{j=2}^{n} \alpha^{j-1} h > \alpha^{n-1} K \right\}, \tag{2.6}$$

which can be reduced to

$$L = \left\lceil \log_\alpha \left[\left(\frac{\alpha}{1-\alpha} \right) \left(\frac{1}{1-\alpha} + \frac{r}{\alpha} \right)^{-1} \right] \right\rceil. \tag{2.7}$$

Table 2.5 shows the optimal production points for the incrementally longer finite horizon problems, until satisfying Step 2 of Algorithm 2.1. The lines within the columns denoting the optimal production points indicate the beginning and end of the interval over which the stopping rule is satisfied. The length of the interval depends upon the value of L, and L varies according to r. When the number of

Period	Demand	**Optimal setups** $r = 5$ ($L = 5$)	**Optimal setups** $r = 10$ ($L = 11$)	**Optimal setups** $r = 20$ ($L = 20$)
1	3	1	1	1
2	1	1	1	1
3	4	1	1, 3	1
4	1	1, 3	1, 3	1
5	5	1, 3, 4, 5	1, 3, 5	1, 5
6	9	1, 3, 4, 5, 6	1, 3, 5	1, 5
7	2		1, 3, 5, 6	1, 5
8	6		1, 3, 5, 8	1, 5
9	5		1, 3, 5, 8	1, 5, 8
10	3		1, 3, 5, 8	1, 5, 8
11	5		1, 3, 5, 8, 10	1, 5, 8
12	8		1, 3, 5, 8, 11	1, 5, 8, 11
13	9			1, 5, 8, 12
14	7			1, 5, 8, 11, 13
15	9			1, 5, 8, 12, 14
16	3			1, 5, 8, 11, 13, 15
17	2			1, 5, 8, 11, 13, 15
18	3			1, 5, 8, 11, 13, 15
19	8			1, 5, 8, 11, 13, 15, 18
20	4			1, 5, 8, 11, 13, 15, 19
21	6			1, 5, 8, 11, 13, 15, 19
22	2			1, 5, 8, 11, 13, 15, 19
23	6			1, 5, 8, 11, 13, 15, 19, 21
24	4			
25	3			

Table 2.5: Forecast Horizon Results for Production Planning Example

finite horizon problems in agreement equals L for that column, then the stopping rule algorithm is satisfied. One observation from the above solutions is that for every finite horizon problem solved for each value of r, the optimal production plan is unique. As we have seen, infinite horizon problems with unique optimal production plans are well-posed, and so the experimental results here are encouraging. Moreover, storage and comparison of the set of optimal initial decisions in step 2 of Algorithm 2.1 is facilitated by having unique optimal initial decisions.

2.4.4 Experimental Results for Algorithm 2.1

We conducted a series of simulation experiments to investigate the sensitivity of Algorithm 2.1 to (i) changes in the magnitude of L and (ii) changes in the variability of demand. To do this, we first derived a general formula for L, which assumes knowledge of the minimum and maximum marginal values of the production and holding costs. Specifically, if $\underline{c} > 0$ and $\bar{c} < \infty$ are upper and lower bounds on the marginal production costs and $\underline{h}$ is a lower bound on the marginal holding cost, then L is the minimum integer N satisfying

$$\underline{c} + \sum_{n=0}^{N-1} \alpha^n \underline{h} > \alpha^N \bar{c}, \tag{2.8}$$

which can be calculated as

$$L = \left\lceil \log_\alpha \left[\frac{\underline{c} + \frac{\underline{h}}{1-\alpha}}{\bar{c} + \frac{\underline{h}}{1-\alpha}} \right] \right\rceil + 1 \tag{2.9}$$

Notice that L is not sensitive to the actual values of demand. For simplicity, we generated the costs as fixed-plus-linear for production and linear for holding. In this case, $\bar{c}$ can be expressed as $\bar{c} = \tilde{c} + \bar{K}$, where $\tilde{c}$ is a bound on the unit production cost and $\bar{K}$ is a bound on the setup costs. All cost parameters and demands were correlated from period to period, but independent of one another, with a 50% chance of increasing or decreasing and an allowable increase or decrease of up to 20% (uniformly distributed) of the distance to the upper or lower bound. There were four combinations of cost parameters and two levels of demand variability, making eight different cases in total, with 1000 runs for each. The summary statistics for these cases lie in Table 2.6, where the last two columns show the average and standard deviation of the calculated forecast horizons for the 1000 runs of each case.

As can be seen in Table 2.6, the average forecast horizon for the cases considered closely follows L, plus a modest linear function of the ratio of the maximum to minimum production cost plus another modest linear function of the minimum holding cost. This is intuitive: smaller marginal production and holding costs tend to encourage producing in larger batches, so the optimal initial production quantities will

Max demand	Prod costs	Setup costs	Holding costs	L	Avg FH	St Dev FH
16	[15, 30]	[50, 100]	[4, 8]	29	34.3	3.13
16	[5, 40]	[50, 100]	[4, 8]	34	38.6	2.01
16	[15, 30]	[50, 100]	[8, 16]	16	19.3	1.08
16	[15, 30]	[50, 100]	[2, 4]	55	62.5	3.24
8	[15, 30]	[50, 100]	[4, 8]	29	37.2	4.07
8	[5, 40]	[50, 100]	[4, 8]	34	41.7	3.46
8	[15, 30]	[50, 100]	[8, 16]	16	21.1	2.80
8	[15, 30]	[50, 100]	[2, 4]	55	66.3	5.11

Table 2.6: Forecast Horizon Results for Production Planning Simulations

tend to cover a larger number of periods than they would if those costs were higher. However, the number of periods required for a forecast horizon in excess of L is small compared to L on average in all cases.

As for the second hypothesis, lower demand variability tends to increase the forecast horizons. Again, this is an intuitive result: higher volatility in demand would accelerate the rate of setups as higher demand quantities become more expensive to build ahead in inventory. This agrees with the experimental results reported in Lundin and Morton [32]. Figure 2.2 shows a comparison of the average forecast horizon versus L for the high demand variability (max_d = 16) and for the low demand variability (max_d = 8) and suggests that increasing demand variability modestly decreases the average forecast horizon. Perhaps the most significant observation from this study is that actual forecast horizons computed with Algorithm 2.1 depend primarily upon the value of L and less upon problem-specific parameters.

2.4.5 Solving for an Entire Infinite Horizon Optimal Production Schedule

The reader may observe that Algorithm 2.1 only yields one decision at a time, and that well-posedness of a problem instance only ensures that at most a few periods' worth of optimal decisions can be determined immediately. This is because in general,

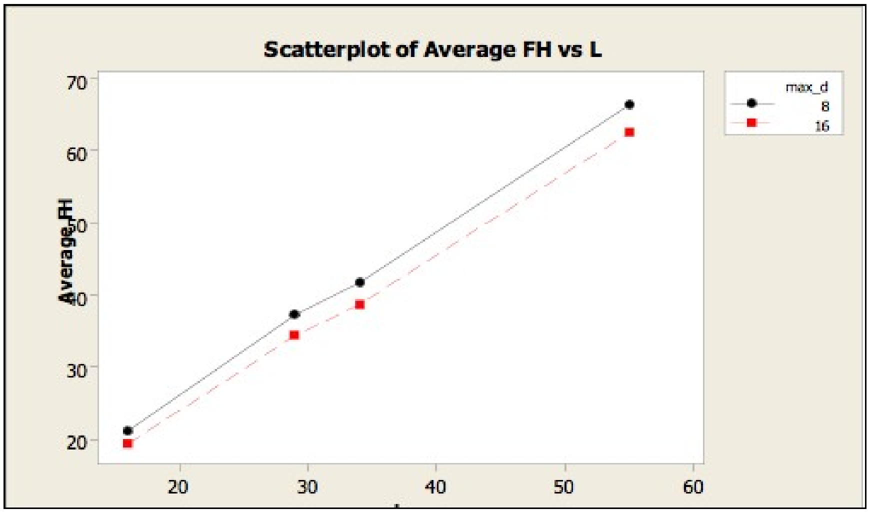

Figure 2.2: Graph of Average Forecast Horizon Versus L for High and Low Demand Variability

only one production decision needs to be made at a time. However, after each decision is made, time rolls forward to the next period and the next decision confronts the decision maker. In effect, the problem is inherently different once time rolls forward. The problem parameters for the periods whose optimal production decisions have been made become irrelevant once those decisions are fixed. In this way, well-posedness can be tested each time the problem rolls forward. Clearly, if a forecast has the property that a forecast horizon exists for *each* of its optimal production decisions, then the forecast is infinite horizon solvable in an algorithmic sense, where repeated calls to Algorithm 2.1 uncover the optimal decisions sequentially, but not all at once. Since the problem class $\mathcal{P}$ is, in general, nonstationary, this algorithmic sense of infinite horizon solvability is the best we can hope for.

2.5 Conclusions

We have considered the class of single-item, deterministic production planning problems under concave costs, discounting, and bounded demands. We focused on solving such problems when the study horizon is infinite. Because solving such problems requires in general an infinite amount of data, we identified coalescence as a necessary and sufficient condition under which an optimal initial (or initial few) decisions can be found using only finitely many periods worth of parameters. When this occurs, we call the problem well-posed. We gave a new solution algorithm that will solve any well-posed problem. Conversely, any non-well-posed problem will fail to be solved by this algorithm or by *any* other exact algorithm. By recursively applying the algorithm, problems possessing a forecast horizon for each period's optimal production decision can also be solved in their entirety.

Chapter 3

An Infinite Horizon Equipment Replacement Model

3.1 Introduction

The equipment replacement problem has a rich literature dating back to the mid-20th century. See Fraser and Posey [25] and Pierskalla and Voelker [34] for more detailed reviews. A firm requires a certain type of equipment in order to operate. That equipment must be purchased and maintained, it may produce revenues or have operating costs or both, and it may have a salvage value (either an income when the equipment is sold or an expense to junk the equipment) at the end of its useful lifetime. All of these cost and revenue sources may vary over time, or by the brand of equipment purchased, or both. At each time step, there is a set of challengers available for purchase. The objective is to schedule machine purchases and disposals in such a way as to minimize total discounted cost.

Classical work on the topic either assumed the problem to extend over only a finite horizon, or to have stationary decision alternatives and cash flows over time, or both. Among the first attempts to break from that tradition is the paper of Sethi and Chand [39], in which they consider a dynamic, infinite horizon equipment replacement problem in an improving technological environment. Because of the impossibility of capturing the entire sequence of decision alternatives and cash flows for an infinite horizon, nonstationary problem, they develop methods for finding only the optimal *initial* decision (to replace or not to replace the existing equipment) while knowing the problem parameters only through some finite horizon, called a *forecast horizon*. In [12], they generalize upon the model in [39] by allowing for multiple replacement alternatives at each replacement period and obtain similar results. Goldstein et al [26] improve upon the computational burden of the algorithm in [12].

Without assuming a particular technological environment, Bean, Lohmann, and Smith [3] generalize the finite horizon equipment replacement model of Oakford, Lohmann, and Salazar [33] to an infinite horizon model, allowing for time-varying sets of challengers and cash flows, as long as the maximum equipment lifetime is known and finite. They discuss a procedure that will uncover an optimal current decision, regardless of potential future cash flows and challengers, whenever it is possible to do so, and call the time at which that occurs an *equivalent finite horizon*. In [4], under the assumption of a non-degrading technological environment, the same authors formulate the infinite horizon, nonstationary, equipment replacement problem with m challengers available as a math program using vector notation. Rather than looking for an exact forecast horizon, they give error bounds on using a finite horizon truncation of the infinite horizon forecast.

3.2 Problem Description and Notation

The model under consideration offers a single choice of replacement equipment (challenger) in each period, with the decision consisting of the duration of the equipment's lifetime before replacement. Since the firm expects to be in business indefinitely long, an infinite horizon approach is appropriate for the model. However, the firm should account for the possibility of either ceasing its operations, or at least, of the obsolescence of the type of equipment. This consideration will appear explicitly in the model parameters below. We also adopt the convention that the equipment does not generate revenue, though this convention could just as easily be discarded through a simple modification of the operating costs.

3.2.1 Notation

$$\alpha \quad = \quad \text{discount factor, where } 0 < \alpha < 1.$$

$$L \quad = \quad \text{maximum physical lifetime of equipment in number of periods.}$$

$o_n(m)$ = undiscounted operating cost ($o_n(m) \geq 0$) of equipment in its mth period of operation ($1 \leq m \leq L$) that was purchased in period n.

p_n = undiscounted purchase cost of equipment purchased in period n.

$s_n(m)$ = undiscounted salvage value ($s_n(m) \geq 0$) of equipment purchased in period n and salvaged after lifetime of m periods.

$c_n(m)$ = $p_n - \alpha^m s_n(m) + \sum_{k=n}^{n+m} \alpha^{k-n} o_n(k-n+1)$, total cost of purchasing new equipment in period n and retaining it for m periods.

b_n = binary variable denoting whether the equipment is required in period n. If $b_n = 1$, then equipment is required in period n. If $b_n = 0$, then equipment is not required in period n, and $b_m = 0$ for all $m > n$.

b = $(b_1, b_2, \ldots)^T$, the business requirement vector.

b^n = $(b_1, b_2, \ldots, b_n, 0, 0, \ldots)^T$: the business requirement vector in which the equipment becomes obsolete at the conclusion of period n.

$$c = \left(\begin{bmatrix} c_1(1) & \cdots & c_1(L) \\ c_2(1) & \cdots & c_2(L) \\ \vdots & \vdots & \vdots \end{bmatrix}, b \right) : \text{ the } \textit{problem} \text{ (or equivalently, } \textit{forecast}\text{).}$$

$\mathcal{C}$ = class of all permissible forecasts.

$\mathcal{C}^N(c)$ = class of all permissible forecasts with the added restriction that for $c' \in \mathcal{C}^N(c)$, $c'_n(m) = c_n(m)\ \forall n = 1, \ldots, N$ and $m = 1, \ldots, L$; i.e., the class of permissible forecasts identical to c through period N.

$$\underline{c}^N = \left(\begin{bmatrix} c_1(1) & \cdots & c_1(L) \\ c_2(1) & \cdots & c_2(L) \\ \vdots & \vdots & \vdots \end{bmatrix}, b^N \right) : \text{ a forecast identical to } c\text{, but in which}$$

the equipment becomes obsolete beginning period $N + 1$.

$$x_n = \begin{cases} 0, & \text{if keeping current equipment in period n} \\ j, & \text{if purchasing equipment for scheduled lifetime } j \text{ in period } n \end{cases}$$

x = $(x_n)_{n=1}^{\infty}$: a *replacement strategy.* If $x_n > 0$, then period n is called a *replacement period* under strategy x.

x^N = $(x_n)_{n=1}^{N}$: a N-horizon replacement strategy.

$V(c, x)$ = $\sum_{n=1}^{\infty} \alpha^n c_n(x_n)$: total cost of replacement strategy x for forecast c.

$\mathcal{X}^*(c)$ = set of all infinite horizon optimal replacement sequences for forecast c.

$\mathcal{X}_1^*(c)$ = set of all infinite horizon optimal initial equipment purchases for forecast c.

$\mathcal{X}_1^{*n}(c)$ = $\bigcap_{c' \in \mathcal{C}^n(c)} X_1^*(c')$: the set of initial decisions that are infinite horizon optimal for all $c' \in \mathcal{C}^n(c)$

Note that if there were multiple challengers in each period, we could reduce the possible selections to a singleton for each combination of period n and feasible lifetime m. Let the cost of purchasing equipment j in period n for lifetime m be $c_n^j(m)$. If there are K equipment choices in period n with feasible lifetimes of m periods, then remove from consideration all equipment $j, 1 \leq j \leq K$ such that $c_n^j(m) > \min_k c_n^k(m)$ and set $c_n(m) = \min_k c_n^k(m)$. We are then indifferent as to the equipment choice as long as it attains the minimum for its combination of replacement period and lifetime. Thus, the results in this chapter apply equally well to the case of multiple replacement equipment options. Also, if it is infeasible to have an operation lifetime of m or more for equipment purchased in period n, then set $o_n(m) = \infty$.

3.2.2 Assumptions and Preliminaries

The following two assumptions are fundamental to the rest of this chapter.

Assumption 3.1. *$s_n(m) \geq s_{n+1}(m)$ for all n and m; i.e. salvage value is nonincreasing in the age of a piece of equipment. In other words, there is no antiquing of equipment.*

Assumption 3.2. *For all $c \in \mathcal{C}$, every feasible replacement strategy has finite total cost.*

We define $\mathcal{C}$ as the set of all equipment replacement forecasts c satisfying Assumptions 3.1-3.2. We now present the discounted equipment replacement problem with no antiquing, as described above, as an optimization problem.

$$\begin{array}{rll} \text{minimize} & \sum_{n=1}^{\infty} \alpha^n c_n(x_n) & \\ \text{subject to} & z_n(x) = \max_{m \le n}[m + x_m - n], & n = 1, 2, \ldots \\ & z_n(x) \ge b_n, & n = 1, 2, \ldots \\ & x_n \in \{0, 1, \ldots, L\}, & n = 1, 2, \ldots \end{array} \tag{3.1}$$

We will call $z_n(x)$ the *remaining lifetime in period n under replacement strategy x.* Thus, the first two constraints in (3.1) represent the requirement that there is a piece of equipment on hand at all required times under business requirement forecast b. Assumption 3.1 results in the first lemma below, while Assumption 3.1 leads directly to the second.

Lemma 3.1. *Suppose that in forecast c, $b_{n-1} = 1$ and $b_n = 0$ (equipment becomes obsolete beginning period n). Then the total cost obtained by salvaging the terminal equipment beginning period n is at most that of keeping it for any additional time beyond period n.*

Proof. Because $b_n = 0$, one need not have any equipment on hand beginning period n. Thus, it would only prove optimal to retain the terminal equipment (purchased in some period m) into period n or beyond if $c_m(n-m) < c_m(n-m-1)$. But by the nonnegative operating costs and Assumption 3.1, $c_n(j) \le c_n(j+1)$ for all $j \le L$ and n. We conclude that it is at least as cheap to salvage the terminal equipment at the close of business as it is to retain it any further. □

By Lemma 3.1, an optimal solution to problem $\underline{c}^N$ has total cost at most that of any optimal solution to any other $c' \in \mathcal{C}^N(c)$ (with business requirement vector b') with $b'_N = 1$. We also have the following corollary, given without proof.

Corollary 3.2. *An optimal solution for an N-horizon problem whose cost and business requirement forecasts are the same as those of the first N periods of c is also optimal for the first N periods of problem $\underline{c}^N$.*

Lemma 3.3. *Without loss of optimality, we can restrict consideration to those strategies for c in which $x_n > 0$ if and only if $z_{n-1}(x) \cdot b_n = 1$ for all $n \ge 1$. In other words,* *one need only consider replacement strategies in which new equipment is purchased if and only if the current equipment has completed its scheduled lifetime (and the equipment is still required).*

Proof. First, observe that if $z_{n-1}(x) \cdot b_n = 1$, then there will be no equipment on hand beginning period n. Since the equipment is still required at period n, it is then necessary to purchase equipment with some lifetime $k \in \{1, \ldots, L\}$. On the other hand, consider the case where $z_{n-1}(x) > 1$. This implies that the current equipment has at least one more period of remaining scheduled lifetime. Let $l < n$ be the period in which the current equipment was purchased and m' its scheduled lifetime where $m' + l \geq n$. Then the cost due to the current equipment is $c_l(m')$.

Suppose the total cost over the first $l+m'$ periods is minimized by setting $x_k > 0$ for some $k \in \{n \ldots, m+m'\}$. By nonnegativity of the operating costs and by Assumption 3.1, total cost could be reduced further by salvaging in period k (rather than in period $l + m'$) the equipment purchased in period l. This contradicts the assertion that the total cost over the first $l+m'$ periods was minimized by setting $x_k > 0$, and the proof is complete. □

3.3 Problem Solvability

In this section, we establish the meaning of coalescence and well-posed for equipment replacement problems and show that the two properties are equivalent.

3.3.1 Additional Definitions

Definition 3.1. *Let x^* and x^{**} be optimal replacement strategies for problem c. We say that x^** reaches *x^{**} if there exists $y^* \in \mathcal{X}^*(c)$ such that $z_M(y^*) = z_M(x^*)$ and $z_N(y^*) = z_N(x^{**})$ for some M and N, $M \leq N$.*

In other words, one optimal replacement strategy reaches another if there is some way of optimally following the first strategy for some initial number of periods, then following some optimal continuation to share a common replacement period with the second strategy. By iteratively reaching, it may be possible to reach many different optimal replacement strategies from a single optimal initial decision. In the special case that all optimal replacement strategies can be reached from one or more strategies sharing an optimal initial decision, the problem has the following characteristic.

Definition 3.2. *We say that c satisfies the* coalescence condition *(or is* coalescent*) if there exists an optimal* source strategy $x^* \in \mathcal{X}^*(c)$ *such that any other optimal replacement strategy for c can be reached from x^*.*

As we did in Chapter 2 for production planning, we define the terms well-posed and forecast horizon as they pertain to the infinite horizon equipment replacement problem.

Definition 3.3. *Forecast $c \in \mathcal{C}$ is* well-posed *if there exists some period N^* such that $\mathcal{X}_1^{*N^*}(c) \neq \emptyset$.*

Well-posedness enables us to directly identify infinite horizon problems which cannot be solved by any finite horizon techniques, as we will soon see. It is also important to formalize the definition of a forecast horizon for the equipment replacement problem.

Definition 3.4. *Period N^* is a* forecast horizon *for class $\mathcal{C}^{N^*}(c)$ if there exists an infinite horizon optimal initial decision x_1^* such that $x_1^* \in \mathcal{X}_1^{*N}(c)$ for all $N \geq N^*$.*

In other words, existence of a forecast horizon N^* implies that one can find an initial optimal solution to any problem in agreement with c through the first N^* periods by solving the finite horizon problem c^{N^*}.

3.3.2 Necessary and Sufficient Conditions for Well-Posed Problems

The proof of the main theorem to follow will flow more easily by first presenting a technical lemma. See Appendix B.1 for a proof.

Lemma 3.4. *Consider a forecast c and a sequence of forecasts $\{c^{(n)}\}_{n=1}^{\infty}$ where $c^{(n)} \in \mathcal{C}^n(c)$ $\forall n$. Then there is some subsequence $\{n_k\}_{k=1}^{\infty}$, a corresponding sequence of optimal replacement strategies $\{\hat{x}^{*(k)} \in \mathcal{X}^*(c^{(n_k)})\}_{k=1}^{\infty}$, and $x^* \in \mathcal{X}^*(c)$ such that $\hat{x}^{*(k)} \to x^*$.*

We now present the main result of this section.

Theorem 3.5. *Problem $c \in \mathcal{C}$ is well-posed if and only if it satisfies the coalescence condition.*

Proof. Suppose that c is well-posed. Choose $x^* \in \mathcal{X}^*(c)$ and N such that x_1^* is an initial optimal decision for all $c' \in \mathcal{C}^N(c)$. Then in particular, x_1^* is an optimal initial decision for $\underline{c}^n$ (and hence c^n) for all $n \geq N$. By Lemma 3.3 and the definition of L, every optimal strategy for c must have a replacement period in $\{N, N+1, \ldots, N+L\}$. Since x_1^* is an initial optimal decision for each of the problems $\underline{c}^N, \underline{c}^{N+1}, \ldots, \underline{c}^{N+L}$, it is straightforward how to construct an optimal replacement strategy beginning with x_1^* to any other optimal plan for c. Thus, well-posedness implies coalescence.

Now suppose that c is coalescent but is not well-posed. Let x^{**} be a source strategy. Then there exists a subsequence $\{n_j\}_{j=1}^{\infty}$ and $c^{(j)} \in \mathcal{C}^{n_j}(c)$ such that $x_1^{**} \notin \mathcal{X}_1^*(c^{(j)})$ for all j. But Lemma 3.4 states that, resorting to a further subsequence $\{n_{j_k}\}_{k=1}^{\infty}$, there exists some sequence of optimal replacement strategies $x^{*(k)} \in \mathcal{X}^*(c^{(n_{j_k})})$ for all k such that $x^{*(k)} \to \tilde{x}^* \in \mathcal{X}^*(c)$.

By assuming that c is coalescent, $\tilde{x}^*$ is reached at some period M by some optimal strategy with initial decision x_1^{**}. Let $K(M)$ be large enough that $x^{*(k)}$ is in agreement with $\tilde{x}^*$ up through and including period M for all $k \geq K(M)$. Note that such a $K(M)$ exists in light of Lemma 3.4. By the Principle of Optimality, there now exists a $c^{(k)}$-optimal plan $\forall k \geq K(M)$ with initial decision x_1^{**}, and which is in agreement with $x^{*(k)}$ thereafter. Hence, $x_1^{**} \in \mathcal{X}_1^*(c^{(k)})$ $\forall k \geq K(M)$. We've reached a contradiction and thus conclude that satisfying the coalescence condition is sufficient for well-posedness. □

3.3.3 Solution Algorithm

If $c \in \mathcal{C}$ is not well-posed, then it should be highlighted that *no* exact solution algorithm will detect an infinite horizon optimal initial decision. It is impossible to amass in finite time all parameters of an infinite horizon equipment replacement problem belonging to $\mathcal{C}$. Thus, a solution algorithm must rely on making decisions based on finite horizon problems.

