

A THEORY OF ROLLING HORIZON DECISION MAKING *

Suresh SETHI and Gerhard SORGER

Faculty of Management, University of Toronto, Toronto, Ontario, M5S 1A4 Canada

In this paper, we develop a theoretical framework for the common business practice of rolling horizon decision making. The main idea of our approach is that the usefulness of rolling horizon methods is, to a great extent, implied by the fact that forecasting the future is a costly activity. We, therefore, consider a general, discrete-time, stochastic dynamic optimization problem in which the decision maker has the possibility to obtain information on the uncertain future at given cost. For this non-standard optimization problem with optimal stopping decisions, we develop a dynamic programming formulation. We treat both finite and infinite horizon cases. We also provide a careful interpretation of the dynamic programming equations and illustrate our results by a simple numerical example. Various generalizations are shown to be captured by straightforward modifications of our model.

1. Introduction

Rolling horizon decision making is a common business practice for making decisions in a dynamic stochastic environment. In essence, this practice involves making the most immediate decisions, i.e., decisions that must be made in the first period, based on a forecast (deterministic or stochastic) of relevant information for a certain number of periods in the future. Clearly, the decision maker would be interested in knowing how far into the future he must forecast in order to make optimal first period decisions. At the beginning of the second period, the second period decisions become most immediate. In order to make these decisions, forecasts for additional periods in the future may be required. In addition, existing forecasts may, in some cases, also be revised or updated. This procedure repeats every period justifying the term rolling horizon decision making for the practice. Here, the term “horizon” refers to the number of periods in the future for which the forecast is made. It is this horizon, that gets “rolled over” each period.

* This research is supported in part by NSERC Grant A4619, SSHRC Grant 410-87-0524, and Manufacturing Research Corporation of Ontario. Comments and suggestions from Qing Zhang and the referees are gratefully acknowledged.

Much of the literature concerning rolling horizon decision making involves production planning problems. While the present paper deals with multiperiod stochastic optimization problems in general, it is desirable for expository purposes to limit our initial discussion to the production planning context.

In the production planning context, the decision maker is concerned with making production decisions in each period in order to satisfy demand at a minimum cost. The relevant costs that are usually considered are cost of production, cost of carrying inventories, and cost of shortages or stockouts. Demand may be either deterministic or stochastic. If deterministic, it is assumed to be given. If stochastic, then its probability law is assumed to be known. It is not usual in the production planning models to include the cost of forecasting demand. It is incorporating forecast costs explicitly in the model, that we part company from the existing literature reviewed in section 2.

Our reasons for incorporating the forecast costs is our belief that these costs along with the accuracy of forecasts (or lack thereof) lies at the heart of the practice of rolling horizon decision making. Indeed, the forecast of the future is either expensive or unreliable or both. Also, the more distant the future, the more expensive, less reliable, or both, the forecast. Moreover, in some cases, the forecast beyond a certain period may simply be unavailable. This is tantamount to saying that the forecast beyond this period is either infinitely expensive or completely uncertain in the sense that no probability law is known for it, or both. In the case of complete uncertainty beyond a certain period in the future, the problem lies in the domain of the theory of forecast and decision horizons; see, e.g. section 2 and section 6.3.

The purpose of this paper is to develop a general theoretical framework for rolling horizon decision making. Every framework must begin somewhere. In the production planning context, we begin by specifying a probability space over which a stochastic demand process is defined *a priori*. Characteristics of this process may be known from related activities, it may be apparent from the past history, or it may simply be educated guesswork on the part of the decision maker. This demand, along with the production decisions in each period, determines the dynamics of the inventory process. Given various costs associated with production and inventories/shortages, a traditional stochastic production planning problem becomes specified.

For our framework, we further assume that the firm has at its disposal some forecasting capabilities. These capabilities may reside within the firm or they may take the form of external forecasting services that the firm may hire. In either case, it costs money to forecast. It would be reasonable to assume that a forecast will represent a partial or full resolution of uncertainty in demand. Furthermore, the cost of forecasting will be assumed to depend on the distance in the future of the first period for which a forecast is not known, and the number of periods for which the forecast is made or requested. Cost may also depend on when the forecast is made. This can, of course, include the discounting of forecasting costs

incurred in the future. In addition, it is also reasonable to assume that forecasting cost will depend on the reliability (or quality) of the forecast. By this, we mean that a forecast that resolves a greater amount of uncertainty than some other forecast will be more expensive than the latter, *ceteris paribus*. Of course, this remark does not apply when the two forecasts are not comparable in the sense just defined.

In practice, sophisticated econometric forecasts will cost more than moving average forecasts or simple regression forecasts. Also, a combination of various forecasts will certainly cost more than each of the individual forecasts. However, there is evidence that a combination of forecasts yields more reliable forecasts than individual forecasts; see, e.g., Newbold and Granger [25], Granger and Newbold [15], Mahmoud [20], and Bopp [12]. Even when it comes to subjective forecasts, aggregating them over several individuals improves the quality of the forecast (Aston and Aston [3]).

Given that there are several possible forecasts to choose from and given their costs, the decision maker faces another decision problem than just making the production decision in each period. He may want to forecast a part of the future at some cost and then make the production decision. Making the production decision this way may result in cost savings that might more than offset the cost of forecasting.

Thus, it is clear that the forecasting decision may have to be made in each period along with the production decision. What is nontraditional about the problem is that, in each period, some future uncertainties may be resolved at some cost and only after this is a production decision made. In a sense, the state of the system changes twice in each period. At the beginning of the period, the decision maker knows the state of his inventory, the past realizations of demand, and previous forecasts made. On this basis, he makes a new forecasting decision, i.e., how many additional periods in the future to forecast and what quality of forecast to obtain. Once this is done, the information part of the state of the system changes. He then bases his production decision on this new state. After he makes the production decision, the demand for the period materializes, he counts the costs in the period and then moves on to the next period. Then the whole process repeats.

Repeating the process also allows the decision maker, as mentioned earlier briefly, to possibly update or revise the forecasts obtained in prior periods. In other words, it may be desirable in a particular period to obtain a less expensive and perhaps less reliable forecast, while in the next period, it may be desirable to revise the forecast at some cost along with possibly obtaining forecasts for additional periods in the future. Such a situation occurs frequently in practice. As we shall see in section 6.2, the possibility of revising past forecasts can be easily incorporated in our model.

Some important remarks from the point of view of decision making practice are in order. Since forecasting for the distant future may be prohibitively

expensive, it would be optimal in each period to forecast for a number of periods in the near future and then repeat the procedure next period. What we see therefore, is that introducing forecasting costs may result in an optimal forecasting policy, which resembles the rolling horizon policy. In fact, in an infinite horizon problem with the stationarity assumption, one would expect that the *policy* for the number of future periods forecasted in each period will be the same as in the previous period. In other words, we will have a stationary rolling horizon policy, as is often the case in practice. Of course, in a finite horizon problem, with or without the stationarity assumption, an optimal forecasting policy cannot be expected to be the same from period to period, unless it is constrained to be so.

What we have just described captures the essential features of the rolling horizon decision making procedure. As we shall see in the next section, the procedure has only been dealt with heuristically in the literature. This paper formalizes the procedure in a rigorous optimization framework.

The plan of the paper is as follows. A brief review of the literature is provided in section 2. In section 3, we precisely state the problem under consideration. The problem specification will be kept simple for convenience in exposition without losing the essential features of the rolling horizon decision making procedure. Another expository device is the assumption that an oracle (or oracles) will substitute for various forecasting methods or forecasting services. Thus, the cost of forecasting will be referred to as the cost charged by the oracle to provide some future information. Cost of going to the oracle, if any, can be considered as the fixed cost of the forecasting activity. It will be seen that our use of the oracle device results in no loss of generality, while it clearly simplifies the exposition. The problem defined in section 3 turns out to be a multi-period stochastic optimization problem with usual stochastic controls as well as stopping time decisions. Such problems have, to our knowledge, not been treated previously in the literature. Section 4 develops a dynamic programming for the finite horizon version of the problem. Interpretations are given in the production planning context for the recursion formulas of the dynamic programming. Boundedness of rolling horizon in each period is also proved. The infinite horizon version of the problem with the stationarity assumption is treated in section 5. Generalizations of the problem to incorporate more realistic as well as more general cases are treated in section 6. Here, we shall also describe the relationship of our paper to the theory of forecast and decision horizons. While the purpose of this paper is not computational, a numerical example is solved in section 7 to illustrate the theory developed in the paper. Section 8 concludes the paper with suggestions for further research on the problem.

2. A brief review of the literature

In 1977, Baker [4] conducted an experimental study of the effectiveness of rolling horizon decision making in production planning. Baker noted that, while

most existing formulations in the production planning literature are finite horizon models (see, e.g., Holt et al. [17] and Wagner and Whitin [29]), the production planning problems themselves occur in systems that will operate indefinitely. He suggested that there are two principal reasons why finite horizon models might be appropriate for decision-making in infinite horizon problems. First, the forecasts for the remote future tend to be unreliable and are, therefore, of limited usefulness. Second, the decisions must for practical reasons be based on limited information about the future. The purpose of the Baker study was to use simulation to investigate the efficiency of decisions obtained from optimizing a finite multiperiod problem with concave costs and implementing those decisions on a rolling basis. The study suggested, with exceptions however, that rolling schedules are quite efficient.

McClain and Thomas [22] examined a linear-cost model using simulation and concluded that good terminal conditions for a finite horizon model might be better than additional periods' worth of information as a way of summarizing future requirements. Baker and Peterson [6] and Baker [5] developed an analytical framework for evaluating rolling schedules. They used a simplified version of the HMMS quadratic cost model to study the effects of the length of the rolling horizon, the terminal conditions, the uncertainty in forecasts, and the periodicity of demand on the rolling horizon decisions as seen in comparison to the optimal infinite horizon policy.

Recognizing the complexity of developing a general analytical framework to evaluate rolling horizon procedures, researchers conducted empirical examinations of the effect of such factors as the length of the rolling horizon, demand variations over time, forecast errors, and lead times on the suboptimality of rolling schedules. Blackburn and Millen [11], Chand [14], Huang and Ong [18], and Wemmerlöv and Whybark [30] studied the behavior of various lot sizing heuristics in the rolling horizon environment. An interesting conclusion was that various lot sizing techniques produce similar results in the presence of uncertainty.

