

INVENTORY MODELS WITH CONTINUOUS, STOCHASTIC DEMANDS¹

BY SIDNEY BROWNE AND PAUL ZIPKIN

Columbia University¹

This article is concerned with the (r, q) inventory model, where demand accumulates continuously, but the demand rate at each instant is determined by an underlying stochastic process. The primary result is the demonstration of a certain insensitivity property, which characterizes the limiting behavior of the model. This property drastically simplifies the computation of performance measures for the system.

1. Introduction. This article is concerned with an inventory model which in most ways is quite simple and standard: There is a single product and a single location. Time is modeled as continuous, and the data are stationary. Orders are placed with an outside supplier, and they arrive after a leadtime, which may be constant or stochastic. All stockouts are backordered.

Also, we restrict attention to a simple, familiar class of control policies, the reorder-point/order-quantity or (r, q) policies: When the *inventory position* (stock on hand plus stock on order minus backorders) reaches the *order point* r , an order is placed for the fixed amount q , the *batch size*.

What is novel here is the demand process: We assume demand is driven by an underlying, exogenous, continuous-time stochastic process, the state of the world, or *world* for short, denoted

$$\mathbf{x} = \{x(t) : t \geq 0\},$$

with state space X . This process may model the economy or conditions in a particular industry, for instance, as well as purely random noise. In addition, we specify a function $\lambda: X \rightarrow \mathcal{R}^+$, the *demand rate*. The process \mathbf{x} and the function λ work together to determine demand as follows: At time t , if $x(t) = \mathbf{x}$, then demand occurs at the rate $\lambda(\mathbf{x})$. That is, if we denote

$$D(t) = \text{cumulative demand in the interval } (0, t],$$

then

$$(1.1) \quad D(t) = \int_0^t \lambda[x(s)] ds.$$

Later, we shall impose specific assumptions on \mathbf{x} and λ , but for now we mention only the most basic of these: We model the world \mathbf{x} as a time-homogeneous Markov process. Moreover, \mathbf{x} is ergodic or regular in the same sense as

Received April 1990; revised October 1990.

¹Partially supported by NSF Grant DDM-89-20660.

AMS 1980 *subject classifications*. 60J25; 90B05.

Key words and phrases. Markov processes, inventory theory, clearing processes, uniform distribution, insensitivity.

the best-behaved Markov chains; in particular, it has a unique stationary probability density π . The demand-rate function λ is sufficiently smooth that the integral in (1.1) is well defined, and also sufficiently variable that $D(t)$ is "truly" stochastic. In addition, we shall need a key irreducibility assumption, described below.

We shall refer often to two special cases of this model: In the first, X is discrete, so \mathbf{x} is a continuous-time Markov chain. The second case is where \mathbf{x} is a (multivariate) diffusion process.

There are several reasons for studying demand processes of this form. First, there is now a substantial body of knowledge concerning the use of Markov processes in continuous time, especially diffusion processes, to model various economic phenomena. See, for example, Sethi and Thompson (1981) and Malliaris and Brock (1982). Specifically, many familiar, widely used demand forecasting models can be cast in this form. So, our model extends traditional inventory analysis to encompass a very rich and flexible class of demand processes.

The second reason is pedagogical: There is a gap in inventory theory between the deterministic EOQ model and the various models with stochastic demand. The Poisson process is by far the most widely studied demand model, but here $D(t)$ and all the associated inventory processes are integer-valued. Thus, the calculation of performance measures involves discrete instead of continuous mathematics. See Hadley and Whitin (1963), for example. When demand follows a compound-renewal process with a continuous batch-size distribution, as in Sahin (1979), for example, the state space does become continuous, but the sample paths of $D(t)$ itself are still piecewise constant with jumps, quite unlike the smoothly evolving world of the EOQ model. In the continuous realm the only model that has received careful attention is one where $D(t)$ itself is Brownian motion with positive drift. See Bather (1966) and Puterman (1975), for example. The problem here, of course, is that demand increments can be negative. Negative demands do occur in practice, but only rarely, and they substantially complicate the analysis.

Third, our model enjoys certain computational advantages over these more familiar stochastic models, as explained below.

There is one disadvantage to our model: Typically, an (r, q) policy is *not* optimal in this setting. We are assuming, in effect, that $x(t)$ is not observed, nor is information about it inferred from observing $D(t)$. Still (r, q) policies are simple and widely used, so it is worthwhile to study their performance.