Bean, Lohmann, and Smith [3] present a solution procedure for the infinite horizon equipment replacement model. We give it here under the notation of this chapter, and show that the procedure will solve any well-posed equipment replacement problem. It will also be shown that the solution procedure will run with effort that is constant, even as the horizon length grows arbitrarily long. As before, let L represent the maximum feasible lifetime.

Algorithm 3.1.

1. *Set* $M = 1$.
2. *If* $\bigcap_{n=M}^{M+L} \mathcal{X}_1^*(c^n) \neq \emptyset$, *stop. Return* $N^* = M$ *and any* $x^* \in \bigcap_{n=M}^{M+L} \mathcal{X}_1^*(c^n)$.
3. *Otherwise, set* $M \leftarrow M + 1$. *Return to 2.*

One must simply store the set of optimal initial purchases for each finite horizon problem and look at the intersection of those sets over $L + 1$ consecutive finite horizon problems. The exact method of determining and storing $\mathcal{X}_1^*(c^n)$ in Step 2 may be determined by the implementor. In practice, it is generally the case that when multiple optimal solutions exist for some finite horizon problem c^n, they are relatively few in number. This algorithm is only intended to be a framework. Nevertheless, it can be shown to have desirable properties.

Theorem 3.6. *Algorithm 3.1 will finitely terminate if and only if c is well-posed.*

Proof. Suppose that c is well-posed. Then there exists some period N^* such that $\mathcal{X}_1^{*N^*}(c) \neq \emptyset$. In particular, since the problems $\underline{c}^n, n = N^*, \ldots, N^* + L$ are in $\mathcal{C}^{N^*}(c)$, there is at least one x^* which is optimal for each of the $\underline{c}^n$ which is also optimal for all $c' \in \mathcal{C}^{N^*}(c)$. Moreover, by Lemmas 3.1 and 3.3, any optimal solution to the problem c^n also is an optimal strategy for the first n periods of the problem $\underline{c}^n$ and vice versa. Thus, x^* is optimal for $c^n, n = N^*, \ldots, N^* + L$ as well, and it will be detected by Algorithm 3.1.

Now suppose that Algorithm 3.1 stops at some period N^* in solving problem c. Then there is an initial optimal solution x^* which is optimal for each of the problems c^n, $n = N^* - L, \ldots, N^*$. By definition of L, any forecast of horizon length greater than N^* must have at least one purchase point in the interval $\{N^* - L, \ldots, N^*\}$. This includes the infinite horizon forecast c, as well as all forecasts $c' \in \mathcal{C}^{N^*}(c)$. By the Principle of Optimality, x_1^* is an optimal initial decision for all $c' \in \mathcal{C}^{N^*}(c)$, so $\mathcal{X}_1^{*N^*}(c) \neq \emptyset$ and c is well-posed. $\square$

If c is well-posed, it is possible to find an arbitrary number of periods' worth of optimal decisions by iteratively applying Algorithm 3.1. If x_1^*, the initial infinite horizon optimal decision as found by the algorithm, purchases an equipment for the first n_1 periods, then one can roll forward to period $n_1 + 1$ and apply the algorithm anew.

By definition, a non-well-posed problem instance always has a future period whose parameters will render the current set of finite horizon optimal decisions non-optimal when solving the finite horizon problems corresponding to that period. In other words, a solution algorithm must make repeated calls to an oracle to discover future problem parameters. When the problem is not well-posed, no matter how many calls to the oracle are made, the algorithm will always find some period whose set of optimal solutions has none in common with those of the previous period. The algorithm will run interminably in this case.

3.3.4 Comments on Solvability of Special Problem Structures

The problem class $\mathcal{C}$ contains a vast number of special problem structures, including various forms of technological improvement, technological degradation, and technological stationarity. Researchers have investigated the *form* of optimal policies under special cases of technological change. This is discussed somewhat in Cheevaprawatdomrong and Smith [15], who prove the paradoxical result that under geometrically improving technology costs, the effect of increasing the geometric factor is to *decrease* the optimal rate of equipment replacement.

It would be desirable to make a claim about the solvability of some dynamic subclass of problems. Unfortunately, without making specific structural assumptions, the lack of knowledge about the entire stream of problem parameters precludes the ability to prove that the problem is well-posed. However, for the subclass of problems that have stationary cash flows and business requirements for a sufficiently long initial period of time, it is possible to show that such problems are well-posed, even in the presence of the possibility of nonstationary parameters further in the future. Solution procedures for the infinite horizon, discounted, purely stationary problem are discussed at length in [25], section 4.2.

That initially stationary problems are well-posed can be seen by reformulating $\underline{c}^N$ as a variant of the famous knapsack problem, where the objective is to minimize total cost with the constraint that $z_n(x) \geq 1$ for $1 \leq n \leq N$, i.e., there is equipment on hand at all times. Such a problem fits within the framework of Shapiro and Wagner [40], who show that for each set of knapsack items (in this case, equipment lifetimes), there exists a time horizon N^* sufficiently long that any truncation of length N^* or longer has an optimal initial decision which is the item which has the maximum (minimum) equated annual reward (cost).

To see that this N^* is a forecast horizon, we note that every optimal replacement strategy for every $c' \in \mathcal{C}^{N^*}(c)$ has at least one replacement period in the time interval $[N^*, N^* + L]$. Thus, there exists an optimal replacement strategy for c' that has the same optimal initial decision as for $\underline{c}^{N^*}$. *The strength of this result is that the cash flows for c need not remain stationary indefinitely in order to claim that a forecast horizon exists.* In general, the forecast horizon can be quite short, especially when ties do not exist among equated annual cost of different equipment lifetimes. Also, even though the sets of decision alternatives and their *total* lifetime costs are assumed to be stationary, this does not imply that the operating costs or salvage values need remain stationary. Rather, they can vary arbitrarily, so long as they are constant for fixed equipment lifetimes.

3.4 Conclusions

We considered the infinite horizon nonstationary equipment replacement problem. The model explicitly considered just a single challenger, although implicitly multiple challengers could be considered. A finite maximum feasible lifetime was assumed, costs were bounded, antiquing of equipment was forbidden, and the equipment was permitted to become obsolete at any point in time. Under these assumptions, we showed that the ability to find an optimal initial decision (lifetime of the first piece of equipment) is equivalent to the *coalescence condition*, which roughly states that the network of optimal replacement strategies must be connected. We also gave a simple solution procedure that will yield an infinite horizon optimal initial decision in finite time whenever it is possible to do so.

Chapter 4

An Infinite Horizon Capacitated Inventory Planning Problem

4.1 Introduction

Frequently, it is not realistic when modeling the operations of a manufacturing firm to assume that inventory levels can be set arbitrarily high ahead of future demands. Liu and Tu [31] indicate that process industries such as paper, petrochemical, or pharmaceutical manufacturing, as well as food processing, are constrained more by inventory capacity than by production capacity, although the latter may also exist. They also report that the inventory capacity case has not been well-studied. Moreover, the existing literature on inventory capacity-constrained production planning has been dominated by algorithmic advances for the finite horizon problem. The reader is referred to [31] for an up-to-date literature review for algorithmic advances for the finite horizon problem.

Sandbothe and Thompson [37] obtain forecast horizon results for the infinite horizon capacitated lot size model with constant cost parameters and variable production and inventory costs under the stockout cost option. They provide sufficient conditions for the existence of a forecast horizon along with an efficient solution procedure to detect one. We generalize their model by allowing arbitrary cost functions and include revenue functions as well, while holding the inventory capacity constant over time. While we do not offer any algorithmic improvements over theirs for a given *finite* horizon problem, we do give necessary and sufficient conditions for finite solvability of any infinite horizon forecast instance. We also give a general algorithm that will compute in finite time an infinite horizon optimal first period decisions for any solvable problem.

4.2 Model and Assumptions

In this application, integer-valued customer demands must be met from existing inventory, and inventory has at all times bounded capacity $\bar{S}$. Unmet demand is lost and not backlogged, so that this is a lost-sales model without stockout penalty. Thus, this is a *conservation model* in the language of [31]. There exist costs $c_n(x)$ and $h_n(y)$ associated with replenishing x units of inventory and of holding y units of inventory, respectively, in period n. Also, the revenue generated from satisfying z units of demand in period n is $r_n(z)$. The objective is to maximize total profit while obeying the inventory capacity. All parameters are deterministic.

With these parameters in place, we define the problem as $\phi = (c, h, r, d)$, where c, h, and r are the infinite sequences of replenishment cost, holding cost, and revenue functions, and d is the infinite sequence of demands. As each ϕ is an infinite stream of functions and integer demands, and we have made no restriction that c, h, r or d are constant sequences, it becomes important to define carefully the class of all permissible forecasts Φ. Specifically, we only require the following two assumptions.

Assumption 4.1. *For all $\phi \in \Phi$, demand is strictly positive at least once every M periods.*

Assumption 4.2. *For all $\phi \in \Phi$, $\sum_{n=1}^{\infty} \alpha^{n-1} r_n(\bar{S}) < \infty$. In other words, total profit is always finite since total revenue is always finite.*

With Φ defined, we can define the subclass function $\Phi^n(\phi)$, which is the class of all $\phi' \in \Phi$ whose parameters in periods 1 through n are identical to those of the fixed problem ϕ. The decisions made in any period should depend upon the on-hand inventory quantity beginning that period and upon all future demands, costs, and revenues. Thus, it is reasonable to formulate a dynamic programming recursion for solving a fixed problem $\phi \in \Phi$ with the state in each period n as the inventory quantity on hand beginning period n. This leads to the following functional equation.

$$f_n(y;\phi) = \max_{0 \le z_n \le d_n, 0 \le y_n \le \bar{S} - z_n} \left[r_n(z_n) - c_n((z_n + y_n - y)^+) - h_n(y_n) + \alpha f_{n+1}(y_n;\phi) \right] \tag{4.1}$$

The optimization problem is to solve $f_1(y_0;\phi)$, where y_0 is the initial on-hand inventory. From this recursion, we see that the dynamic programming state at any period n for a fixed forecast ϕ is the inventory quantity ending the previous period, y_{n-1}. Because of limited capacity $\bar{S}$, the set of feasible states is limited to cardinality at most $\bar{S} + 1$ in each period.

The *decisions* in each period n are doubles $x_n = (y_n, z_n)$, where y_n is the *inventory quantity* and z_n is the *fulfillment quantity*. For a given entering inventory quantity y_{n-1} and demand level d_n, and for a desired fulfillment quantity $z_n \le d_n$ and ending inventory quantity $y_n \le \bar{S} - z_n$, there are actually an infinite number of *feasible* replenishment quantities. However, since inventory is bounded at all times by $\bar{S}$, replenishment amounts over and above $\bar{S} - y_{n-1}$ are superfluous and must be discarded. Therefore, it is without loss of optimality that we consider only replenishment amounts in the interval $\{0, \ldots, \bar{S} - y_{n-1}\}$. This fact is reflected in (4.1) by expressing the replenishment quantity as $(z_n + y_n - y_{n-1})^+$, where y is actually y_{n-1}, and $y_n \le \bar{S} - z_n$.

We call an infinite sequence $x = (y_n, z_n)_{n=1}^{\infty}$ of inventory and fulfillment quantity decisions a *fulfillment strategy*. Feasibility of a fulfillment strategy x depends upon the chosen problem ϕ. We say that x is feasible for $\phi = (c, h, r, d)$ if, for all periods

n,

1. $0 \leq z_n \leq d_n$, and

2. $1 \leq y_n \leq \bar{S} - z_n$.

Additionally, we say that x is an *optimal* fulfillment strategy for ϕ if x is feasible for ϕ and if fixing the decisions under x result in total profit $f_1(y_0, \phi)$.

Remark 4.1. *Because demand doesn't have to be satisfied, it is feasible for any problem ϕ to set $x_n = (y_0, 0)$ for all n.*

Remark 4.2. *For any inventory quantity $0 \leq y_n \leq \bar{S}$ in any period n for any problem ϕ, there is a feasible fulfillment strategy for ϕ with inventory quantities y_{n-1} ending period $n-1$ and y_n ending period n for any $y_{n-1} \leq y_n$.*

Proof. Replenish $y_n - y_{n-1}$ units of inventory in period n and set $z_n = 0$. □

Remark 4.3. *For any pair of inventory quantities y_n and y_m with $y_m < y_n$ and any problem $\phi \in \Phi$, there is a feasible fulfillment strategy with inventory levels y_n ending period n and y_m ending period m for any $m \geq n + M(y_n - y_m)$.*

Proof. By definition of M, any problem in ϕ has at least one unit of demand each M periods. Therefore, in any interval of M periods, there exists at least one period n' in which $z_{n'}$ can feasibly be set to 1. By replenishing no inventory between periods n and m for $m = n + M$ and setting $z_{n'} = 1$ at some point in that interval, we can reduce the inventory level by 1 unit every M periods. Thus, it is feasible for any $\phi \in \Phi$ to reduce the inventory quantity by an amount equal to $(y_n - y_m)$ in $M(y_n - y_m)$ periods. □

4.3 Problem Solvability

As with production planning and equipment replacement, the nonstationary nature of an infinite horizon capacitated inventory problem mandates that one rely upon finite horizon truncations of the problem. When an initial fulfillment decision that is optimal after knowing only the parameters an n-period problem is also optimal for any infinite horizon problem identical to the n-period problem through its first n periods, then the problem can be solved in finite time. This prompts us to define well-posed and forecast horizon for this new problem class.

Definition 4.1. *Problem $\phi \in \Phi$ is* well-posed *if there exists a period N^* and initial decision $x_1^* = (y_1^*, z_1^*)$ such that y_1^* and z_1^* are optimal first period inventory and fulfillment quantities, respectively, for any $\phi' \in \Phi^{N^*}(\phi)$.*

Definition 4.2. *Period N^* is a* forecast horizon *for problem $\phi \in \Phi$ if there exist first period inventory and fulfillment quantities y_1^* and z_1^* that are optimal for all $\phi' \in \Phi^{N^*}(\phi)$.*

Again, as with production planning and equipment replacement, we seek general conditions under which a problem $\phi \in \Phi$ is well-posed. To this end, we define reachability and coalescence in the context of this chapter, and show that coalescence is equivalent to well-posed under the construction we have demonstrated.

Definition 4.3. *Let $x^1 = (y^1, z^1)$ and $x^2 = (y^2, z^2)$ be two feasible fulfillment strategies for some problem $\phi \in \Phi$. We say that x_1* reaches *x_2 if there exists a feasible fulfillment strategy $x^3 = (y^3, z^3)$ such that $y_m^1 = y_m^3$ and $y_n^2 = y_n^3$ for some $m \leq n$.*

Definition 4.4. *Problem $\phi \in \Phi$ is* coalescent *if there exists some optimal fulfillment strategy $x^* = (y^*, z^*)$ for ϕ such that any other optimal fulfillment strategy for ϕ can be reached from x^*.*

A fulfillment strategy satisfying the properties of x^* in Definition 4.4 is called a *source strategy*. We proceed to give two technical results that will be required in the proof of the main theorem to follow.

Lemma 4.4. *Let $\phi \in \Phi$ and choose any subsequence of integers $\{n_j\}$ with associated forecasts $\{\phi(n_j)\}, \phi(n_j) \in \Phi^{n_j}(\phi)$ for all n. Then there exist a further sub-*

sequence $\{n_{j_k}\} \subseteq \{n_j\}$ *with associated optimal fulfillment strategies* $\{\tilde{x}^*(n_{j_k})\}$ *optimal for* $\phi(n_{j_k})$ *for each* k *and some optimal fulfillment strategy* $\tilde{x}^*$ *for* ϕ *such that* $\tilde{x}^{*(n_{j_k})} \to \tilde{x}^*$, *where convergence is componentwise.*

Please see Appendix C.1 for a proof.

Lemma 4.5. *For each* $\phi \in \Phi$, *period* n, *and inventory quantity* $0 \leq y_n \leq \bar{S}$, *there exists a forecast* $\phi(y_n)$ *such that* y_n *is the* unique *optimal inventory quantity ending period* n *for* $\phi(y_n)$.

Proof. We will construct a $\phi(y_n)$ that will have y_n as its unique optimal inventory quantity ending period n. Let $\phi(y_n)$ be identical to ϕ through period n. Set $d_{n+1} = y_n$ and $d_{n+2} = 0$. Set $r_{n+1}(\cdot) = c_{n+1}(\cdot) = \frac{1}{\alpha}(1+\epsilon)c_n(\cdot)$ and $h_{n+1}(\cdot) = (1+\epsilon)h_n(\cdot)$ for some $\epsilon > 0$. Also, set $c_{n+2}(y) = \min_{1 \leq m \leq n}(c_m(x+y) - c_m(x)$ for $0 \leq y \leq \bar{S}$ and $0 \leq x \leq \bar{S} - y$. With these parameter and function constructions for $\phi(y_n)$, we observe the following.

First, any demand in periods $n+3$ or later can be fulfilled with strictly greater profit by a replenishment in period $n+2$ or later since $c_{n+2}(y) \leq c_m(x+y) - c_m(x)$ for all y and $m \leq n+1$ and positive inventory costs are incurred in period $n+1$. Second, the profit incurred by holding y_n units forward from period n for demand fulfillment in period $n+1$ is strictly greater than any profit by retaining any of that amount for demand fulfillment in period n. Putting these two facts together, since the demand in period $n+1$ is exactly y_n, the unique optimal inventory quantity ending period n under $\phi(y_n)$ is y_n. Note that the parameters of $\phi(y_n)$ after period $n+2$ are irrelevant to this result. $\square$

Theorem 4.6. *Problem* $\phi \in \Phi$ *is well-posed if and only if it is coalescent.*

Proof. Suppose first that ϕ is well-posed. Then there exists a period N^* and first period decision $x_1^* = (y_1^*, z_1^*)$ such that x_1^* is also optimal for any $\phi' \in \Phi^{N^*}(\phi)$. If

x_1^* is the unique optimal first period decision for ϕ, then ϕ is trivially coalescent. So suppose that there exist multiple optimal first period decisions for ϕ, and in particular, that there are $k^* \geq 2$ distinct optimal period N^* inventory quantities. Call these quantities $y_{N^*}^1, \ldots, y_{N^*}^{k^*}$. By Lemma 4.5, each $y_{N^*}^i$ must also be the unique optimal inventory quantity ending period N^* for some $\phi^i \in \Phi^{N^*}(\phi)$. But x_1^* is an optimal first period decision for each $\phi^i, i = 1, \ldots, k^*$. By the Principle of Optimality, there exists an optimal fulfillment strategy for ϕ that has x_1^* as an optimal first period decision and $y_{N^*}^i$ as its inventory quantity ending period N^* for $i = 1, \ldots, k^*$. Thus, ϕ is coalescent.

Now suppose that ϕ is coalescent but is not well-posed. Let x^{**} be an optimal source strategy. Then there exists a subsequence $\{n_j\}_{j=1}^{\infty}$ and $\phi^{(j)} \in \Phi^{n_j}(\phi)$ such that x_1^{**} is not optimal for $\phi^{(j)}$ for all j. But Lemma 4.4 states that, resorting to a further subsequence $\{n_{j_k}\}_{k=1}^{\infty}$, there exists some sequence of optimal fulfillment strategies $x^{*(k)}$ for forecasts $\phi^{(j_k)}$ for all k such that $x^{*(k)} \to \tilde{x}^*$, which is optimal for ϕ.

By assuming that ϕ is coalescent, $\tilde{x}^*$ is reached at some period M by some optimal fulfillment strategy with initial decision x_1^{**}. Let $K(M)$ be large enough that $x^{*(k)}$ is in agreement with $\tilde{x}^*$ up through and including period M for all $k \geq K(M)$. Note that such a $K(M)$ exists in light of Lemma 4.4. By the Principle of Optimality, there now exists a $\phi^{(j_k)}$-optimal fulfillment strategy for all $k \geq K(M)$ with first period decision x_1^{**}, and which is in agreement with $x^{*(k)}$ thereafter. Hence, x_1^{**} is an optimal fulfillment strategy for $\phi^{(j_k)}$ for all $k \geq K(M)$. We've reached a contradiction and thus conclude that satisfying the coalescence condition is sufficient for well-posedness. $\square$

4.4 Solution Procedure

We have identified coalescence as a necessary and sufficient condition for finite solvability of a given problem. However, because coalescence is a property of an infinite horizon problem, which cannot be known in finite time, to solve a given problem will generally require an algorithm. In this section, we provide an algorithm which will provide in finite time an optimal first period decision for any well-posed $\phi \in \Phi$.

Algorithm 4.1. *Choose $\phi \in \Phi$ and set $n = 1$.*

1. *For $y = 0, \ldots, \bar{S}$, find $X_1^*(\phi^n|y)$, the set of optimal first-period decisions for the n-period truncation of p with the terminal state restricted to be y.*

2. *If $\exists x_1^* \in \cap_{0 \leq y \leq \bar{S}} X_1^*(\phi^n|y)$, then stop. x_1^* is an optimal first-period decision for ϕ and n is a forecast horizon. Otherwise, increase n by 1 and return to step 1.*

Theorem 4.7. *Algorithm 4.1 stops in finite time if and only if ϕ is well-posed.*

Proof. Suppose that ϕ is well-posed. Then there exists a period N^* and first period decision x_1^* such that x_1^* is also optimal for all $\phi' \in \Phi^{N^*}(\phi)$. By Lemma 4.5, for each $y \in \{0, \ldots, \bar{S}\}$, there exists $\phi(y)$ such that x_1^* is the unique optimal first-period decision for $\phi(y)$. Thus, at period N^*, Algorithm 4.1 will stop and return x_1^*.

Now suppose that Algorithm 4.1 stops at period N^* with optimal first-period decision x_1^*. Then x_1^* is an optimal first-period decision along some optimal finite horizon fulfillment strategy to each inventory quantity $y \in \{0, \ldots, \bar{S}\}$. By finiteness of $\{0, \ldots, \bar{S}\}$, and the fact that every optimal fulfillment strategy for any $\phi' \in \Phi^{N^*}(\phi)$ must have an inventory quantity ending period N^* in that set, we see by the Principle of Optimality that $x - 1^*$ is an optimal first-period decision for all $\phi' \in \Phi^{N^*}(\phi)$. Thus, ϕ is well-posed. $\square$

4.5 Conclusions

We have studied the infinite horizon, nonstationary, capacitated inventory planning problem with deterministic demands, costs, and revenues. Under the assumptions that demand is strictly positive at least once every M periods and that maximum total profit is always finite, we showed that the ability to find optimal first-period fulfillment and inventory quantities is equivalent to a condition called coalescence. Coalescence roughly stipulates that if there exist multiple optimal fulfillment strategies, then there must be some optimal first-period decision such that any other optimal fulfillment strategy has a common inventory quantity in some period with another optimal fulfillment strategy that has the specified optimal first-period decision. We also gave a solution procedure that will detect in finite time an infinite horizon optimal first-period decision any time it is possible to do so.

Chapter 5

General Deterministic Infinite Horizon Optimization

5.1 Introduction

Each of the previous three chapters established the equivalence between coalescence and well-posedness for a specific operations management application. One might well suppose that such a relationship persists for more complex problems, or at a higher level of abstraction. In this chapter, we show that indeed, a large class of optimization problems that can be formulated as dynamic programs also satisfy the coalescence - well-posedness equivalence.

Within the literature on the existence and discovery of forecast horizons in general deterministic optimization, Bean and Smith [1] show that a forecast horizon exists for a general class of problems when the optimal strategy is unique. In [2], they show that a weak reachability condition is necessary and sufficient for finite discovery of the optimal initial decision, again under the uniqueness assumption. They also give a solution procedure that will detect the optimal initial decision in finite time when the uniqueness and weak reachability conditions are met, and *stopping sets* of states are appropriately chosen.

In this chapter, we consider a slightly less general formulation than the generic networks used in [1] and [2] in that we restrict decisions and transitions to occur at fixed time periods. However, we maintain the level of generality sufficiently high so as to include many typical applications. We also relax the assumption of uniqueness of the optimal strategy. We give necessary and sufficient conditions for finite detection of the optimal initial condition under two different sets of finiteness and reachability conditions, along with algorithms for each. We show that these finiteness and reachability conditions are satisfied by some well-known applications.

5.2 Model and Assumptions

We consider a class $\mathcal{P}$ of infinite horizon, discrete time, deterministic *problems* or *forecasts* p, each of which can be formulated as a dynamic program. A forecast p consists of all transition functions, constraints, and cost functions that are necessary to give, for any sequence of decisions x, a certificate of feasibility or infeasibility, the resulting state sequence $s(x;p)$ if x is feasible, and the total cost of the decisions and resulting states. The objective is to select, for a given forecast p, a sequence x of decisions that is feasible for problem p and that minimizes total discounted cost. We assume that for each $p \in \mathcal{P}$, there exists a feasible decision sequence which results in total cost bounded in each period by some uniform constant $B < \infty$. Coupling this assumption with a one period discount factor α, where $0 < \alpha < 1$, we see that each p has an optimal solution which has finite total cost, with upper bound $\frac{B}{1-\alpha}$. For each $p \in \mathcal{P}$ and each period n, let $\mathcal{P}^n(p)$ represent the subclass of all forecasts whose parameters in the first n periods are identical to those of p.