More recently, Alden and Smith [1] have derived error bounds for the cost of rolling horizon procedures in the general setting of a finite state Markov control model. More specifically, they have shown that, as the length of the rolling horizon increases, the cost of the corresponding rolling horizon procedure geometrically approaches the infinite horizon cost provided the model includes discounting. In addition, the error bound can be improved if ergodicity is taken into account. These results have been improved and extended to Markov control models with arbitrary state spaces by Hernandez-Lerma and Lasserre [16].

Other related work has been carried out in connection with the maintenance of a stable Master Production Schedule (MPS), a critical issue in many firms that utilize Material Requirement Planning (MRP) systems to plan and control manufacturing operations. It is known that frequent adjustments to the MPS caused by changes in customer order requirements, sales forecasts, or production

plans can induce major changes in the detailed MRP schedules – a phenomenon referred to as nervousness (see, e.g., Mather [21]). Berry et al. [7] examine the data from Abott Laboratories to study the effectiveness of the approach of freezing the MPS, used frequently in practice, in reducing the nervousness of the system. Carlson et al. [13] have modified several standard lot sizing procedures used in a rolling horizon environment to incorporate the cost of changing the MPS; thereby alleviating nervousness by considering its economic effect. In a recent paper, Sridharan et al. [28] have examined the impact on cost of three important decision variables in managing the stability of the MPS within a rolling horizon framework: the method used to freeze the MPS, the proportion of the MPS that is frozen, and the length of the rolling horizon for the MPS; see also other works on the topic cited in Sridharan et al. [28].

Finally, the extensive literature on forecast horizons is intimately related to rolling horizon decision making. We shall not review this literature as it has been surveyed in Morton [24], Aronson and Thompson [2], Bhaskaran and Sethi [10], and Sethi [26]. For our brief discussion, it is sufficient to know that a forecast horizon is a finite horizon that is far enough off that the data beyond it have no effect on the optimal decisions in the current period, see, e.g., Bes and Sethi [9]. Clearly, if we can find a forecast horizon in each period on a real-time basis without any cost, then using these forecast horizons as successive rolling horizons will provide a rolling schedule that is optimal for the infinite horizon problem. At least, in this sense, the theory of forecast horizons provides a partial justification for the practice of rolling horizon decision making.

A related paper for our work is that of Kleindorfer and Kunreuther [19], who for the first time introduced forecasting costs in their consideration of forecast horizons. Importance of these costs had been emphasized earlier by Modigliani and Cohen [23]. Kleindorfer and Kunreuther define a forecast horizon to be a period that is sufficiently distant in the future so that the data beyond it has no larger effect on the optimal decision in the first period than the cost of getting the forecast for one additional period following the period under consideration. Thus, the forecasting cost is considered exogenously. In other words, the forecast itself, except for its length, does not have any impact on other decisions. The distinguishing feature of our paper, on the other hand, is that it integrates the forecasting activity as a decision variable to be determined simultaneously with other decision variables.

3. Statement of the problem

In this section, we introduce some notation and we develop a mathematical model for the following scenario. A decision maker controlling a stochastic system over a finite or infinite planning horizon has access to an oracle that can provide exact realization of a part of the future at some given finite cost. This

assumption is relaxed in section 6.1. The decision maker seeks to minimize his total expected cost consisting of his expenses for the oracle and the cost incurred by controlling the system. Although the oracle is costly, it helps to reduce the uncertainty of the system and thereby reduces the decision maker's total cost. In each period t , two decisions have to be made, namely, how much information on the future to buy and which control action to take. Of course, the choice of the control action may depend on all the information available in that period as well as on the state of the system.

3.1. STATE EQUATION

We denote by N the problem horizon, where N is either a positive integer or ∞ . If $N < \infty$, then $Z^N = \{1, 2, \dots, N\}$ and $Z_0^N = Z^N \cup \{0\}$. On the other hand, if $N = \infty$, then Z^∞ and Z_0^∞ denote the set of positive, respectively, non-negative integers. We consider a probability space $\{\Omega, \mathcal{F}, P\}$ and a stochastic process $(\xi_t)_{t \in Z^N}$ defined on this space with values $\xi_t \in D_t$. The set D_t , $t \in Z^N$, will be termed as the set of all possible disturbances in period t and is assumed to be countable. This countability assumption helps us to avoid the delicate technical details concerned with various measurability questions. In particular, it will not be necessary to assume measurability of the transition function and the cost functions introduced below. Moreover, no structural assumptions are required on the state spaces and control spaces defined in the next paragraph.

We let ξ_0 denote an arbitrary constant and \mathcal{F}_t the sub- σ -algebra of \mathcal{F} generated by the truncated process $\xi^t := (\xi_0, \dots, \xi_t)$. For every $t \in Z^N$, there is a state space S_t and a control space C_t . The state variable at the end of period t is $x_t \in S_t$ and the control action applied in period t is $u_t \in C_t$. Furthermore, we let $S_0 = \{x_0\}$ be a given deterministic initial state. A sequence $(u_t)_{t \in Z^N}$ of controls generates a sequence $(x_t)_{t \in Z^N}$ of states via the transition equation

$$x_t = f_t(x_{t-1}, u_t, \xi_t), \quad t \in Z^N. \tag{1}$$

Here $f_t: S_{t-1} \times C_t \times D_t \rightarrow S_t$ are given functions.

3.2. ROLLING HORIZONS

Let h denote a stopping time with respect to the family $(\mathcal{F}_t)_{t \in Z^N}$ of σ -fields. This means that h is a non-negative random variable such that the event $\{h = t\} \in \mathcal{F}_t$ for all $t \in Z_0^N$. Thus, stopping times can be interpreted as random variables independent of the future. We also define \mathcal{F}_h to be the σ -algebra associated with h ; it consists of the sets $A \in \mathcal{F}$ for which $A \cap \{h = t\} \in \mathcal{F}_t$ for all $t \in Z_0^N$. In other words, \mathcal{F}_h is the totality of events to be observed over the random time h .

We let $H(s, N)$ denote the set of all stopping times h such that $s \leq h \leq N$, if $N < \infty$, and $s \leq h < \infty$, if $N = \infty$. A sequence $(h_t)_{t \in Z_0^N}$ of stopping times

$h_t \in H(t - 1, N)$ will be called a sequence of admissible rolling horizons, if $h_0 = 0$ and if

$$h_{t-1} \vee (t - 1) \leq h_t \quad \text{for all } t \in Z^N. \tag{2}$$

In what follows let us denote by h'_t the stopping time $h'_t = h_t \vee t$.

The interpretation of rolling horizons is as follows. At the beginning of period t , before approaching the oracle, the decision maker knows the realizations of the random variables $\xi^{h_{t-1}(\omega)}(\omega)$. If he decides not to go to the oracle, then those realizations along with the value of x_{t-1} is all the information he has at his disposal to decide which control u_t to implement. In this case $h_t = h'_{t-1}$. Otherwise, he asks the oracle for realizations up to period $h_t(\omega)$ at a cost to be defined in section 3.5. It should be emphasized that, since getting future information from the oracle is expensive, the decision maker may want to obtain the future realizations successively one at a time. The definition of rolling horizons as stopping times captures precisely this sequential character of information gathering.

Alternatively, the decision maker may be required to specify the number of periods for which he must obtain information all at once. This case of nonsequential information gathering can be handled as a special case of our model as described in section 4.1.

Note that the definition of rolling horizons above implies that no information on some future period can be obtained unless all realizations prior to that period are already known. Furthermore, (2) implies that no information gets lost, i.e., that any information available in period $t - 1$ will also be available in period t . Finally, since $h'_{t-1} \geq t - 1$, the information available at the beginning of period t includes at least the realizations of the past periods up to $t - 1$.

3.3. THE AUGMENTED STATE

In a dynamic optimization problem, the state variable at time $t - 1$ usually contains all the information available to make the control decision in period t . From this observation and from section 3.2, it should be clear that x_{t-1} is an insufficient state variable, since it does not include the known realizations $\xi^{h_{t-1}}$. Let us, therefore, define the augmented state spaces

$$\Sigma_t = \{ (x, s, \xi^s) \mid x \in S_t, t \leq s < N + 1, \xi^s \in D_0 \times \dots \times D_s \}. \tag{3}$$

Every augmented state vector $X_t \in \Sigma_t$ consists of the state x_t at the end of period t , the number s_t of realizations known at the end of period t , and these realizations ξ^{s_t} themselves. Given sequences $(u_t)_{t \in Z^N}$ and $(h_t)_{t \in Z^N}$ of controls and admissible rolling horizons, the augmented state evolves according to

$$X_t = F_t(X_{t-1}, u_t, h_t, \xi^N), \quad t \in Z^N, \tag{4}$$

where the function F_t is defined by

$$F_t(X, u, h, \xi^N) = (f_t(x, u, \xi_t), h \vee t, \xi^{h \vee t}) \tag{5}$$

if $X = (x, s, \xi^s)$. Note that the transition function F_t is expressed formally as a function of ξ^N , although only realizations up to period h will affect its value. This is unlike in the usual case when the transition function depends only on the t -th disturbance ξ_t . Since h could conceivably be equal to N , the above property is the main reason that makes our problem a non-standard dynamic optimization problem. In this regard we would also like to point out that the usual state variable x_t depends on the future disturbances ξ_{t+1}, \dots, ξ_h , provided that a forecast has been obtained, i.e., $h_t > t$.

3.4. DECISION LAWS

We have seen in section 3.2 (see eq. (2)) that, for a given augmented state $(x_{t-1}, h'_{t-1}, \xi^{h'_{t-1}}) \in \Sigma_{t-1}$, the set of admissible horizons for period t is given by $H(h'_{t-1}, N)$. Therefore, we call any function $\eta_t: \Sigma_{t-1} \rightarrow H(t-1, N)$ that satisfies

$$\eta_t(x, s, \xi^s) \in H(s, N), \quad \text{for all } (x, s, \xi^s) \in \Sigma_{t-1}, \tag{6}$$

an admissible rolling horizon law for period t . We denote by \mathcal{H}_t the set of all admissible rolling horizon laws for period t . Note that the elements of \mathcal{H}_t are not stopping times but functions, which map into a set of stopping times.