The primary result of this article is the demonstration of a certain *insensitivity property*, which characterizes the limiting behavior of the model. This property drastically simplifies the computation of performance measures for the system.

Specifically, let $\mathbb{P} = \{\mathbb{P}(t): t \geq 0\}$ denote the inventory position. We will assume $r < \mathbb{P}(0) \leq r + q$. It will be convenient to work with a simple transformation of \mathbb{P} :

$$c(t) = (r + q) - \mathbb{P}(t).$$

So, $c(t)$ measures the distance between $\mathbb{P}(t)$ and its maximum value, and we can write

$$(1.2) \quad c(t) = [c(0) + D(t)] \pmod{q}.$$

The process $\mathbf{c} = \{c(t): t \geq 0\}$ has state space C , the real numbers \pmod{q} or the circle with circumference q . Also, let ν denote the uniform density on C .

Now, define the joint process $\mathbf{w} = (\mathbf{c}, \mathbf{x})$ with state space $W = C \times X$. Clearly, \mathbf{w} is a Markov process. We show that \mathbf{w} is well behaved in the same sense \mathbf{x} is. Specifically, \mathbf{w} has the unique stationary probability density $\rho = \nu\pi$, which also describes the limiting behavior of \mathbf{w} as $t \rightarrow \infty$. That is, let $w(\infty) = [c(\infty), x(\infty)]$ denote a random variable having density ρ . Then, $c(\infty)$ is *uniformly* distributed on C , and $c(\infty)$ and $x(\infty)$ are *independent*. The limiting behavior of \mathbf{w} is thus *insensitive* to the specification of \mathbf{x} and λ , except of course for the marginal density π of $x(\infty)$.

The meaning of this property can be elucidated in the following way: As \mathbf{w} evolves, $c(t)$ rotates around the circle C at rate $\lambda[x(t)]$; the *movement* of \mathbf{c} is determined by \mathbf{x} . Nevertheless, after a sufficiently long time, the *position* of $c(t)$ contains negligible information about $x(t)$. In other words, the future of the demand process becomes independent of $c(t)$, and hence of the inventory position.

Similar properties have been demonstrated for several models in which demand is a point process. See Galliher, Morse and Simond (1959) or Hadley and Whitin (1963) for the Poisson case, and Finch (1961), Sivazlian (1974), Sahin (1979, 1983, 1990) and Zipkin (1986a) for more general point processes. Comparable results can be obtained when demand includes both jumps (as in a point process) and continuous accumulation (as in our model), provided both depend on the history of the process only through \mathbf{x} . We shall not pursue such extensions here.

Incidentally, economists have become interested in this kind of result recently; such properties are helpful in describing the behavior of aggregate inventories at the economy-wide level. See Caplin (1985) and Mosser (1986), for example.

Given this characterization of \mathbf{w} , one can derive relatively simple formulas for the most important inventory performance measures. For example, suppose the order leadtime is a fixed constant, L , and let $\mathbb{L}(t)$ denote the *inventory level* (inventory minus backorders) at time t . Using standard arguments and the definition of $\mathbb{P}(t)$, one can show that

$$\mathbb{L}(t + L) = \mathbb{P}(t) - [D(t + L) - D(t)], \quad t \geq 0.$$

Now, if it makes any sense at all to take limits here, we should be able to write

$$\mathbb{L}(\infty) = \mathbb{P}(\infty) - D(L),$$

where $D(L)$ represents the demand during a leadtime under equilibrium conditions, in some sense. The question is, what precisely does $D(L)$ mean,

and what is its relationship to $\mathbb{P}(\infty)$? Our results imply that $\mathbb{L}(\infty)$ can indeed be characterized by this equation, where $D(L)$ is precisely the demand in the interval $(0, L]$ when \mathbf{x} is initialized with $x(0) \sim \pi$, $\mathbb{P}(\infty)$ is uniformly distributed on the interval $(r, r + q]$, and furthermore $D(L)$ and $\mathbb{P}(\infty)$ are independent. (Here and below, “ \sim ” means “is distributed as.”)

Now, let $B(\infty)$ denote the expected average outstanding backorders in equilibrium, and F the complementary cumulative distribution of $D(L)$. Given the equation above, and the fact that $B(\infty) = E[\mathbb{L}(\infty)]^-$, it is not hard to show that

$$(1.3) \quad B(\infty) = [\beta(r) - \beta(r + q)]/q,$$

where

$$\beta(y) = \int_y^\infty (x - y)F(x) dx.$$

The same formula applies when the leadtime is stochastic. Here, the leadtime demand refers to a mixture of the $D(L)$ over L ; specifically, the leadtime demand is the demand over a random interval of time, whose distribution is that of the leadtimes, with $x(0) \sim \pi$. [This extension requires some additional assumptions about how the leadtimes are generated. See Zipkin (1986a).]

