

Solvability in Infinite Horizon Optimization

Timothy D. Lortz* Irina S. Dolinskaya† Archis Ghaté‡ Robert L. Smith§

January 20, 2015

Abstract

We give necessary and sufficient conditions for finite detection of an optimal initial decision for infinite horizon optimization under a broad set of assumptions and provide an algorithm that is guaranteed to solve every solvable problem under these assumptions. We illustrate the theory and algorithms developed with applications in production planning and equipment replacement.

1 Introduction

Planners confronted by the task of making an optimal decision today cannot avoid the challenge that the quality of today's decision depends on the world one will face and the associated decisions one will make in the future. Key to this decision process is the length of the horizon over which decisions will be made. For example, the best current decision will preclude expensive actions whose investment must be recovered by savings over the long run if the study horizon is short. On the other hand, for firms that anticipate being in business for the long run, that horizon is typically indefinite with no predetermined end. In these situations, an infinite horizon model seems appropriate. However, unless one adopts the heroic assumption that the future will bring a world like the present, one is confronted by the task of forecasting an unlimited quantity of data. To attempt to overcome this, planners have traditionally adopted one of two alternatives. One can model the problem with stationary but stochastic data at least providing the illusion of a dynamic and changing world while the probability distributions remain stationary. Alternatively, the planner can adopt a limited lookahead and simply incorporate a nonstationary finite horizon model in place of the infinite horizon problem one actually confronts. The problem that arises here is the tendency of today's decision to be distorted by end of study effects. However, if those end of study distortions eventually stop, one can argue that one has found an optimal decision for today to the true underlying infinite horizon problem. In order to be certain a finite horizon optimal decision has in fact stabilized and will not change were we to lengthen it, we must somehow know that unrevealed data beyond the current horizon will not affect the optimality of the currently optimal decision in hand. Such a horizon is called a *forecast horizon*.

Considerable effort has been expended in the literature to establish conditions under which such a forecast horizon exists. Examples for which no forecast horizon exists can be found in the context

*Booz Allen Hamilton, Inc., Annapolis Junction, MD 20701

†Department of Industrial Engineering and Management Sciences, Northwestern University, Evanston, IL 60208

‡Department of Industrial and Systems Engineering, University of Washington, Seattle, WA 98195

§Department of Industrial and Operations Engineering, University of Michigan, Ann Arbor, Michigan 48109

of general infinite horizon optimization [1], production planning [4], and asset selling [6]. Classes of nonstationary infinite horizon optimization problems for which forecast horizons do exist have generally relied upon either uniqueness of the optimal immediate decision or monotonicity of that decision as a function of horizon length. For example, uniqueness of the first optimal decision has been established as a sufficient condition for existence of a forecast horizon for a very general class of infinite horizon optimization problems in [1]. However, although it is believed that real world problems typically satisfy this condition, we typically do not know how to check for this condition in a specific instance [11]. Indeed, if a forecast horizon does not exist (a situation we will call failure to be well-posed), the problem cannot be solved. In [6], authors give an example of such a problem whereby no matter how many periods of data is forecasted, the data not yet seen can make any decision either optimal or not optimal, thus rendering the task of determining the next best decision impossible. Such a problem is not solvable by forecasts of finite data sets, no matter how large.

This paper formalizes these notions of well-posedness and solvability for a very general class of infinite horizon problems and introduces the new notion of *coalescence* of optimal states which we show to be equivalent to being well-posed and solvable. We give necessary and sufficient conditions for finite detection of an optimal initial decision under a broad set of assumptions involving reachability between states. We show that under these assumptions, coalescence and well-posedness are equivalent, and give a solution procedure that is guaranteed to solve any well-posed problem with these properties. In particular, we provide a finite algorithm that is guaranteed to solve every solvable problem within this class of infinite horizon, deterministic nonstationary problems. We show that these reachability conditions are satisfied in infinite horizon production planning and equipment replacement.

1.1 Related Work

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 can be found in [5]. In light of that article, we will only mention other research as it directly pertains to our work.

The problems studied in this paper all satisfy the condition that there exists a feasible solution with finite total discounted cost, and our optimality criterion will be that of total discounted cost. Under such conditions, one can bound the maximum deviation from optimal total cost by solving to optimality a finite horizon version of the problem. That is, one can obtain optimal solutions over a finite horizon together with bounds on cost-to-go that ensure total cost within error ϵ of the optimal infinite horizon cost for any $\epsilon > 0$. This paper adopts the more challenging objective of *seeking convergence by looking at incrementally longer finite horizon problems not just in terms of cost error, but also in terms of policy error.*