We will follow a construction of decision and state sets and transition functions very similar to that in Schochetman and Smith [38], although we will be precise here in specifying the forecast $p \in \mathcal{P}$ for which each states and decisions are possible and/or feasible. First, we denote the sets of all *possible* decisions and states available in period n by Y_n and S_n, respectively. These sets may be finite or countable, although we will require finiteness of these sets at times in order to obtain results in section 5.3. The sets $Y = \prod_{n=1}^{\infty} Y_n$ and $S = \prod_{n=1}^{\infty} S_n$ contain all possible decision and state sequences, respectively. However, because we are concerned here with ensuring that decisions are optimal for a variety of forecasts, we find it necessary to further specify the sets of all possible decisions and states available in period n *for a given forecast* p by $Y_n(p)$ and $S_n(p)$, respectively, so that $Y_n(p) \subseteq Y_n$ and $S_n(p) \subseteq S_n$ for each n.

We adopt the convention that state transitions occur at the beginning of the period, so the state at the end of each period is a function of the previous period's state and the decision made in the current period. Specifically, let s_0 be the initial state. Then, for a forecast p, the state s_n in period n is dynamically assigned via the equations $s_n = f_n(s_{n-1}, y_n; p)$, where $y_n \in Y_n(p)$. Of course, not all decisions $y_n \in Y_n(p)$ may be feasible for a given s_{n-1} and p, so let $Y_n(s_{n-1}; p)$ be the set of all available decisions when the state is s_{n-1} beginning period n under forecast p. Thus,

$$S_n(p) = \{f_n(s_{n-1}, y_n; p) : s_{n-1} \in S_{n-1}(p), y_n \in Y_n(s_{n-1}; p)\} \tag{5.1}$$

represents all the feasible states in period n under forecast p. We say that $x \in Y$ is a *feasible decision sequence* for p if $x_n \in Y_n(s_{n-1}; p)$ and $s_n = f_n(s_{n-1}, x_n; p)$ for all n. Similarly, we say that $s \in S$ is a *feasible state sequence* for p if $x_n \in Y_n(s_{n-1}; p)$ and $s_n = f_n(s_{n-1}, x_n; p)$ for all n. We define the *feasible decision space for* p $X(p)$ to be the subset of Y consisting of all x which are feasible decision sequences for p ,and the *feasible state space for* p $T(p)$ to be the subset of S consisting of all s which are feasible state sequences for p.

We will assume that for any forecast p, if state s_n is feasible in period n, i.e. $s_n \in S_n(p)$, then $s_n \in S_n(p')$ for all p' in agreement with p through period n. Moreover, $Y_{n+1}(s_n; p)$ is nonempty for any p, $s_n \in S_n(p)$, and all n, so that any finite horizon feasible state or decision sequence can be feasibly extended arbitrarily far beyond period n, for any forecast in agreement with p through period n.

For each forecast p, exercising a feasible decision sequence x results in a unique state sequence $s(p, x)$. The converse is not necessarily true, that is, if s is a sequence of feasible states for p, then there may exist multiple feasible decision sequences $x^1, x^2, \ldots, x^k$ such that $s_n(p, x^1) = \cdots = s_n(p, x^k)$ for all periods j. However, in such a case, we will assume that only a feasible decision sequence $x^i, i \in \{1, \ldots, k\}$ such that $s(p, x)$ achieving minimum total cost among all x^i will be considered, where ties among minimum-cost decision sequences resulting in the same state sequence can be broken arbitrarily. Thus, we will operate under the assumption that for a particular problem p, feasible state and decision sequences have a one-to-one correspondence, although this assumption does not preclude greater generality. Consequently, although the total cost of a decision sequence x for a forecast p is a function of both the decisions and their resulting states, it is sufficient to identify only one of the two sequences in order to know both.

Let $V(p, x)$ denote the total discounted cost of exercising decision sequence x under forecast p, where $V(p, x) = \infty$ if $x \notin X(p)$. Then the optimization problem for p is

$$V^*(p) = \min_{x \in X(p)} V(p, x). \tag{5.2}$$

Let $T^*(p)$ ($X^*(p)$) represent the set of all *optimal state sequences* (*optimal decision sequences*) for forecast p, with $T_n^*(p)$ ($X_n^*(p)$) as the set of optimal states (optimal decisions) in period N. Necessarily, $T^*(p) \subseteq T(p)$ and $X^*(p) \subseteq X(p)$, with analogous inclusions holding for the single-period sets. It is also important to identify, for each forecast, the sets of states and decisions in each period which are optimal for *some* forecast in agreement with the forecast through that period. To this end, let

$$T_n^{**}(p) = \bigcup_{p' \in \mathcal{P}^n(p)} T_n^*(p') \tag{5.3}$$

and

$$X_n^{**}(p) = \bigcup_{p' \in \mathcal{P}^n(p)} X_n^*(p') \tag{5.4}$$

denote the sets of *potentially optimal states* and *potentially optimal decisions* for forecast p, respectively, in period n. If $s_n \in T_n^{**}(p)$, then clearly $s_n \in T_n^{**}(p')$ for all $p' \in \mathcal{P}^n(p)$. We will also assume that $s_n \in T_n^{**}(p)$ implies that $Y_{n+1}(s_n; p') \neq \emptyset$ for all $p' \in \mathcal{P}^n(p)$; that is, if s_n is potentially optimal for p in period n, then s_n lies on some feasible state sequence for any forecast in agreement with p through period n. We now introduce three important assumptions that are properties of the problem class $\mathcal{P}$.

Assumption 5.1 (**Type I Assumptions**).

1. *For all forecasts $p \in \mathcal{P}$, periods n, and $s_n \in T_n^{**}(p)$, there exists a forecast $p' = (p_1, \ldots, p_n, q_{n+1}, q_{n+2}, \ldots) \in \mathcal{P}$ such that s_n is the* unique *infinite horizon optimal state ending period n for forecast p'.*

2. *For all forecasts $p \in \mathcal{P}$, horizons n, $s_n \in T_n^{**}(p)$ and $s_{n+L} \in T_{n+L}^{**}(p)$, there are decisions $x_{n+1}, x_{n+2}, \ldots, x_{n+L}$ and states $t_{n+1}, t_{n+2}, \ldots, t_{n+L-1}$ such that $x_{n+1} \in Y_{n+1}(s_n; p)$, $t_{n+1} = f_{n+1}(s_n, x_{n+1}; p)$, $t_j = f_j(t_{j-1}, x_j; p) \in S_j(p)$ for $j = n+2, \ldots, n+L$ and $f_{n+L}(t_{n+L-1}, x_{n+L}; p) = s_{n+L}$. In other words, for any forecast, there exists a feasible decision sequence that connects any pair of potentially optimal states that are at least L periods apart, where $L < \infty$ and L is independent of p and n.*

3. *For all forecasts $p \in \mathcal{P}$ and horizons n, $T_n^{**}(p)$ and $X_n^{**}(p)$ are finite (though not necessarily uniformly bounded).*

The first Type I Assumption implies that all potentially optimal states in any period n for any forecast p are the unique optimal states in that period for some forecast in $\mathcal{P}^n(p)$, and therefore, every member of the potentially optimal state set must be considered in checking for agreement of the optimal initial decision (making each potentially optimal state set a minimal regeneration set in the language of the literature). The second Type I Assumption impedes an optimal state sequence for one forecast from getting too "far away" in the sense of total cost by using a carefully chosen continuation from any of its potentially optimal state sequences. The third Type I Assumption will become important in obtaining convergence of finite horizon optimal state sequences and in creating a solution procedure.

5.3 Problem Solvability

Suppose that p has multiple optimal decision sequences. Let x^* and x^{**} be two such sequences. We say that x^{**} is *optimally reachable* from x^* if there exists an optimal solution x^{***} (possibly x^* itself) such that $s_m(p, x^*) = s_m(p, x^{***})$ and $s_n(p, x^{**}) = s_n(p, x^{***})$ for some m, n with $m \leq n$.

Definition 5.1. *We say that p is* coalescent *or satisfies the* coalescence condition *if there exists some optimal solution x^* such that every other optimal solution for p is optimally reachable from x^*. In this case, such an x^* is called a* source solution.

Figure 5.1 shows a depiction of coalescence in a general, time-staged, discrete optimization problem. The single first state on each side represents the problem's fixed initial state. The bold paths, solid or dashed, represent the optimal solutions. For the problem on the left, the two optimal solutions do not reach one another in the figure, and assuming that they continue to run separately over the infinite horizon, the problem is not coalescent. On the other hand, for the problem on the right, assuming that there are just two optimal initial decisions, the dashed path is a source solution and the problem is coalescent. We now define well-posed and forecast horizon in the context of general deterministic infinite horizon optimization, and we also formalize the definition of *planning horizon* that frequently appears in the literature.

Definition 5.2. *We say that p is* well-posed *if there exist a period N^* and initial decisions $x_n^*, n = 1, \ldots, N$ for some $1 \leq N \leq N^*$ such that $x_1^*, \ldots, x_N^*$ are optimal*

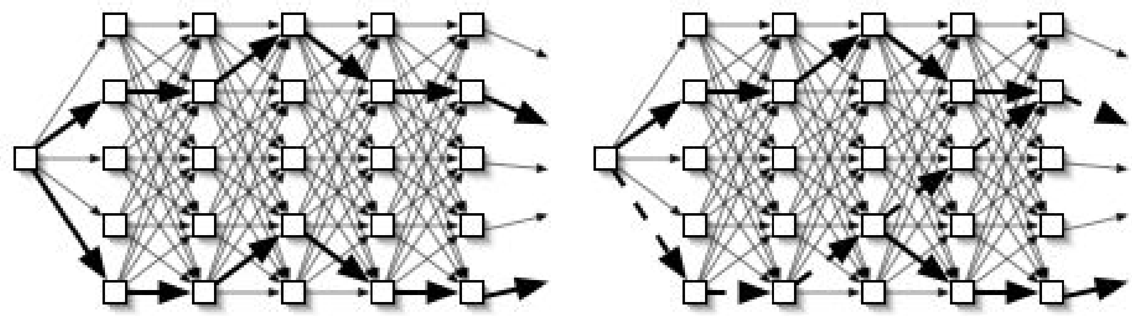

Figure 5.1: Sketch of Coalescence in General Deterministic Optimization

decisions for all $p' \in \mathcal{P}^{N^*}(p)$.

Definition 5.3. *If there exist a period* N^* *and initial decisions* $x_n^*, n = 1, \ldots, N$ *for some* $1 \leq N \leq N^*$ *such that* $x_1^*, \ldots, x_N^*$ *are optimal decisions for all* $p' \in \mathcal{P}^{N^*}(p)$, *then we call* N^* *a* forecast horizon *and* N *a* planning horizon *for problem* p *with respect to class* $\mathcal{P}$.

Lemma 5.1. *Let* $p \in \mathcal{P}$ *and choose any increasing subsequence of positive integers* $\{n_j\}$ *with associated forecasts* $\{p(n_j)\}$, *where* $p(n_j) \in \mathcal{P}^{n_j}(p)$ $\forall n$. *Then, under the Type I Assumptions, there exist a further subsequence* $\{n_{j_k}\} \subseteq \{n_j\}$ *with associated optimal decision sequences* $\{\tilde{x}^*(n_{j_k})\}, \tilde{x}^*(n_{j_k}) \in \mathcal{X}^*\left(p(n_{j_k})\right)$ $\forall n$, *and some* $\tilde{x}^* \in X^*(p)$ *such that* $\tilde{x}^{*(n_{j_k})} \to \tilde{x}^*$, *where convergence is componentwise.*

See Appendix D for a proof. This result prepares us for the main result of this section.

Theorem 5.2. *Under the Type I Assumptions, forecast* $p \in \mathcal{P}$ *is well-posed if and only if it is coalescent.*

Proof. Suppose p is well-posed and the Type I Assumptions hold. Let y^* be an arbitrary infinite horizon optimal decision sequence and let N^* be a period such that x_1^* is an infinite horizon optimal initial decision for all $p' \in \mathcal{P}^{N^*}(p)$. Such an N^* and x_1^* exist by well-posedness of p. Then by the first Type I Assumption, there exists a specific forecast $p' \in \mathcal{P}^{N^*}(p)$ such that $s_{N^*}(p, y^*)$ is the unique infinite horizon optimal state ending period N^* for forecast p'. By well-posedness, x_1^* is an infinite horizon optimal initial decision for p', so that there is some $x^* \in X^*(p')$ which has initial decision x_1^* and state $s_{N^*}(p, y^*)$ ending period N^*. Thus, y^* is optimally reachable from x^*, and having chosen y^* arbitrarily, p is coalescent.

Now suppose that p is coalescent but is not well-posed with respect to $\mathcal{P}$. Let x^{**} be a source solution. Then there exists a subsequence $\{n_j\}_{j=1}^{\infty}$ and $p(n_j) \in \mathcal{P}^{n_j}(p)$ such that $x_1^{**} \notin X_1^*(p(n_j))$ for all j. But Lemma 5.1 states that, resorting to a further subsequence $\{n_{j_k}\}_{k=1}^{\infty}$ if necessary, there exists some sequence of infinite horizon optimal decision sequences $x^*(n_{j_k}) \in X^*(p^(n_{j_k}))$ for all k such that $x^*(n_{j_k}) \to \tilde{x}^* \in X^*(p)$.

By assuming that p is coalescent, $\tilde{x}^*$ is optimally reached at some period M by some infinite horizon decision sequence, optimal for p, with initial decision x_1^{**}. Let $K(M)$ be large enough that $x^*(n_{j_k})$ is in agreement with $\tilde{x}^*$ up through and including period M for all $k \geq K(M)$. By the Principle of Optimality, there now exists a $p(n_{j_k})$-optimal decision sequence $\forall k \geq K(M)$ with initial decision x_1^{**}, and which is in agreement with $x^*(n_{j_k})$ beginning period M. Hence, $x_1^{**} \in X_1^*(p(n_{j_k}))$ $\forall k \geq K(M)$. We've reached a contradiction and conclude that satisfying the coalescence condition is sufficient for well-posedness. $\square$

5.4 Solution Procedures

In this section, we present three different solution procedures. The first is the simplest and can be applied whenever the sets of potentially optimal states are readily available. Recognizing that these sets often require thoughtful derivation that is application-specific, the second procedure incorporates a general method from [24] that both iteratively constructs the sets of potentially optimal states and invokes a stopping rule at each iteration. However, since for many applications the problem-specific functions required to successfully implement the second procedure may not exist, we introduce a third solution procedure. This procedure requires a stronger assumption (introduced in [38]) of reachability than in the second Type I Assumption that, when satisfied, gives rise to an implementation as simple as the first procedure.

5.4.1 A Simple Solution Procedure

While coalescence can occasionally aid in identifying well-posed problems, the nature of infinite horizon nonstationary optimization generally precludes a priori knowledge of all problem parameters. Rather, these parameters are more frequently uncovered sequentially, and conclusions regarding well-posedness of a problem instance must be determined algorithmically. A primary difficulty in constructing such an algorithm lies in ensuring that it yields an infinite horizon optimal initial decision any time a problem is well-posed without stopping prematurely and yielding an initial decision that is not optimal for some potential future parameters. That is, the algorithm should stop finitely if and only if the problem is well-posed.

Clearly, p is well-posed if and only if there exists some optimal first period state s_1^* and period N such that for each $s_N \in T_N^{**}(p)$, s_1^* lies on a minimum cost solution to s_N. By the third Type I Assumption, we can hope to verify this condition in finite time. However, to verify this condition requires a complete characterization of $T_n^{**}(p)$ for all n and $p \in \mathcal{P}$. In general, this can be extremely difficult, although we shall show that it is possible for both the concave cost production planning and equipment replacement problems, and we will give a method in the next section to generate $T_n^{**}(p)$ for a general class of problems. For now, however, we will assume knowledge of $T_n^{**}(p)$ to give the following solution procedure.

Assumption 5.2. *For all $p \in \mathcal{P}$ and periods n, $T_n^{**}(p)$ is known.*

Algorithm 5.1. *Consider a fixed $p \in \mathcal{P}$. Set $n = 1$.*

1. *For all $s_n \in T_n^{**}(p)$, find $X_1^*(p^n|s_n)$, the set of all optimal initial decisions for*

the n-horizon truncation of p with the state ending period n restricted to be s_n.

2. *If there exists $x_1^* \in \cap_{s_n \in T_n^{**}(p)} X_1^*(p^n|s_n)$, then return x_1^* and stop. Otherwise, increment n by 1 and return to step 1.*

Theorem 5.3. *Under the Type I Assumptions and Assumption 5.2, Algorithm 5.1 stops finitely if and only if $p \in \mathcal{P}$ is well-posed.*

Proof. Suppose that the Type I Assumptions hold. Suppose that p is well-posed. Then there exists some period N^* and initial decision x_1^* such that x_1^* is an optimal initial decision for all $p' \in \mathcal{P}^{N^*}(p)$. By the third Type I Assumption, the set $T_{N^*}^{**}(p)$ is finite. Furthermore, by definition, every optimal state sequence for every $p' \in \mathcal{P}^{N^*}(p)$ passes through $T_{N^*}^{**}(p)$. But Algorithm 5.1 will find the set of optimal initial decisions for each partial state sequence with an ending state in $T_{N^*}^{**}(p)$. Since x_1^* is an optimal initial decision for each of those partial state sequences by well-posedness of p, and each of the potentially optimal ending states is known by Assumption 5.2, Algorithm 5.1 will stop at period N^* and return x_1^*.

Now suppose that Algorithm 5.1 stops at period N^* with optimal initial decision x_1^*. Then x_1^* is an optimal initial decision for any optimal state sequence passing through $T_{N^*}^{**}$. Since every optimal solution for every $p' \in \mathcal{P}^{N^*}(p)$ must pass through there, the Principle of Optimality says that x_1^* is an optimal initial decision for all $p' \in \mathcal{P}^{N^*}(p)$. Thus, p is well-posed. □

5.4.2 An Integrated Procedure to Derive Potentially Optimal State Sets and Find Optimal First Decisions

Since a choice for $T_n^{**}(p)$ is often not obvious, we seek a general method for constructing these sets. Following Federgruen and Tzur [24], we can efficiently generate $T_n^{**}(p)$ for a subset of the general class of problems we have chosen in this chapter. First, we need to express the forecast as a shortest path problem in an acyclic network. Then, we need to develop difference functions to compare two nodes as optimal predecessor nodes to a fixed, third node. Finally, we can use the difference functions to characterize regions of potential future parameters for which a given node is optimal; in other words, we can generate sets of potentially optimal states. The method used in this section to construct $T_n^{**}(p)$ will follow that in [24], being adapted to the problem class $\mathcal{P}$ defined in this chapter.

Reformulation as a Shortest Path Problem

Let p^N represent the N-period truncation of forecast $p \in \mathcal{P}$. Under the additional assumption that there exist *finite, known* sets $\bar{X}_n^{**}(p)$ and $\bar{T}_n^{**}(p)$ such that $X_n^{**}(p) \subseteq \bar{X}_n^{**}(p)$ and $T_n^{**}(p) \subseteq \bar{T}_n^{**}(p)$ for $n = 1, \ldots, N$, we can express p^N as a shortest path problem on an acyclic network. The node set $\mathbb{N} = \{0, 1, \ldots, n\}$ contains all states in $\cup_{n \in \{1,\ldots,N\}} T_n^{**}(p) \cup s_0$. For all nodes $i, j, i < j$, the cost of arc (i, j) is given by $c(i, j)$, where $c(i, j) = +\infty$ if it is infeasible to transition from i to j. In general, $c(i, j)$ is finite if and only if $i \in S_n(p), j \in S_{n+1}(p)$ and there exists some decision y such that $j = f_{n+1}(i, y; p)$, although there may exist feasible decisions that incur infinite cost. We observe that in this construction, the node sets and arc costs are completely determinable from the forecast p, which contains all potentially optimal decision and state sets, transition functions, and costs.

Let $F(j)$ denote the cost of the shortest path from node 0 to node $j \in \mathbb{N}$, and

$$F(l, j) = F(l) + c(l, j)$$

is the minimum cost of a path from node 0 to node j with the restriction that node $l < j$ is the next-to-last node on the path (the *predecessor* of j). Also, let

$$L(j) = \{l : F(l, j) = F(j)\}$$

represent the set of optimal predecessor nodes to node j and
$Q(j) = \{l : l \text{ is the } \textit{first} \text{ node after node 0 on some optimal path to node } j\}$. The Q

sets can be expressed recursively as

$$Q(j) = \begin{cases} \cup_{i \in L(j)} Q(i) & \text{if } 0 \notin L(j) \\ \cup_{i \in L(j)} Q(i) \cup \{j\} & \text{otherwise} \end{cases} \quad (j = 1, \ldots, n).$$

A Finite Solution Procedure for Well-Posed Problems

Continuing with the model derived in Federgruen and Tzur, we now define *difference functions* $\Delta_{k,l}(j) : \mathbb{N} \to \mathbb{R}$ by

$$\Delta_{k,l}(j) = F(k,j) - F(l,j),$$

which determines which of a given pair of nodes k and l is preferable as the predecessor node to node $j \in \mathbb{N}$. We now include the following key assumption, which Federgruen and Tzur have shown is satisfied by a number of important applications, such as multi-item joint replenishment systems, combined inventory routing, machine scheduling issues, and single item stochastic inventory settings.

Definition 5.4. *Problem $p \in \mathcal{P}$ is said to satisfy the* Difference Function Assumption *if each node in $\mathbb{N}$ can be characterized by a finite set of indicators $\mathbf{X} = (X_1, X_2, \ldots, X_m)$ such that all difference functions can be written in the form*

$$\Delta_{k,l}(j) = \delta_{k,l}(\mathbf{X}(j)) = \delta_{k,l}(X_1(j), X_2(j), \ldots, X_m(j)),$$

where the function $\delta_{k,l}(\cdot)$ can be evaluated knowing only the forecast parameters through node l and $\mathbf{X}(j)$ can be evaluated knowing only the forecast parameters through node j.

In this manner, upon reaching node l, the difference function can be evaluated for any future node $j \geq l$ given the indicator values $X_1(j), X_2(j), \ldots, X_m(j)$.

We now proceed to show how the difference functions can be used to construct the sets of potentially optimal predecessor nodes, and hence, the sets $T_N^{**}(p)$ of potentially optimal states. Let $\Omega(j)$ be the set of all nodes in $\{0, \ldots, j\}$ that are optimal predecessor nodes for *some* later node with appropriate indicator values. Here, we will follow the approach of Federgruen and Tzur for the general case of possible non-linear difference functions.

For each $k = 1, \ldots, r$, let $R_{i_k}(j)$ denote the region of all vectors of indicator values $\mathbf{X} \in \mathbb{R}^m$ for which i_k is the best predecessor node among all nodes in $\Omega(j)$. Then

$$R_{i_k} = \{\mathbf{X} = (X_1, X_2, \ldots, X_m) \in \mathbb{R}^m : \delta_{i_k,i_l}(\mathbf{X}) \leq 0 \; \forall l = 1, \ldots, r, l \neq k\}.$$

To obtain the set Ω and the R sets each time the node set is augmented (i.e. the horizon is increased), we introduce here a simple update procedure. Initially, $\Omega(0) = \{0\}$ and $R_0(0) = \mathbb{R}^m$. Assume at iteration j that $\Omega(j) = \{i_1, i_2, \ldots, i_r\}$ with the corresponding R regions are known. By definition of Ω, $\Omega(j+1) \subseteq \{\Omega(j) \cup \{j+1\}\}$. To check whether to include node $j+1$ or not in $\Omega(j+1)$, it is sufficient to see whether the (possibly) nonlinear program

$$R_{j+1}(j+1) = \{\mathbf{X} \in \mathbb{R}^m : \delta_{j+1,i_l}(\mathbf{X}) \leq 0, l = 1, \ldots, r\}$$

has a feasible solution. If it does, then node $j+1$ should be added. Likewise, to check whether a node $i_k \in \Omega(j)$ needs to be eliminated in iteration $j+1$, see if

$$R_{i_k}(j+1) = R_{i_k}(j) \cap \{X \in \mathbb{R}^m : \delta_{i_k,j+1}(\mathbf{X}) \leq 0\}$$

is nonempty. If it is not, then node i_k may be eliminated as an optimal predecessor for some potential future node. With this update procedure established, an algorithm to solve any well-posed instance immediately follows.

Algorithm 5.2.

1. *Set* $N = 0$, $\Omega(0) = \{0\}$ *and* $R_0(0) = \mathbb{R}^m$.