Formula (1.3) with β and F defined as above has been used as an approximation for some time; see Hadley and Whitin (1963), for instance. To our knowledge, our model is the first for which this formula is exact.

As shown by Zipkin (1986b), $B(\infty)$ in (1.3) is a convex function of the policy parameters (r, q) for *any* complementary distribution F . Indeed, if the leadtime demand has a positive density on \mathcal{R}^+ , $B(\infty)$ is strictly convex for $r \geq 0$. The average inventory has the same properties, as does the frequency of orders. Thus, all the components of the standard average cost function are convex in (r, q) . To compute an optimal policy within the (r, q) class, therefore, one need only submit this cost function to any standard nonlinear-program solver.

The situation for discrete demand is markedly different. Amazingly, until quite recently there was *no* reliable, straightforward method for computing an optimal (r, q) policy, even in the simple case of Poisson demand. The first such algorithm, to our knowledge, was presented in Zipkin's (1987) class notes; this procedure is based on an approach developed by Sahin (1982). Federgruen and Zheng (1988) have since substantially refined and clarified the algorithm.

Still, this is a special-purpose algorithm. Using our continuous model, all the joys and sorrows of implementation and testing can be dispensed with. This is the computational advantage of our model mentioned above.

The process \mathbf{c} above is sometimes called a *clearing process*, and similar processes have been studied by Stidham (1974, 1977), Serfozo and Stidham (1978), Whitt (1981) and Schmidt (1986). In these papers the focus is on long-run frequency distributions (i.e., time averages) instead of limiting distributions. Several of them show, under various conditions, that \mathbf{c} is asymptoti-

cally uniform on C in this sense. (Generally, a slightly different definition of \mathbf{c} from (1.2) is used, and then in many cases \mathbf{c} is *not* uniform on C .)

Results like these concerning the behavior of \mathbf{c} itself, without broader results concerning \mathbf{w} or something similar, are insufficient to characterize the inventory system except in special circumstances. Basically, we need other reasons to treat the leadtime demand as independent of the inventory position. We do not need to worry about the leadtime demand, of course, when the leadtime is identically 0, but this condition severely restricts the scope of the model. Otherwise, we need to assume the demand process has stationary, independent increments. If we also wish $D(t)$ to be nondecreasing (and nonexplosive), then demand can only be a compound-Poisson process, while if we require continuous $D(t)$ we are left with Brownian motion; these too are quite special cases.

Apart from this qualification, the frequency approach and our distributional approach each have their own advantages. The frequency models include the EOQ model as a special case, whereas ours cannot; the resulting periodicity rules out a limiting distribution. Also, several of the papers cited above assume quite general ergodic demand processes without requiring the Markov property. On the other hand, distributional results are generally stronger than frequency results. That is, a distributional result often implies the corresponding frequency result, but not conversely. (However, see the discussion at the end of Section 4.)

Also, the frequency approach typically presumes that the demand process has stationary increments [which in our terms means $x(0) \sim \pi$], and that $c(0)$ is a fixed constant, and in particular independent of $x(0)$. This approach may seem natural, but it also masks a critical distinction among models concerning the significance of initial conditions. As we shall see, there are some models which behave well when initialized in this way, but otherwise they behave badly, and in particular \mathbf{c} is *not* asymptotically uniform in any sense.

This distinction is expressed below by a key irreducibility condition, Assumption 3.5. This condition means, essentially, that \mathbf{x} must be exogeneous to the inventory-control system in a specific sense. Only with this assumption can we ensure that initial conditions do not affect limiting behavior.

The rest of the article is organized as follows: Section 2 treats the special case where \mathbf{x} is a discrete-state, continuous-time Markov chain. This case requires much simpler assumptions and proofs than the general case.

The next two sections deal with the general case. The assumptions for our model are presented and discussed in Section 3. Section 4 proves the insensitivity property.

To use formula (1.3), we still need to compute the functions F and β . Section 5 shows how to do this for certain special cases of the model.

2. Insensitivity for continuous-time Markov chains. We first consider the special case where \mathbf{x} is a countable-state, continuous-time Markov chain. This model has interesting and important applications, and its analysis requires only elementary methods.

Here are the assumptions we need in this case:

ASSUMPTION 2.1. The chain \mathbf{x} is regular, that is, irreducible and positive-recurrent.