It is often the case that there exists a sequence of selections of finite horizon optimal initial decisions that agrees, in finite time, with an infinite horizon optimal initial decision. Algorithms to determine such selections typically require the existence of forecast horizons. In general deterministic optimization, Bean and Smith [1] show that a forecast horizon exists for a very general class of problems when the optimal solution is unique. Later Bean and Smith [2] extending an algorithm

in [8] show that a weak reachability condition is necessary and sufficient for finite discovery of the optimal initial decision whenever it is unique. They also give a solution procedure that will detect an optimal initial decision in finite time when the uniqueness and weak reachability conditions are met, and stopping sets of states are appropriately chosen. In subsequent work, Bes and Sethi [3] adopted a more abstract approach to exploring general conditions for the existence of forecast horizons although these remained difficult to check in practice. However, the assumption of uniqueness of the optimal initial decision is only a sufficient condition for the solvability of a particular problem instance. Since uniqueness is not a necessary condition for discovery of an optimal initial decision, we will depart from [1, 2, 8], in that *we explicitly drop the assumption that the optimal initial decision (or any optimal decision, for that matter) is unique, allowing for the presence of multiple optima*. Moreover, we establish coalescence properties that are not only sufficient but, unlike uniqueness and weak reachability, are also necessary to solve an infinite horizon optimization problem. We follow [6] which seeks a finite algorithm that can solve every well-posed Markov Decision Process problem (a problem is well-posed if it is solvable). Such an algorithm is presented in [6] for solving well-posed problems whose optimal policies are monotone in horizon. In this paper, we show well-posed problems are those satisfying a coalescence property for optimal states and we provide an algorithm that solves every such problem satisfying some mild structural properties.

1.2 Paper Outline

The rest of the paper is organized as follows. Section 2 presents the problem statement, including the notation used, necessary assumptions and an illustrative application. In Section 3, we establish definitions of the terminology used throughout the paper, including that of a well-posed problem, forecast horizon and coalescence, and establish the key relationship between problem *solvability*, *coalescence* and *well-posed* properties. The solution procedure for our class of problems is described in Section 4. We come back to the motivating example in Section 5 and demonstrate application of our key results to infinite horizon single-item production planning (Section 5.1), as well as an additional example of equipment replacement problem (Section 5.2). The paper is summarized and concluded in Section 6 and the Appendix contains some supporting proofs.

2 Notation and Problem Statement

We consider a class of infinite horizon, deterministic, discrete time, non-stationary problems parameterized by data *forecast* $\phi \in \Phi$, where $\phi = (\phi_1, \phi_2, \dots)$, ϕ_i is a finite set of problem data associated with period i for $i = 1, 2, \dots$, and Φ is the set of all possible forecasts ϕ . We will often refer to Problem ϕ by which we mean the problem with forecast $\phi \in \Phi$. Let $\phi^n = (\phi_1, \dots, \phi_n)$, for $n \geq 1$, denote the problem data for the first n periods from a forecast ϕ in Φ . We call ϕ^n a *truncated forecast at n* , and Φ^n the class of all truncated forecasts ϕ^n from $\phi \in \Phi$. Then, $\Phi(\phi^n)$ is the set of all forecasts $\theta \in \Phi$ whose data for the first n periods match that of ϕ^n , that is, $\theta^n = \phi^n$.

For a given forecast $\phi \in \Phi$, the underlying system is observed at the beginning of periods $n = 1, 2, \dots$ to be in state $s_n \in S_n(\phi^{n-1})$, where $S_n(\phi^{n-1})$ is the set of all feasible states associated with truncated forecast ϕ^{n-1} at the beginning of period n . Then, an action $y_n \in Y_n(s_n; \phi^n)$ is chosen and a cost $c_n(y_n, s_n; \phi^n)$ is incurred, where $Y_n(s_n; \phi^n)$ is the *feasible decision space at period n* , corresponding to the set of all feasible actions when the system is in state s_n beginning period n under truncated forecast ϕ^n . The system then transitions to state s_{n+1} at the end of period n ,

following the state equation:

$$s_{n+1} = f_n(s_n, y_n; \phi^n), \quad \phi^n \in \Phi^n, \quad y_n \in Y_n(s_n; \phi^n), \quad \forall n \geq 1, \quad (1)$$

where f_n is the given state transition function for period n .

Then,

$$S_n(\phi^{n-1}) = \{f_{n-1}(s_{n-1}, y_{n-1}; \phi^{n-1}) : s_{n-1} \in S_{n-1}(\phi^{n-2}), y_{n-1} \in Y_{n-1}(s_{n-1}; \phi^{n-1})\}$$

represents all the *feasible states* in period n under truncated forecast ϕ^{n-1} . Finally, s_1 denotes the initial state at the beginning of period 1.

Note that implicit from the notation provided above, we assume the following.

Remark 2.1. $S_n(\phi^{n-1})$ only depends on the data parameters for the first $n - 1$ periods from a forecast ϕ . Therefore, for any forecast ϕ , if state s_n is feasible in period n (i.e., $s_n \in S_n(\phi^{n-1})$), then $s_n \in S_n(\theta^{n-1})$ for all $\theta \in \Phi(\phi^{n-1})$.