2. *Set* $N \leftarrow N + 1$. *For* $j \in \bar{T}_N^{**}(p)$,

    (a) *add node* j *to* $\mathbb{N}$, *with arc costs* $c(i, j)$ *for all* $i < j$;

(b) *determine whether* $\{j\} \in \Omega(j)$*;*

(c) *determine whether to remove any* $i \in \Omega(j-1)$ *in updating* $\Omega(j)$*.*

3. *Let* $\Omega\left(\sum_{k=1}^{N}\left|\bar{T}_k^{**}(p)\right|\right) = \{i_1, \ldots, i_m\} \subseteq \bar{T}_N^{**}(p)$*. If there exists a node* $q^* > 0$ *with* $q^* \in Q(i_r)$ *for all* $r = 1, \ldots, m$*, then stop. The initial decision resulting in the state corresponding to node* q^* *is an optimal initial decision for all* $p' \in \mathcal{P}^N(p)$*, and* p *is well-posed. Otherwise, return to 2.*

We note that following the complete Update step for period N, the Ω set is precisely $T_N^{**}(p)$, as desired. Thus, any state in that set must be optimal for some forecast in agreement with p through period N. If every one of those states lies on an optimal path (for problem p^N) passing through a common initial state, then the initial decision resulting in that initial state must be optimal for any $p' \in \mathcal{P}^N(p)$, and p is well-posed. Conversely, if p is well-posed, then there must exist some period N at which Algorithm 5.2 will stop. We can thus state the following result.

Theorem 5.4. *Under the Type I Assumptions and the Difference Function Assumption, Algorithm 5.2 stops in finite time if and only if* $p \in \mathcal{P}$ *is well-posed.*

5.4.3 A Solution Procedure for Problems with Finite Feasible State Spaces

To this point, we have avoided assuming that the sets of *feasible* states are finite, rather requiring that the sets of potentially optimal states are finite. We also required in the second Type I Assumption that any pair of potentially optimal states at least L periods apart can be feasibly connected. In many applications, such as equipment replacement, it is also reasonable to expect that any pair of *feasible* states can be feasibly connected after some uniformly bounded number of periods. Schochetman and Smith call this property Bounded Reachability in Definition 4.1 of [38], which we repeat here in slightly different language as one of the second set of assumptions.

Assumption 5.3 (**Type II Assumptions**).

1. *For all forecasts* $p \in \mathcal{P}$*, periods* n*, and* $s_n \in S_n(p)$*, there exists a forecast* $p' = (p_1, \ldots, p_n, q_{n+1}, q_{n+2}, \ldots) \in \mathcal{P}$ *such that* s_n *is the* unique *infinite horizon optimal state ending period* n *for forecast* p'*.*

2. *For all forecasts* $p \in \mathcal{P}$*, periods* n*,* $s_n \in S_n(p)$ *and* $s_{n+R} \in S_{n+R}(p)$*, there are decisions* $x_{n+1}, x_{n+2}, \ldots, x_{n+R}$ *and states* $t_{n+1}, t_{n+2}, \ldots, t_{n+R-1}$ *such that* $x_{n+1} \in Y_{n+1}(s_n; p)$*,* $t_{n+1} = f_{n+1}(s_n, x_{n+1}; p)$*,* $t_j = f_j(t_{j-1}, x_j; p) \in S_j(p)$ *for* $j = n + 2, \ldots, n+R$ *and* $f_{n+R}(t_{n+R-1}, x_{n+R}; p) = s_{n+R}$*. In other words, for any forecast, there exists a feasible decision sequence that connects any pair of feasible states that are at least* R *periods apart, where* $R < \infty$ *and* R *is independent of* p *and* n*.*

3. *For all forecasts* $p \in \mathcal{P}$ *and periods* n*,* $S_n(p)$ *and* $Y_{n+1}(s_n; p)$ *for all* $s_n \in S_n(p)$ *are finite (though not necessarily uniformly bounded).*

Proposition 5.5. *If the second Type II Assumption is satisfied and* $R \leq L$*, then the second Type I Assumption is also satisfied. Moreover, the first and third Type II Assumption imply their Type I counterparts.*

Proof. To show that the first and third Type II Assumptions imply their Type I counterparts is straightforward since all potentially optimal states and decisions are also feasible for any p. Suppose that the second Type II Assumption is satisfied. Then consider any pair of potentially optimal states at least L periods apart. Since $L \geq R$ and potentially optimal states are also feasible, we conclude that the second Type I Assumption is also satisfied. □

Note that even if $L > R$, the second Type II Assumption still implies that a modified form of the second Type I Assumption holds by setting $L = R$. For the sake of simplicity, we will assume in this section that $R \leq L$.

In section 5.3, the second Type I Assumption was invoked only in the proof of Lemma 5.1. Thus, since the second Type II Assumption implies the second Type I Assumption, all the results in that section will still hold here, and we claim the following without having to reconstruct the proofs.

Corollary 5.6. *When all of the Type I assumptions are met or if their Type II counterparts are met, $p \in \mathcal{P}$ is well-posed if and only if it is coalescent.*

We could also follow the same solution procedure here and claim that Algorithm 5.1 stops finitely if and only if $p \in \mathcal{P}$ is well-posed. However, we have assumed in this section a more favorable problem structure; namely, that Bounded Reachability holds. Note that Schochetman and Smith [38] assumed a uniformly bounded feasible state space in each period in order to obtain their solvability results for the same class of problems. When the Type II Assumptions are satisfied, the only additional requirement to apply a solution procedure analogous to Algorithm 5.1 is a knowledge of the feasible state sets. We know assume such a knowledge.

Assumption 5.4. *For all $p \in \mathcal{P}$ and periods n, $S_n(p)$ is known.*

The following solution procedure exploits the finiteness and complete characterization of the feasible decision and state spaces as well as the Bounded Reachability Property.

Algorithm 5.3. *Consider a fixed $p \in \mathcal{P}$. Set $n = 1$.*

1. *For all $s_n \in S_n(p)$, find $X_1^*(p^n|s_n)$, the set of all optimal initial decisions for the n-horizon truncation of p with the state ending period n restricted to be s_n.*

2. *If there exists $x_1^* \in \cap_{s_n \in S_n(p)} X_1^*(p^n|s_n)$, then return x_1^* and stop. Otherwise, increment n by 1 and return to step 1.*

Theorem 5.7. *Under the Type II Assumptions and Assumption 5.4, Algorithm 5.3 stops finitely if and only if $p \in \mathcal{P}$ is well-posed.*

Proof. Suppose that p is well-posed. Then there exists some period N^* and initial decision x_1^* such that x_1^* is an optimal initial decision for all $p' \in \mathcal{P}^{N^*}(p)$. By the third Type II Assumption, the set $S_{N^*}(p)$ is finite. Furthermore, by the first Type II Assumption, every member of $S_{N^*}(p)$ is the unique optimal state in period N^* for some $p' \in \mathcal{P}^{N^*}(p)$. But Algorithm 5.3 will find the set of optimal initial decisions for each partial state sequence with an ending state in $S_{N^*}(p)$, which set is known by Assumption 5.4. Since x_1^* is an optimal initial decision for each of those partial state sequences by well-posedness of p, Algorithm 5.3 will stop at period N^* and return x_1^*.

Now suppose that Algorithm 5.3 stops at period N^* with optimal initial decision x_1^*. By the third Type II Assumption and the Principle of Optimality, x_1^* is an optimal initial decision for all $p' \in \mathcal{P}^{N^*}(p)$. Thus, p is well-posed. □

Theorem 5.7 tells us that any time the sets of feasible states are known, finite, and identical to the sets of potentially optimal states for all problems in the problem class, then any well-posed problem will be solved by Algorithm 5.3. Although the hypothesis for this result may appear strong, it is not vacuous, as we will show in section 5.5. By Corollary 5.6 and Theorem 5.7, we can recast Theorem 5.7 as follows.

Corollary 5.8. *Under the Type II Assumptions and Assumption 5.4, Algorithm 5.3 stops finitely if and only if $p \in \mathcal{P}$ is coalescent.*

This claim concerning Algorithm 5.3 now becomes stronger than that made for the corresponding solution procedure in [38] in that they only gave a sufficient condition for its stopping rule to be satisfied.

5.5 Example Applications

A number of applications can be shown to satisfy the Type I or Type II Assumptions or some combination thereof, and can therefore be solved by one or more of the solution procedures in the previous section. We show that the applications in chapters 2-4 fit in one or more of these sets of assumptions.

5.5.1 Concave Cost, Single-Item Production Planning

Under the assumption that there exist uniform upper bounds on demand ($\bar{d}$) and on the number of periods between optimal production points (L), we see that the second Type I Assumption is satisfied. Since production capacity is unbounded and a unit of inventory is never optimally carried more than L periods, we also see that the third Type I Assumption is satisfied. Finally, since we allow the possibility that no demand occurs beyond period N of any N-period problem, and any finite horizon problem always optimally terminates with zero inventory, we see that the first Type I Assumption is also satisfied. The sets $T_n^{**}(p)$ for any period n and any problem p are the inventory levels $\{0, \ldots, L\bar{d}\}$. See Appendix D.2 for a full justification of this claim.

When the Type I Assumptions are met, problem $p \in \mathcal{P}$ is well-posed if and only if it is coalescent (by Theorem 5.2), and a direct implementation of Algorithm 5.1 will find an optimal initial decision for p if and only if it is well-posed (by Theorem 5.3). These are precisely the claims made in Chapter 2. Algorithm 2.1 runs differently. Rather than requiring the inventory ending period n to take on all possible values in $T_n^{**}(p)$, it instead looks over the optimal initial production decisions to the finite horizon problems for each of the previous L periods. Either method may be used.

5.5.2 Equipment Replacement

The application in Chapter 3 differs from the production planning example in that there exists a uniform upper bound on the maximum *feasible* lifetime. The state in this application is the remaining scheduled lifetime of the current equipment. By allowing the machine to become obsolete at any time, the first Type II Assumption is satisfied through careful construction of a forecast extension. Moreover, because of the bound on feasible lifetimes and the freedom to purchase machines at any time, the second Type II Assumption is met. Finally, since it is optimal to have only one piece of equipment on hand at any given time, and replacement decisions are limited to the lifetime of the next piece of equipment, the third Type II Assumption is met.

With the Type II Assumptions satisfied, Corollary 5.6 holds and problems are well-posed if and only if they are coalescent. Additionally, since the feasible state and decision sets are always known, an implementation of Algorithm 5.3 will stop finitely if and only the problem is well-posed, as established by Theorem 5.7. As with the production planning example in Chapter 2, these claims were proven for the equipment replacement problem in Chapter 3.

5.5.3 Capacitated Inventory Planning

Because of limited inventory capacity $\bar{S}$, both the set of potentially optimal decisions and the set of feasible states are finite in each period. Thus, the third Type I Assumption is satisfied. The only reason why the third Type II Assumption is not satisfied is that the sets of feasible decisions are not finite. However, the sets of potentially optimal decisions *are* finite and known since it is never optimal to replenish inventory above the capacity level, so in effect, the third Type II Assumption is satisfied for this problem.

We know from Remarks 4.2 and 4.3 that the second Type II Assumption is satisfied, and by Remark 4.5, the first Type II Assumption is also satisfied. With the Type II Assumptions satisfied, it follows from Corollary 5.6 that p is well-posed if and only if it is coalescent. Since the sets of potentially optimal states and decisions are known in all periods for any given problem, Assumption 5.4 is satisfied. Thus, from Theorem 5.7, an implementation of Algorithm 5.3 stops finitely if and only if

p is well-posed. As with the previous two examples, these claims were previously established in Chapter 4.

5.6 Conclusions

We have studied a general class of deterministic, nonstationary, infinite horizon optimization problems. We defined a condition called coalescence for these problems, and showed that it is equivalent to the property that an infinite horizon optimal initial decision can be found by knowing only finitely many periods' worth of problem parameters, regardless of future parameters. This relationship between coalescence and solvability was shown to hold for two different sets of assumptions: a weaker one in which potentially optimal states are finite in each time period and satisfy a reachability property, and a stronger one in which feasible states are finite in each time period and satisfy a similar reachability property.

We also gave solution procedures for these two sets of assumptions, and included a general method originating from [24] to construct the sets of potentially feasible states. Finally, we showed that the two sets of assumptions are, in fact, satisfied by some meaningful applications.

Chapter 6

Solvability in a Stochastic Environment: An Infinite Horizon Asset Selling Problem

6.1 Introduction

The asset selling problem is a well-known optimal stopping problem. An asset owner, faced with a sequence of nonnegative random offers, must choose in real time whether to accept or reject each offer. One offer arrives per period, and if that offer is not accepted, the owner incurs a penalty for continuing to hold the asset. Moreover, the capital invested in the asset could accrue interest at a rate $r \geq 0$, so it is appropriate to discount future rewards by a factor $\alpha = \frac{1}{1+r}$. Once an offer is accepted, no further offers may be accepted. Thus, the problem is to determine, a priori, a decision rule for each period that indicates which offers should be accepted and which should be rejected.

We assume that the asset owner can forecast the penalty costs and the probability distributions of offers in each period arbitrarily far into the future. Unlike Bertsekas [7] and Hayes [27], we do not assume that the offers are identically distributed. Derman et al [20] assume that the offers are stationary, but that they come from one of $k \geq 2$ underlying probability distributions, and that that distribution is unknown when the problem is initialized. In that work, observing a finite sequence of offers provides some learning about their underlying distribution. Rosenfield et al [36] generalize upon that by assuming that the family of distributions is unknown and learning about the distribution occurs over time, although the offers all come from the same distribution.

In all of these works mentioned above, it is shown that the optimal policy for a sequence of offer distributions (whether finite or infinite) is to set a threshold c_n in each period n such that offers below c_n are rejected, the owner is indifferent between acceptance and rejection for offers of value c_n, and offers exceeding c_n are accepted. All the aforementioned approaches to the asset selling problem have placed the set of possible offers on an interval (possibly infinite) in $\mathbb{R}$. Breaking from tradition, we will follow the approach of Cheevaprawatdomrong et al [16], who cast the infinite horizon asset selling problem within a larger class of nonhomogeneous MDPs on a bounded, discrete state space. In this formulation, all reward and transition functions can vary arbitrarily with time, so that no two periods need necessarily have the same parameters. As we shall show, this approach complicates the analysis in that two or more thresholds may be optimal for a given period, and it precludes the possibility of knowing the entire forecast simultaneously, but it also facilitates the construction of a forecast horizon and encompasses a large class of potential problems.

This chapter will proceed as follows. In section 6.2, we formally present the infinite horizon asset selling problem, as well as its finite horizon truncation. Under discounting and a finite upper bound on the offer distributions and penalty costs in each period, we show in section 6.3 that optimal thresholds in each period of the N-horizon truncation are monotone nondecreasing in N. Moreover, we show that uniqueness of the optimal initial threshold is sufficient to guarantee that there exists a forecast horizon, and comment on the likelihood that this condition is met.

In section 6.4, we give a deterministic reformulation of the asset selling problem and analyze the properties of optimal state sequences in the deterministic formulation to give both necessary and sufficient conditions for the existence of a forecast horizon. As in Chapters 2-5, this necessary and sufficient condition is coalescence. However, coalescence in the context of asset selling problems is easier to verify and has a richer meaning than in the other applications in this thesis. In section 6.5, we present an algorithm that will find a forecast horizon any time one exists, and therefore, that can solve for *all* optimal decisions for any forecast in which, looking forward from each period, there exists a forecast horizon for that period's optimal threshold. We also show the results of simulations designed to test the performance of the algorithm. Finally, in section 6.6, we give some concluding remarks.

6.2 Model, Assumptions, and the Form of an Optimal Policy

We begin by formulating the infinite horizon, nonstationary asset selling problem as a Markov Decision Process, and characterize the form of its optimal solutions.

6.2.1 General Model Definition and Description

An asset owner, faced with a sequence of nonnegative random offers, must choose in real time whether to accept or reject each offer for the asset. One offer arrives per period, and if that offer is not accepted, the owner incurs a deterministic penalty, possibly zero, for continuing to hold the asset. Future offers are discounted at a single-period rate α. Offers are assumed to come from the set of integers $\{l, \ldots, u\}$, where $l \geq 0$. We will use the convention that offers arrive at the beginning of each period and decisions are made at the end of the same period. The penalty immediately follows a decision to retain the asset and is not discounted to the next period.

More formally, the infinite horizon asset selling problem can be expressed as a Markov Decision Process, which we now proceed to do, following the development in [16].

In each period n, the state i is an element of the set $\{l, \ldots, u+1\}$ for some integers $0 \leq l < u$. When $l \leq i \leq u$, i represents the random offer received in period n, while the state $u+1$ represents the state of the system after the sale of the asset. We thus denote the state space $\mathcal{S} = \{l, \ldots, u+1\}$. We assume that the offer in period n follows the probability mass function q_n, so that for each n, $\sum_{i=l}^{u} q_n(i) = 1$ and $q_n(u+1) = 0$. Let F_n be the cumulative distribution function corresponding to q_n.

Corresponding to each state $i \in \mathcal{S}$, the decision set D_i is given by

$$D_i = \begin{cases} \{0,1\}, & \text{if } l \leq i \leq u, \\ \{1\}, & \text{if } i = u+1, \end{cases}$$

where the decision $k = 0$ represents rejection and the decision $k = 1$ represents acceptance. We define the transition structure as

$$p_n(i,j;k) = \begin{cases} q_n(j), & \text{if} \quad l \leq j \leq u, \quad k = 0 \\ 0, & \text{if} \quad j = u+1, \quad k = 0; \quad \text{or} \quad l \leq j \leq u, \quad k = 1 \\ 1, & \text{if} \quad j = u+1, \quad k = 1. \end{cases}$$

Furthermore, for each period n, let the reward function be

$$\rho_n(i,k) = \begin{cases} -h_n, & \text{if} \quad l \leq i \leq u \quad k = 0, \\ i, & \text{if} \quad l \leq i \leq u, \quad k = 1, \\ 0, & \text{if} \quad i = u+1, \quad k = 1, \end{cases}$$

where $h_n \geq 0$ is the holding cost in period n.

With the transition and reward functions defined, we now define the *forecast* or *problem* as $\phi = (\phi_n)_{n=1}^{\infty} = (p_n, \rho_n)_{n=1}^{\infty}$. Let Φ be the collection of all forecasts. We will assume that for all $\phi \in \Phi$ and all periods n, $0 \leq h_n \leq \alpha u$. Let $\mu_n(\phi) \equiv \sum_{i=l}^{u} i \cdot q_n(i)$ be the expected value of the offer in period n under forecast ϕ.

Because of the impossibility of capturing an entire infinite horizon forecast at once, solution methods for the infinite horizon asset selling problem must rely upon incrementally longer finite horizon problems. To that end, let ϕ^N represent the N-horizon truncation of ϕ, and let $\Phi^N(\phi)$ be the class of all forecasts that are equivalent to ϕ through period N.

By definition of the decision sets $D_i, i \in \mathcal{S}$, it is clear that a *strategy* π specifies, for each state i in each period n, some $k \in D_i \subseteq \{0,1\}$. Let Π denote the set of all feasible strategies, and let

$$\beta_n(\pi;\phi) = \prod_{k=1}^{n-1} \sum_{i:\pi_k(i)=0} q_k(i)$$

for $n \geq 2$, where $\beta_1(\pi;\phi) = 1$. Thus, $\beta_n(\pi;\phi)$ is the probability of rejecting the offers in periods 1 through $n-1$ under forecast ϕ and strategy π. Then

$$R_n(\pi;\phi) = \sum_{k=n}^{\infty} \alpha^{k-n} \frac{\beta_k(\pi;\phi)}{\beta_n(\pi;\phi)} \left(\sum_{i:\pi_k(i)=1} i \cdot q_k(i) - \sum_{i:\pi_k(i)=0} h_k \cdot q_k(i) \right)$$

represents the total expected discounted rewards beginning period n under forecast ϕ and strategy π. The optimization problem can be expressed as

$$\langle \phi \rangle \equiv \max_{\pi \in \Pi} R_1(\pi;\phi).$$

Let $V_n(\phi) = \max_{\pi \in \Pi} R_n(\pi;\phi)$ be the maximum total expected discounted rewards beginning period n under forecast ϕ.

6.2.2 Form of an Optimal Policy

Cheevaprawatdomrong et al show that one can optimally restrict consideration of strategies to those for which in each period n, a value c_n is chosen such that for $i < c_n$, the decision for state i is 0, while for $i \geq c_n$, the decision for state i is 1 (including, of course, state $u+1$). In particular, the value for c_n can be determined from the forecast ϕ as the function

$$c_n(\phi) = \alpha V_{n+1}(\phi) - h_n.$$

That is, c_n is the maximum total expected discounted rewards incurred when the period n offer is rejected. Clearly, if the offer received in period n meets or exceeds $c_n(\phi)$, then it should be accepted. Let

$$S_n^*(\phi) = \{i \in \mathcal{S} : i \geq c_n(\phi), q_n(i) > 0\}$$

and let $i_n^*(\phi)$ be the smallest element of $S_n^*(\phi)$. Then $i_n^*(\phi)$ is called the *optimal period n threshold* under ϕ. From this point forward, we will characterize a policy π as a sequence of thresholds, i.e., $\pi = (i_n)_{n=1}^{\infty}$. Note that we have restricted optimal thresholds to take on offer levels that have positive probability.

6.3 Problem Solvability

In solving an asset selling problem, the objective is to determine the optimal thresholds prior to observing the offers so that one can act optimally when faced with offers. Again, however, because of the impossibility of capturing an entire infinite horizon forecast at once, one must rely upon the convergence of the optimal first period threshold of increasingly longer finite horizon problems to determine the optimal first period threshold of the infinite horizon problem. This section will give conditions under which an infinite horizon asset selling problem can be solved by finite horizon techniques. We begin by establishing the monotonicity of the optimal initial threshold with respect to the problem horizon.

We note that a policy monotonicity result was already obtained in [16], section 3, for optimal policies in *all* periods, in a more general class of Markov Decision Processes. This problem fits within that class of problems, so the monotonicity results we obtain are not novel. However, they will be used at subsequent points in this chapter in a way particular to the asset selling problem, so the derivation here is necessary for our use.

6.3.1 Monotonicity of the Optimal Thresholds

In this section, we will discuss the monotonicity of the *initial* optimal threshold. However, the results for the initial threshold do, in fact, hold for the thresholds in all periods. Please refer to E.1 for further analysis on this matter.

Let ϕ^N represent the N-horizon truncation of the infinite horizon forecast ϕ, so that only N offers are available and $\langle\phi\rangle$ is an N-period optimization problem. Then

$$f(i_1, \phi^N) = \left[\alpha V_2(\phi^N) - h_1\right] \sum_{i=l}^{i_1 - 1} q_1(i) + \sum_{i=i_1}^{u} i \cdot q_1(i) \tag{6.1}$$

is the maximum expected discounted revenue for the N-period problem ϕ^N when setting the initial threshold to i_1. We wish to analyze the behavior of $f(i_1, \phi^N)$ with respect to changes in both i_1 and N. Let

$$\delta(i_1, \phi^N) = f(i_1 + 1, \phi^N) - f(i_1, \phi^N), \tag{6.2}$$

so that $\delta(i_1, \phi^N)$ represents the marginal value of increasing the initial threshold from i_1 to $i_1 + 1$ and continuing optimally in the N-period problem ϕ^N. The next result shows that $\delta(i_1, \phi^N)$ is monotonic with respect to N. This is important because for fixed N and ϕ, zeroes of $\delta(i_1, \phi^N)$ can identify maximizing values of i_1, which are the optimal initial thresholds for the N-period problems.

Lemma 6.1. $\delta(i_1, \phi^{N+1}) \geq \delta(i_1, \phi^N)$.

Proof. We first note that, for any $\phi \in \Phi$,

$$V_2(\phi^{N+1}) \geq V_2(\phi^N) \tag{6.3}$$

because all rewards are nonnegative (we will discuss this in more detail in section

6.3.2). Then

$$\begin{aligned}
\delta(i_1, \phi^N) &= f(i_1+1, \phi^N) - f(i_1, \phi^N) && (6.4)\\
&= \big[\alpha V_2(\phi^N) - h_1\big] \sum_{i=l}^{i_1} q_1(i) + \sum_{i=i_1+1}^{u} i \cdot q_1(i) \\
&\quad - \big[\alpha V_2(\phi^N) - h_1\big] \sum_{i=l}^{i_1-1} q_1(i) - \sum_{i=i_1}^{u} i \cdot q_1(i) && (6.5)\\
&= \big[\alpha V_2(\phi^N) - h_1 - i_1\big] \cdot q_1(i_1). && (6.6)
\end{aligned}$$

Similarly, we obtain $\delta(i_1, \phi^{N+1}) = \big[\alpha V_2(\phi^{N+1}) - h_1 - i_1\big] \cdot q_1(i_1)$. Applying (6.3), the desired result follows. □

Lemma 6.2. *For $l \leq i_1 < i_1^*(\phi^N)$, $\delta(i_1, \phi^N) \geq 0$.*

Proof. Choose any $l \leq i_1 < i_1^*(\phi^N)$. Then

$$\begin{aligned}
\delta(i_1, \phi^N) &= f(i_1+1, \phi^N) - f(i_1, \phi^N) \\
&= q_1(i_1)\big[\alpha V_2(\phi^N) - h_1 - i_1\big] \\
&\geq 0
\end{aligned}$$

since $\alpha V_2(\phi^N) - h_1 > i_1$. Otherwise, there would be no loss of optimality in setting $i_1^*(\phi^N) = i_1$. □

Clearly, if $\delta(i_1^*(\phi^N), \phi^N) > 0$, then one could increase the expected total revenue by increasing i_1 by one unit. Thus, we state the following.