Therefore [see, e.g., Ross (1970)] \mathbf{x} has a unique stationary probability density π . Also, π is the limiting density for \mathbf{x} .

ASSUMPTION 2.2. Define $\lambda[\pi] = \sum_{x \in X} \pi(x)\lambda(x)$. Then $0 < \lambda[\pi] < \infty$.

This assumption merely prohibits explosive demand.

ASSUMPTION 2.3. There exist two positive-recurrent states x_1 and x_2 , such that $\lambda(x_1) \neq \lambda(x_2)$.

Thus, demand is “truly” stochastic.

We now proceed to the results. We prove two lemmas, followed by the main result of this section, Theorem 2.6.

LEMMA 2.4. *Suppose Assumptions 2.1 and 2.2 (but not necessarily 2.3) hold. Then the probability density ρ is stationary for \mathbf{w} .*

PROOF. Suppose $w(0) = [c(0), x(0)] \sim \rho$, and choose any $t > 0$. We may write

$$[c(t)|x(t)] = \{c(0) + [D(t)|x(t)]\} \pmod{q}.$$

Notice that $[D(t)|x(t)]$ is independent of $c(0)$, and by Assumption 2.2 it has a proper distribution. So, in the expression above, we have a uniformly distributed random variable $c(0)$, plus another, independent random variable, all mod(q). Any such combination also has a uniform distribution on C . [See, e.g., Feller (1971), page 64.] Thus, $[c(t)|x(t)] \sim \nu$, and since $x(t) \sim \pi$, we have $w(t) \sim \rho$. \square

Our next assertion speaks about the irreducibility of \mathbf{w} , among other properties. We have not said what this means for processes like \mathbf{w} , whose state spaces are not discrete. (A precise definition will come in the next section.) For now we shall use this concept loosely, to mean that the entire state space W is accessible (in an intuitive sense) from every starting point. Actually, the proof demonstrates that \mathbf{w} is irreducible in the precise sense.

LEMMA 2.5. *Given Assumptions 2.1–2.3, the process \mathbf{w} is regenerative and irreducible, and the regeneration cycle time has a nonlattice distribution.*

PROOF. We can choose as a regeneration state any $w = (c, x)$ having $\pi(x) > 0$ and $\lambda(x) > 0$. If we start with $w(0) = w$, then \mathbf{x} will stay at x for small t (with probability 1), so \mathbf{c} will immediately move away from c , hence \mathbf{w}

will move away from w . Now, \mathbf{x} may stay at x until $t = t[x] \equiv q/\lambda(x)$, when \mathbf{c} "comes around" again to c , in which case the cycle time is certainly finite. Let $p(x) > 0$ denote the probability of this event. Otherwise, \mathbf{x} will return again to x at some finite time, and on each such return there is a positive probability that \mathbf{x} will stay at x long enough for \mathbf{c} to come around to c . These probabilities, moreover, are all bounded below by $p(x)$. Thus, \mathbf{w} will return to w in finite time (with probability 1), and indeed the mean cycle time is finite. That is, \mathbf{w} is regenerative.

A similar argument shows that \mathbf{w} can reach any state $w' = (c', x')$ with $\pi(x') > 0$ and $\lambda(x') > 0$ from any starting state w in finite time. (Indeed, every such state *will* be reached in finite time with probability 1.) In case $\lambda(x') = 0$ [but still $\pi(x') > 0$], we have to be a bit more careful: Let W' be a subset of W with elements $w' = (c', x')$, where x' is fixed and the c' comprise some interval in C . Then any such subset W' can (and will) be reached in finite time from any starting state. Thus, \mathbf{w} is irreducible.

Now, let w be any regeneration state as above, and define

$$\begin{aligned} T_1 &= \text{time until } \mathbf{x} \text{ first leaves } x, \\ T_2 &= \text{time from } T_1 \text{ until } \mathbf{x} \text{ first returns to } x, \\ t_2 &= T_1 + T_2, \\ T_3 &= \text{time from } t_2 \text{ until } \mathbf{x} \text{ again leaves } x, \\ T &= \text{cycle time.} \end{aligned}$$

Condition on the event $E = \{T_1 < t[x] \text{ and } T_3 \geq t[x]\}$. This joint event has probability $[1 - p(x)]p(x) > 0$. In this case, $t_2 \leq T < t_2 + T_3$. Note that even conditional on E , T_1 and T_2 are independent, and both have densities, so t_2 also has a density.