Remark 2.2. Feasible decision space $Y_n(s_n; \phi^n)$ is nonempty for any $\phi^n \in \Phi^n, s_n \in S_n(\phi^{n-1})$, 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 $\Phi(\phi^n)$.

Let $x \in Y(\phi) := \prod_{n=1}^{\infty} Y_n(\phi^n)$, where $Y_n(\phi^n) = \bigcup_{s_n \in S_n(\phi^{n-1})} Y_n(s_n; \phi^n)$, denote a *feasible decision sequence* for ϕ , if for $x = (x_1, x_2, \dots)$, we have $x_n \in Y_n(s_n; \phi^n)$ and $s_{n+1} = f_n(s_n, x_n; \phi^n)$ for all n . Then, the *feasible decision space* $X(\phi) \subseteq Y(\phi)$ is the subset of $Y(\phi)$ consisting of all feasible decision sequences for ϕ . Similarly, let $s \in S(\phi) := \prod_{n=1}^{\infty} S_n(\phi^{n-1})$ denote a *feasible state sequence* for ϕ if for $s = (s_1, s_2, \dots)$, we have $x_n \in Y_n(s_n; \phi^n)$ and $s_{n+1} = f_n(s_n, x_n; \phi^n)$ for all n . And then the *feasible state space* $T(\phi) \subseteq S(\phi)$ is the subset of $S(\phi)$ consisting of all feasible state sequences for ϕ .

Finally, let $V(\phi, x)$ denote the total discounted cost of implementing decision sequence x under forecast ϕ , i.e., $V(\phi, x) = \sum_{n=1}^{\infty} \alpha^{n-1} c_n(x_n, s_n; \phi^n)$ where s is the state sequence associated with decision sequence x under ϕ and where $V(\phi, x) = \infty$ if $x \notin X(\phi)$.

We now can state our **infinite horizon optimization problem**: for a given forecast ϕ , choose a feasible decision sequence $x \in X(\phi)$ that minimizes total discounted cost. That is, find an x^* that achieves

$$V^*(\phi) = \min_{x \in X(\phi)} V(\phi, x). \quad (2)$$

We adopt three basic assumptions about our problem class Φ that are listed below.

Assumption 2.1. Decision space $Y_n(\phi^n)$ and state space $S_n(\phi^{n-1})$ are finite for all n and all $\phi \in \Phi$.

Assumption 2.2. For a particular problem $\phi \in \Phi$, feasible state and decision sequences have a one-to-one correspondence.

Note that Assumption 2.2 may be considered without loss of generality. For each forecast ϕ , exercising a feasible decision sequence x results in a unique state sequence $s(\phi, x) = (s_n(\phi^{n-1}, x))_{n=1}^{\infty}$. The converse is not necessarily true, that is, if s is a sequence of feasible states for ϕ , then there may

exist multiple feasible decision sequences x^1, x^2, \dots, x^k such that $s_n(\phi^{n-1}, x^1) = \dots = s_n(\phi^{n-1}, x^k)$ for all periods n . However, in such a case, we will assume that only a feasible decision sequence $x^i, i \in \{1, \dots, k\}$ such that $s(\phi, x)$ 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 Assumption 2.2 without loss of generality. Consequently, although the total cost of a decision sequence x for a forecast ϕ 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.

Assumption 2.3. *For each $\phi \in \Phi$, there exists a feasible decision sequence which results in non-discounted cost bounded in each period by some uniform constant $B < \infty$.*

Let $X^*(\phi) \subseteq X(\phi)$ denote the set of all optimal decision sequences for ϕ , and let $X_n^*(\phi)$ be the set of all optimal decisions for period n for ϕ . Note that $X^*(\phi) \neq \emptyset$ since $V(\phi, x)$ is continuous in x over the compact set $X(\phi)$. Then, we call $X_n^{**}(\phi^n) = \bigcup_{\theta \in \Phi(\phi^n)} X_n^*(\theta)$ the set of *candidate optimal decisions* for forecast ϕ in period n . Similarly, let $T^*(\phi) \subseteq T(\phi)$ denote the set of all optimal state sequences for ϕ , $T_n^*(\phi)$ the set of all optimal states at the beginning of period n , and $T_n^{**}(\phi^{n-1}) = \bigcup_{\theta \in \Phi(\phi^{n-1})} T_n^*(\theta)$ the set of *potentially optimal states* for forecast ϕ at the beginning of period n . The states in $T_n^{**}(\phi^{n-1})$ are potentially optimal in the sense that they may be infinite horizon optimal depending on how the forecast evolves after period $n - 1$.

2.1 Key Assumptions

The following two assumptions are key to establishing the main results of this paper. In Section 2.2 below we provide some of the well-known applications that satisfy these assumptions.