Lemma 6.3. *For fixed N and ϕ, the marginal value of increasing the initial threshold by one from its optimal value is nonpositive, i.e. $\delta(i_1^*(\phi^N), \phi^N) \leq 0$.*

From Lemmas 6.2 and 6.3, we can draw the following useful corollary.

Corollary 6.4. *If there are multiple values of i_1 that maximize $f(i, \phi^N)$ (i.e. multiple optimal initial thresholds for problem ϕ^N), then they are all consecutive integers.*

In the event that there are multiple maximizers of $f(i_1, \phi^N)$, we say that $i_1^*(\phi^N)$ is not unique. This property becomes important when investigating the solvability of a forecast, as we shall see in the next section. We are now prepared to present the most important result of this section.

Theorem 6.5 (Initial Threshold Monotonicity). $i_1^*(\phi^{N+1}) \geq i_1^*(\phi^N)$ *for all* $\phi \in \Phi$ *and for all periods* N.

Proof. By Lemma 6.3, in order for $i_1 < i_1^*(\phi^N)$ to be a maximizer of $f(i, \phi^{N+1})$, it must satisfy $\delta(i_1^*(\phi^N), \phi^N) \leq 0$. However, for all $i_1 < i_1^*(\phi^N)$

$$\begin{aligned} \delta(i_1, \phi^{N+1}) &\geq \delta(i_1, \phi^N) \\ &\geq 0, \end{aligned}$$

where the first inequality is from Lemma 6.1 and the second is from Lemma 6.2. So consider the case where $\delta(i_1, \phi^{N+1}) = 0$. Then clearly, $\delta(i_1, \phi^N) = 0$, but i_1 was not a maximizer of $f(i, \phi^N)$. Therefore, $\exists i_1', i_1 < i_1' < i_1^*(\phi^N)$ such that $\delta(i_1', \phi^N) > 0$. Moreover, $\delta(i_1', \phi^{N+1}) \geq \delta(i_1', \phi^N)$, so that $f(i_1, \phi^{N+1}) < f(i_1^*(\phi^{N+1}), \phi^{N+1})$, and i_1 is not an optimal initial threshold for ϕ^{N+1}. Having chosen i_1 arbitrarily, we conclude that $i_1^*(\phi^{N+1}) \geq i_1^*(\phi^N)$ and the proof is complete. $\square$

Figure 6.1 shows the properties of $i_1^*(\phi^N)$ and $\delta(i_1, \phi^N)$ that we have just proved. For fixed values of i_1, $\delta(i_1, \phi^N)$ is monotone nondecreasing in N, and therefore, so is $i_1^*(\phi^N)$.

We note that an elegant alternate proof of Theorem 6.5 follows from the application of Theorem 6.1 of Topkis [41]. See E.2 for details. Again, Theorem 6.5 follows from Corollary 3.1 of [16]; we have simply recast it in the context of asset selling and obtained insight and useful results in the process.

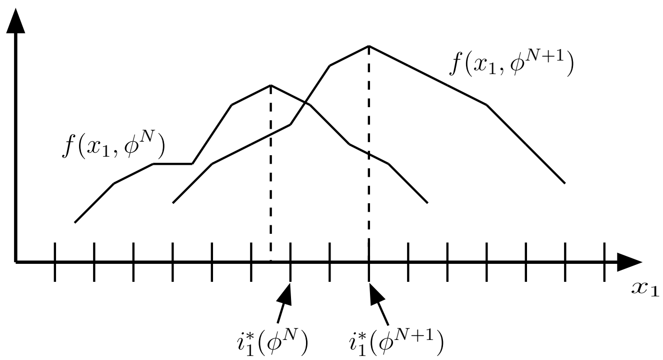

Figure 6.1: Monotonicity of Optimal Initial Thresholds

6.3.2 Characterizing Solvability of the Infinite Horizon Asset Selling Problem

We now direct our attention to conditions under which the infinite horizon problem ϕ has a common optimal initial threshold with every problem in agreement with at least the first N periods of ϕ. First, we define well-posed and forecast horizon as they pertain to the asset selling problem.

Definition 6.1. *Forecast $\phi \in \Phi$ is called* well-posed (with respect to Φ) *if, for some period N^*, $\exists i_1^*(\phi^{N^*})$ such that $i_1^*(\phi^{N^*})$ is an optimal initial threshold for all $\phi' \in \Phi^{N^*}(\phi)$.*

Definition 6.2. *Period N^* is called a* forecast horizon (for forecast ϕ in problem class Φ) *if $\exists i_1^*(\phi^{N^*})$ such that $i_1^*(\phi^{N^*})$ is an optimal initial threshold for all $\phi' \in \Phi^{N^*}(\phi)$.*

The next result follows immediately from the two preceding definitions and the results in the previous section. Recall that F_n represents the cumulative distribution function of the offers in period n under the forecast ϕ.

Remark 6.6. *Suppose that $F_{N'}(\lfloor \alpha u \rfloor) = 0$. Then N' is a forecast horizon for any problem ϕ having offer distribution forecast $q_{N'}$ in period N'.*

Proof. $F_{N'}(\lfloor \alpha u \rfloor) = 0$ implies that the offer arriving in period N' will certainly exceed the discounted value of any future offer, no matter the sequence of future offer distributions. Thus, any offer in period N' will be accepted under any $\phi' \in \Phi^{N'}(\phi)$, thereby terminating the optimal stopping problem. Thus, any optimal initial threshold for problem $\phi^{N'}$ is optimal for any $\phi' \in \Phi^{N'}(\phi)$, and N' is a forecast horizon for problem ϕ with respect to Φ. □

We cannot expect in general that an unbeatable offer will arrive with certainty. Thus, it is desirable to have more general conditions under which a forecast horizon exists for forecast $\phi \in \Phi$. In [16], the authors present an instance of an asset selling problem in which there are two values of i_1 that maximize $f(i_1, \phi)$. They then show that by making a slight perturbation in the holding penalty arbitrarily many periods into the future, one can render one of the two infinite horizon optimal initial thresholds non-optimal. Because this perturbation can be done at any time in the future and still have the effect of changing the set of optimal initial thresholds, they conclude that the given instance is not well-posed. It would seem, then, that if ϕ has a unique infinite horizon optimal initial threshold, then one might be able to avoid these types of issues. The second result below confirms the truth of this observation. The first result is a technical lemma used in the proof of the second.

Lemma 6.7.

$$V_1(\phi^{N+1}) \leq V_1(\phi^N) - \alpha^{N-1}h_N + \alpha^N \mu_{N+1}(\phi),$$

and

$$V_1(\phi) \leq V_1(\phi^N) - \alpha^{N-1}h_N + \alpha^N u.$$

If $h_N \geq \alpha\mu_{N+1}(\phi)$, remove the $-\alpha^{N-1}h_N + \alpha^N \mu_{N+1}(\phi)$ from the first claim and change its inequality to equality.

Proof. We begin by noting that $i^*_N(\phi^N) = l$. If $i^*_N(\phi^{N+1}) = i^*_N(\phi^N)$ (which is the case if $\alpha^{N-1}h_N \geq \alpha^N \mu_{N+1}(\phi)$), then no offer will be accepted in period $N+1$ and hence, $V(\phi^{N+1}) = V(\phi^N)$ and the result holds trivially. So suppose that $i^*_N(\phi^{N+1}) > i^*_N(\phi^N)$. Since the $i^*_k(\phi^N), k = 1, \ldots, N$ were chosen to maximize expected revenue for ϕ^N, it must be the case that the expected revenue over the first N periods of ϕ^{N+1} is at most

$V(\phi^N)$. Also, $i^*_{N+1}(\phi^{N+1}) = l$, giving the result that the (undiscounted) expected revenue in period $N+1$ of ϕ^{N+1}, conditional upon rejecting all previous offers, is $\mu_{N+1}(\phi)$. However, one incurs a penalty of $\alpha^{N-1}h_N$ when rejecting the offer in period N in order to accept the offer in period $N+1$. Thus, $V(\phi^{N+1}) \leq V(\phi^N) - \alpha^{N-1}h_N + \alpha^N \mu_{N+1}(\phi)$, as desired.

To prove the second claim, it suffices to note that the total rewards incurred over periods $N+1$ onward, for any forecast, are at most $\alpha^N u$. □

If one wished to give a bound on the increase in the rewards to go from the second period ($V_2(\phi^{N+1})$ vs. $V_2(\phi^N)$), Lemma 6.7 could be slightly modified by changing the $V_1(\cdot)$ terms to $V_2(\cdot)$ and multiplying the latter by α. We now give the main result of this section. This same result is established for a more general class of problems by Theorems 6.6-6.7 of [16]. The hope in recasting it here is to make the result more intuitive and transparent for asset selling problems in particular.

Theorem 6.8. *$\phi \in \Phi$ is well-posed if $i^*_1(\phi)$ is unique.*

Proof. Suppose that $i^*_1(\phi)$ is unique. Then $i^*_1(\phi^N) \to i^*_1(\phi)$ monotonically in N by Theorem 6.5. Since $i^*_1(\phi^N)$ is integer-valued, it must attain $i^*_1(\phi)$ by some finite time; call this time N'. Then, for all $N > N', i^*_1(\phi^N) = i^*_1(\phi^{N'})$. To see that $i^*_1(\phi'^N) = i^*_1(\phi^{N'}) = i^*_1(\phi)$ for all $\phi' \in \Phi^{N'}(\phi)$ and $N \geq N'$ (implying that ϕ is well-posed), let

$$N^*(\phi) = \min\left[N | i^*_1(\phi^N) > \alpha V_2(\phi) - h_1 - \alpha^{N-1}h_N + \alpha^N u\right]. \tag{6.7}$$

Since $i^*_1(\phi)$ is unique, $\alpha V_2(\phi) - h_1$ is not integer-valued, and therefore $N^*(\phi)$ is finite. By Lemma 6.7, for any $\phi' \in \Phi^{N^*(\phi)}(\phi)$, we have that

$$\alpha V_2(\phi') \leq \alpha V_2(\phi^N) - \alpha^{N-1}h_N + \alpha^N u,$$

and therefore

$$\alpha V_2(\phi') - h_1 \leq \alpha V_2(\phi^N) - h_1 - \alpha^{N-1}h_N + \alpha^N u.$$

However, for all $N \geq N^*(\phi)$,

$$\begin{aligned}
\alpha V_2(\phi'^N) - h_1 &\leq \alpha V_2(\phi') - h_1 \\
&\leq \alpha V_2(\phi^N) - h_1 - \alpha^{N-1} h_N + \alpha^N u \\
&\leq \alpha V_2(\phi) - h_1 - \alpha^{N-1} h_N + \alpha^N u \\
&< i_1^*(\phi^N) \\
&= i_1^*(\phi)
\end{aligned}$$

But $i_1^*(\phi'^N)$ is the minimum integer which is at least $\alpha V_2(\phi'^N) - h_1$ and $i_1^*(\phi'^N)$ is nondecreasing in N, so we see that $i_1^*(\phi'^N) = i_1^*(\phi)$ for all $N \geq N^*(\phi)$. This completes the proof. □

It would be very desirable to be able to place a bound on $N^*(\phi)$ independent of ϕ, for this would limit the number of periods of parameters to forecast in order to find an optimal initial threshold. However, since $\alpha V_2(\phi) - h_1$ can approach an integer arbitrarily close from below depending upon ϕ, this is not possible. We will give in §6.5 a solution procedure that will solve for the optimal threshold for any problem in which it is possible to do so.

6.3.3 Examples

First, a well-posed asset selling forecast with multiple optimal initial thresholds. Let $l = 1, u = 2$, and $q_n(i)$ represents the probability of receiving an offer of value i in period n. Let $\alpha, 0 < \alpha < 1$ be the discount factor. Suppose that $h_1 = 0.9$, and the first two periods have offer distribution forecasts as follows.

$$\begin{aligned}
q_1(1) &= a \\
q_1(2) &= 1 - a \\
q_2(1) &= 0 \\
q_2(2) &= 1
\end{aligned} \tag{6.8}$$

That is, an offer of value 2 arrives with probability 1 in period 2. Clearly, no offer after period 2 can have a present worth greater than 2α, the present worth of accepting the offer of value 2 in period 2. By Remark 6.6, an optimal policy over the first two periods is optimal for any longer horizon problem.

There are three possible thresholds in period 1: 1 (accept any offer), 2 (accept only an offer of value 2), and 3 (don't accept any offers). Let $V(1), V(2)$ and $V(3)$ be the expected discounted total revenue of using initial thresholds 1, 2, and 3, respectively. Then

$$\begin{aligned} V(1) &= a + 2(1-a) = 2 - a \\ V(2) &= 2 - 2a + a(2\alpha - 0.9) = 2a\alpha + 2 - 2.9a \\ V(3) &= 2\alpha - 0.9 \end{aligned}$$

In order to have multiple optimal initial thresholds, we must have $V(1) = V(2) > V(3), V(1) = V(3) > V(2)$, or $V(2) = V(3) > V(1)$. Choosing the first of the three options, we must satisfy $2a\alpha + 2 - 2.9a = 2 - a$, or $\alpha = 0.95$. Substituting $\alpha = 0.95$ into the equation for $V(3)$, we find that $V(3) = 1$. Now, $V(1) = 2 - a > 1$ for any $0 < a < 1$, and since $V(1) = V(2)$ by design, we conclude that $V(1)$ and $V(2)$ are both optimal initial thresholds for any asset selling forecast with the first two periods satisfying (6.8) when $\alpha = 0.95$. Thus, any ϕ with q_1, q_2, h_1, and α as given above is well-posed with multiple initial optimal thresholds.

Now we present a non-well-posed asset selling forecast that will appear similar to the example in section 1 of [16]. As before, let $l = 1, u = 2$, and fix α. To simplify the analysis, suppose that there are no holding penalties. For all periods n, let the offer distribution forecast be

$$\begin{aligned} q_n(1) &= a, \\ q_n(2) &= 1 - a. \end{aligned}$$

Thus, ϕ is a stationary forecast. The two stationary policies are to set the threshold equal to 1 and to 2 in all periods (the third policy, setting the threshold equal to 3, has zero expected revenue and need not be considered). Denote the total expected revenue of these two policies by $V(1)$ and $V(2)$, respectively. We can express these two quantities as

$$\begin{aligned} V(1) &= a + 2(1-a) \\ &= 2 - a, \\ V(2) &= a\alpha V(2) + 2(1-a), \end{aligned}$$

so that $V(2) = \frac{2-2a}{1-a\alpha}$. We can make the total expected revenue the same under the two policies by setting $V(1) = V(2)$ and finding that a must satisfy

$$a = \frac{2\alpha - 1}{\alpha}. \tag{6.9}$$

Note that $V(1)$ depends only upon the period 1 forecast, whereas $V(2)$ depends upon all forecasts. Further, it can be shown that for all $0 < a, \alpha < 1$, $V(2)$ is increasing in a. Thus, if at some period $n \geq 1$, $q_n(1) = a + \epsilon$ and $q_n(2) = 1 - a - \epsilon$, the total expected revenue from period n forward increases from that under the original stationary forecast. The opposite effect can be achieved by decreasing $q_n(1)$ by ϵ. By performing this perturbation on future probabilities, the tie between $V(1)$ and $V(2)$ can be broken. Moreover, because we can perform the perturbation for any n and achieve the desired result (i.e. breaking the tie), thus rendering one of the two initial thresholds non-optimal, we conclude that ϕ is not well-posed with respect to Φ. Note that we could achieve a similar result by introducing a holding penalty at some point in the future, which would have the effect of encouraging the asset owner to lower (as compared to the thresholds under the forecast without any penalties) the acceptance thresholds prior to the appearance of the penalty. The fundamental source of the unsolvability of this problem instance is the fact that Φ consists of, in general, nonstationary forecasts, so that it can never be known if the observed forecast is truly stationary.

6.3.4 Likelihood That the Optimal Initial Threshold is Not Unique

We shall show that the conditions under which the infinite horizon optimal initial threshold is not unique are quite rare. By Corollary 6.4, multiple infinite horizon optimal initial thresholds are consecutive integers. If $i_1^*(\phi)$ and $i_1^*(\phi)+1$ are both optimal initial thresholds for ϕ, then they result in the same total expected discounted revenue, i.e.

$$\begin{aligned}
\sum_{i=1}^{i_1^*(\phi)-1} \left[\alpha V_2(\phi) - h_1\right] q_1(i) + \sum_{i=i_1^*(\phi)}^{u} i \cdot q_1(i) &= \sum_{i=1}^{i_1^*(\phi)} \left[\alpha V_2(\phi) - h_1\right] q_1(i) \\
&\quad + \sum_{i=i_1^*(\phi)+1}^{u} i \cdot q_1(i) \\
i_1^*(\phi) \cdot q_1(i_1^*(\phi)) &= \left[\alpha V_2(\phi) - h_1\right] q_1(i_1^*(\phi)) \\
i_1^*(\phi) &= \alpha V_2(\phi) - h_1 \qquad (6.10)
\end{aligned}$$

In order for (6.10) to occur, there must be a careful balancing of the offer distributions and the discount factor, which is difficult to achieve except by intentional construction.

Lemma 6.9. *For fixed ϕ, $\alpha V_2(\phi) - h_1$ is strictly increasing in α.*

Proof. Observe that

$$V_2(\phi) = \sum_{n=2}^{\infty} \alpha^{n-2} \frac{\prod_{m=1}^{n-1} F_m(i_m^*(\phi)-1)}{F_1(i_1^*(\phi)-1)} \left[-F_n(i_n^*(\phi)-1) \cdot h_n + \sum_{i=i_k^*(\phi)}^{u} q_n(i) \cdot i \right].$$

If the $i_n^*(\phi)$ are fixed in α for each n, then the proof follows immediately. However, the $i_n^*(\phi)$ are chosen to maximize $V_2(\phi)$, so that increasing α can only increase $V_2(\phi)$ and therefore $\alpha V_2(\phi) - h_1$ is strictly increasing in α. □

Lemma 6.10. *For fixed ϕ, there exists at most one $\alpha_i \in [0,1]$ such that $\alpha_i V_2(\phi) - h_1 = i$, for $i \in \{l, \ldots, u\}$.*

Proof. This follows immediately from Lemma 6.9. □

The following theorem unifies our discussion of this section and shows that practically all asset-selling instances can be solved exactly using finite horizon methods.

Theorem 6.11. *Each $\phi \in \Phi$ is well-posed for almost every discount factor.*

Proof. Choose $\phi \in \Phi$. By Lemma 6.10, there are at most $u-l+1$ values of α, call them $\alpha_l, \ldots, \alpha_u$, such that $\alpha V_2(\phi) - h_1 = i$ is integer-valued, resulting in multiple optimal initial thresholds. Since α is continuous on the interval $(0, 1)$, the set $\{\alpha_l, \ldots, \alpha_u\}$ has (Lebesgue) measure zero, so that ϕ has a unique optimal initial threshold is therefore well-posed for almost every discount factor. □

6.4 Analyzing Solvability Via a Deterministic Reformulation

In this section, we present the asset selling problem as a deterministic dynamic program and show that its optimal solutions are the probability distributions over the stochastic states generated by choosing a sequence of optimal thresholds. With the characterization of optimal deterministic state sequences, we can more easily determine necessary and sufficient conditions for well-posedness of an asset selling forecast.

6.4.1 The Deterministic Formulation of an Asset Selling Forecast

Any asset selling forecast $\phi \in \Phi$ can be recast as a deterministic dynamic program. The states, rather than representing the current offer value or the position of having already sold the asset, become *distribution functions* over the possible offer levels and the position of having already sold the asset. The decision tree has, for each node (state), $u - l + 1$ arcs emanating from it, one for each possible threshold level. The distribution functions are determined by the sequence of thresholds lying along its path from the root node, which is null, along with the offer distribution of the current period. We note that the sum of the probabilities over offer levels l through u in each node is, in general, strictly less than one, due to the increasing probability over time of having already sold the asset (this probability is included in the distribution function).

More precisely, states are determined as follows. Let t be a node in the decision tree, with the period of t denoted by $T(t)$ and the predecessor node of t by $A(t)$. Let the thresholds that generate the state sequence leading to t be given by $i_1^*(t), i_2^*(t), \ldots, i_{T(t)}^*(t)$. Define $\beta(t) = \prod_{n=1}^{T(t)-1} F_n(i_n^*(t)-1)$, so that $\beta(t)$ represents the probability of rejecting the offers in periods 1 through $T(t)-1$ given the threshold sequence on the path to t. For any node t such that $T(t) = 1$, let $\beta(t) = 1$. The reward associated with node t is

$$R(t) = \alpha^{T(t)-1}\beta(t)\left[\sum_{j=i_{T(t)}^*(t)}^{u} q_{T(t)}(j) \cdot j - F_{T(t)}(i_{T(t)}^*(t)-1) \cdot h_{T(t)}\right]. \tag{6.11}$$

Thus, $R(t)$ accounts for the effects of monetary discounting, as well as the diminishing probability of holding the asset until time $T(t)$, so that $R(t)$ approaches zero as $T(t)$ increases. Moreover, if at any node along the path to t, a threshold is set so low that any offer is accepted, $\beta(t) = 0$ and no rewards are incurred in node t. Aside from the effects of $\beta(t)$ and $\alpha^{T(t)-1}$, the expected rewards in node t are the scalar product of the vector of offer levels that are at least $i_{T(t)}^*(t)$ with their corresponding vector of probabilities in period $T(t)$, minus the probability of rejecting the offer times the value of the holding penalty in period $T(t)$. With a concise description of the period, predecessor node, previous threshold sequence, survival probability, and rewards associated with each node in the decision tree, we now give a graphical illustration of the state of the node. As each node has a single entering arc, these parameters are unique for each t. Again, the state is a distribution function. Notice that the probability of *not* being in $u+1$ when exiting node t decreases by a factor of $\frac{1}{F_{T(t)}(i_{T(t)}^*(t)-1)}$ from the same probability when exiting node $A(t)$, as the conditional probability of rejecting the asset in node t given rejection in all of node t's predecessors is $F_{T(t)}(i_{T(t)}^*(t)-1)$. In general, the state description for node t approaches a distribution with zero probability in all of the offer levels and probability

Value	Probability
l	$\beta(t)q_{T(t)}(l)$
$l+1$	$\beta(t)q_{T(t)}(l+1)$
$\vdots$	$\vdots$
$i^*_{T(t)}(t)-1$	$\beta(t)q_{T(t)}(i^*_{T(t)}(t)-1)$
$i^*_{T(t)}(t)$	0
$\vdots$	$\vdots$
u	0
$u+1$	$1-\beta(t)F_{T(t)}(i^*_{T(t)}(t)-1)$

Table 6.1: Generating a Deterministic State from a Threshold Sequence

one in $u+1$ as $T(t)$ grows large. See Appendix E.3 for a sketch of a tree of deterministic states generated by setting threshold levels for a given asset selling forecast.

Proposition 6.12. *Choose any forecast $\phi \in \Phi$ and any sequence of thresholds $\{i^*_n\}$, and let $\{t_n\}$ be the corresponding sequence of nodes in the deterministic formulation. Then*

$$\frac{1}{\beta(t_N)} \sum_{n=N}^{\infty} R(t_n)$$

is the total expected discounted reward beginning period N of choosing threshold sequence $\{i^*_n\}$ for ϕ, conditional upon rejecting the offers in periods 1 through $N-1$.

Proof. Proof by induction. Conditional upon rejecting the offers in periods 1 through $N-1$ for ϕ, one expects to incur a reward of

$$\alpha^{N-1}\left[\sum_{j=i^*_N}^{u} q_N(j)\cdot j - F_N(i^*_N-1)\cdot h_N\right] = \frac{R(t_N)}{\beta(t_N)}$$

in period N. The probability of rejecting the period N offer is $F_N(i^*_N-1)$. Then the expected discounted reward in period $N+1$ conditional upon rejecting the first

$N-1$ offers is

$$\alpha^N F_N(i_N^* - 1)\left[\sum_{j=i_{N+1}^*}^{u} q_{N+1}(j)\cdot j - F_{N+1}(i_{N+1}^* - 1)\cdot h_{N+1}\right] = \frac{R(t_{N+1})}{\beta(t_N)}.$$

So the total expected discounted rewards in periods N and $N+1$ conditional upon rejecting the first $N-1$ offers is $\frac{1}{\beta(t_N)}\left[R(t_N) + R(t_{N+1})\right]$. Now suppose that the result is true for periods N through some $M > N$. The probability of rejecting the offers in periods N through M is $\frac{\beta(M+1)}{\beta(N)}$, so the expected discounted reward in period $M+1$ is

$$\alpha^M \frac{\beta(t_{M+1})}{\beta(t_N)}\left[\sum_{j=i_{M+1}^*}^{u} q_{M+1}(j)\cdot j - F_{M+1}(i_{M+1}^* - 1)\cdot h_{M+1}\right] = \frac{R(t_{M+1})}{\beta(t_N)}.$$

Thus, for any $M > N$, the total expected rewards in periods N through M is given by $\frac{1}{\beta(t_N)}\sum_{n=N}^{M} R(t_n)$. Since total expected rewards are bounded and converge as $M \to \infty$, they are equal to $\frac{1}{\beta(t_N)}\sum_{n=N}^{\infty} R(t_n)$. □

We can immediately draw the two following useful corollaries.