Consider $(T|t_2, E)$. Conditional on t_2 (and E), the only remaining source of uncertainty in T is the demand during t_2 . This demand itself has a density, because of Assumption 2.3 and the remaining uncertainty in $(T_1|t_2, E)$. Thus, $(T|t_2, E)$ has a density. Now, when we decondition over t_2 (still conditioning on E), recalling that t_2 itself has a density, we find that $(T|E)$ has a density.

Finally, deconditioning on E , we conclude that T has a density on some interval. Specifically, T has a nonlattice distribution. \square

Lemmas 2.4 and 2.5, along with the basic theorems on regenerative processes [see, e.g., Kingman (1972)], immediately imply the insensitivity property:

THEOREM 2.6.

- (a) ρ is the unique stationary density for \mathbf{w} ;
- (b) ρ is the limiting density for \mathbf{w} ;
- (c) the long-run frequency distribution of \mathbf{w} has density ρ (with probability 1).

3. Assumptions and examples. We now turn to the general case where \mathbf{x} is a Markov process. In this section we set forth the assumptions we require of the model. For purposes of accessibility and motivation, we also review certain key elements of Markov-process theory, discuss our assumptions at some length and illustrate them with examples.

We first need to deal with some technical preliminaries: Our first assumption ensures that \mathbf{x} is well behaved in a rather mild sense and that the integral (1.1) is well defined; it is likely to be satisfied by all models of practical interest.

ASSUMPTION 3.1. The world \mathbf{x} is a time-homogeneous, continuous-time, Markov process, specifically \mathbf{x} is a *Feller–Dynkin* process in the sense of Williams (1979), pages 114–127. Also, the demand-rate function λ is measurable.

Let \mathcal{X} denote the sigma-field upon which each $x(t)$ is defined, the collection of subsets of X to which probabilities can be legitimately assigned. Assumption 3.1 implies, among other things, that \mathcal{X} is *separable* in the sense of Orey (1971), page 5. Let \mathcal{C} be the standard (Borel) sigma-field on C , and let \mathcal{W} be the sigma-field on W generated by $\mathcal{C} \times \mathcal{X}$. It is straightforward to verify, then, that \mathcal{W} too is separable in this sense. (This fact will be needed below.) In addition, with one more assumption below (Assumption 3.4), ensuring that $D(t)$ is finite, one can show that \mathbf{w} is also a Feller–Dynkin process.

A *measure* ϕ on (X, \mathcal{X}) is a nonnegative, countably additive set function defined on \mathcal{X} , and a *probability measure* is a measure ϕ with $\phi(X) = 1$. In this context we may not be able to speak of probability densities, so π will denote a probability measure on (X, \mathcal{X}) . Likewise, ν is the uniform probability measure on (C, \mathcal{C}) , and $\rho = \nu\pi$ is a probability measure on (W, \mathcal{W}) .

The *transition probability function* for \mathbf{x} is given by

$$P^t(x, Y) = \Pr\{x(t) \in Y | x(0) = x\}, \quad x \in X, Y \in \mathcal{X}, t \geq 0.$$

Similarly, for \mathbf{w} we shall denote

$$Q^t(w, Z) = \Pr\{w(t) \in Z | w(0) = w\}, \quad w \in W, Z \in \mathcal{W}, t \geq 0.$$

Note that Assumption 3.1 implies that, for fixed (x, Y) , P^t is a continuous function of t ; see Williams (1979), page 115. Likewise, Q^t is continuous in t . We shall use this fact below.

We shall need to consider the discrete-time process obtained by observing \mathbf{x} at equally spaced points in time. Specifically, using n as the discrete time index, for any $\Delta > 0$, define $x^\Delta(n) = x(n\Delta)$, and the process $\mathbf{x}^\Delta = \{x^\Delta(n) : n \geq 0\}$. Define \mathbf{w}^Δ similarly.

We now introduce concepts of irreducibility and recurrence for discrete-time processes of this sort; see Orey (1971), for example. Define the *accessibility function*

$$R^\Delta(x, Y) = \Pr\{x^\Delta(n) \in Y \text{ for some } n > 0 | x^\Delta(0) = x\}, \quad x \in X, Y \in \mathcal{X}.$$

For any measure ϕ on (X, \mathcal{X}) , we say the chain \mathbf{x}^Δ is ϕ -irreducible if, for all

$x \in X$ and all $Y \in \mathcal{Y}$ with $\phi(Y) > 0$, we have $R^\Delta(x, Y) > 0$. If, in addition, all such $R^\Delta(x, Y) = 1$, we say \mathbf{x}^Δ is ϕ -recurrent. We can extend these definitions to continuous-time processes in the following way: We shall say that \mathbf{x} is ϕ -irreducible if \mathbf{x}^Δ is ϕ -irreducible for all $\Delta > 0$. Likewise, \mathbf{x} is ϕ -recurrent if \mathbf{x}^Δ is ϕ -recurrent for all $\Delta > 0$.