Assumption 2.4. *For all forecasts $\phi \in \Phi$, periods n , and potentially optimal states $s_n \in T_n^{**}(\phi^{n-1})$, there exists a forecast $\theta \in \Phi(\phi^{n-1})$ such that s_n is the unique infinite horizon optimal state at the beginning of period n for forecast θ .*

Assumption 2.4 provides that all potentially optimal states in any period n for any forecast ϕ are the unique optimal states in that period for some forecast in $\Phi(\phi^{n-1})$ in agreement with ϕ through period $n - 1$, and therefore, these and only these need to be considered in checking for possible agreement and therefore identification of an optimal initial decision (making each potentially optimal state set a *minimal regeneration set* in the language of related literature).

Assumption 2.5. *There exists an $L < \infty$, such that for all forecasts $\phi \in \Phi$, horizons n , states $s_n \in T_n^{**}(\phi^{n-1})$ and $s_{n+L} \in T_{n+L}^{**}(\phi^{n+L-1})$, there are decisions $x_n, x_{n+1}, \dots, x_{n+L-1}$ and states $t_{n+1}, t_{n+2}, \dots, t_{n+L-1}$ such that $x_n \in Y_n(s_n; \phi^n)$, $t_{n+1} = f_n(s_n, x_n; \phi^n)$, $x_j \in Y_j(s_j; \phi^j)$, $t_j = f_{j-1}(t_{j-1}, x_{j-1}; \phi^{j-1}) \in S_j(\phi^{j-1})$ for $j = n+2, \dots, n+L-1$ and $f_{n+L-1}(t_{n+L-1}, x_{n+L-1}; \phi^{n+L-1}) = s_{n+L}$. In other words, for every forecast ϕ , there exists a feasible decision sequence that connects any pair of potentially optimal states that are at least L periods apart, where L is independent of ϕ .*

2.2 Infinite Horizon Single-Item Production Planning with Concave Costs Example

In this subsection, we provide an example in production planning that satisfies Assumptions 2.1-2.5.

Consider a production planning problem parameterized by forecast $\phi = (c, h, d)$ where

- $c = (c_n)_{n=1}^{\infty}$ and $c_n(\cdot)$ is a non-negative concave production cost function for period n for $n = 1, 2, \dots$ with $c_n(0) = 0$.
- $h = (h_n)_{n=1}^{\infty}$ and h_n is a non-negative concave inventory holding cost for period n for $n = 1, 2, \dots$, with $h_n(0) = 0$. We assume $h_n(i_n) \rightarrow \infty$ as $i_n \rightarrow \infty$.
- $d = (d_n)_{n=1}^{\infty}$ with d_n equal to the integer demand for period n , $n = 1, 2, \dots$ and with $0 \leq d_n \leq \bar{d} < \infty$ for all n .

The optimization problem is to determine a discounted cost-minimizing sequence of *production decisions* $x = (x_n)_{n=1}^{\infty}$ that satisfies demand over the entire horizon, or more formally, for $0 < \alpha < 1$

$$\text{minimize } \sum_{n=1}^{\infty} \alpha^{n-1} [c_n(x_n) + h_n(i_n)] \quad (3)$$

$$\text{subject to } i_n + x_n - i_{n+1} = d_n, \quad n = 1, 2, \dots \quad (4)$$

$$x_n, i_n \geq 0 \text{ integer}, \quad n = 1, 2, \dots \quad (5)$$

$$i_1 = 0 \quad (6)$$

where i_n is the inventory level at the beginning of period n , corresponding to the state of the system. Without loss of optimality, we limit consideration of production plans to those that satisfy the *zero-inventory optimal ordering policy*, which stipulates that one need only contemplate production in periods with zero beginning inventory [10].

Clearly the discounted cost of a production schedule that builds no inventory is feasible while a production schedule that builds unlimited inventory is suboptimal so that:

- there exists a finite cost feasible solution for ϕ (which implies Assumption 2.3); and
- there exists a uniform upper bound L on the number of periods between optimal production points.

Observe that under the assumptions that there exist uniform upper bounds on demand (\bar{d}) and on the number of periods between optimal production points (L), we see that Assumptions 2.1 and 2.5 are satisfied. Furthermore, it is evident from equation (4) that Assumption 2.2 is satisfied. In addition, since we allow the possibility that no demand occurs beyond period N of any N -period problem, and any finite horizon problem can optimally terminate with zero inventory, we see that Assumption 2.4 is also satisfied. That is, for an arbitrary inventory level i_n at time n , there exists a forecast $\theta \in \Phi(\phi^{n-1})$ for which i_n is the unique optimal state at time n . For example, consider a forecast with the cost of production in period n prohibitively large so one needs to carry inventory into period n ; at the same time, by setting cost of carrying inventory in period $n + 1$ large enough, the optimal inventory carried into period n must be precisely d_n .

3 Well-Posed Problem and Coalescence

In this section, we introduce the notion of state coalescence and show that coalescence is a necessary and sufficient condition for the infinite horizon problem to be solvable.