Now we assume that, for every $x \in S_{t-1}$, there exists a non-empty subset $U_t(x) \subseteq C_t$ called the set of admissible controls. Any function $\mu_t: \Sigma_{t-1} \rightarrow C_t$ that satisfies

$$\mu_t(x, s, \xi^s) \in U_t(x), \quad \text{for all } (x, s, \xi^s) \in \Sigma_{t-1}, \tag{7}$$

is called an admissible control law at time t . We denote by \mathcal{M}_t the set of all admissible control laws for period t .

We have assumed that the set of admissible horizons depends only on the number s of periods already known and not on the state x or the realizations ξ^s . Also, the set of admissible controls is assumed to depend only on the state x and neither on s nor on ξ^s . It should be mentioned that both of these assumptions are only imposed for simplicity in exposition. They could be relaxed without affecting the results of this paper.

We can now summarize the stochastic dynamics of the problem. At the beginning of period $t \in Z_0^N$, the augmented state is a realization of the vector $X_{t-1} = (x_{t-1}, h'_{t-1}, \xi^{h'_{t-1}})$. The decision maker applies a rolling horizon law $\eta_t \in \mathcal{H}_t$ to obtain the rolling horizon

$$h_t = \eta_t(x_{t-1}, h'_{t-1}, \xi^{h'_{t-1}}). \tag{8}$$

This rolling horizon is a stopping time and it defines a future up to which the oracle is asked for information. After receiving this information, the decision maker faces a new augmented state vector, say \bar{X}_t , defined as $\bar{X}_t = (x_{t-1}, h_t, \xi^{h_t})$. He now applies a control law $\mu_t \in \mathcal{M}_t$ to obtain

$$u_t = \mu_t(x_{t-1}, h_t, \xi^{h_t}). \tag{9}$$

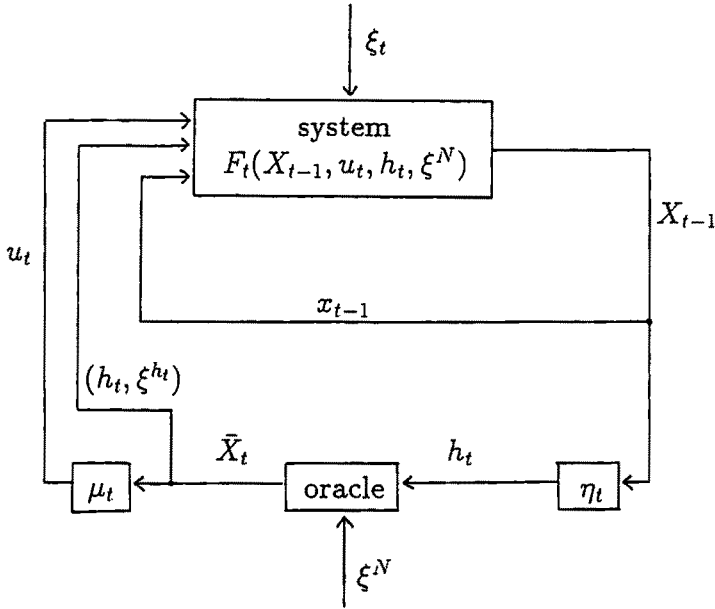


Fig. 1. Schematic of the process of information gathering and system control.

By implementing this control and by observing the realization ξ_t (if $h_t = t - 1$), the augmented state finally changes to $X_t = (x_t, h'_t, \xi^{h'_t})$. For convenience of reference, the process of information gathering and controlling the system is summarized in fig. 1.

From the above discussion (in particular, eq. (5)) and the fact that x_0 and ξ_0 are deterministic, i.e., \mathcal{F}_0 -measurable, it should be clear that x_t and X_t are \mathcal{F}_{h_t} -measurable random variables with values in S_t and Σ_t , respectively.

3.5. COST FUNCTION

The total cost incurred in period t consists of two parts. First, getting information from the oracle involves the cost $\alpha^{t-1}c_t(h'_{t-1}, h_t)$, where $\alpha \in (0, 1]$ is a constant discount factor and $c_t: Z_0^N \times Z_0^N \rightarrow [0, \infty)$ is a given function. This part of the cost will be referred to as the oracle charge. The price of the future information may depend on the number h'_{t-1} of realizations already known, on the forecast horizon h_t required, and on the period t itself. It should be noted that this formulation of the oracle charge is very general and is capable of describing various phenomena (e.g., fixed costs of going to the oracle). Reasonable assumptions on c_t include (i) that storing information is free ($c_t(h'_{t-1}, h_t) = 0$, for $h_t = h'_{t-1}$) and (ii) that forecasting for a larger number of periods is more costly than forecasting for a smaller number of periods ($c_t(h'_{t-1}, h_t)$ is non-de-

creasing in h_t). All results except for those in section 4.3, however, hold without these assumptions.

The second component of cost is the one incurred by the implementation of the control action. Since this cost will also depend on the state, it will be referred to as the running cost rather than merely the control cost. As usual, we assume that the running cost is given by $\alpha^{t-1}g_t(x_{t-1}, u_t, \xi_t)$, where $g_t: S_{t-1} \times C_t \times D_t \rightarrow [0, \infty]$ is a given extended real-valued function. Note that by allowing g_t to take on the value $+\infty$ we can easily include state constraints in our model. With the definition of the total t -period cost k_t as

$$k_t(x, u, \xi, h'_{t-1}, h_t) = g_t(x, u, \xi) + c_t(h'_{t-1}, h_t)$$

the cost functional for the N -period problem is given by

$$J^N(x_0, (\eta_t)_{t \in Z^N}, (\mu_t)_{t \in Z^N}) = E \sum_{t=1}^N \alpha^{t-1} k_t(x_{t-1}, u_t, \xi_t, h'_{t-1}, h_t), \quad (10)$$

if $N < \infty$, and

$$J^\infty(x_0, (\eta_t)_{t \in Z^\infty}, (\mu_t)_{t \in Z^\infty}) = \lim_{n \rightarrow \infty} E \sum_{t=1}^n \alpha^{t-1} k_t(x_{t-1}, u_t, \xi_t, h'_{t-1}, h_t), \quad (11)$$

if $N = \infty$. Here the rolling horizon h_t and the control u_t at time t are expressed by their feedback form (8) and (9), respectively, and E denotes the expectation with respect to the probability measure P . The decision maker seeks to minimize (10) (or (11) if $N = \infty$) over the set of all sequences $(\mu_t)_{t \in Z^N}$ and $(\eta_t)_{t \in Z^N}$ with $\mu_t \in \mathcal{M}_t$ and $\eta_t \in \mathcal{H}_t$, $t \in Z^N$, and subject to the constraints (1), (4), (8), and (9). We denote by $J^{N*}(x_0)$ the optimal cost of the N -period problem, i.e., the infimum of $J^N(x_0, \cdot, \cdot)$ subject to the above-stated constraints. Note that we assume that initially no information on the future realizations is available, i.e., $h'_0 = 0$. This allows us to write the objective functional and the optimal value function in terms of x_0 only. In a more general setting with arbitrary initial information $X_0 = (x_0, h'_0, \xi^{h_0})$ and $h'_0 > 0$, we would have to write $J^N(X_0, \cdot, \cdot)$ and $J^{N*}(X_0)$ in place of $J^N(x_0, \cdot, \cdot)$ and $J^{N*}(x_0)$, respectively. Moreover, the expectations on the right hand side of (10) and (11) would be conditional expectations given X_0 .

Finally, we would like to mention that it is straightforward to include a terminal value $S(x_N, N, \xi^N)$ in the cost functional provided $N < \infty$. This would be the case, for example, in production-inventory models, where a salvage value for the ending inventory in period N has to be taken into account.

We should note that, if the oracle is not accessible, then $h_t = t - 1$ for all $t \in Z^N$ and $u_t = \mu_t(x_{t-1}, h_t, \xi^{h_t})$ is non-anticipative. In this case the problem defined above reduces to the standard stochastic control problem. It should be obvious that access to the oracle at any cost can only decrease the optimal cost.

3.6. REMARK ON NOTATION

If $s \in Z_0^N$, then the conditional expectation with respect to the σ -algebra \mathcal{F}_s will be written as E_s , i.e. $E(Y | \mathcal{F}_s) = E_s Y$. Similarly, if h is a stopping time and \mathcal{F}_h the associated σ -algebra, then $E(Y | \mathcal{F}_h) = E_h Y$. Finally, if (h, ξ^h) is a particular realization of $(h(\omega), \xi^{h(\omega)}(\omega))$, then the conditional expectation $E(Y | \xi^{h(\omega)}(\omega) = \xi^h) = (E_h Y)(\omega)$ will be denoted by $E(Y | \xi^h)$.

4. Finite horizon problem

In this section, we consider the problem defined in section 3 for $N < \infty$. We first derive two different but equivalent dynamic programming formulations. The first one is a single recursion but it contains the optimization over the set of stopping times. Since this formulation is not suitable for computational purposes, we resolve the part of the problem concerned with the optimization over stopping times by another recursion. After the formal derivation of the dynamic programming equations, we give an intuitive explanation. Finally, we show that in a discounted model with bounded cost, the rolling horizons are uniformly bounded.

4.1. DYNAMIC PROGRAMMING

THEOREM 1

Define the value function

$$V^N: \{(t, x, h, \xi^h) | t \in Z^{N+1}, (x, h, \xi^h) \in \Sigma_{t-1}\} \rightarrow \mathbb{R}$$

recursively by

$$V^N(N + 1, x, h, \xi^h) = 0 \tag{12}$$

and, for $t \in Z^n$,

$$\begin{aligned} &V^N(t, x, h, \xi^h) \\ &= \inf_{\sigma \in H(h, N)} E \left\{ \inf_{u \in U_t(x)} E_\sigma [k_t(x, u, \xi_t, h, \sigma) \right. \\ &\quad \left. + \alpha V^N(t + 1, f_t(x, u, \xi_t), \sigma \vee t, \xi^{\sigma \vee t})] | \xi^h \right\}. \end{aligned} \tag{13}$$

Then, it holds that $J^{N*}(x_0) = V^N(1, x_0, 0, \xi^0)$.