ASSUMPTION 3.2. For some measure ϕ on (X, \mathcal{X}) the process \mathbf{x} is ϕ -recurrent. Also, \mathbf{x} has a stationary probability measure π .

Now, there is also a concept of periodicity in this context [again, see Orey (1971)], which we shall not state. We point out, however, that since \mathbf{x}^Δ is ϕ -irreducible for all Δ , every \mathbf{x}^Δ is also aperiodic.

We can now invoke the basic limit theorem for Markov chains [Orey (1971), pages 30–34] to conclude: (a) π is the unique stationary probability measure for \mathbf{x}^Δ ; (b) π is the limiting probability measure for \mathbf{x}^Δ (in the total-variation norm); and (c) \mathbf{x}^Δ is π -recurrent. As for the continuous process \mathbf{x} , results (a) and (c) carry over immediately. Furthermore, (b) above and the continuity of P^t imply that (b) holds for \mathbf{x} also; see Kingman (1963), for example. We record this conclusion as Lemma 3.3.

LEMMA 3.3.

- (a) π is the unique stationary probability measure for \mathbf{x} ;
- (b) π is the limiting probability measure for \mathbf{x} (in the total-variation norm);
- (c) \mathbf{x} is π -recurrent.

Before proceeding further, we pause to mention some examples. Suppose \mathbf{x} is a multivariate diffusion process. Specifically, $x(t)$ is a vector of k dimensions, and \mathbf{x} satisfies a stochastic differential equation of the form

$$(3.1) \quad dx = \mu(x) dt + \sigma(x) dw.$$

The parameters here are the k -vector $\mu(x)$ and the $k \times k$ matrix $\sigma(x)$. Also, w indicates a k -dimensional Wiener process. The equation is *linear* when $\mu(x) = Mx$ and $\sigma(x) = \sigma$ for constant $k \times k$ matrices M and σ . (This condition is sometimes referred to as linearity in the narrow sense; we shall use the term linear for short.) The equation is *nonsingular* when $\sigma(x)$ is nonsingular for all x .

The linear, nonsingular case is the most basic one, and it has been studied extensively. Here, irreducibility is immediate, because $[x(t)|x(0)]$ has a multivariate normal distribution with a full-rank covariance matrix. The conditions required for the existence of a stationary measure π are quite simple: All the eigenvalues of M must have negative real parts. In this case π corresponds to a normal distribution with mean vector 0; computation of the covariance matrix requires the solution of a matrix equation, in general. Given π , recurrence is also straightforward, so Assumption 3.2 is satisfied, and Lemma 3.3 applies. In the scalar case ($k = 1$) this condition reduces to $M < 0$; here, \mathbf{x}

is the classic Ornstein–Uhlenbeck process. See Arnold (1974), pages 133 and 134, for example.

For nonlinear equations (3.1) the scalar case is well understood. In this case irreducibility must be assumed, but the property is usually easy to recognize; the state space X cannot be divided into mutually inaccessible regions. Also, the conditions needed for a stationary measure π are fairly simple to verify; see, for example, Karlin and Taylor (1981).

We now continue with our assumptions. Our next assumption is the analog of Assumption 2.2.

ASSUMPTION 3.4. Define $\lambda[\pi] = \int_X \pi(dx) \lambda(x)$. Then $0 < \lambda[\pi] < \infty$.

Evidently, this is a growth condition on the function λ . For example, suppose \mathbf{x} satisfies a linear, nonsingular version of (3.1). Then Assumption 3.4 requires that the mean of λ with respect to a certain normal distribution be finite; this condition is usually quite easy to check.

Now, recall the accessibility function R^Δ above, and let S^Δ be the analog for \mathbf{w}^Δ . That is,

$$S^\Delta(w, Z) = \Pr\{w^\Delta(n) \in Z \text{ for some } n > 0 | w^\Delta(0) = w\}, \quad w \in W, Z \in \mathscr{Z}.$$

As above, \mathbf{w} is ϕ -irreducible when all $S^\Delta(w, Z) > 0$ whenever $\phi(Z) > 0$.