Suppose that a problem with forecast ϕ 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




Figure 1: Sketch of Coalescence in General Deterministic Optimization

solution x^{***} (possibly x^* itself) such that $s_m(\phi^{m-1}, x^*) = s_m(\phi^{m-1}, x^{***})$ and $s_n(\phi^{n-1}, x^{**}) = s_n(\phi^{n-1}, x^{***})$ for some m, n with $m \leq n$.

Definition 3.1. We say that a problem with forecast ϕ is **coalescent** if there exists some optimal solution x^* such that every other optimal solution is optimally reachable from x^* . In this case, such an x^* is called a **source solution**.

Figure 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 fail to intersect 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.

Definition 3.2. We say that a problem with forecast ϕ is **solvable** if there exists a procedure that uses a finite collection of data and requires a finite number of computations to deliver an optimal action x_1^* .

Definition 3.3. We say that a problem with forecast ϕ is **well-posed** if there exists a period N^* and subsequent decisions $x_k^*, k = 1, \dots, N$ for some $1 \leq N \leq N^*$ such that x_1^*, \dots, x_N^* are optimal decisions for all $\theta \in \Phi(\phi^{N^*})$. When such N^* and N exist, we shall call N^* a **forecast horizon** and N a **planning horizon** for problem ϕ with respect to class Φ .

Intuitively, a problem is well-posed when one is able to find an optimal action by forecasting its parameters only finitely far into the future. Applied recursively, one can repeatedly select the next optimal action without requiring the entire infinite forecast. The following lemma shows that well-posedness is a necessary condition for a problem to be solvable.

Lemma 3.1. A problem is solvable only if it is well-posed.

Proof. Suppose ϕ is not well-posed. Then, for any $N^* < \infty$, there exist $\theta_1 \in \Phi(\phi^{N^*})$ such that for every optimal action $x_1 \in X_1^*(\theta_1)$ we have a $\theta_2 \in \Phi(\phi^{N^*})$ such that $x_1 \notin X_1^*(\theta_2)$. That is the information in $\Phi(\phi^{N^*})$ is insufficient to establish the optimality of a candidate initial decision x_1^* regardless of how large N^* is. Therefore, the problem with forecast ϕ is not solvable. \square

We will show in the next section that under the assumption that potentially optimal states are known, a problem is solvable if and only if it is well-posed. We will now establish that coalescence is a criterion for a problem to be well-posed.

Lemma 3.2. *Let $\phi \in \Phi$ and choose any increasing subsequence of positive integers $\{n_j\}$ and associated forecasts $\{\phi(n_j)\}$, where $\phi(n_j) \in \Phi(\phi^{n_j}) \forall n_j$. Then, under Assumptions 2.1 through 2.5, there exists a further subsequence $\{n_{j_k}\}$ and associated optimal decision sequences $\{\tilde{x}^*(n_{j_k})\}$, where $\tilde{x}^*(n_{j_k}) \in X^*(\phi(n_{j_k})) \forall n_{j_k}$, and some $\tilde{x}^* \in X^*(\phi)$ such that $\tilde{x}^*(n_{j_k}) \rightarrow \tilde{x}^*$, 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 decision sequences $\{\tilde{x}^*(n_{j_k})\}$ converging to a limit point $\tilde{x}^* \in X(\phi)$.
2. $\tilde{x}^* \in X^*(\phi)$.

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 $\{\phi(n_j)\}$. By Assumption 2.1, there are finitely many candidate optimal decisions in period n for any $\theta \in \Phi(\phi^n)$, so that $X_n^{**}(\phi^n)$ is compact. The Tychonoff Theorem gives that $\prod_{n=1}^{\infty} X_n^{**}(\phi^n)$ is compact in the product topology of componentwise convergence. Thus, we can construct subsequences $\{n_{j_k}\} \subseteq \{n_j\}$ and $\{\phi(n_{j_k})\} \subseteq \{\phi(n_j)\}$ 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 ϕ , we note that the first n_{j_k} decisions of $\tilde{x}^*(n_{j_k})$ are candidate optimal and therefore feasible decisions and hence feasible for ϕ , for each k .
2. Now we will show that $\tilde{x}^* \in X^*(\phi)$. Let y^* be an infinite horizon optimal decision sequence for ϕ , with corresponding state sequence s^* . Let \tilde{s}^* and $\tilde{s}^*(n_{j_k})$ be the state sequences resulting from the decision sequences \tilde{x}^* and $\tilde{x}^*(n_{j_k})$, respectively. We first make the observation that for any k , by Assumption 2.5, since $s_{n_{j_k}-L}^*$ and $\tilde{s}_{n_{j_k}}^*(n_{j_k})$ are both potentially optimal for $\phi(n_{j_k})$ there exists a decision sequence $\tilde{y}(n_{j_k})$ such that

$$\tilde{y}(n_{j_k}) \equiv (y_1^*, \dots, y_{n_{j_k}-L}^*, \tilde{y}_{n_{j_k}-L+1}(n_{j_k}), \tilde{y}_{n_{j_k}-L+2}(n_{j_k}), \dots)$$