Corollary 6.13. *Choose any forecast $\phi \in \Phi$ and any sequence of thresholds $\{i_n^*\}$, and let $\{t_n\}$ be the corresponding sequence of nodes in the deterministic formulation. Then*

$$\sum_{n=1}^{\infty} R(t_n)$$

is the total expected discounted reward of choosing threshold sequence $\{i_n^*\}$ for ϕ.

Proof. Apply Proposition 6.12 with $N = 1$. Since $\beta(t_1) = 1$, the result follows. □

Corollary 6.14. *For a given forecast $\phi \in \Phi$, a sequence of thresholds is optimal in the stochastic formulation if and only if its corresponding node sequence is optimal in the deterministic formulation.*

Proof. We have already shown that the total discounted expected rewards of a threshold sequence in the stochastic formulation is equivalent to the total rewards of the node sequence generated by the threshold sequence in the deterministic formulation.

Furthermore, any threshold sequence can be represented by a node sequence in the deterministic formulation, so the maximum expected discounted total rewards in the stochastic formulation forms a lower bound on the maximum reward node sequence in the deterministic formulation. It remains only to show that there is no node sequence that is not generated by some threshold sequence. But this is implied by the decision tree structure, wherein each node bears the cumulative effects of all thresholds that generated its predecessors. $\square$

6.4.2 Characterizing Solvable Asset Selling Forecasts Using the Deterministic Formulation

We would like to know when to expect to find an optimal initial threshold in finite time. Having shown that the total rewards of optimal paths in the decision tree are equal to the total expected rewards of optimal threshold sequence, we can now exploit the structure of the deterministic formulation to characterize problems whose optimal initial threshold can be solved in finite time.

Definition 6.3. *An optimal path $\{t_k\}$ in the decision tree of the deterministic formulation of $\phi \in \Phi$ is said to be* optimally reachable *from an optimal path $\{t'_k\}$ if there exists an optimal path $\{\hat{t}_k\}$ such that $\hat{t}_m = t_m$ and $\hat{t}_n = t'_n$ for some periods m and n, $n \geq m$.*

Although the decision tree structure forbids the intersection of any two paths whose initial arcs are distinct, two such paths *could* intersect after a state aggregation process, transforming the decision tree into a dynamic programming network. This is possible because the entire pertinent history - the survival probability after $N-1$ periods - of a threshold sequence is contained in $\beta(t_N)$. No information about past rewards and costs is necessary because only one reward can be incurred, and each cost is a one-time expense. Alternatively, this can be seen by recalling that the asset selling problem is a Markov decision process, so the problem is inherently forward-looking and there exist Markovian optimal policies.

Definition 6.4. $\phi \in \Phi$ *is* coalescent *if, in the decision tree of its deterministic formulation, there exists an optimal path* $\{t_k^*\}$ *such that any other optimal path is optimally reachable from* $\{t_k^*\}$*.*

We now explicitly identify coalescence as a structural property of an asset selling problem.

Lemma 6.15. $\phi \in \Phi$ *is coalescent if and only if one or both of the following is true.*

(a) $\alpha V_2(\phi) - h_1$ *is not integer-valued; i.e. the optimal initial threshold is unique.*

(b) For some period N'*,* $F_{N'}(i^*_{N'}(\phi) - 1) = 0$ *and* $\alpha V_{N'+1}(\phi) - h_{N'} < i^*_{N'}(\phi)$*; i.e. it is strictly optimal to accept an offer with probability one in period* N'*.*

Proof. If the optimal initial threshold is unique, then there is a unique optimal first-period node in the decision tree. Thus, any optimal path has the same first-period node, and ϕ is coalescent. If it is strictly optimal to accept an offer with probability one in some period N', then any optimal path has $\beta(t_N) = 0$ for all $N > N'$. Thus, all optimal paths have the same state for all periods beyond N', which has zero probability for all offer levels l through u and probability one for $u + 1$. In this case, any of the optimal paths can optimally reach any other one in period $N' + 1$, and again, ϕ is coalescent.

Now suppose that neither condition (a) nor condition (b) holds. Then the optimal initial threshold is not unique, and for any optimal node sequence $\{t_n\}$, $\beta(t_n) > 0$ for all k. For ease of exposition, we assume that "multiple optimal thresholds" means there are two optimal thresholds, although the argument to follow can easily be shown to hold for higher numbers of thresholds. If no other period has multiple optimal thresholds, then there are just two optimal node sequences $\{t_n^1\}$ and $\{t_n^2\}$, and since $\beta(t_2^1) \neq \beta(t_2^2)$, $\beta(t_n^1) \neq \beta(t_n^2)$ for all $n \geq 2$, so $t_n^1 \neq t_n^2$ for all $n \geq 1$ and ϕ is not coalescent.

Value	State $t^1_{N'}$	State $t^2_{N'}$	State $t^3_{N'}$	State $t^4_{N'}$
l	$\beta(t^1_{N'})q_{N'}(l)$	$\beta(t^1_{N'})q_{N'}(l)$	$\beta(t^2_{N'})q_{N'}(l)$	$\beta(t^2_{N'})q_{N'}(l)$
$\vdots$	$\vdots$	$\vdots$	$\vdots$	$\vdots$
$i^{*1}_{N'}-1$	$\beta(t^1_{N'})q_{N'}(i^{*1}_{N'}-1)$	$\beta(t^1_{N'})q_{N'}(i^{*1}_{N'}-1)$	$\beta(t^2_{N'})q_{N'}(i^{*1}_{N'}-1)$	$\beta(t^2_{N'})q_{N'}(i^{*1}_{N'}-1)$
$i^{*1}_{N'}$	0	$\beta(t^1_{N'})q_{N'}(i^{*1}_{N'})$	0	$\beta(t^2_{N'})q_{N'}(i^{*1}_{N'})$
$i^{*1}_{N'}+1$	0	0	0	0
$\vdots$	$\vdots$	$\vdots$	$\vdots$	$\vdots$
u	0	0	0	0
$u+1$	$1-\beta(t^1_{N'})F_{N'}(i^{*1}_{N'}-1)$	$1-\beta(t^1_{N'})F_{N'}(i^{*1}_{N'})$	$1-\beta(t^2_{N'})F_{N'}(i^{*1}_{N'}-1)$	$1-\beta(t^2_{N'})F_{N'}(i^{*1}_{N'})$

Table 6.2: Multiple Optimal States for Non-Coalescent Asset Selling Forecast

On the other hand, suppose there exists some period N' with multiple optimal thresholds $i^{*1}_{N'}$ and $i^{*2}_{N'} = c^1_{N'}+1$. Then there are four optimal period N' nodes, displayed in Table 6.2 in their tabular form.

Since $\beta(t^1_{N'}) \neq \beta(t^2_{N'})$, it can easily be seen that none of the four states are identical. Furthermore, any optimal state $t_N, N > N'$ must have one of these four states along its path from the root node. If all optimal thresholds beyond period N' are unique, then there are four optimal paths, and four optimal states in period $N'+1$, say $t^1_{N'+1}, \ldots, t^4_{N'+1}$. Let $t^1_{N'+1}$ and $t^2_{N'+1}$ have first-period state t^1_1 in their paths and let the other two have first period state t^2_1. Even though $t^1_{N'+1}$ and $t^2_{N'+1}$ have the same state in period 1, the fact that one arises from using threshold $i^{*1}_{N'}$ and the other from threshold $i^{*2}_{N'} = i^{*1}_{N'}+1$ in period N' means that $\beta(t^1_{N'+1}) \neq \beta(t^2_{N'+1})$. A similar argument shows that $\beta(t^3_{N'+1}) \neq \beta(t^4_{N'+1})$. While it may be true that $\beta(t^2_{N\prime+1}) = \beta(t^3_{N\prime+1})$ (if, for example, $F_1(i^*_1-1) = F_{N'}(i^*_{N'}-1)$ and $F_1(i^*_1) = F_{N'}(i^*_{N'})$), it is not true that $\beta(t^1_{N\prime+1}) = \beta(t^4_{N\prime+1})$; in fact, $\beta(t^1_{N\prime+1}) < \beta(t^4_{N\prime+1})$. Thus, if $\{t^1_n\}$ and $\{t^4_n\}$ are the optimal node sequences passing through states $t^1_{N'+1}$ and $t^4_{N'+1}$, respectively, we see that there is no optimal node sequence that optimally reaches both $\{t^1_n\}$ and $\{t^4_n\}$.

Applying the same argument to any quadruplets of optimal states in any future period, or even to higher numbers of optimal states if there are other periods with multiple optimal thresholds, we can construct at least one pair of optimal state sequences, neither of which can be optimally reached by a common optimal path. We conclude that ϕ is not coalescent. $\square$

We now present the main result of this section.

Theorem 6.16. $\phi \in \Phi$ *is well-posed if and only if it is coalescent.*

Proof. First, suppose that ϕ is coalescent. Then, by Lemma 6.15, either the optimal initial threshold is unique or it is strictly optimal to accept an offer with probability one in some period N' (or both). We will show that either of these conditions establishes well-posedness of ϕ.

Case 1: $i_1^*(\phi)$ is unique. That uniqueness of the optimal initial threshold implies well-posedness was already established in Theorem 6.8.

Case 2: $F_{N'}(i_{N'}^*(\phi) - 1) = 0$. In this case, we can determine N^* such that $i_1^*(\phi)$ is also an optimal initial threshold for any $\phi' \in \Phi^{N^*}(\phi)$. Define

$$M^*(\phi) = \min\left[N | i_{N'}^*(\phi^N) > \alpha V_{N'+1}(\phi) - h_{N'} - \alpha^{N-1} h_N + \alpha^N u\right]. \tag{6.12}$$

Observe that $M^*(\phi)$ is finite since, by hypothesis, $\alpha V_{N'+1}(\phi) - h_{N'} < i_{N'}^*(\phi)$. Then, following the same idea used in the proof of Theorem 6.8, for all $\phi' \in \Phi^{N'+M^*(\phi)}(\phi)$, $i_{N'}^*(\phi') = i_{N'}^*(\phi)$. Thus, for each $\phi' \in \Phi^{N'+M^*(\phi)}(\phi)$, it is strictly optimal to accept any offer in period N'. Since problem $\phi^{N'}$ also optimally accepts any offer in period N', the optimal initial threshold for $\phi^{N'}$ is also optimal for any $\phi' \in \Phi^{N'+M^*(\phi)}(\phi)$, and ϕ is well-posed.

Now suppose that ϕ is well-posed but not coalescent. Then there exists a period N^* and an initial threshold i_1^* such that i_1^* is also an optimal initial threshold for any $\phi' \in \Phi^{N^*}(\phi)$. Since ϕ is not coalescent, it must have multiple optimal initial thresholds, and in no period is it strictly optimal to accept every offer. Due to this latter fact, for any optimal path $\{t_n\}$ in the decision tree of the deterministic formulation (and there must be at least two such paths since there are multiple optimal initial thresholds), $\beta(t_n) > 0$ for all n, i.e., the probability of survival to an arbitrary future state is always strictly positive.

It must also be true that for all n, $F_n\left(\lfloor \alpha u - h_n \rfloor\right) > 0$; i.e. no period guarantees an unbeatable offer. Consider the two forecasts ${}^n\bar{\phi}, {}^n\underline{\phi} \in \Phi^n(\phi)$, with the distinction that ${}^N\bar{\phi}$ satisfies $q_{N+1}(u) = 1$ and for ${}^N\underline{\phi}$, $q_n(l) = 1$ for all $n > N$. Also, let the holding penalty in period n of both ${}^N\bar{\phi}$ and ${}^N\underline{\phi}$ be identical to that under ϕ, in all periods $n > N$, so that ${}^N\bar{\phi}$ and ${}^N\underline{\phi}$ constitute best- and worst-case forecasts, respectively, among all those in agreement with ϕ through period N. Then, for each N, the maximum expected rewards beginning period $N+1$ under forecast ${}^N\bar{\phi}$ are strictly greater than those under forecast ϕ, and the maximum expected rewards beginning period $N+1$ under forecast ${}^N\underline{\phi}$ are strictly less than those under forecast ϕ. We conclude that

$$\alpha V_2({}^N\underline{\phi}) < \alpha V_2(\phi) < \alpha V_2({}^N\bar{\phi}),$$

and since $i_1^*(\phi) = \alpha V_2(\phi) - h_1$, it is true that $i_1^*({}^N\underline{\phi}) < i_1^*({}^N\bar{\phi})$, so that there is no optimal initial threshold that is optimal for all $\phi' \in \Phi^N(\phi)$. Having chosen N arbitrarily, we see that ϕ is not well-posed. We have achieved a contradiction and this completes the proof. $\square$

6.5 Solving the Infinite Horizon Asset Selling Problem

In [16], §6, the authors give a solution procedure for a general class of MDPs, including the asset selling problem formulated in this chapter, which will find a forecast horizon any time one exists for a given forecast with respect to its class of potential forecasts. The solution procedure suggested in 6.5.1, while differing in nomenclature, operates in much the same manner. It is also shown that the forecast horizon calculated by the solution procedure is, in fact, the minimum forecast horizon. Section 6.5.2 discusses the average forecast horizon, as well as factors that contribute to the variability in the forecast horizon. Section 6.5.3 analyzes the performance of the solution procedure in a series of randomly generated asset selling forecasts.

6.5.1 Solution Procedure

As in [16], §6, we first define the finite horizon best- and worst-case problems $\bar{\phi}^N$ and $\underline{\phi}^N$ (similar to ${}^N\bar{\phi}$ and ${}^N\underline{\phi}$ in the proof of Theorem 6.16 of [16]). Let $\bar{\phi}^N$ be equivalent to the finite horizon problem ϕ^N, with the only difference being that the reward associated with rejecting an offer in period N is $\alpha u - h_N$ rather than simply $-h_N$. The quantity $\alpha u - h_N$ represents the maximum possible rewards beginning period N, discounted to period N, conditional upon rejecting the offers in periods 1 through N, over all $\phi' \in \Phi^N(\phi)$. On the other hand, let $\underline{\phi}^N = \phi^N$, since the maximum total expected reward of ϕ is at least that of ϕ^N. Then $\langle\bar{\phi}^N\rangle$ and $\langle\underline{\phi}^N\rangle$ represent the problems $\left\langle\underline{\phi}^N\right\rangle$ and $\left\langle\bar{\phi}^N\right\rangle$, respectively, in the Iterate step of the Forecast Horizon Algorithm in [16].

The Forecast Horizon Algorithm looks for policy convergence in the optimal policies to the increasingly longer best- and worst-case problems. Because we know that optimal policies to the asset selling problem take the form of thresholds, we can simplify the Terminate phase of the algorithm. To do this, we need only to keep track of $V_2(\bar{\phi}^N)$ and $V_2(\underline{\phi}^N)$ since the optimal initial thresholds to $\bar{\phi}^N$ and $\underline{\phi}^N$ depend solely upon these quantities and on their first-period parameters (ϕ_1).

Algorithm 6.1.

1. *Set* $N = 2$.

2. *Solve* $\langle \bar{\phi}^N \rangle$ *and* $\langle \underline{\phi}^N \rangle$ *to get* $V_2(\bar{\phi}^N)$ *and* $V_2(\underline{\phi}^N)$.

3. *If one of the following is satisfied, stop and go to step 4. Otherwise, increment* N *by 1 and return to step 2.*

   (a) *For some* $l \leq k \leq u-1$, $k < \alpha V_2(\underline{\phi}^N) - h_1 < k+1$ *and* $k < \alpha V_2(\bar{\phi}^N) - h_1 \leq k+1$.

   (b) *For some* $l \leq k \leq u-1$, $k \leq \alpha V_2(\underline{\phi}^N) - h_1 < k+1$ *and* $k < \alpha V_2(\bar{\phi}^N) - h_1 < k+1$.

4. *Return* $i_1^*(\phi) = i_1^*(\phi^N)$.

Theorem 6.17. *Algorithm 6.1 will stop in finite time if and only if* ϕ *is well-posed.*

Proof. Suppose first that ϕ is well-posed. Then by Theorem 6.16, either $k < \alpha V_2(\phi) - h_1 < k+1$ for some $l \leq k \leq u-1$ or $F_{N'}(i^*_{N'}(\phi) - 1) = 0$ for some $N' \geq 1$. If the first condition is true, then since $V_2(\underline{\phi}^N) \leq V_2(\phi) \leq V_2(\bar{\phi}^N)$ for all N, and $V_2(\underline{\phi}^N)$ and $V_2(\bar{\phi}^N)$ converge monotonically to $V_2(\phi)$ from below and above, respectively, there necessarily exists some period N^* such that either

$$k < \alpha V_2(\underline{\phi}^{N^*}) - h_1 \leq \alpha V_2(\phi) - h_1 \leq \alpha V_2(\bar{\phi}^{N^*}) - h_1 \leq k+1$$

or

$$k \leq \alpha V_2(\underline{H}^{N^*}, N^*) - h_1 \leq \alpha V_2(\phi) - h_1 \leq \alpha V_2(\bar{\phi}^{N^*}) - h_1 < k+1.$$

Thus, either condition (a) or condition (b) of Algorithm 6.1 will be satisfied at period N^* and the algorithm will stop.

If the second condition above is true, i.e. $F_{N'}(i^*_{N'}(\phi) - 1) = 0$ for some $N' \geq 1$, then $V_2(\phi^N) = V_2(\phi^{N'})$ for all $N \geq N'$. This is true because $i^*_{N'}(\phi^{N'}) = l$ and $i^*_{N'}(\phi^N)$ is nondecreasing in N, so that there is always zero probability of optimally rejecting the offer in period N' under problem $\underline{\phi}^N$. If $k < \alpha V_2(\phi^{N'}) - h_1 < k + 1$ for some $l \leq k \leq u - 1$, then there exists some period $M^*(\phi) > N'$ such that $\alpha V_2(\phi^{N'}) + \alpha^{M^*(\phi)} u - h_1 < k + 1$, and thus both condition (a) and condition (b) will be satisfied at period $M^*(\phi)$ and the algorithm will stop. On the other hand, if $k = V_2(\phi^{N'}) - h_1$ for some $l \leq k \leq u - 1$, then after $M^{**}(\phi) = N' + \log_\alpha \lceil \frac{1}{u} \rceil$ periods, it is true that $\alpha V_2(\phi^{N'}) + \alpha^{M^{**}(\phi)} u - h_1 < k + 1$, and condition (b) of the algorithm will cause it to stop. We see that any well-posed problem will be detected by the algorithm.

Now suppose that the algorithm stops at some time M. Our investigation will proceed casewise, depending on whether condition (a) or condition (b) caused the algorithm to stop.

Case 1: Condition (a). For some $l \leq k \leq u - 1$, $k < \alpha V_2(\underline{\phi}^M) - h_1 < k + 1$ and $k < \alpha V_2(\bar{\phi}^M) - h_1 \leq k + 1$. If the last inequality is strict, i.e. $\alpha V_2(\bar{\phi}^M) - h_1 < k + 1$, then $k < \alpha V_2(\underline{\phi}^M) - h_1 \leq \alpha V_2(\phi) - h_1 \leq \alpha V_2(\bar{\phi}^M) - h_1 < k + 1$, so that the infinite horizon optimal initial threshold is unique and ϕ is well-posd. If the inequality is not strict, it is still possible that the infinite horizon optimal initial threshold is unique. The other possibility is that $\alpha V_2(\phi) - h_1 \leq \alpha V_2(\bar{\phi}^M) - h_1 = k + 1$. But this implies that an offer of value u arrives with probability one in period $M + 1$, so that offer is always optimally accepted with probability one and ϕ is well-posed.

Case 2: Condition (b). For some $l \leq k \leq u - 1$, $k \leq \alpha V_2(\underline{\phi}^M) - h_1 < k + 1$ and $k < \alpha V_2(\bar{\phi}^M) - h_1 < k + 1$. Similar to Case 1, if the *first* inequality is strict, then the infinite horizon optimal initial threshold is unique and ϕ is well-posed. If, on the other hand, $k = \alpha V_2(\underline{\phi}^M) - h_1$ and $\alpha V_2(\bar{\phi}^M) - h_1 < k+1$, then $k \leq \alpha V_2(\phi) - h_1 < k+1$, so that $k+1$ is an infinite horizon optimal initial threshold for any $\phi' \in \Phi^M(\phi)$, and ϕ is well-posed. □

We have just shown that Algorithm 6.1 will solve any well-posed problem. In other words, if a forecast horizon exists, Algorithm 6.1 will detect it in finite time. Keeping in mind that a primary consideration in solving the infinite horizon asset selling problem is to minimize the required number of periods' worth of forecast, Algorithm 6.1 would ideally stop at the *minimum* forecast horizon, which is defined as follows.

Definition 6.5. *Period $N^{**}(\phi)$ is called the* minimum forecast horizon *for problem ϕ if it is a forecast horizon for ϕ and, for any $n < N^{**}(\phi)$, there exists some $\phi' \in \Phi^n(\phi)$ such that $i_1^*(\phi') > i_1^*(\phi)$.*

Theorem 6.18. *For any well-posed $\phi \in \Phi$, Algorithm 6.1 will stop after exactly $N^{**}(\phi)$ periods.*

Proof. For any $n < N^{**}(\phi)$, it must be that $i_1^*(\bar{\phi}^n) > i_1^*(\underline{\phi}^n)$ since $\langle \underline{\phi}^n \rangle$ and $\langle \bar{\phi}^n \rangle$ must possess the minimum and maximum optimal initial thresholds over all $\phi' \in \Phi^n(\phi)$. But under this condition, Algorithm 6.1 will not stop at period n.

On the other hand, since $N^{**}(\phi)$ is a forecast horizon for ϕ, for every $\phi' \in \Phi^{N^{**}(\phi)}(\phi)$, $i_1^*(\phi)$ must be an optimal initial threshold for ϕ'. The forecasts that achieve the minimum and maximum values of the $V_2(\cdot)$ function in this set are $\underline{\phi}^{N^{**}(\phi)}$ and $\bar{\phi}^{N^{**}(\phi)}$. But since these both have $i_1^*(\phi)$ as a common optimal initial threshold, Algorithm 6.1 will surely stop at period $N^{**}(\phi)$. □

6.5.2 Algorithm Performance

We are interested in predicting the performance of Algorithm 6.1 with respect to the problem class Φ. In particular, to avoid placing any restrictions on the offer distributions or holding penalties, we will analyze the algorithm's performance as a function of u. Implicitly, we will assume that $l = 0$, though the results here will hold no matter the value of l. Indeed, it is the magnitude of u which will be seen to primarily dictate the algorithm's performance.

We begin by introducing the problem-dependent quantities

$$d(\phi) = i_1^*(\phi) - (\alpha V_2(\phi) - h_1) \tag{6.13}$$

$$g(\phi) = \min\left[N : i_1^*(\phi) - 1 < \alpha V_2(\bar{\phi}^N) - h_1 \leq i_1^*(\phi)\right] \tag{6.14}$$

$$h(\phi) = \min\left[N : i_1^*(\phi) - 1 < \alpha V_2(\phi^N) - h_1 \leq i_1^*(\phi)\right] \tag{6.15}$$

For ease of exposition, we will momentarily assume that $\alpha V_2(\phi) - h_1$ is not integer-valued. Under this assumption, $N^{**}(\phi) = \max[g(\phi), h(\phi)]$. Thus, by Theorem 6.18, Algorithm 6.1 will solve problem ϕ after $\max[g(\phi), h(\phi)]$ periods. In general, because $g(\phi)$ is explicitly dependent upon the magnitude of u due to the maximum terminal reward applied in $\bar{\phi}^N$, we can expect that $g(\phi) > h(\phi)$. Of course, in a minority of problems, this may not be true. In those problems, $\alpha V_2(\phi^N) - h_1$ approaches $i_1^*(\phi) - 1$ from below more slowly than $\alpha V_2(\bar{\phi}^N) - h_1$ approaches $i_1^*(\phi)$ from above. It is not immediately clear what role u plays, if any, in delaying the convergence of $\alpha V_2(\phi^N) - h_1$ (this is difficult to investigate without imposing some structure upon the offer distributions). However, the following is true for the majority of problems.

Proposition 6.19. *For all ϕ such that $g(\phi) > h(\phi)$, the average minimum forecast horizon is proportional to* $\log(u)$.