Proof

By definition, we have

$$\begin{aligned}
 J^{N*}(x_0) &= \inf_{(\eta_t)_{t \in Z^N}, (\mu_t)_{t \in Z^N}} J^N(x_0, (\eta_t)_{t \in Z^N}, (\mu_t)_{t \in Z^N}) \\
 &= \inf_{(\eta_t)_{t \geq 1}, (\mu_t)_{t \geq 1}} E \left\{ \sum_{t=1}^N \alpha^{t-1} k_t [x_{t-1}, \mu_t(x_{t-1}, \eta_t, \xi^{n_t}), \right. \\
 &\quad \left. \xi_t, h'_{t-1}, \eta_t(x_{t-1}, h'_{t-1}, \xi^{h'_{t-1}})] \middle| \xi^0 \right\}, \tag{14}
 \end{aligned}$$

where $\eta_t(x_{t-1}, h'_{t-1}, \xi^{h'_{t-1}})$, when it occurs inside the function μ_t , has been abbreviated to η_t , and where the infimum is taken over all sequences $(\eta_t)_{t \in Z^N}$ and $(\mu_t)_{t \in Z^N}$ which satisfy $\eta_t \in \mathcal{H}_t$, $\mu_t \in \mathcal{M}_t$ for all $t \in Z^N$. Because $k_t \geq 0$, the expectation and summation in the above expression can be interchanged. Moreover, k_1 does not depend on η_t and μ_t , for $t = 2, \dots, N$. This shows that $J^{N*}(x_0)$ can be written as

$$J^{N*}(x_0) = \inf_{\eta_1 \in \mathcal{H}_1, \mu_1 \in \mathcal{M}_1} \left[E(k_1 | \xi^0) + \inf_{(\eta_t)_{t \geq 2}, (\mu_t)_{t \geq 2}} E \left(\sum_{t=2}^N \alpha^{t-1} k_t \middle| \xi^0 \right) \right], \tag{15}$$

where we have omitted the arguments of k_t , η_t , and μ_t . Both $\eta_1(x_0, 0, \xi^0)$ and $\eta'_1(x_0, 0, \xi^0) = \eta_1(x_0, 0, \xi^0) \vee 1$ are stopping times greater or equal to zero so that

$$\begin{aligned}
 E(k_1 | \xi^0) &= E[(E_{\eta_1} k_1) | \xi^0], \\
 E \left(\sum_{t=2}^N \alpha^{t-1} k_t \middle| \xi^0 \right) &= \alpha E \left[\left(E_{\eta'_1} \sum_{t=2}^N \alpha^{t-2} k_t \right) \middle| \xi^0 \right].
 \end{aligned}$$

Now observe that the decision laws η_t and μ_t for $t \geq 2$ are not applied before period two when the information $(x_1, \eta'_1(x_0, 0, \xi^0), \xi^{\eta'_1})$ is known. This means that

$$\inf_{(\eta_t)_{t \geq 2}, (\mu_t)_{t \geq 2}} E \left[\left(E_{\eta'_1} \sum_{t=2}^N \alpha^{t-2} k_t \right) \middle| \xi^0 \right] = E \left[\inf_{(\eta_t)_{t \geq 2}, (\mu_t)_{t \geq 2}} \left(E_{\eta'_1} \sum_{t=2}^N \alpha^{t-2} k_t \right) \middle| \xi^0 \right].$$

Therefore, we obtain from (15)

$$\begin{aligned}
 J^{N*}(x_0) &= \inf_{\eta_1 \in \mathcal{H}_1, \mu_1 \in \mathcal{M}_1} \left\{ E[(E_{\eta_1} k_1) | \xi^0] \right. \\
 &\quad \left. + \alpha E \left[\inf_{(\eta_t)_{t \geq 2}, (\mu_t)_{t \geq 2}} \left(E_{\eta'_1} \sum_{t=2}^N \alpha^{t-2} k_t \right) \middle| \xi^0 \right] \right\} \\
 &= \inf_{\eta_1 \in \mathcal{H}_1, \mu_1 \in \mathcal{M}_1} E \left\{ E_{\eta_1} k_1 + \alpha \inf_{(\eta_t)_{t \geq 2}, (\mu_t)_{t \geq 2}} \left(E_{\eta'_1} \sum_{t=2}^N \alpha^{t-2} k_t \right) \middle| \xi^0 \right\} \\
 &= \inf_{\eta_1 \in \mathcal{H}_1, \mu_1 \in \mathcal{M}_1} E \left\{ E_{\eta_1} \left[k_1 + \alpha \inf_{(\eta_t)_{t \geq 2}, (\mu_t)_{t \geq 2}} \left(E_{\eta'_1} \sum_{t=2}^N \alpha^{t-2} k_t \right) \right] \middle| \xi^0 \right\}.
 \end{aligned}$$

The last step is true because of $0 \leq \eta_1 \leq \eta'_1$. A similar argument as before shows that we can interchange the infimum with respect to $\mu_1 \in \mathcal{M}_1$ and the expectation operator $E(\cdots | \xi^0)$. Indeed, the control law μ_1 is not applied before the information $(x_0, \eta_1(x_0, 0, \xi^0), \xi^{\eta_1})$ becomes known. Thus, we arrive at

$$J^{N*}(x_0) = \inf_{\eta_1 \in \mathcal{X}_1} E \left\{ \inf_{\mu_1 \in \mathcal{M}_1} E_{\eta_1} \left[k_1 + \alpha \inf_{(\eta_t)_{t \geq 2}, (\mu_t)_{t \geq 2}} \left(E_{\eta'_1} \sum_{t=2}^N \alpha^{t-2} k_t \right) \right] \middle| \xi^0 \right\}.$$

The innermost infimum in this expression has exactly the same structure as the definition of $J^{N*}(x_0)$ in (14). We can, therefore, apply the same argument $N - 2$ more times to obtain

$$J^{N*}(x_0) = \inf_{\eta_1 \in \mathcal{X}_1} E \left\{ \inf_{\mu_1 \in \mathcal{M}_1} E_{\eta_1} \left[k_1 + \alpha \inf_{\eta_2 \in \mathcal{X}_2} E_{\eta'_1} \inf_{\mu_2 \in \mathcal{M}_2} E_{\eta_2} \left[k_2 + \cdots + \alpha \inf_{\eta_N \in \mathcal{X}_N} E_{\eta'_{N-1}} \inf_{\mu_N \in \mathcal{M}_N} E_{\eta_N} k_N \right] \cdots \right] \middle| \xi^0 \right\}.$$

Using the simple fact that, for any functions G_1 and G_2 defined on appropriate spaces, we have

$$\begin{aligned} \inf_{\mu \in \mathcal{M}} G_1(x, s, \xi^s, \mu(x, s, \xi^s)) &= \inf_{u \in U_1(x)} G_1(x, s, \xi^s, u), \\ \inf_{\eta \in \mathcal{X}} G_2(x, s, \xi^s, \eta(x, s, \xi^s)) &= \inf_{h \in H(s, N)} G_2(x, s, \xi^s, h), \end{aligned}$$

we can also write

$$J^{N*}(x_0) = \inf_{h_1 \in H(0, N)} E \left\{ \inf_{u_1 \in U_1(x_0)} E_{h_1} \left[k_1 + \alpha \inf_{h_2 \in H(h'_1, N)} E_{h'_1} \inf_{u_2 \in U_2(x_1)} E_{h_2} \left[k_2 + \cdots + \alpha \inf_{h_N \in H(h'_{N-1}, N)} E_{h'_{N-1}} \inf_{u_N \in U_N(x_{N-1})} E_{h_N} k_N \right] \cdots \right] \middle| \xi^0 \right\}. \tag{16}$$

From (12) and (13) follows

$$\begin{aligned} \left[\inf_{h_N \in H(h'_{N-1}, N)} E_{h'_{N-1}} \inf_{u_N \in U_N(x_{N-1})} E_{h_N} k_N \right](\omega) &= V^N(N, x_{N-1}(\omega), \\ &h'_{N-1}(\omega), \xi^{h'_{N-1}(\omega)}(\omega)). \end{aligned}$$

Using this equation, (13), and (16), we obtain recursively $J^{N*}(x_0) = V^N(1, x_0, 0, \xi^0)$. This proves the theorem. \square

Let us, just for a moment, assume that information gathering is non-sequential, i.e., that the decision maker is required to specify the number of periods for which he can obtain the realizations before getting any information from the oracle. This situation can be captured by theorem 1, if we replace the constraint

$\sigma \in H(h, N)$ in (13) by $\sigma \in \{h, \dots, N\}$, that is, if we replace the optimization over stopping times by an optimization over integer numbers. In that case (12) and (13) are a set of equations that can be used to compute $J^{N*}(x_0) = V^N(1, x_0, 0, \xi^0)$. Note that such a situation arises often in practice when the decision maker hires a consulting organization to do a forecasting project for him. In this case, sequential information gathering is far too expensive to be practical.

If, on the other hand, the original formulation with $\sigma \in H(h, N)$ being a stopping time is considered, then eq. (13) is not suitable for computational purposes. In this case, we replace the optimization over stopping times on the right hand side of (13) by a recursion. This step is based on the following result (see Shirayev [27], ch. 2.2).

LEMMA 1

Let (Y, B) be a measurable space and let, for every $y \in Y$, P_y be a probability measure on (Ω, \mathcal{F}) . Moreover, let $(y_s, \mathcal{G}_s, P_y)_{s \in Z_0^N}$ be a homogeneous Markov chain and $G: Y \rightarrow [0, \infty)$ a Borel-measurable function. We denote by E_y the expectation with respect to P_y and by $H(s, N)$ the set of all stopping times σ with respect to $(\mathcal{G}_s)_{s \in Z_0^N}$ that satisfy $s \leq \sigma \leq N$. Then the function

$$A^N(y_s, s) = \inf_{\sigma \in H(s, N)} E_{y_s} G(y_\sigma)$$

satisfies the recursion

$$A^N(y_N, N) = G(y_N), \tag{17}$$

$$A^N(y_s, s) = \min\{G(y_s), E_{y_s} A^N(y_{s+1}, s+1)\}. \tag{18}$$

Moreover, the optimal stopping time in the class $H(s, N)$ is

$$\sigma_{s, N}^* = \min\{s \leq z \leq N \mid A^N(y_z, z) = G(y_z)\}.$$

We are now in a position to derive the second dynamic programming formulation suitable for computational purposes.