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

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

Then, for all k ,

$$V(\phi(n_{j_k}), \tilde{y}(n_{j_k})) \geq V(\phi(n_{j_k}), \tilde{x}^*(n_{j_k})). \quad (7)$$

We are now in a position to show that \tilde{x}^* is infinite horizon optimal for ϕ . Since $\phi(n_{j_k}) \rightarrow \phi$ and $\tilde{y}(n_{j_k}) \rightarrow y^*$, continuity of $V(\cdot, \cdot)$ gives us that $V(\phi(n_{j_k}), \tilde{y}(n_{j_k})) \rightarrow V(\phi, y^*)$. A similar argument shows that $V(\phi(n_{j_k}), \tilde{x}^*(n_{j_k})) \rightarrow V(\phi, \tilde{x}^*)$. Thus, by inequality (7), $V(\phi, \tilde{x}^*) \leq V(\phi, y^*)$. 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.

□

Theorem 3.3. *Under Assumptions 2.1 through 2.5, a problem with forecast $\phi \in \Phi$ is well-posed if and only if it is coalescent.*

Proof. Suppose ϕ is well-posed and Assumptions 2.4 and 2.5 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 $\theta \in \Phi(\phi^{N^*})$. Then by Assumption 2.4, there exists a specific forecast $\theta \in \Phi(\phi^{N^*})$ such that $s_{N^*+1}(\phi, y^*)$ is the unique infinite horizon optimal state ending period N^* for forecast θ . By well-posedness of ϕ , x_1^* is an infinite horizon optimal initial decision for θ , so that there is some $x^* \in X^*(\theta)$ which has initial decision x_1^* and state $s_{N^*+1}(\phi, y^*)$ ending period N^* . Thus, y^* is optimally reachable from x^* , and having chosen y^* arbitrarily, ϕ is coalescent.

Now suppose that ϕ is coalescent but is not well-posed with respect to Φ . Let x^{**} be a source solution. Then there exists a subsequence $\{n_j\}_{j=1}^\infty$ and $\phi(n_j) \in \Phi(\phi^{n_j})$ such that $x_1^{**} \notin X_1^*(\phi(n_j))$ for all j . But Lemma 3.2 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 $x^*(n_{j_k}) \in X^*(\phi(n_{j_k}))$ for all k such that $x^*(n_{j_k}) \rightarrow \tilde{x}^* \in X^*(\phi)$.

By assuming that ϕ is coalescent, \tilde{x}^* is optimally reached at some period M by some infinite horizon decision sequence, optimal for ϕ , 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 $\phi(n_{j_k})$ -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, $x_1^{**} \in X_1^*(\phi(n_{j_k})) \forall k \geq K(M)$. Contradiction. □

4 Solution Procedure

We now present a solution procedure for any well-posed problem ϕ , building on the results presented above. Our procedure relies on the set of potentially optimal states $T_n^{**}(\phi^{n-1})$ to be known for all periods n .

Assumption 4.1. *$T_n^{**}(\phi^{n-1})$ is known for all periods n for problem $\phi \in \Phi$.*

Assumption 4.1 is not as restrictive as it might seem. One can often show through the problem structure that every feasible state at some time n is a potentially optimal state, i.e., $S_n(\phi^{n-1}) = T_n^{**}(\phi^{n-1})$. See Section 5.2 for an example. Alternatively, one can use a procedure based on work by Federgruen and Tzur [7] to construct $T_n^{**}(\phi^{n-1})$ when only a finite superset, $\bar{T}_n^{**}(\phi^{n-1})$, of potentially optimal states, $T_n^{**}(\phi^{n-1}) \subseteq \bar{T}_n^{**}(\phi^{n-1})$, is known.

Our solution procedure (Algorithm 4.1) finds an optimal first period action x_1^* and period N^* , such that $s_2^* = f_1(s_1, x_1^*; \phi^1)$ lies on a minimum cost solution to s_{N^*} for each of the potentially optimal states $s_{N^*} \in T_{N^*}^{**}(\phi^{N^*-1})$. Since for all forecasts $\phi \in \Phi$ and horizons n , $T_n^{**}(\phi^{n-1})$ and $X_n^{**}(\phi^n)$ are finite, we can (for example by complete enumeration of feasible policies) verify this in finite time. Note that this verification requires a complete characterization of $T_n^{**}(\phi^{n-1})$ for all n and $\phi \in \Phi$, which we have by Assumption 4.1.

Algorithm 4.1.

Consider a fixed $\phi \in \Phi$. Set $n = 2$.

1. For all $s_n \in T_n^{**}(\phi^{n-1})$, find $X_1^*(\phi^{n-1}|s_n)$, the set of all optimal initial decisions for the $(n-1)$ -horizon truncation of ϕ with the state at the beginning of period n constrained to be s_n .
2. If there exists $x_1^* \in \bigcap_{s_n \in T_n^{**}(\phi^{n-1})} X_1^*(\phi^{n-1}|s_n)$, then return x_1^* and stop. Otherwise, increment n by 1 and return to step 1.