Proof. Since $V_2(\phi^N) \leq V_2(\phi) \leq V_2(\bar{\phi}^N)$ for all N, and by the definition of $d(\phi)$, we see that for ϕ such that $g(\phi) > h(\phi)$,

$$\begin{aligned} g(\phi) &\leq \min[N : \alpha V_2(\bar{\phi}^N) - \alpha V_2(\phi^N) \leq d(\phi)] \\ &\leq \min[N : \alpha^N u \leq d(\phi)] \end{aligned}$$

Taking the log (with base > 1) of both sides, we see that

$$\begin{aligned} N\log(\alpha) + \log(u) &\leq \log(d(\phi)) \\ N &\leq \frac{\log(d(\phi)) - \log(u)}{\log(\alpha)} \end{aligned}$$

Now, $d(\phi)$ is an unknown constant that varies as a function of ϕ, for fixed u, between zero and one. Thus, the *average* value of $\frac{\log(d(\phi))}{\log(\alpha)}$ is some positive constant, and since $\log(\alpha) < 0$, $\frac{-\log(u)}{\log(\alpha)}$ is also a positive constant, dependent upon u. We conclude that the average minimum forecast horizon, for all problems ϕ in which $g(\phi) > h(\phi)$, is proportional to $\log(u)$. $\square$

6.5.3 Numerical Studies

Although we know that for the majority of problems, the minimum forecast horizon is, on average, proportional to $\log u$, it is not clear what the magnitudes of forecast horizons are. Also, since we have seen that the forecast horizon for well-posed problem $\phi \in \Phi$ is dependent upon the unknown quantities $g(\phi)$ and $h(\phi)$, we cannot yet say anything about the *distribution* of the forecast horizon for a given u, where ϕ is chosen randomly. To better approximate the distribution of the forecast horizon, we ran a series of numerical experiments where asset selling forecasts were randomly generated and the minimum forecast horizon was computed for each, for a variety of values of u: 10, 30, 50, 70, 90, 100, 200, 400, 600, 800, and 1000.

For simplicity, the asset selling forecasts consisted solely of the offer distributions, i.e., there were no holding penalties. The offer distributions were discrete, triangular distributions on the interval $[0, u]$. The peak of the triangle shifted randomly from period to period, with a 50% chance of increasing and a 50% chance of decreasing, and was permitted to shift at most 80% of the distance to u (if increasing) or to 0 (if decreasing).

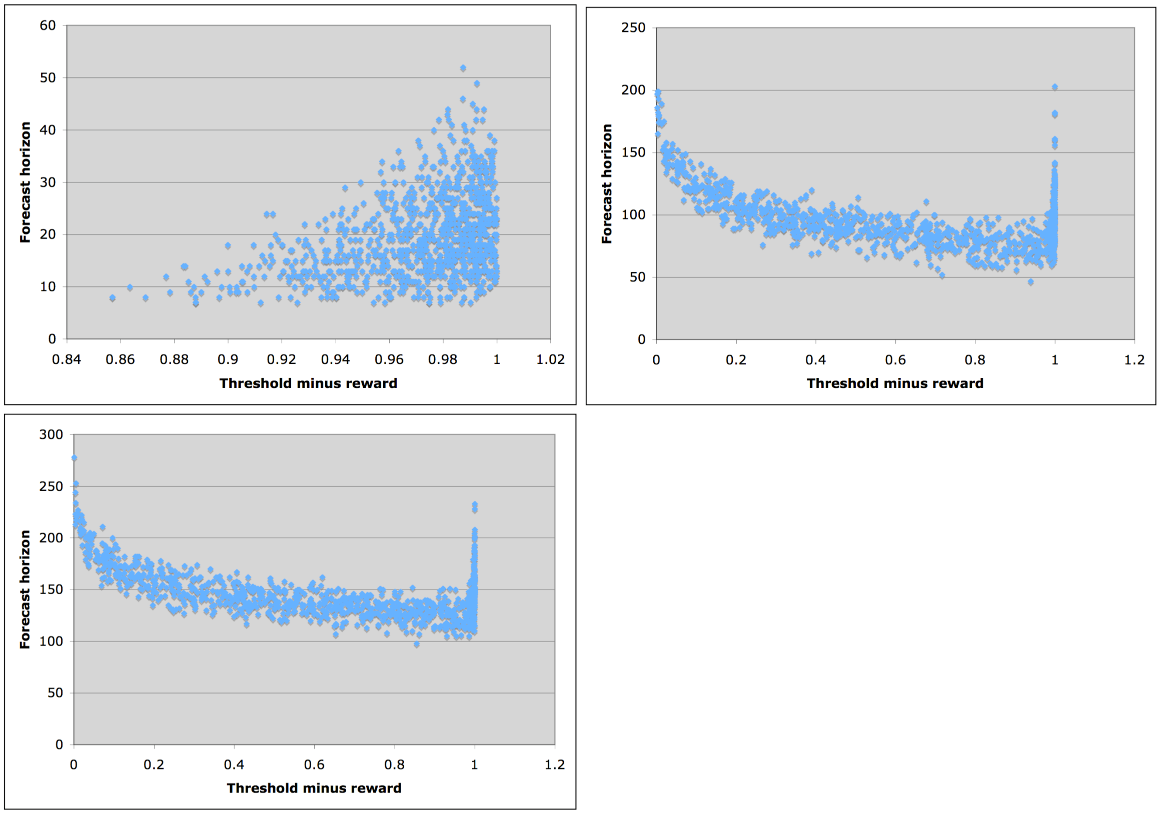

Figure 6.2: Plots of Forecast Horizon Versus $i_1^*(\phi) - \alpha V_2(\phi^{N^{**}(\phi)})$

For each value of u, 1000 streams of offer distributions were generated, and each was stopped as soon as the stopping rule in the algorithm was satisfied. Thus, for each value of u, the study found 1000 forecast horizons, and also calculated (for each ϕ) the quantity $i_1^*(\phi) - \alpha V_2(\phi^{N^{**}(\phi)})$. Plots of $N^{**}(\phi)$ versus $i_1^*(\phi) - \alpha V_2(\phi^{N^{**}(\phi)})$ are shown in Figure 6.2 for $u = 10$ (top left), $u = 100$ (top right), and $u = 1000$ (bottom).

Confirming Proposition 6.19, the average forecast horizon increased logarithmically in u. The results were also encouraging in that the maximum forecast horizon (out of 1000 trials) and the standard deviation also increased logarithmically in u, and the standard deviation even appeared to asymptotically approach a constant. These trends can be seen in Figure 6.3, where the maximum is represented by triangles, the average by diamonds, and the standard deviation by squares, for 1000 randomly

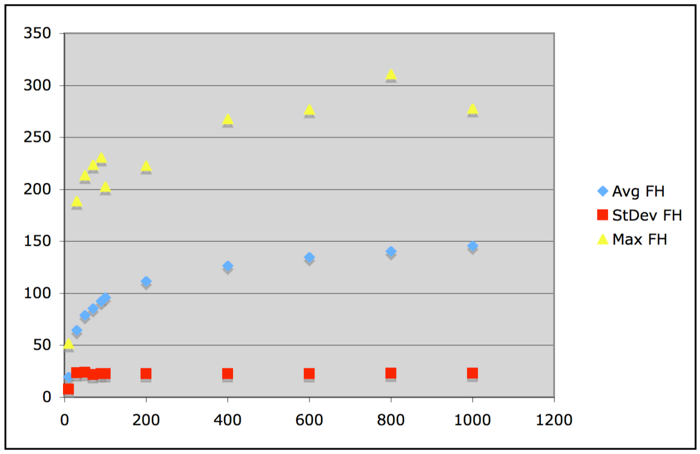

Figure 6.3: Plots of Performance Statistics for Algorithm 6.1

generated problems for each value of u. Values of u are on the x-axis.

We also see that as $i_1^*(\phi) - \alpha V_2(\phi^{N^{**}(\phi)})$ approaches 0 (implying that $g(\phi) \ll h(\phi)$ probably) or 1 ($h(\phi) \gg g(\phi)$ probably), forecast horizons are longer in general. It appears, in fact, that for larger values of u, there are relatively fewer forecasts which result in $d(\phi)$ is close to zero, so that it is more likely that a given problem will have a minimum forecast horizon proportional to $\log(u)$. Lastly, it should be noted that none of the 11,000 forecasts (as well as several thousand additional trial runs) failed to be solved, which supports the claim in Theorem 6.11.

6.6 Conclusions

Optimal solutions to the discrete-valued, infinite horizon asset selling problem are characterized by a threshold acceptance policy in each period. The thresholds in each period of the finite horizon truncations of any asset selling problem are monotonically increasing in the length of the truncation. Using a rigorous definition of solvability called well-posedness, we showed that any problem whose infinite horizon optimal initial threshold is unique is well-posed. Practically all problem instances satisfy the uniqueness condition, so that practically all problems are well-posed.

Under a deterministic reformulation of the asset selling problem, it can be seen that a forecast is well-posed if and only if either the optimal initial threshold is unique or it is optimal to accept any offer in some period. Although it may not be possible to bound a priori the number of periods worth of forecast required to solve for an optimal initial threshold, we provided a solution procedure which will detect in finite time an optimal initial threshold for any well-posed problem. Moreover, the solution procedure can be applied recursively to find in finite time each of the optimal thresholds for a given infinite horizon forecast.

Chapter 7

Conclusions

This thesis approached the topic of solving discrete, nonstationary, infinite horizon optimization problems - problems that are in general unsolvable unless (i) the decision maker is content to make at most a few decisions at a time, and (ii) the dynamic programming formulation of the problem possesses some key characteristics. We have assumed the first condition and have given a characteristic which is both necessary and sufficient in support of the second condition. This characteristic is called *coalescence*, and it roughly implies that there exists an optimal sequence of states for the given problem (called a source path) such that any other optimal sequence of states can be optimally reached from some leading segment of the source path. Naturally, when the problem possesses a unique optimal solution, it is coalescent, but we have expressly avoided assuming uniqueness of any optimal decisions. Indeed, coalescence can be satisfied even in the presence of multiple optima.

Coalescence is shown to be necessary and sufficient for finite solvability of an infinite horizon nonstationary version of the following problems: uncapacitated production planning without backlogging, single equipment replacement, and capacitated inventory planning with revenues and lost sales. The common problem characteristics that allowed coalescence to be equivalent to solvable were

- the possibility that each feasible (alternatively, potentially optimal) state in each period for each problem is strictly optimal for some problem in agreement with the given problem through the given period;

- the existence of a partial feasible decision sequence to connect any pair of feasible (potentially optimal) states for any problem within a uniformly bounded number of periods;

- the finiteness of the feasible (potentially optimal) states in each period for each problem.

Using these common problem properties as assumptions, we showed that coalescence is equivalent to finite solvability for more generic classes of deterministic optimization problems. Interestingly, the same relationship holds for a simple stochastic optimization problem - the asset selling problem. Moreover, coalescence has a straightforward meaning in the context of that problem - either the optimal solution is unique (almost surely met) or an unbeatable offer is forecast to arrive with certainty in some period (trivializing the problem). Thus, the solvability of an infinite horizon nonstationary asset selling problem is all but assured.

Future work could include establishing that other well-known or significant applications fall within the generic classes of optimization problems defined in Chapter 5. In general, the most difficult part of making these establishments is likely to be showing that the first Type I or Type II Assumption - that each potentially optimal (or finite) state in each period for each problem is the unique optimal state in that period for some other problem. Another major complication that has posed difficulties in the development of stopping rule algorithms for the dynamic lot size model (among other applications) in the literature is determining the *minimal* sets of potentially optimal states in each period for each problem. Discussions on this topic can be found in [32], [23], and [24].

Another area of fruitful work could be determining a more general class of stochastic problems for which coalescence is equivalent to finite solvability. A difficulty here becomes the *state space*, as a deterministic reformulation may be required in order to describe the state of the system after invoking a strategy over a finite number of periods.

Appendices

Appendix A

An Infinite Horizon Concave Cost Production Planning Problem

A.1 Justification of Zero-Inventory Ordering Policy (Equation (2.2))

If there were an optimal production plan for the infinite horizon problem that did not satisfy (2.2), a unit of inventory would be carried over either an infinite or a finite interval in which there is at least one production point. If the former case (carried over an infinite horizon) is true, then clearly the extra unit of inventory is never used to satisfy demand. By Assumption 2.1, total cost remains the same or decreases by dropping the extra unit of inventory and hence not paying production or holding costs for it. In the latter case, there is some interval $\{a, a+1, \ldots, b\}$ such that a and b are both production points and $c \in \{a+1, a+2, \ldots, b-1\}$ is also a production point, but i_c, the inventory level beginning period c, is strictly positive. By the Principle of Optimality, with optimal production points occurring in periods a and b, the decisions made in periods $\{a, a+1, \ldots, b-1\}$ should be optimal for the subproblem on that interval. However, as noted in Denardo [19], any finite horizon concave cost production planning problem has at least one optimal solution satisfying the *zero-inventory ordering policy.* Thus, there exists an optimal solution to the subproblem (and hence the infinite horizon problem) such that $i_c = 0$.

A.2 A DLSM Instance With Non-Monotonic Last Production Points

Let p be a stationary problem with $d_n = 2$ for all n, $h_n(x) = \frac{x}{4}$ for all n, and

$$c_n(x) = \begin{cases} 1+x, & n=1 \\ 5, & n=2 \\ x, & n \geq 3 \end{cases} .$$

Then for the first six periods, for any $.77 \leq \alpha \leq 1.0$, the optimal finite horizon production plans are unique, and the last optimal production point is non-monotonic. As can be seen in the table below, the optimal last production point actually decreases from period 3 to period 2 when going to the 5-period problem from the 4-period problem. Moreover, note that all of the stopping rules based on monotonicity in the last production point given by the authors referenced in this chapter would have stopped after three periods because an optimal production point occurred in the terminal period of that problem. We see here that those stopping rules do not hold when costs are general concave, or even when the DLSM is given in its general form.

Period	**Optimal setups**
1	1
2	1
3	1,3
4	1,3
5	1,2
6	1,2

A.3 Proof of Lemma 2.1

Lemma 2.1. *Let $p \in \mathcal{P}$ and choose any subsequence of integers $\{n_j\}$ with associated forecasts $\{p(n_j)\}, p(n_j) \in \mathcal{P}^{n_j}(p)\ \forall n$. Then there exist a further subsequence $\{n_{j_k}\} \subseteq \{n_j\}$ with associated optimal production plans $\{\tilde{x}^*(n_{j_k})\}, \tilde{x}^*(n_{j_k}) \in \mathcal{X}^*(p(n_{j_k}))\ \forall n$ and some $\tilde{x}^* \in \mathcal{X}^*(p)$ such that $\tilde{x}^{*(n_{j_k})} \to \tilde{x}^*$, where convergence is componentwise.*

Proof. It will suffice to show the following:

1. Some subsequence $\{p(n_{j_k})\}$ of $\{p(n_j)\}$ must have optimal production plans $\{\tilde{x}^*(n_{j_k})\}$ converge to a limit point $\tilde{x}^*$ which is feasible for p.

2. $\tilde{x}^*$ is optimal for p.

We now proceed with the proof, claim by claim.

1. Let $\{\tilde{x}^*(n_j)\}$ be any subsequence of optimal production plans corresponding to the forecast subsequence $\{p(n_j)\}$. By Assumption 2.2 and the definition of $\bar{d}$, there are at most $L\bar{d}+1$ potentially optimal production decisions in each period when solving any $p' \in \mathcal{P}$. Thus, each period's potentially optimal decision space $Y = \{0, 1, 2, \ldots, L\bar{d}\}$ is finite and therefore compact. The Tychonoff Theorem gives that $\prod_{n=1}^{\infty} Y$ is compact in the product topology of componentwise convergence. Thus, we can construct subsequences $\{n_{j_k}\} \subseteq \{n_j\}$ and $\{p(n_{j_k})\} \subseteq \{p(n_j)\}$ of periods and associated forecasts, respectively, such that the optimal production plans $\{\tilde{x}^*(n_{j_k})\}$ over the forecasts must converge to a limit point $\tilde{x}^*$, as claimed. To see that $\tilde{x}^*$ is feasible for p, we note that the first n_{j_k} decisions of $\tilde{x}^*(n_{j_k})$ are feasible for p, for each k.

2. Now we will show that $\tilde{x}^* \in \mathcal{X}^*(p)$. Let y^* be an infinite horizon optimal production plan for p, with corresponding inventory plan i^*. Let $\tilde{i}^*$ and $\tilde{i}^*(n_{j_k})$ be the optimal inventory plans resulting from the optimal production plans $\tilde{x}^*$ and $\tilde{x}^*(n_{j_k})$, respectively. Also, let $d(n_{j_k})$ be the demand forecast belonging to $p(n_{j_k})$. We first make the observation that for any k, the production plan

$$y(n_{j_k}) \equiv (y_1^*, \ldots, y_{n_{j_k}}^*, d_{n_{j_k}+1}(n_{j_k}), d_{n_{j_k}+2}(n_{j_k}), \ldots)$$

is feasible (although not necessarily optimal) for $p(n_{j_k})$. Then, for all k,

$$V(p(n_{j_k}), y(n_{j_k})) \geq V(p(n_{j_k}), \tilde{x}^*(n_{j_k})). \tag{A.1}$$

We are now in a position to show that $\tilde{x}^*$ is infinite horizon optimal for p. Since $p(n_{j_k}) \to p$ and $y(n_{j_k}) \to y^*$, finiteness and continuity of $V(\cdot, \cdot)$ gives us that

$V(p(n_{j_k}), y(n_{j_k})) \to V(p, y^*)$. A similar argument shows that $V(p(n_{j_k}), \tilde{x}^*(n_{j_k})) \to V(p, \tilde{x}^*)$. Thus, by (A.1), $V(p, \tilde{x}^*) \leq V(p, y^*)$. But $\tilde{x}^*$ is feasible for p and y^* is optimal for p, so that $\tilde{x}^*$ must also be optimal for p, as desired.

□

Appendix B

An Infinite Horizon Equipment Replacement Problem

B.1 Proof of Lemma 3.4

Lemma 3.4. *Consider a forecast c and a sequence of forecasts $\{c^{(n)}\}_{n=1}^{\infty}$ where $c^{(n)} \in \mathcal{C}^n(c)\ \forall n$. Then there is some subsequence $\{n_k\}_{k=1}^{\infty}$, a corresponding sequence of optimal replacement strategies $\{\hat{x}^{*(k)} \in \mathcal{X}^*(c^{(n_k)})\}_{k=1}^{\infty}$, and $x^* \in \mathcal{X}^*(c)$ such that $\hat{x}^{*(k)} \to x^*$.*

Proof. It will suffice to show the following:

1. Some subsequence $\{c(n_{j_k})\}$ of $\{c(n_j)\}$ must have optimal replacement strategies $\{\tilde{x}^*(n_{j_k})\}$ converging to a limit point $\tilde{x}^*$ which is feasible for c.

2. $\tilde{x}^*$ is optimal for c.

We now proceed with the proof, claim by claim.

1. Let $\{\tilde{x}^*(n_j)\}$ be any subsequence of optimal replacement strategies corresponding to the forecast subsequence $\{c(n_j)\}$. By definition of the maximum feasible lifetime L, there are at most L potentially optimal replacement decisions in each period when solving any $c' \in \mathcal{C}$. Thus, each period's potentially optimal decision space $Y = \{0, 1, 2, \dots, L\}$ is finite and therefore compact. The Tychonoff Theorem gives that $\prod_{n=1}^{\infty} Y$ is compact in the product topology of componentwise convergence. Thus, we can construct subsequences $\{n_{j_k}\} \subseteq \{n_j\}$ and

$\{c(n_{j_k})\} \subseteq \{c(n_j)\}$ of periods and associated forecasts, respectively, such that the optimal replacement strategies $\{\tilde{x}^*(n_{j_k})\}$ over the forecasts must converge to a limit point $\tilde{x}^*$, as claimed. To see that $\tilde{x}^*$ is feasible for c, we note that the first n_{j_k} decisions of $\tilde{x}^*(n_{j_k})$ are feasible for c, for each k, and a new piece of equipment can be purchased for any lifetime up to L at any point in the future.

2. Now we will show that $\tilde{x}^* \in \mathcal{X}^*(c)$. Let y^* be an infinite horizon optimal replacement strategy for c. Let $b(n_{j_k})$ be the business requirement forecast belonging to $c(n_{j_k})$. We first make the observation that for any k, the replacement strategy

$$y(n_{j_k}) \equiv (y^*_1, \ldots, y^*_{n_{j_k}}, b(n_{j_k})_{n_{j_k}+1}, b(n_{j_k})_{n_{j_k}+2}, \ldots)$$

is feasible (although not necessarily optimal) for $c(n_{j_k})$. Then, for all k,

$$V(p(n_{j_k}), y(n_{j_k})) \geq V(p(n_{j_k}), \tilde{x}^*(n_{j_k})). \tag{B.1}$$

We are now in a position to show that $\tilde{x}^*$ is infinite horizon optimal for c. Since $c(n_{j_k}) \rightarrow c$ and $y(n_{j_k}) \rightarrow y^*$, finiteness and continuity of $V(\cdot,\cdot)$ gives us that $V(c(n_{j_k}), y(n_{j_k})) \rightarrow V(c, y^*)$. A similar argument shows that $V(c(n_{j_k}), \tilde{x}^*(n_{j_k})) \rightarrow V(c, \tilde{x}^*)$. Thus, by (B.1), $V(c, \tilde{x}^*) \leq V(c, y^*)$. But $\tilde{x}^*$ is feasible for c and y^* is optimal for c, so that $\tilde{x}^*$ must also be optimal for c, as desired.

$\square$

Appendix C

An Infinite Horizon Capacitated Inventory Planning Problem

C.1 Proof of Lemma 4.4

Lemma 4.4. *Let $\phi \in \Phi$ and choose any subsequence of integers $\{n_j\}$ with associated forecasts $\{\phi(n_j)\}, \phi(n_j) \in \Phi^{n_j}(\phi)$ for all n. Then there exist a further subsequence $\{n_{j_k}\} \subseteq \{n_j\}$ with associated optimal fulfillment strategies $\{\tilde{x}^*(n_{j_k})\}$ optimal for $\phi(n_{j_k})$ for each k and some optimal fulfillment strategy $\tilde{x}^*$ for ϕ such that $\tilde{x}^{*(n_{j_k})} \to \tilde{x}^*$, where convergence is componentwise.*

Proof. It will suffice to show the following:

1. Some subsequence $\{\phi(n_{j_k})\}$ of $\{\phi(n_j)\}$ must have optimal fulfillment strategies $\{\tilde{x}^*(n_{j_k})\}$ converge to a limit point $\tilde{x}^*$ which is feasible for ϕ.

2. $\tilde{x}^*$ is optimal for ϕ.

We now proceed with the proof, claim by claim.

1. Let $\{\tilde{x}^*(n_j)\}$ be any subsequence of optimal fulfillment strategies corresponding to the forecast subsequence $\{\phi(n_j)\}$. By the definition of $\bar{S}$, there are at most $(\bar{S}+1)^2$ (the product of the maximum number of feasible fulfillment and inventory quantities) potentially optimal fulfillment decisions in each period when solving any $\phi' \in \Phi$. Thus, each period's potentially optimal decision space Y

is finite and therefore compact. The Tychonoff Theorem gives that $\prod_{n=1}^{\infty} Y$ is compact in the product topology of componentwise convergence. Thus, we can construct subsequences $\{n_{j_k}\} \subseteq \{n_j\}$ and $\{p(n_{j_k})\} \subseteq \{p(n_j)\}$ of periods and associated forecasts, respectively, such that the optimal fulfillment strategies $\{\tilde{x}^*(n_{j_k})\}$ over the forecasts must converge to a limit point $\tilde{x}^*$, as claimed. To see that $\tilde{x}^*$ is feasible for ϕ, we note that the first n_{j_k} decisions of $\tilde{x}^*(n_{j_k})$ are feasible for ϕ, for each k, and by Remarks 4.2 and 4.3, one can feasibly connect any pair of feasible states for any problem in a uniformly bounded amount of time.