THEOREM 2

Define the value functions

$$V^N: \{(t, x, h, \xi^h) \mid t \in Z^{N+1}, (x, h, \xi^h) \in \Sigma_{t-1}\} \rightarrow \mathbb{R}$$

and

$$A^N: \{(t, h, x, s, \xi^s) \mid t \in Z^{N+1}, t-1 \leq h \leq s \leq N, (x, s, \xi^s) \in \Sigma_{t-1}\} \rightarrow \mathbb{R}$$

recursively by

$$V^N(N + 1, x, h, \xi^h) = 0, \tag{19}$$

$$V^N(t, x, h, \xi^h) = A^N(t, h, x, h, \xi^h), \tag{20}$$

$$A^N(t, h, x, N, \xi^N) = \inf_{u \in U_t(x)} \left[k_t(x, u, \xi_t, h, N) + \alpha V^N(t + 1, f_t(x, u, \xi_t), N, \xi^N) \right], \tag{21}$$

and, for $h \leq s \leq N - 1$,

$$A^N(t, h, x, s, \xi^s) = \min \left\{ E \left[A^N(t, h, x, s + 1, \xi^{s+1}) \middle| \xi^s \right], \right. \\ \left. \inf_{u \in U_t(x)} E \left[k_t(x, u, \xi_t, h, s) + \alpha V^N(t + 1, f_t(x, u, \xi_t), s \vee t, \xi^{s \vee t}) \middle| \xi^s \right] \right\}. \tag{22}$$

Then it holds that $J^{N*}(x_0) = V^N(1, x_0, 0, \xi^0)$. Moreover, the optimal rolling horizon h_t in period t with augmented state $X_{t-1} = (x, h, \xi^h)$ is given by

$$h_t = \min \left\{ h \leq s \leq N \mid A^N(t, h, x, s, \xi^s) \right. \\ = \left. \inf_{u \in U_t(x)} E \left[k_t(x, u, \xi_t, h, s) + \alpha V^N(t + 1, f_t(x, u, \xi_t), s \vee t, \xi^{s \vee t}) \middle| \xi^s \right] \right\}. \tag{23}$$

Proof

From theorem 1 and (20) it is clear that it is sufficient to show that the right hand side of (13) is given by $A^N(t, h, x, h, \xi^h)$. We do this by applying lemma 1. We specify $Y = \{(s, d_0, \dots, d_s) \mid s \in Z_0^N, d_i \in D_i\}$, $\mathcal{G}_s = \mathcal{F}_s$, and $y_s = (s, \xi^s)$. For given $y = (s, \xi^s) \in Y$, the probability measure P_y on \mathcal{F} is defined by $P_y(A) = P(A \mid \xi^s(\omega) = \xi^s)$ for each $A \in \mathcal{F}$. Of course, with these specifications, the process $(y_s, \mathcal{G}_s, P_y)$ is a homogeneous Markov chain. The function $G(y)$ of lemma 1 is given by

$$G(y) = G(s, \xi^s) = \inf_{u \in U_t(x)} E_s \left[k_t(x, u, \xi_t, h, s) + \alpha V^N(t + 1, f_t(x, u, \xi_t), s \vee t, \xi^{s \vee t}) \right].$$

It is straightforward to check that (21) and (22) are identical to (17) and (18) with $A^N(y_s, s)$ replaced by $A^N(t, h, x, s, \xi^s)$. This proves the theorem. \square

4.2. INTERPRETATION OF D.P. EQUATIONS (19)–(22)

Equation (19) is obvious, since we are concerned with an N period problem with zero terminal value. $A^N(t, h, x, s, \xi^s)$ denotes the optimal cost from period t to N , if we begin period t with state x and with the information about the demands up to period h and ask the oracle for further information about the demands up to period N (if $s = N$) or up to at least period s (if $s \leq N - 1$) with these demands denoted as ξ^s . With this meaning of A^N , eq. (20) is obvious. That is, $V^N(t, x, h, \xi^h)$ is the optimal cost from period t to N , if the beginning inventory of period t is x and if the demands from period 0 to h are known to be ξ^h . To interpret (21), we note that the first quantity inside the bracket on the RHS is the total t -period cost including the oracle charge in period t to obtain the demands from period h to N . The second quantity represents the value function in period $t + 1$ with the inventory $f_t(x, u, \xi_t)$ and known demands ξ^N . Clearly, if we minimize the sum of the quantities in the bracket by a control with the full knowledge of demand and the beginning inventory level in period t , then we should obtain the LHS of (21).

Finally, we interpret the general recursion term in (22), i.e., when $s \leq N - 1$. Because of $s \leq N - 1$, the RHS is the minimum of two terms inside $\{\dots\}$. The second term has nearly the same meaning as the one above for (21) except that now we have $s \leq N - 1$ instead of $s = N$, which requires some additional care in writing this term. There are two important differences. The demand information argument inside the value function at $t + 1$, represented by the second quantity inside $[\dots]$, is $\xi^{s \vee t}$ instead of ξ^s (i.e., ξ^N) in (21). This is because even if $s = t - 1$, demand ξ_t becomes realized in period t and will, therefore, be known in period $t + 1$. The other difference is that we need an expectation of the quantities inside $[\dots]$ conditioned on (x, s, ξ^s) , since $s < N$, before we take the infimum with respect to the control in period t . All told, the second term inside $\{\dots\}$ represents the optimal cost from period t to N if we begin period t with inventory level x and with the information about the demand up to period h and ask the oracle for further information about the demands up to period $s \leq N - 1$ with these demands denoted as ξ^s . The first term inside $[\dots]$, on the other hand, represents the same as above except that for this quantity, we ask the oracle for demands up to *at least* period $s \leq N - 1$. Therefore, to obtain the LHS, we must compare the second term inside $\{\dots\}$ with the first term $E[A^N(t, h, x, s + 1, \xi^{s+1}) | \xi^s]$ inside $\{\dots\}$ representing the expected value of the LHS with $s + 1$ replacing s . Now it should be clear why the LHS equals the minimum of the two terms inside $\{\dots\}$ on the RHS.

Finally, we note that the observations made just above also explain (23). It says that, as long as the second term inside $\{\dots\}$ on the RHS of (22) exceeds the first term inside $\{\dots\}$, we continue to ask the oracle for the demand in one additional period. We stop asking the oracle for any further demand when the first term inside $\{\dots\}$ on the RHS of (22) exceeds the second term for the first time.

We can now define the domain of continued additional information

$$C_t^N(x, h, \xi^h) = \left\{ \xi_{h+1}^s \mid A^N(t, h, x, s, \xi^s) < \inf_{u \in U_t(x)} E \left[k_t(x, u, \xi_t, h, s) + \alpha V^N(t+1, f_t(x, u, \xi_t), s \vee t, \xi^{s \vee t}) \mid \xi^s \right] \right\},$$

where $\xi_{h+1}^s = (\xi_{h+1}, \dots, \xi_s)$. Thus, at time t with information (x, h, ξ^h) , if we obtain ξ_{h+1}^s from the oracle, then we continue asking the oracle for additional information if, and only if, $\xi_{h+1}^s \in C_t^N(x, h, \xi^h)$. We stop asking, if the received information is not in $C_t^N(x, h, \xi^h)$.

4.3. UNIFORM BOUNDEDNESS OF HORIZONS IN THE DISCOUNTED MODEL

In this subsection, we assume that the cost functions g_t and c_t satisfy

$$0 \leq g_t(x, u, \xi) \leq \bar{g},$$

$$c_t(h, s) = \begin{cases} K + \sum_{i=h+1}^s c(i-t) & \text{for } s \geq h+1, \\ 0, & \text{otherwise,} \end{cases} \tag{24}$$

where $c(\cdot)$ is non-negative and \bar{g} and K are nonnegative constants. Let us also assume that the discount factor α is strictly less than 1. Finally, assume that forecasting the distant future is sufficiently expensive; specifically we impose the condition

$$\liminf_{i \rightarrow \infty} c(i) > \bar{g} \frac{1 + \alpha}{1 - \alpha}. \tag{25}$$

THEOREM 3

Under the assumptions stated above, there exists a constant \bar{h} independent of the problem horizon N such that an optimal sequence of rolling horizon laws $(\eta_t)_{t \in Z^N}$ satisfies

$$\eta_t(x, s, \xi^s) \in H(s, t + \bar{h}), \quad \text{for all } (x, s, \xi^s) \in \Sigma_{t-1}.$$

More specifically,

$$\bar{h} = 1 + \inf \left\{ j \in Z: c(i) > \bar{g} \frac{1 + \alpha}{1 - \alpha}, \forall i \geq j \right\}. \tag{26}$$

Proof

Consider the problem defined in section 2 with the additional constraint $\eta_t(X) \in H(t-1, t-1)$ for all $X \in \Sigma_{t-1}$. This corresponds to a situation where

no forecasting is possible. Because of the additional constraint, the optimal cost of this new problem is higher than $J^{N*}(x_0)$. More generally, we have

$$V^N(t, x, h, \xi^h) \leq \inf \sum_{\tau=t}^N \alpha^{\tau-t} g_{\tau}(x_{\tau-1}, u_{\tau}, \xi_{\tau}), \tag{27}$$

where the infimum is taken over all admissible sequences $(\eta_t)_{t \in \mathbb{Z}^N}, (\mu_t)_{t \in \mathbb{Z}^N}$ with $\eta_t \in \mathcal{H}_t, \mu_t \in \mathcal{M}_t$ and subject to the constraints (1), (4), (8), (9), $\eta_{\tau}(X) \in H(\tau - 1, \tau - 1)$, and $x_{t-1} = x$. Note that, because of $h_{\tau} = \tau - 1$ and (25), $k_{\tau}(x_{\tau-1}, u_{\tau}, \xi_{\tau}, h'_{\tau-1}, h_{\tau}) = g_{\tau}(x_{\tau-1}, u_{\tau}, \xi_{\tau})$. From (24) and (27) follows

$$V^N(t, x, h, \xi^h) \leq \frac{\bar{g}}{1 - \alpha}. \tag{28}$$

Now consider eq. (21). The non-negativity of g_t and V^N implies

$$A^N(t, h, x, N, \xi^N) \geq c_t(h, N).$$

Using this inequality and (22) and noting that (25) and $c(i) \geq 0$ imply that $c_t(h, s)$ is non-decreasing in s , we can see by a simple induction argument that

$$A^N(t, h, x, s, \xi^s) \geq c_t(h, s). \tag{29}$$

From (22) and (23), it follows that a sufficient condition for $h_t \leq s$ is

$$\begin{aligned} & E[A^N(t, h, x, s + 1, \xi^{s+1}) | \xi^s] \\ & \geq \inf_{u \in U_t(x)} E[k_t(x, u, \xi_t, h, s) \\ & \quad + \alpha V^N(t + 1, f_t(x, u, \xi_t), s \vee t, \xi^{s \vee t}) | \xi^s]. \end{aligned}$$