Theorem 4.1. *Under Assumptions 2.1 through 2.5 and 4.1, Algorithm 4.1 stops finitely if and only if $\phi \in \Phi$ is well-posed. Hence the following properties are equivalent:*

1. forecast ϕ is coalescent,
2. forecast ϕ is well-posed
3. forecast ϕ is solvable.

Proof. From Lemma 3.1 and Theorem 3.3, it suffices to show that Algorithm 4.1 solves every well-posed problem. Suppose that ϕ 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 $\theta \in \Phi(\phi^{N^*})$. Since S_n is finite for all n , the set $T_{N^*+1}^{**}(\phi^{N^*})$ is also finite. Furthermore, by definition, every infinite horizon optimal state sequence for every $\theta \in \Phi(\phi^{N^*})$ passes through $T_{N^*+1}^{**}(\phi^{N^*})$. But Algorithm 4.1 will find the set of optimal initial decisions for each partial state sequence with an ending state in $T_{N^*+1}^{**}(\phi^{N^*})$. By the principle of optimality, each of these sets must contain x_1^* and Algorithm 4.1 will stop at the end of period N^* (beginning of period $N^* + 1$) and return x_1^* . \square

5 Applications

We return here to the infinite horizon single-item production planning illustration introduced in Section 2.2 and demonstrate application of our key results. We also briefly demonstrate applicability of our results to an equipment replacement problem. The reader is referred to [9] for a more detailed discussion of these applications, as well as an additional illustration of a capacitated inventory planning problem.

5.1 Infinite Horizon Single-Item Production Planning with Concave Costs

Recall from Section 2.2 the infinite horizon single-item production planning problem, where forecast ϕ corresponds to infinite sequences of production cost c , holding cost h and demand forecasts d , while $\Phi = \{(c, h, d)\}$. Since production capacity is unbounded and a unit of inventory is never optimally carried more than L periods, we see that the sets $T_n^{**}(\phi^{n-1})$ for any period n and any problem $\phi \in \Phi$ are the inventory levels $\{0, \dots, L\bar{d}\}$, where \bar{d} is the uniform upper bound on demand. Note that Assumptions 2.4 and 2.5 are satisfied by all $\phi \in \Phi$ and set $T_n^{**}(\phi^{n-1})$ is known for any period n .

Due to the zero-inventory optimal ordering policy, it is straightforward to characterize coalescence in the context of concave cost production plans. For a fixed demand forecast, an optimal zero-inventory ordering policy can be completely described by the periods in which production occurs. That is, once production occurs, the amount is depleted until there is no inventory left, at which




Figure 2: Illustration of coalescence in a production planning problem with three optimal production plans

time production occurs again. If there exist two optimal production plans for a given forecast, then one is said to *optimally reach* another if at some period, they have the same inventory level. By the Principle of Optimality, in such a case, it is optimal to follow one plan until a period at which they share an inventory level, and then to follow the other. When all optimal production plans can be reached from a single optimal production plan, then the problem is said to be coalescent. For a graphical illustration of coalescence, see Figure 2, 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 four, marked by the first arrow, and plans 1 and 3 join at period five, marked by the second arrow. Thus, plan 1 is a source plan and this problem is coalescent.

With the Assumptions 2.4 and 2.5 met, we conclude the following by Theorem 3.3.

Corollary 5.1. *A problem with forecast $\phi \in \Phi$ under equations (3)-(6) is well-posed if and only if it is coalescent.*

A direct implementation of Algorithm 4.1 takes the following form for a concave cost production planning problem. Let $X_1^*(\phi^{n-1}|i)$ be the set of optimal first period decisions for the $(n-1)$ -period problem ϕ^{n-1} with the inventory level at the beginning of period n restricted to be i .

Algorithm 5.1.

1. Set $n = 2$.

2. If $\bigcap_{i=0}^{L\bar{d}} X_1^*(\phi^{n-1}|i) \neq \emptyset$, stop. Return $N^* = n$ and any $x^* \in \bigcap_{i=0}^{L\bar{d}} X_1^*(\phi^{n-1}|i)$.

3. Otherwise, set $n \leftarrow n + 1$. Return to 2.

By Theorem 4.1, we conclude the following.

Corollary 5.2. *Algorithm 5.1 will find an optimal initial decision for ϕ under equations (3)-(6) if and only if it is well-posed.*

The algorithm constructs, for each period n , the set $T_n^{**}(\phi^{n-1}) = \{0, \dots, L\bar{d}\}$ representing all potentially optimal inventory levels at the beginning of period n . When there exists an initial decision that lies on an optimal path to any potentially optimal inventory level at the beginning of period n , then the algorithm terminates.