ASSUMPTION 3.5. The joint process \mathbf{w} is ρ -irreducible.

This simple-sounding condition plays a critical role in our analysis: First, it rules out the case of a constant demand-rate function λ , for in this case $c(t)$ is a deterministic, periodic function with period q/λ , so the discrete-time process \mathbf{w}^Δ with $\Delta = q/\lambda$ is *not* ρ -irreducible. Thus, the assumption requires that demand be “truly” stochastic, and so plays a role analogous to that of Assumption 2.3.

However, Assumption 3.5 also prohibits certain anomalies which can never arise in the discrete-state case. For example, suppose \mathbf{x} is a bivariate diffusion process $(\mathbf{x}_1, \mathbf{x}_2)$, where \mathbf{x}_2 happens to have the state space C and dynamics $d\mathbf{x}_2 = \lambda(\mathbf{x}) dt$. Thus, regardless of the behavior of \mathbf{x}_1 , \mathbf{x}_2 is essentially a copy of \mathbf{c} . If $c(0) = x_2(0)$, then $c(t) = x_2(t)$ for all t ; more generally, the difference between $c(0)$ and $x_2(0)$ will be preserved forever. Thus, \mathbf{w} cannot possibly be ρ -irreducible.

This example is pathological, of course, but the principle it illustrates applies to certain more realistic models: Researchers in marketing have developed models describing the introduction of new products, in which current demand depends on cumulative demand to date, among other factors. The idea is to model the possible saturation of the market for the product. See Mahajan and Wind (1986), for example. In our terms $D(t)$ would be one of the components of $x(t)$, so \mathbf{w} cannot be ρ -irreducible.

Certain variants of these new-product models, however, do satisfy our assumptions. For example, \mathbf{x} may include a component that is influenced by

the current demand rate, like \mathbf{x}_2 above, as long as there are other influences as well, such as decay or stochastic factors. Such constructs could be used to model markets with limited capacity to absorb *short-term* surges in demand, provided the market adjusts so that current demand is “forgotten” in the long run.

In sum, Assumption 3.5 requires that the world \mathbf{x} be *exogenous* to the inventory-control system in a rather precise sense. A similar assumption is employed by Zipkin (1986a) in the context of point-process demand for similar reasons. For the examples above where this assumption is violated, indeed, one can show that our main results in the next section fail.

On a positive note, Assumption 3.5 is satisfied for most plausible models of interest. For example, suppose \mathbf{x} satisfies a linear, nonsingular version of (3.1). (Note, nonsingularity rules out most of the badly behaved models above.) Also, suppose there are two compact subsets $X_1, X_2 \in \mathcal{X}$ with $\pi(X_1) > 0$ and $\pi(X_2) > 0$, such that $\lambda(x_1) < \lambda(x_2)$ for all $x_1 \in X_1, x_2 \in X_2$. Then one can show that Assumption 3.5 holds. The argument is similar to the proof of Lemma 2.5 above.

4. Insensitivity: The general case. In this section we prove the insensitivity property for a general state space X .

LEMMA 4.1. *The process \mathbf{w} is ρ -recurrent.*

PROOF. We need to show that, for every fixed $\Delta' > 0$, and for all $w_0 \in W$ and $Z' \in \mathcal{W}$ with $\rho(Z') > 0, S^{\Delta'}(w_0, Z') = 1$. To do this, we invoke a fundamental result about ϕ -irreducible Markov chains: Because \mathcal{W} is separable, $\mathbf{w}^{\Delta'}$ is ρ -irreducible, and $\rho(Z') > 0$, there exist a subset $Z'' \subseteq Z'$ with $\rho(Z'') > 0$, a constant $\gamma > 0$ and a positive integer m , such that, setting $\Delta = m\Delta'$,

$$(4.1) \quad Q^{\Delta}(w, Z) \geq \gamma\rho(Z) \quad \text{for all } w \in Z'' \text{ and } Z \in \mathcal{W}.$$

This result is due to Orey (1971); the version stated here can be found in Nummelin (1984), page 19.