Because of (28) and (29), this holds if

$$c_t(h, s + 1) \geq \bar{g} + c_t(h, s) + \alpha \frac{\bar{g}}{1 - \alpha}.$$

From (25), it follows that this is true whenever

$$c(s + 1 - t) \geq \bar{g} \frac{1 + \alpha}{1 - \alpha}.$$

Hence, the result follows from (26). \square

5. Infinite horizon problem

We now treat the problem defined in section 3 with $N = \infty$. We assume throughout this section that the problem is autonomous (time homogeneous or stationary). This means that the functions f_t, g_t , and c_t as well as the sets D_t, S_t, C_t and $U_t(x)$ do not depend explicitly on the time argument t (so that we can omit the subscript t). Moreover, for the autonomous problem, the disturbances $(\xi_t)_{t \in \mathbb{Z}^{\infty}}$ are independent and identically distributed random variables.

The independence assumption for $(\xi_t)_{t \in \mathbb{Z}^+}$ implies that past disturbances are not relevant for decision making in a given period. Let us therefore define the set of augmented state vectors by

$$\Sigma = \{ (x, s, \xi^s) \mid x \in S, 0 \leq s < \infty, \xi_i \in D \},$$

with the understanding that s denotes the number of periods in the future (counting the current period as the first period in the future), for which the future realizations ξ^s are known. If $s = 0$, then $\xi^s = \xi^0$ denotes the null element; in other words, no future realizations are known. Any function $\eta : \Sigma \rightarrow H(0, \infty)$ that satisfies

$$\eta(x, s, \xi^s) \in H(s, \infty), \quad \text{for all } (x, s, \xi^s) \in \Sigma,$$

is called an admissible rolling horizon law and any function $\mu : \Sigma \rightarrow C$ satisfying

$$\mu(x, s, \xi^s) \in U(x), \quad \text{for all } (x, s, \xi^s) \in \Sigma,$$

is an admissible control law. Note that the decision laws η and μ are stationary, i.e., they are time independent functions of the augmented state (x, s, ξ^s) . It can be shown (see theorem 6 below) that permitting the decision laws to depend explicitly on time t does not improve the optimal cost.

Our first aim is to prove a counterpart of theorem 1, i.e., the existence of a value function that satisfies the dynamic programming recursion. To obtain this result, we apply the monotone operator approach developed in Bertsekas and Shreve ([8], ch. 5). Let us, therefore, denote by F the set of all extended non-negative real-valued functions $V : \Sigma \rightarrow [0, \infty]$. The basic mapping $H : \Sigma \times C \times H[0, \infty) \times F \rightarrow [0, \infty]$ is defined by

$$H(x, h, \xi^h, u, \sigma, V) = E \left[k(x, u, \xi_1, h, \sigma) + \alpha V(f(x, u, \xi_1), (\sigma - 1) \vee 0, \xi_2^\sigma \mid \xi^h) \right],$$

where $\xi_2^\sigma = (\xi_2, \dots, \xi_\sigma)$ if $\sigma \geq 2$ and $\xi_2^\sigma = \xi^0$ if $\sigma \leq 1$. Moreover, let us define the function $V^0 : \Sigma \rightarrow [0, \infty]$ by $V^0(X) = 0$ for all $X \in \Sigma$. The following properties are true for all $X = (x, h, \xi^h) \in \Sigma$, $u \in U(x)$, $\sigma \in H(h, \infty)$, $V \in F$, $\bar{V} \in F$, and $r > 0$:

$$V^0 \leq V \leq \bar{V} \Rightarrow H(X, u, \sigma, V) \leq H(X, u, \sigma, \bar{V}), \tag{30}$$

$$V^0(X) \leq H(X, u, \sigma, V^0), \tag{31}$$

$$V \geq V^0 \Rightarrow H(X, u, \sigma, V) \leq H(X, u, \sigma, V + r) \leq H(X, u, \sigma, V) + \alpha r. \tag{32}$$

The validity of (30)–(32) follows trivially from the definitions of H and V^0 and from the non-negativity of the cost function k . Finally, the monotone convergence theorem implies

$$\lim_{n \rightarrow \infty} H(X, u, \sigma, V^n) = H(X, u, \sigma, \lim_{n \rightarrow \infty} V^n) \tag{33}$$

for any sequence $(V^n)_{n \in \mathbb{Z}^+}$ satisfying $V^0 \leq V^n \leq V^{n+1}$. The following theorem is an immediate consequence of (30)–(33) and proposition 5.2 in Bertsekas and Shreve [8].

THEOREM 4

The optimal cost satisfies $J^{\infty*}(x_0) = V^{\infty}(x_0, 0, \xi^0)$, where $V^{\infty}: \Sigma \rightarrow [0, \infty]$ is a solution of the dynamic programming equation

$$\begin{aligned}
 &V^{\infty}(x, h, \xi^h) \\
 &= \inf_{\sigma \in H(h, \infty), u \in U(x)} E \left[k(x, u, \xi_1, h, \sigma) + \alpha V^{\infty}(f(x, u, \xi_1), \right. \\
 &\quad \left. (\sigma - 1) \vee 0, \xi_2^{\sigma} \mid \xi^h \right]. \tag{34}
 \end{aligned}$$

In the same way as in the finite horizon case, we can resolve the optimization over stopping times on the right hand side of (34). The proof of the following result is an obvious modification of the one for theorem 2 and it is therefore omitted.

THEOREM 5

Define the function

$$A^{\infty}: \{(h, x, s, \xi^s) \mid 0 \leq h \leq s < \infty, (x, s, \xi^s) \in \Sigma\} \rightarrow [0, \infty]$$

as the largest function satisfying

$$\begin{aligned}
 A^{\infty}(h, x, s, \xi^s) = \min \left\{ E \left[A^{\infty}(h, x, s + 1, \xi^{s+1}) \mid \xi^s \right], \right. \\
 \left. \inf_{u \in U(x)} E \left[k(x, u, \xi_1, h, s) + \alpha A^{\infty}((s - 1) \vee 0, \right. \right. \\
 \left. \left. f(x, u, \xi_1), (s - 1) \vee 0, \xi_2^s \mid \xi^s \right] \right\}. \tag{35}
 \end{aligned}$$

Then, $J^{\infty*}(x_0) = A^{\infty}(0, x_0, 0, \xi^0)$.

We conclude this section by the following important remark. We have shown that the problem under consideration is a particular instance of the abstract sequential optimization problem treated in Bertsekas and Shreve [8]. Therefore, other results proved for the abstract problem are also valid for our model. In particular, the value function V^N of the autonomous finite horizon problem converges to V^{∞} as N approaches infinity if, and only if, the limit function $(\lim_{n \rightarrow \infty} V^N)$ satisfies eq. (34). Moreover, we have the following result on the stationarity of the optimal policies.

THEOREM 6

Under the assumptions stated above and $\alpha < 1$, the optimal cost $J^{\infty*}(x_0)$ of the given problem is the same as the optimal cost of a relaxed problem in which the decision laws η and μ are allowed to be explicit functions of time.

Proof

See proposition 5.1 in Bertsekas and Shreve [8].

6. Generalizations

In this section, we briefly outline how our model can be modified in order to describe more general and more realistic situations. Moreover, we point out the connection with the theory of forecast horizons.

6.1. ORACLE WITHOUT PERFECT FORESIGHT

In most realistic situations, it is not possible to resolve the future uncertainty entirely regardless of the effort devoted to forecasting research. This means that the oracle can only provide some information on the future but not a perfect forecast. Also, in some cases, the interpretation of the forecast might not be unambiguous so that some uncertainty remains unresolved. Both of these cases can easily be captured by our model, if we allow for two stochastic processes $(\xi_t)_{t \in \mathbb{Z}^N}$, $(\delta_t)_{t \in \mathbb{Z}^N}$ instead of only $(\xi_t)_{t \in \mathbb{Z}^N}$. The oracle can only reveal the realizations of $(\xi_t)_{t \in \mathbb{Z}^N}$ but not those of $(\delta_t)_{t \in \mathbb{Z}^N}$. The functions f_t and g_t will now depend on both ξ_t and δ_t , and the augmented state at the end of period $t-1$ will be of the form $X_{t-1} = (x_{t-1}, h'_{t-1}, \xi^{h_{t-1}}, \delta^{t-1})$, since the $(t-1)$ -th realization of the δ -process becomes known automatically at the beginning of the t -th period. Without going into details, we just note that the dynamic programming equations of theorem 2 have to be modified in the following way

$$\begin{aligned} V^N(N+1, X_N) &= 0, \\ V^N(t, x, h, \xi^h, \delta^{t-1}) &= A^N(t, h, x, h, \xi^h, \delta^{t-1}), \\ A^N(t, h, x, N, \xi^N, \delta^{t-1}) \\ &= \inf_{u \in U_t(x)} E \left[k_t(x, u, \xi_t, \delta_t, h, N) \right. \\ &\quad \left. + \alpha V^N(t+1, f_t(x, u, \xi_t, \delta_t), N, \xi^N, \delta^t) \middle| \xi^N, \delta^{t-1} \right] \end{aligned}$$

and, for $h \leq s \leq N-1$,

$$\begin{aligned} A^N(t, h, x, s, \xi^s, \delta^{t-1}) \\ &= \min \left\{ E \left[A^N(t, h, x, s+1, \xi^{s+1}, \delta^{t-1}) \middle| \xi^s, \delta^{t-1} \right], \right. \\ &\quad \left. \inf_{u \in U_t(x)} E \left[k_t(x, u, \xi_t, \delta_t, h, s) \right. \right. \\ &\quad \left. \left. + \alpha V^N(t+1, f_t(x, u, \xi_t, \delta_t), s \vee t, \xi^{s \vee t}, \delta^t) \middle| \xi^s, \delta^{t-1} \right] \right\}. \end{aligned}$$

Note that the equation corresponding to (21) now has also an expectation because δ_t is a random variable.