5.2 Equipment Replacement

Suppose a piece of equipment can be replaced at the beginning of any period, and the cost of replacement may vary with time. The total equipment cost incorporates purchase costs, operating and maintenance costs, and salvage value. All costs are subject to discounting. Let $c_n(x_n)$ be the cost of acquiring a piece of equipment in period n and retiring it x_n periods later. We assume $c_n(\cdot)$ is uniformly bounded over n and that the equipment may become obsolete at any point in time, eliminating the need for further possession of equipment. Let $b = (b_n)_{n=1}^{\infty}$ be a *business requirement forecast*, where $b_n = 1$ if the equipment is required in period n , and $b_n = 0$, otherwise. We assume there exists an upper bound $R < \infty$ on the maximum feasible equipment lifetime in any period. Note that since $c_n(\cdot)$ is uniformly bounded, the discounted cost of retiring and replacing equipment in any period is finite. Finally, we assume that if there is a salvage value, it is non-increasing in the age of equipment, i.e., there is no value in antiquing of equipment. (See [9] for a more formal statement of the problem.)

The equipment replacement problem class is $\Phi = \{(c, b)\}$, where $c = (c_n)_{n=1}^{\infty}$ is a cost forecast, and $b = (b_n)_{n=1}^{\infty}$ is the business requirement forecast. In [9], it is shown that it is never strictly optimal to hold two pieces of equipment simultaneously and that equipment is optimally salvaged as soon as it becomes obsolete. Then, since there exists a finite cost feasible solution to the problem, we conclude that for all $\phi \in \Phi$ there exists an optimal solution that satisfies the condition $x_n \cdot z_n(x) = 0$ in all periods n , where $z_n(x)$ represents the remaining lifetime in period n under replacement strategy x . This leads to an interpretation of coalescence highly similar to that of concave cost production planning.

This application 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, through careful construction of a forecast extension, we see that the set of all feasible states in period n is equal to the set of all potentially optimal states in period n (i.e., $S_n(\phi^{n-1}) = T_n^{**}(\phi^{n-1})$) and Assumption 2.4 is satisfied. Moreover, because of the bound on feasible lifetimes and the freedom to purchase machines at any time, the Assumption 2.5 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, $S_n(\phi^{n-1})$ and $Y_n(s_n; \phi^n)$ for all periods n and $s_n \in S_n(\phi^{n-1})$ are finite. This is sufficient for proving that problem $\phi \in \Phi$ presented here is well-posed if and only if it is coalescent.

Let $X_1^*(\phi^{n-1}|i)$ be the set of optimal first period decisions for the $(n-1)$ -period problem ϕ^{n-1} with the remaining scheduled lifetime at the beginning of period n restricted to be i . Then we can implement Algorithm 4.1, similarly to Algorithm 5.1. The algorithm should construct, for each period n , the set $T_n^{**}(\phi^{n-1}) = \{0, \dots, R\}$ representing all potentially optimal remaining scheduled lifetimes at the beginning of period n . When there exists an initial decision that lies on an optimal path to any potentially optimal remaining scheduled lifetime at the beginning of period n , then the algorithm should terminate.

6 Conclusion

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, independent of future parameters. This relationship between coalescence and solvability was shown to hold for a set of assumptions in which potentially optimal states are finite in each time period and satisfy a reachability property. We also provided a solution procedure for this set of problems. Finally, we showed that this set of assumptions are, in fact, satisfied by some commonly studied applications.

Acknowledgment

This work was supported in part by the National Science Foundation through the grant CMMI-1333260.

References

- [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] Bès, C. and S. Sethi. Concepts of Forecast and Decision Horizons: Applications to Dynamic Stochastic Optimization Problems. *Mathematics of Operations Research*, **13(2)**, 295-310 (1988).
- [4] Bhaskaran, S. and Sethi, S. The Dynamic Lot Size Model with Stochastic Demands: a Decision Horizon Study. *INFOR* **26(3)**, 213-224 (1988).
- [5] 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).
- [6] 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).
- [7] Federgruen, A. and Tzur, M. Detection of Minimal Forecast Horizons in Dynamic Programs with Multiple Indicators of the Future. *Naval Research Logistics* **43**, 169-189 (1996).
- [8] Lasserre, J.B. Detecting Planning Horizons in Deterministic Infinite Horizon Optimal Control. *IEEE Transactions on Automatic Control* **AC-31(1)** (1986).
- [9] Lortz, T.D. Solvability in Discrete, Nonstationary, Infinite Horizon Optimization. Unpublished Ph.D. dissertation, Department of Industrial and Operations Engineering, University of Michigan, 2008.
- [10] Romeijn, E., Sharma, D., and Smith, R.L. Extreme Point Characterizations for Infinite Network Flow Problems. *Networks* **48(4)**, 209222 (2006).

- [11] Ryan, S.H., Bean, J.C., and Smith, R.L. A Tie-Breaking for Discrete Infinite Horizon Optimization. *Operations Research* **40(S1)**, S117-S126 (1992).