2. Now we will show that $\tilde{x}^*$ is optimal for ϕ. Let $x^* = (y^*, z^*)$ be an infinite horizon optimal fulfillment strategy for ϕ. Let $\tilde{y}^*$ and $\tilde{y}^*(n_{j_k})$ be the optimal inventory quantities resulting from the optimal fulfillment strategies $\tilde{x}^*$ and $\tilde{x}^*(n_{j_k})$, respectively. Also, let $d(n_{j_k})$ be the demand forecast belonging to $\phi(n_{j_k})$. We first make the observation that for any k, the fulfillment strategy

$$\begin{aligned} x(n_{j_k}) \quad \equiv \quad & ((y_1^*, z_1^*), \ldots, (y_{n_{j_k}}^*, z_{n_{j_k}}^*), ((y_{n_{j_k}}^* - d_{n_{j_k}+1}(n_{j_k}))^+, \min(d_{n_{j_k}+1}(n_{j_k}), \bar{S}), \\ & ((y_{n_{j_k}}^* - d_{n_{j_k}+1}(n_{j_k}) - d_{n_{j_k}+2}(n_{j_k}))^+, \min(d_{n_{j_k}+2}(n_{j_k}), \bar{S}), \ldots) \end{aligned}$$

is feasible (although not necessarily optimal) for $\phi(n_{j_k})$. Then, for all k, the total profit of using fulfillment strategy $x(n_{j_k})$ under forecast $\phi(n_{j_k})$ is no more than that of using fulfillment strategy $\tilde{x}^*(n_{j_k})$ for forecast $\phi(n_{j_k})$. We are now in a position to show that $\tilde{x}^*$ is infinite horizon optimal for ϕ. Since $\phi(n_{j_k}) \to \phi$ and $x(n_{j_k}) \to x^*$, by Assumption 4.2, the total profit of fulfillment strategy $x(n_{j_k})$ for forecast $\phi(n_{j_k})$ must converge in k to the total profit of fulfillment strategy x^* for forecast ϕ. A similar argument shows that the total profit of fulfillment strategy $\tilde{x}^*(n_{j_k})$ for forecast $\phi(n_{j_k})$ must converge in k to the total profit of fulfillment strategy $\tilde{x}^*$ for forecast ϕ. Thus, the total profit of fulfillment strategy $\tilde{x}^*$ for forecast ϕ is at least that of fulfillment strategy x^* for foercast ϕ. But $\tilde{x}^*$ is feasible for ϕ and y^* is optimal for ϕ, so that $\tilde{x}^*$ must also be optimal for ϕ, as desired.

$\square$

Appendix D

General Deterministic Infinite Horizon Optimization

D.1 Proof of Lemma 5.1

Lemma 5.1. *Let $p \in \mathcal{P}$ and choose any increasing subsequence of positive integers $\{n_j\}$ with associated forecasts $\{p(n_j)\}$, where $p(n_j) \in \mathcal{P}^{n_j}(p)$ $\forall n$. Then, under the Type I Assumptions, there exist a further subsequence $\{n_{j_k}\} \subseteq \{n_j\}$ with associated optimal decision sequences $\{\tilde{x}^*(n_{j_k})\}, \tilde{x}^*(n_{j_k}) \in \mathcal{X}^*(p(n_{j_k}))$ $\forall n$, and some $\tilde{x}^* \in X^*(p)$ such that $\tilde{x}^{*(n_{j_k})} \to \tilde{x}^*$, where convergence is componentwise.*

Proof. It will suffice to show the following:

1. Some subsequence $\{p(n_{j_k})\}$ of $\{p(n_j)\}$ must have optimal decision sequences $\{\tilde{x}^*(n_{j_k})\}$ converging to a limit point $\tilde{x}^* \in X(p)$.

2. $\tilde{x}^* \in X^*(p)$.

We now proceed with the proof, claim by claim.

1. Let $\{\tilde{x}^*(n_j)\}$ be any subsequence of optimal decision sequences corresponding to the forecast subsequence $\{p(n_j)\}$. By the third Type I Assumption, there are finitely many potentially optimal decisions in each period for any $p' \in \mathcal{P}$, so that $X_N^{**}(p)$ compact. The Tychonoff Theorem gives that $\prod_{N=1}^{\infty} X_N^{**}(p)$ is compact in the product topology of componentwise convergence. Thus, we can construct

subsequences $\{n_{j_k}\} \subseteq \{n_j\}$ and $\{p(n_{j_k})\} \subseteq \{p(n_j)\}$ of periods and associated forecasts, respectively, such that the optimal decision sequences $\{\tilde{x}^*(n_{j_k})\}$ over the forecasts must converge to a limit point $\tilde{x}^*$, as claimed. To see that $\tilde{x}^*$ is feasible for p, we note that the first n_{j_k} decisions of $\tilde{x}^*(n_{j_k})$ are potentially optimal and therefore feasible for p, for each k.

2. Now we will show that $\tilde{x}^* \in X^*(p)$. Let y^* be an infinite horizon optimal decision sequence for p, with corresponding state sequence s^*. Let $\tilde{s}^*$ and $\tilde{s}^*(n_{j_k})$ be the optimal state sequences resulting from the optimal decision sequences $\tilde{x}^*$ and $\tilde{x}^*(n_{j_k})$, respectively. We first make the observation that for any k, by the second Type I Assumption, since $s^*_{n_{j_k}-L}$ and $\tilde{s}^*_{n_{j_k}}(n_{j_k})$ are both potentially optimal for $p(n_{j_k})$ there exists a decision sequence $\tilde{y}(n_{j_k})$ such that

$$\tilde{y}(n_{j_k}) \equiv (y^*_1, \ldots, y^*_{n_{j_k}-L}, \tilde{y}_{n_{j_k}-L+1}(n_{j_k}), \tilde{y}_{n_{j_k}-L+2}(n_{j_k}), \ldots)$$

is feasible (although not necessarily optimal) for $p(n_{j_k})$, and

$$s_{n_{j_k}}(p(n_{j_k}), \tilde{y}(n_{j_k})) = \tilde{s}^*_{n_{j_k}}(n_{j_k}).$$

Then, for all k,

$$V(p(n_{j_k}), \tilde{y}(n_{j_k})) \geq V(p(n_{j_k}), \tilde{x}^*(n_{j_k})). \tag{D.1}$$

We are now in a position to show that $\tilde{x}^*$ is infinite horizon optimal for p. Since $p(n_{j_k}) \to p$ and $y(n_{j_k}) \to y^*$, finiteness and continuity of $V(\cdot,\cdot)$ gives us that $V(p(n_{j_k}), y(n_{j_k})) \to V(p, y^*)$. A similar argument shows that $V(p(n_{j_k}), \tilde{x}^*(n_{j_k})) \to V(p, \tilde{x}^*)$. Thus, by (D.1), $V(p, \tilde{x}^*) \leq V(p, y^*)$. But $\tilde{x}^*$ is feasible for p and y^* is optimal for p, so that $\tilde{x}^*$ must also be optimal for p, as desired.

□

D.2 Justification for Setting $T_n^{**}(p) = \{0, \ldots, L\bar{d}\}$ for Production Planning

For any period n and any production planning problem p, it is straightforward to show that $T_n^{**}(p) = \{0, \ldots, L\bar{d}\}$. Clearly, for the forecast $\underline{p}^n$ in $\mathcal{P}^N(p)$, it is optimal to have zero inventory ending period n. Now suppose that demand in period $n+1$ is $d_{n+1} \in \{1, \ldots, \bar{d}\}$. If, for example, $c_{n+1}(d_{n+1}) = \frac{1}{\alpha}\left(c_n(d_{n+1}) + h_n(d_{n+1}) + \epsilon\right)$ for some $\epsilon > 0$, then $c_{n+1}(d_{n+1}) \geq \frac{1}{\alpha}\left(c_n(x) - c_n(x - d_{n+1}) + h_n(d_{n+1}) + \epsilon\right)$ for any $x \in \{0, \ldots, \bar{d}\}$, so that no matter what the amount produced in period n to satisfy demand in period n, it would be strictly optimal to carry inventory forward to satisfy demand in period $n+1$ as well.

Similarly, let d_{n+m} be the demand in period $n+m$ for $2 \leq m \leq L$. By setting $c_{n+m}(d_{n+m}) = \frac{1}{\alpha^m}\left(c_n(d_{n+m}) + h_n(d_{n+m}) + \epsilon\right)$ for some $\epsilon > 0$, then $c_{n+1}(d_{n+m}) \geq \frac{1}{\alpha}\left(c_n(x) - c_n(x - d_{n+m}) + h_n(d_{n+m}) + \epsilon\right)$ for any $x \in \{0, \ldots, m\bar{d}\}$, so that no matter what the amount produced in period n to satisfy demand in periods n through $n+m-1$, it would be strictly optimal to carry inventory forward to satisfy demand in period $n+m$ as well. In this way, it can be strictly optimal to have any inventory level in the set $\{1, \ldots, L\bar{d}\}$ ending period n.

Since L is a bound on the time between optimal production periods and $\bar{d}$ is a bound on feasible demands, it is never optimal to have more than $L\bar{d}$ units of inventory on hand at the end of period n. Thus, $T_n^{**}(p)$ is precisely the set of inventory levels $\{0, \ldots, L\bar{d}\}$.

To verify that the production cost functions over periods $n+1$ through $n+L$ as generated above are concave, we note the following. If, for $1 \leq m \leq L$, we define $c_{n+m}(d) = \frac{1}{\alpha^m}\left(c_n(x) - c_n(x - d) + h_n(x) + \epsilon\right)$ for some $0 \leq x \leq (L-1)\bar{d}$ and $\epsilon > 0$, then

$$\begin{aligned}
& c_{n+m}(d+2) - c_{n+m}(d+1) \\
& - (c_{n+m}(d+1) - c_{n+m}(d)) \quad = \quad c_{n+m}(d+2) - 2c_{n+m}(d+1) + c_{n+m}(d) \\
& \qquad = \quad c_n(x-d-2) - 2c_n(x-d-1) + c_n(x-d) \\
& \qquad = \quad c_n(x-d-2) - c_n(x-d-1) \\
& \qquad \qquad - (c_n(x-d-1) - c_n(x-d)) \\
& \qquad \leq \quad 0
\end{aligned}$$

since $c_n(\cdot)$ is a concave function. Thus, $c_{n+m}(\cdot)$ as defined is concave.

Appendix E

An Infinite Horizon Asset Selling Problem

E.1 Monotonicity of the Optimal Thresholds for All Periods

Let $f_j(x, \phi^N)$ represent the maximum expected discounted (to period j) rewards beginning period j with incoming state x under problem ϕ^N, and $\delta_j(x, \phi^N) = f_j(x + 1, \phi^N) - f_j(x, \phi^N)$ for $j \leq N$ and $l \leq x \leq u$. Also, let $V_j(\phi^N), j \leq N$ be the maximum expected discounted revenue over periods j through N under problem ϕ^N, given that the offers in periods 1 through $j-1$ were rejected. Using the fact that $V_j(\phi^{N+1}) \geq V_j(\phi^N)$ by non-negativity of the offers, it is straightforward to show that results analogous to Lemmas 6.1 and 6.2, and Lemma 6.3 hold for all optimal thresholds, i.e.

$$\delta_j(x, \phi^{N+1}) \geq \delta_j^*(x, \phi^N) \tag{E.1}$$

for all $1 \leq j \leq N$ and $\phi \in \Phi$,

$$\delta_j(x, \phi^N) \geq 0 \tag{E.2}$$

for $l \leq x \leq i_j^*(\phi^N)$, and

$$\delta_j(i_j^*(\phi^N)\phi^N) \leq 0. \tag{E.3}$$

This gives us all we need to prove the following result.

Theorem E.1 (Threshold Monotonicity). *For all $\phi \in \Phi$, N and $1 \leq k \leq N$, $i_k^*(\phi^{N+1}) \geq i_k^*(\phi^N)$.*

The proof would proceed similarly to that of Theorem 6.5.

E.2 Alternate Proof of Theorem 6.5

An elegant proof of Theorem 6.5 follows by an application of the following theorem from Topkis. That main question in that paper concerns the collection of optimization problems

$$\min g(x,t), x \in S_t, \tag{E.4}$$

where the variable is x and both the constraint set S_t and the objective function $g(x,t)$ depend upon the parameter t, with T being a member of the parameter set T. Let S_t^* be the set of optimal solutions for (E.4) given t in T.

Theorem E.2 (Theorem 6.1, Topkis). *If S is a lattice, T is a poset, $S_t \subseteq S$ is ascending in t on T, $g(x,t)$ is submodular in x on S for each $t \in T$, and $g(x,t)$ has antitone differences in (x,t) on $S \times T$, then S_t^* is ascending in t on T^*.*

Here, for a fixed forecast $\phi \in \Phi$, we wish to minimize $-f(x,\phi^N)$. We let $t = N$ so that $T = \mathbb{Z}_+$. Also, $S_1 = S_2 = \cdots = S = \{l, \ldots, u\}$. Clearly, S is a lattice, T is a poset, and $S_t \subseteq S$ is ascending in t on T. To see that $g(x,t) = -f(x,H,t)$ is submodular in x for each $t \in T$, note that for any $x, y \in S$, $x \leq y$ implies that $x \vee y = y$ and $x \wedge y = x$, and $x > y$ implies that $x \vee y = x$ and $x \wedge y = y$. Then for any t,

$$g(x \vee y, t) + g(x \wedge y, t) = -f(x,H,t) - f(y,H,t) = g(x,t) + g(y,t),$$

so that $g(x,t)$ is submodular in x for each $t \in T$.

To show that $g(x,t)$ has antitone differences in (x,t), by Lemma 6.1, $\delta^*(x,\phi^N) \leq \delta^*(x,\phi^{N+1})$. Expanding the terms and rearranging, we obtain

$$-f(x+1,\phi^{N+1}) + f(x+1,\phi^N) \leq -f(x,\phi^{N+1}) + f(x,\phi^N), \tag{E.5}$$

which can be simplified to $g(x+1,t+1) - g(x+1,t) \leq g(x,t+1) - g(x,t)$. Thus, $g(x,t)$ has antitone differences in (x,t), and we have satisfied all the hypotheses of Topkis' theorem. The result should be interpreted here as meaning that the minimizing arguments (in x) of $-f(x, \phi^N)$ (and consequently, the maximizing arguments of $f(x, \phi^N)$) are monotonically increasing in N, which is precisely what we claimed in Theorem 6.5.

E.3 Deterministic Decision Tree for Asset Selling Problem

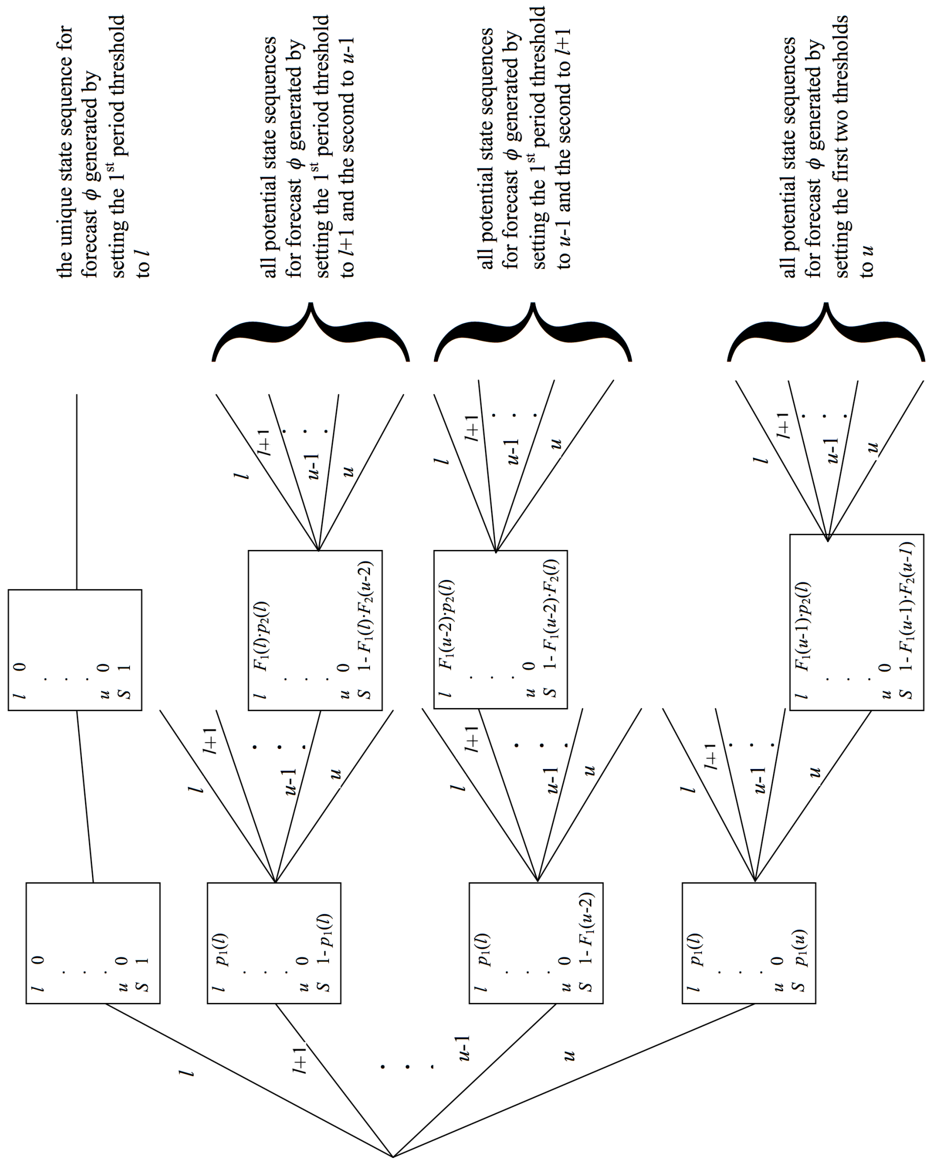

Figure E.1: Deterministic Decision Tree for Asset Selling Problem

Bibliography

[1] Bean, J., and Smith, R. L. Conditions for the Existence of Planning Horizons. *Mathematics of Operations Research* **9**, 391-401 (1984).

[2] Bean, J.C., and Smith, R.L. Conditions for the Discovery of Solution Horizons. *Mathematical Programming* **59**, 215-229 (1993).

[3] Bean, J.C., Lohmann, J.R., and Smith, R.L. A Dynamic Infinite Horizon Replacement Economy Decision Model. *The Engineering Economist* **30**, 99-120 (1985).

[4] Bean, J.C., Lohmann, J.R. and Smith, R.L. Equipment Replacement Under Technological Change. *Naval Research Logistics* **41**, 117-128 (1994).

[5] Bean, J., Smith, R.L. and Yano, C. Forecast Horizons for the Discounted Dynamic Lot-Size Model Allowing Speculative Motive. *Naval Research Logistics* **34**, 761-774 (1987).

[6] Bensoussan, A., Proth, J.-M. and Queyranne, M. A Planning Horizon Algorithm for Deterministic Inventory Management with Piecewise Linear Concave Costs. *Naval Research Logistics* **38**, 729-742 (1991).

[7] Bertsekas, D.P. "Dynamic Programming and Stochastic Control." Academic Press, New York, NY, 1976.

[8] Bhaskaran, S. and Sethi, S. The Dynamic Lot Size Model with Stochastic Demands: a Decision Horizon Study. *INFOR* **26(3)**, 213-224 (1988).

[9] Bylka, S., Sethi, S. and G. Sorger. Minimal Forecast Horizons in Equipment Replacement with Multiple Technologies and General Switching Costs. *Naval Research Logistics* **39**, 487-507 (1992).

[10] Chand, S., Hsu, V.N., and Sethi, S. Forecast, Solution, and Rolling Horizons in Operations Management Problems: A Classified Bibliography. *Manufacturing & Service Operations Management* **4(1)**, 25-43 (2002).

[11] Chand, S. and Morton, T. Minimal Forecast Horizon Procedures for Dynamic Lot Size Models. *Naval Research Logistics Quarterly* **33**, 111-122 (1986).

[12] Chand, S., and Sethi, S. Planning Horizon Procedures for Machine Replacement Models with Several Possible Replacement Models. *Naval Research Logistics Quarterly* **29(3)**, 483-493 (1982).

[13] Chand, S., Sethi, S. and Proth, J.-M. Existence of Forecast Horizons in Undiscounted Discrete-Time Lot Size Models. *Operations Research* **38(5)**, 884-892 (1990).

[14] Chand, S., Sethi, S. and Sorger, G. Forecast Horizons in the Discounted Dynamic Lot Size Model. *Management Science* **38(7)**, 1034-1048 (1992).

[15] Cheevaprawatdomrong, T., and Smith, R.L. A Paradox in Equipment Replacement Under Technological Change. *Operations Research Letters* **31**, 77-82 (2003).

[16] Cheevaprawatdomrong, T., Schochetman, I.E., Smith, R.L. and A. Garcia. Solution and Forecast Horizons for Infinite-Horizon Nonhomogeneous Markov Decision Processes. *Mathematics of Operations Research* **32(1)**, 51-72 (2007).

[17] Chen, H.-D. and Lee, C.-Y. Error Bound for the Dynamic Lot Size Model Allowing Speculative Motive. *IIE Transactions* **27**, 683-688 (1995).

[18] Dawande, S., Gavirneni, S., Naranpanawe, S., and Sethi, S. Forecast Horizons for a Class of Dynamic Lot-Size Problems Under Discrete Future Demand. *Operations Research* **55(4)**, 688-702 (2007).

[19] Denardo, E.V. *Dynamic Programming: Theory and Applications*. Dover Publications, Mineola, NY (2003).

[20] Derman, C., Lieberman, G.J. and Ross, S.M. Adaptive Disposal Models. *Naval Research Logistics Quarterly*, **26(1)**, 33-40 (1979).

[21] Eppen, G. D., Gould, F. J., and Pashigian, B. P. Extensions of the Planning Horizon Theorem in the Dynamic Lot Size Model. *Management Science* **15(5)**, 268-277 (1969).

[22] Federgruen, A. and Tzur, M. Minimal Forecast Horizons and a New Planning Procedure for the General Dynamic Lot Sizing Model: Nervousness Revisited. *Operations Research* **42(3)**, 456-468 (1994).

[23] Federgruen, A. and Tzur, M. Fast Solution and Detection of Minimal Forecast Horizons in Dynamic Programs with a Single Indicator of the Future: Applications to Dynamic Lot-Sizing Models. *Management Science* **41(5)**, 874-893 (1995).

[24] Federgruen, A., Tzur, M. Detection of Minimal Forecast Horizons in Dynamic Programs with Multiple Indicators of the Future. *Naval Research Logistics* **43**, 169-189 (1996).

[25] Fraser, J.M., and Posey, J.W. A Framework for Replacement Analysis. *European Journal of Operational Research* **40**, 43-57 (1989).

[26] Goldstein, T., Ladany, S.P., and Mehrez, A. A Dual Machine Replacement Model: A Note on Planning Horizon Procedures for Machine Replacements. *Operations Research* **34(6)**, 938-941 (1986).

[27] Hayes, R.H. Optimal Strategies for Divestiture. *Operations Research* **17(2)**, 292-310 (1969).

[28] Hopp, W.J., and Spearman, M.L. *Factory Physics: Foundations of Manufacturing Management (Second Edition)*. McGraw-Hill, New York, NY (2001).

[29] Lasserre, J.B. Detecting Planning Horizons in Deterministic Infinite Horizon Optimal Control. *IEEE Transactions on Automatic Control* **AC-31(1)** (1986).

[30] Lee, C.-Y. and Denardo, E. Rolling Planning Horizons: Error Bounds for the Dynamic Lot Size Model. *Mathematics of Operations Research* **11**, 423-432 (1986).

[31] Liu, X., and Tu, Y. Production Planning with Limited Inventory Capacity and Allowed Stockout. *International Journal of Production Economics* **111**, 180-191 (2008).

[32] Lundin, R. and Morton, T.. Planning Horizons for the Dynamic Lot Size Model: Zabel vs. Protective Procedures and Computational Results. *Operations Research* **23(4)**, 711-734 (1975).

[33] Oakford, R., Lohmann, J., and Salazar, A. A Dynamic Replacement Economy Decision Model. *IIE Transaction* **16**, 65-72 (1984).

[34] Pierskalla, W.P., and Voelker, J.A. A Survey of Maintenance Models: The Control and Surveillance of Deteriorating Systems. *Naval Research Logistics Quarterly* **23(3)**, 353-388 (1976).

[35] Puterman, M.L. *Markov Decision Processes: Discrete Stochastic Dynamic Programming.* John Wiley & Sons, New York, 1994.

[36] Rosenfield, D.B., Shapiro, R.D. and Butler, D.A. Optimal Strategies for Selling an Asset. *Management Science* **29 (9)**, 1051-1061 (1983).

[37] Sandbothe, R.A., and Thompson, G.L. Decision Horizons for the Capacitated Lot Size Model with Inventory Bounds and Stockouts. *Computers & Operations Research* **20(5)**, 455-465 (1993).

[38] Schochetman, I.E., Smith, R.L. Existence and Discovery of Average Optimal Solutions in Deterministic Infinite Horizon Optimization. *Mathematics of Operations Research* **23(2)**, 416-432 (1998).

[39] Sethi, S. and S. Chand. Planning Horizon Procedures for Machine Replacement Models. *Management Science* **25(2)**, 140-151 (1979).

[40] Shapiro, J. and Wagner, H.. A Finite Renewal Algorithm for the Knapsack and Turnpike Models. *Operations Research* **15(2)**, 319-341 (1967).

[41] Topkis, D.M. Minimizing a Submodular Function on a Lattice. *Operations Research* **26(2)**, 305-321 (1978).

[42] Wagner, H.M., and Whitin, T.M.. Dynamic Version of the Economic Lot Size Model. *Management Science* **5**, 89-96 (1958).

[43] Zabel, E. Some Generalizations of an Inventory Planning Horizon Theorem. *Management Science* **10**, 465-471 (1964).