An important special case arises when the distribution of the random variable δ_t is known up to some parameters, for which forecasts can be obtained from the oracle. For example, demand δ_t in a production planning problem may be geometrically distributed with mean ξ_t . This implies that $\delta_t(\omega) = d_t(\xi_t(\omega), \omega)$, where $d_t(\xi, \cdot)$ is a geometrically distributed random variable with mean ξ . Moreover, we have $f_t(x_{t-1}, u_t, \xi_t, \delta_t) = f_t(x_{t-1}, u_t, \delta_t)$ and similarly for the cost functions.

Another example for an imperfect oracle arises when the oracle can rule out a future disturbance with certainty. Let, for example, d_1, d_2, \dots be the possible values of the disturbance δ_t in period t and let $A_l = \{\delta_t = d_l\}$, $l \in Z^\infty$, be the corresponding events. Moreover, let $(A_{lk})_{k \in Z^\infty}$ be a given partition of A_l , $l \in Z^\infty$, into mutually disjoint sets. Each partition that satisfies $A_{ll} = \emptyset$, $l \in Z^\infty$, defines an oracle via

$$\xi_t(\omega) = d_l, \text{ for all } \omega \in \bigcup_{k \neq l} A_{kl}.$$

The interpretation of this is that the oracle announces that d_l will not happen in period t whenever $\omega \in A_{kl}$ for some $k \in Z^\infty$. Note that the oracle cannot give an incorrect forecast, because we have assumed that the sets A_{ll} are empty. The extension of this example to oracles which rule out more than one of the possible disturbances d_1, d_2, \dots is straightforward.

Finally, let us demonstrate how to model an unreliable oracle that gives the correct answer only in $\beta \cdot 100\%$ of all cases. As in the above example assume that the possible values of δ_t are given by d_l , $l \in Z^\infty$, and that a partition of each of the sets A_l , $l \in Z^\infty$, is given. If this partition satisfies $P(A_{ll}) = \beta P(A_l)$ for all $l \in Z^\infty$, then the oracle defined by

$$\xi_t(\omega) = d_l, \text{ for all } \omega \in \bigcup_{k \in Z^\infty} A_{kl}$$

has the following property. In $\beta \cdot 100\%$ of all cases in which the the state of nature at time t is going to be d_l , the forecast of the oracle for period t is $\xi_t = d_l$.

We conclude this subsection by remarking that more complicated oracles as well as combinations of the oracles defined above can easily be modelled in our framework with two processes $(\delta_t)_{t \in Z^N}$ and $(\xi_t)_{t \in Z^N}$. Moreover, the perfect oracle defined in section 3 is the special case in which these two processes are identical.

6.2. MORE THAN ONE ORACLE

Now consider the situation where the uncertainty of the model is described by m different random processes $(\xi_t^{(1)})_{t \in Z^N}, \dots, (\xi_t^{(m)})_{t \in Z^N}$. As an example, consider a firm that produces for markets in m different countries. Let $\xi_t^{(i)}$ denote the

demand in country i during period t . Research institutes may conduct surveys in different countries to obtain forecasts for the future demands in these countries. There might be one research institute for each country (parallel oracles) or there might be some institutes which can conduct surveys in more than one country (possibility of hierarchical oracles). The firm (i.e., the decision maker) has now an additional decision variable describing which oracles to ask for information. For each of the oracles chosen, a sequential information gathering procedure is applied as described in section 3. It should be noted that the presence of hierarchical oracles can incorporate the realistic case in which the decision maker in a given period can decide to update or revise the past forecasts (i.e., forecasts requested in prior periods) along with obtaining forecasts for additional future periods. It should be clear that this generalization of our model to incorporate several oracles can also be formulated as a dynamic programming problem similar to the one in theorem 2. The notation, however, becomes very cumbersome and, therefore, we have chosen not to present the relevant equations here.

6.3. INTEGRATION WITH THEORY OF FORECAST HORIZON

As mentioned in section 1, there are situations when forecasts beyond a certain period in the future are simply unavailable. In these cases, the probability measure is not defined for events that are related to these future periods and our model will not apply. However, we can use the theory of forecast horizons in conjunction with our model. To be more specific, then, let us assume that we have a measure space $\{\Omega, \mathcal{F}\}$ over which a measurable process $\{\xi_t\}_{t \in \mathbb{Z}^\infty}$ is defined and that \mathcal{F}_t is the σ -field generated by the process ξ^t . Moreover, we let \mathcal{P} denote a set of probability measures on $\{\Omega, \mathcal{F}\}$ for which the only event with zero probability is the empty set. As in Bes and Sethi [9], we define a space of forecasts D . Let $\{\Omega, \mathcal{F}_t, P_t^d\}_{t \in \mathbb{Z}^\infty}$ denote a sequence of probability spaces for $d \in D$ with the requirement that $P_t^d \in \mathcal{P}$ for all $d \in D$ and all $t \in \mathbb{Z}^\infty$ and that P_{t+1}^d agrees with P_t^d on the events in $\mathcal{F}_t \subseteq \mathcal{F}_{t+1}$.

Furthermore, we define the sets

$$I_t(d) = \{e \in D \mid P_s^d = P_s^e, 1 \leq s \leq t\} \subseteq D$$

and the partitions

$$\mathcal{I}_t = \{I_t(d), d \in D\}, \quad t \in \mathbb{Z}^\infty$$

in D . We can now define the concept of a forecast horizon. Let $T \in \mathbb{Z}^\infty$ and $I \in \mathcal{I}_T$. We say that T is an I -forecast horizon, if there exist policies $\gamma_1^I \in \mathcal{H}_1$ and $\theta_1^I \in \mathcal{M}_1$ such that, for all $N \geq T$ and for all $d \in I$, there is at least one optimal pair of strategies $\eta_1^{N,d} \in \mathcal{H}_1$ and $\mu_1^{N,d} \in \mathcal{M}_1$ for the N -period problem with forecast d such that

$$\eta_1^{N,d} = \gamma_1^I \quad \text{and} \quad \mu_1^{N,d} = \theta_1^I.$$

Of course, as in Bes and Sethi [9], other parameters such as oracle costs etc. could also be made dependent on $d \in D$.

7. Numerical example

In this section, we apply the algorithm developed in section 4 to a simple production planning model with quadratic holding/backordering cost and linear production cost. The model is constructed so as to illustrate simply the possible cases that can occur and is not an attempt to describe any realistic situation. More specifically, we assume the running cost function

$$g_t(x, u, \xi) = x^2/2 + u$$

and the transition function

$$f_t(x, u, \xi) = x + u - \xi.$$

Demand is assumed to be a finite state stationary Markov chain with values in $D_t = \{0, 1, 2, 3, 4\}$ and

$$P(\xi_1 = i) = \begin{cases} 0.4 & \text{if } i = 1, \\ 0.6 & \text{if } i = 3, \\ 0 & \text{otherwise.} \end{cases}$$

The transition probabilities $P(\xi_{t+1} = j | \xi_t = i) = q_{ij}$, $i, j = 0, \dots, 4$, are given by the matrix

$$Q = \begin{pmatrix} 0.2 & 0.8 & 0 & 0 & 0 \\ 0.2 & 0 & 0.8 & 0 & 0 \\ 0 & 0.5 & 0 & 0.5 & 0 \\ 0 & 0 & 0.7 & 0 & 0.3 \\ 0 & 0 & 0 & 0.6 & 0.4 \end{pmatrix}.$$

We specify the oracle charge by

$$c_t(h'_{t-1}, h_t) = \sum_{i=h'_{t-1}+1}^{h_t} 0.01(i-t+1).$$

This means that forecasting demand for a period that is τ time units ahead incurs the cost 0.01τ . By specifying the parameter values $N = 5$, $\alpha = 1$, $x_0 = 0$ and the control set $U_t(x) = \{0, 1, 3.5\}$ for all $t \in \{1, \dots, 5\}$ and all $x \in \mathbb{R}$, the model is completely determined.

Applying the algorithm of theorem 2 we obtain, for the first period, the optimal rolling horizon

$$h_1 = \eta_1(1, 0, \xi_0) = \begin{cases} 1, & \xi_1 = 3, \\ 2, & \xi_1 = 1 \text{ and } \xi_2 = 0, \\ 3, & \xi_1 = 1 \text{ and } \xi_2 = 2. \end{cases}$$

This shows that it is optimal to go to the oracle in period one and ask for a forecast of at least the first period. The first period demand can be 3 (with probability 0.6) or 1 (with probability 0.4). If the oracle tells us that the first

period demand will be 3 units, then we stop asking for more information and produce 3.5 units of the product since

$$u_1 = \mu_1(0, 1, (\xi_0, 3)) = 3.5.$$

Let us denote this as case A.

Now consider the case that the oracle tells us that demand in period 1 will be one unit. In this situation, the optimal decision is to ask also for the demand in period 2. According as the oracle announces $\xi_2 = 0$ or $\xi_2 = 2$ we stop asking and produce 1 unit (case B) or we continue asking for the demand of period 3 (case C), respectively. Whatever information on the demand ξ_3 we obtain in case C, it turns out to be optimal to stop the sequential information gathering process. The optimal amount of units to produce, however, depends on ξ_3 in the following way: $u_1 = 0$ if $\xi_3 = 1$ and $u_1 = 1$ if $\xi_3 = 3$. Formally, the above optimal decisions can be written as

$$u_1 = \mu_1(0, 2, (\xi_0, 1, 0)) = 1,$$

$$u_1 = \mu_1(0, 3, (\xi_0, 1, 2, 3)) = 1,$$

$$u_1 = \mu_1(0, 3, (\xi_0, 1, 2, 1)) = 0.$$

Note that the above mentioned cases constitute a complete discussion of the optimal decisions in period 1. The “curse of dimensionality” leads to a tremendous increase of possible cases in period 2. We, therefore, restrict the following discussion to one particular scenario (one realization ω). Let us assume that case A has occurred in period 1. This means that, at the beginning of period 2, we know only the demand $\xi_1 = 3$. The inventory contains $x_1 = 0.5$ units, because we have produced 3.5 units and demand has been 3 units. From theorem 2, we obtain

$$h_2 = \eta_2(0.5, 1, (\xi_0, 3)) = 2, \quad \text{for all } \omega \in \Omega.$$

This means that it is optimal to ask the oracle for a forecast of the second period’s demand but nothing more (independently of the information obtained). Moreover,

$$u_2 = \mu_2(0.5, 2, (\xi_0, 3, 2)) = 1,$$

$$u_2 = \mu_2(0.5, 2, (\xi_0, 3, 4)) = 3.5.$$

Let us assume that the oracle tells us that demand in period 2 will be 4 units and, therefore, we produce 3.5 units. This implies that, at the beginning of period 3, we have an empty inventory, $x_2 = 0$, and we know the demands $\xi_1 = 3$, $\xi_2 = 4$. We obtain for the third period

$$h_3 = \eta_3(0, 2, (\xi_0, 3, 4)) = 2, \quad \text{for all } \omega \in \Omega,$$

which shows that it is optimal to make the production decision without any information on the demands in the current and the future periods. The solution of this (usual) stochastic dynamic programming problem is to produce 3.5 units

again. The reason for the suboptimality of asking for any information in this case is that, from the known probability distribution of demand and from $\xi_2 = 4$, we can conclude that the demand in the following periods will remain relatively high. This knowledge is enough to make the optimal decision of producing at the highest level possible, i.e., 3.5 units. Formally, we have

$$u_3 = \mu_3(0, 2, (\xi_0, 3, 4)) = 3.5.$$

Let us assume that the demand actually realizing in period 3 is 3 units, i.e., $\xi_3 = 3$. Then the augmented state vector at the beginning of period 4 is

$$X_3 = (0.5, 3, (\xi_0, 3, 4, 3)).$$

We obtain

$$h_4 = \eta_4(0.5, 3, (\xi_0, 3, 4, 3)) = 4, \quad \text{for all } \omega \in \Omega,$$

and

$$u_4 = \mu_4(0.5, 4, (\xi_0, 3, 4, 3, 2)) = 0,$$

$$u_4 = \mu_4(0.5, 4, (\xi_0, 3, 4, 3, 4)) = 3.5.$$

The result $h_4 = 4$ says that it is optimal to go to the oracle but to ask only for the current period's demand. Assuming that we obtain the information $\xi_4 = 2$, the control law μ_4 implies that we should produce nothing at all. This leads to a backlog of $x_4 = -1.5$ units at the beginning of period 5 and because of

$$h_5 = \eta_5(-1.5, 4, (\xi_0, 3, 4, 3, 2)) = 4,$$

$$u_5 = \mu_5(-1.5, 4, (\xi_0, 3, 4, 3, 2)) = 0,$$

we neither go to the oracle in the last period nor do we produce anything. Of course, this follows from the fact that we did not include a salvage value function evaluating a terminal inventory or a terminal backlog.

8. Concluding remarks

In this final section, we point out several directions for further research. First of all, we would like to mention that the purpose of this paper was to develop a general theoretical framework for stochastic dynamic optimization problems in which the future uncertainty can partly be resolved at some cost. An obvious next step would be to evaluate various existing heuristic methods of rolling horizon decision making (like those mentioned in section 2) in this framework. Since we believe that costly forecasts are one of the major reasons for the application of rolling horizon methods, we expect that such a performance evaluation of heuristics should yield good results under reasonable assumptions. We hope that our framework will further clarify these assumptions and probably lead to new heuristic methods.

The second research topic follows from the obvious computational difficulties of solving our optimization problem. Since the augmented state vector X_t contains the usual state x_t , as well as the known realizations ξ^{h_t} , its dimension is quite large even for simple problems. It might be possible to reduce the computational complexity of our dynamic programming algorithm by taking into account the special structure of the problem under consideration. Moreover, it might be useful to develop other algorithms like policy iteration, value iteration, or even heuristic methods to solve the non-standard optimization problem.

Another direction of research is concerned with sensitivity analysis with respect to various parameters. Let us mention just two particular questions. First, we would like to know under what conditions the optimal rolling horizon h_1^N of the N -period problem is monotonic with respect to the problem horizon N . Together with a boundedness result like the one in section 4.3, such a monotonicity condition would ensure finite convergence of the optimal rolling horizon as N tends to infinity. The second question concerns the influence of fixed cost of going to the oracle on the structure of the optimal solution. Of course, increasing the fixed part of $c_t(h, s)$, if any, will likely reduce the number of periods in which the oracle is approached for information. The exact relation between the fixed costs and the number and lengths of forecasts would be worthwhile to explore.

In some special cases, it might be possible to obtain an explicit solution of our problem. Consider a simple model with stationary demand and simple cost structure (e.g., only fixed costs) as an extension of the lot size model. Is it possible to derive a production policy of the (s, S) type? What is the structure of the optimal forecast policy?

Finally, we would like to mention that a continuous time version of our problem, especially of the forecasting cost, might be difficult, if not impossible, to obtain.

References

- [1] J.M. Alden and R.L. Smith, Rolling horizon procedures in nonhomogeneous Markov decision processes, Technical report 87-25, Dept. of Industrial and Operations Engineering, Univ. of Michigan (1987).
- [2] J. Aronson and G.L. Thompson, Forward methods in mathematical programming, *J. Large Scale Syst.* 7 (1984) 1–16.
- [3] A.H. Aston and R.H. Aston, Aggregating subjective forecasts in some empirical results, *Man. Sci.* 31 (1985) 1499–1508.
- [4] K.R. Baker, An experimental study of rolling schedules in production planning, *Decision Sci.* 8 (1977) 19–27.
- [5] K.R. Baker, An analysis of terminal conditions in rolling schedules, *Eur. J. Oper. Res.* 7 (1981) 355–361.
- [6] K.R. Baker and D.W. Peterson, An analytic framework for evaluating rolling schedules, *Man. Sci.* 25 (1979) 341–351.
- [7] W.L. Berry, T.E. Vollman and D.C. Whybark, *Master Production Scheduling: Principles and Practice* (American Production and Inventory Control Society, 1979).

- [8] D.P. Bertsekas and S.E. Shreve, *Stochastic Optimal Control: The Discrete Time Case* (Academic Press, New York, 1978).
- [9] C. Bes and S. Sethi, Concepts of forecast and decision horizons: Application to dynamic stochastic optimization problems, *Math. Oper. Res.* 13 (1988) 295–310.
- [10] S. Bhaskaran and S. Sethi, Decision and forecast horizons in stochastic environment: A survey, *Opt. Control Appl. Meth.* 8 (1987) 201–217.
- [11] J.D. Blackburn and R.A. Millen, A heuristic lot-sizing performance in a rolling-schedule environment, *Decision Sci.* 11 (1980) 691–701.
- [12] A.E. Bopp, On combining forecasts: Some extensions and results, *Man. Sci.* 31 (1985) 1492–1498.
- [13] R.C. Carlson, J.V. Jucker and D.H. Kropp, Less nervous MRP systems: A dynamic economic lot-sizing approach, *Man. Sci.* 25 (1979) 754–761.
- [14] S. Chand, A note on dynamic lot sizing in a rolling-horizon environment, *Decision Sci.* 13 (1982) 113–119.
- [15] C. Granger and P. Newbold, *Forecasting Economic Time Series* (Academic Press, New York, 1977).
- [16] P. Hernandez-Lerma and J.B. Lasserre, Error bounds for rolling horizon policies in general Markov control processes, Report 88298, Laboratoire d'Automatique et d'Analyse des Systèmes, Toulouse (1989).
- [17] C. Holt, F. Modigliani, J. Muth and H. Simon, *Planning Production Inventories and Work Force* (Prentice-Hall, Englewood Cliffs, NJ, 1960).
- [18] H.C. Huang and H.L. Ong, Expected performance of some lot-sizing heuristics in a rolling schedule environment, *IIE Trans.* 16 (1984).
- [19] P. Kleindorfer and H. Kunreuther, Stochastic horizons for the aggregate planning problems, *Man. Sci.* 24 (1978) 485–497.
- [20] E. Mahmoud, Accuracy in forecasting: A survey, *J. Forecasting* 3 (1984) 1319–159.
- [21] H. Mather, Reschedule the reschedules you just rescheduled – Way of life for MRP?, *Prod. Inventory Man.* 18 (1977) 60–79.
- [22] J.O. McClain and J.L. Thomas, Horizon effects in aggregate production planning with seasonal demand, *Man. Sci.* 23 (1977) 728–736.
- [23] F. Modigliani and K. Cohen, *The Role of Anticipation and Plans in Economic Behavior and Their Use in Economic Analysis and Forecasting* (University of Illinois Bureau of Economic and Business Research, 1961).
- [24] T.E. Morton, Forward algorithms for forward-thinking managers, *Appl. Man. Sci.* 1 (1981).
- [25] P. Newbold and C. Granger, Experience with forecasting univariate series and the combination of forecasts, *J. Roy. Statist. Soc. Ser. A* 137 (1974) 131–165.
- [26] S. Sethi, Decision and forecast horizon in dynamic optimization, in: *The Systems and Control Encyclopedia Supplement*, ed. M. Singh (Pergamon, London, 1989).
- [27] N. Shirayev, *Optimal Stopping Rules* (Springer, New York, 1978).
- [28] V. Sridharan, W.L. Berry and V. Udayabhanu, Freezing the master production schedule under rolling planning horizons, *Man. Sci.* 33 (1987) 1137–1149.
- [29] H. Wagner and T. Whitin, Dynamic version of the economic lot size model, *Man. Sci.* 5 (1958) 89–96.
- [30] U. Wemmerlöv and D.C. Whybark, Lot-sizing under uncertainty in a rolling schedule environment, *Int. J. Prod. Res.* 22 (1984) 467–484.