Now, since $\rho(Z'') > 0$, we can find a subset $Z_1 \subseteq Z''$ of the form $Z_1 = B_1 \times A$, where $B_1 \in \mathcal{C}, A \in \mathcal{X}, \nu(B_1) > 0$ and $\pi(A) > 0$. Indeed, we can choose B_1 to be an interval of C of length q/K for some positive integer K , an interval closed at one end and open at the other end. Now, partition C into the K equal intervals $\{B_k: k = 1, \dots, K\}$, each of them semi-open like B_1 , and set $Z_k = B_k \times A$. By construction, (4.1) holds with Z_1 replacing Z'' , and by symmetry the same is true of each of the Z_k . In particular, for all $j, k = 1, \dots, K$,

$$(4.2) \quad Q^{\Delta}(w, Z_j) \geq \gamma\rho(Z_j) \quad \text{for all } w \in Z_k.$$

Consider the chain \mathbf{w}^{Δ} , and write $w_0 = (c_0, x_0)$. Because \mathbf{x}^{Δ} is π -recurrent, \mathbf{x}^{Δ} will reach A infinitely often starting from any $x_0 \in X$ (with probability 1); see the corollary in Orey (1971), page 22. Each time \mathbf{x}^{Δ} reaches A , \mathbf{w}^{Δ} reaches one of the Z_k , so *some* Z_k is reached infinitely often (with probability 1). Now,

(4.2) and Proposition 5.1 in Orey (1971) imply that *all* the Z_k are reached infinitely often (with probability 1). In particular, Z_1 , hence Z'' , hence Z' , are reached infinitely often by \mathbf{w}^Δ , hence also by \mathbf{w}^Δ . Thus, $S^\Delta(w_0, Z') = 1$. \square

THEOREM 4.2.

- (a) ρ is the unique stationary probability measure for \mathbf{w} ;
- (b) ρ is the limiting probability measure for \mathbf{w} (in the total-variation norm).

PROOF. Notice first that, with only minor modifications, the proof of Lemma 2.4 can be extended to the current model. Thus, ρ is a stationary probability measure for \mathbf{w} . This fact and Lemma 4.1 are all we need to invoke the limit theorem from Orey (1971), pages 30–34, used for Lemma 3.3 above. This result describes the limiting behavior of the \mathbf{w}^Δ , but it can be extended to the continuous-time process \mathbf{w} , as described just before Lemma 3.3. \square

Notice that Theorem 4.2 does *not* include the analog of Theorem 2.6(c) concerning limiting frequencies. This would require showing that \mathbf{w} , initialized with $w(0) \sim \rho$, is an ergodic process. We expect this is true for most cases of interest, but we are unaware of results that would allow us to prove it in general.

Specifically, it would suffice to show that \mathbf{w} is regenerative in the extended sense developed by Athreya, McDonald and Ney (1978). Their results apply to continuous-time Markov chains with general state spaces. We expect this approach can be extended to more general processes, including \mathbf{w} , but such a theory has yet to be fully developed.

5. Computation of functions describing leadtime demand. In this section we describe how to compute the functions F and β describing the leadtime demand for an important special case of the model. Specifically, we assume \mathbf{x} is a continuous-time Markov chain, as in Section 2, with a *finite* state space X . Let the matrix Q denote the infinitesimal generator of \mathbf{x} . Here, the stationary density π can be viewed as a row vector with $\pi Q = 0$. In addition, we suppose all $\lambda(x) > 0$.

Also, we assume the leadtime L has a *phase-type distribution*. This means that there is another finite-state, continuous-time Markov chain \mathbf{j} with a single absorbing state, where \mathbf{j} is independent of \mathbf{x} , such that L is the time until \mathbf{j} reaches its absorbing state. The data describing \mathbf{j} and L comprise the pair (α, A) , where α is a row vector and A a square matrix. The vector α is substochastic and gives the initial probabilities of \mathbf{j} for the nonabsorbing, transient states. The matrix A is the generator of \mathbf{j} , restricted to the transient states; see Neuts (1981).

We shall use I to denote an identity matrix and e a column vector of 1's. The dimensions will be clear from the context. We assume that $\alpha e = 1$. This means $L > 0$ with probability 1.

Operationally, the assumption above on L means the following: Let $G(t)$ be the complementary cumulative distribution of L . We can compute G in the following manner: Let $g(t)$ be a row vector which satisfies the system of ordinary linear differential equations

$$(5.1) \quad \begin{aligned} g(0) &= \alpha, \\ g_t &= gA. \end{aligned}$$

(The subscript t here denotes differentiation.) Thus, $g(t)$ is the probability density of $j(t)$, restricted to the transient states, so

$$(5.2) \quad G(t) = g(t)e = \alpha[\exp(At)]e.$$

Our main result is that *the leadtime demand also has a phase-type distribution*, so F can be computed by solving differential equations of the form (5.1), followed by a sum like (5.2). Furthermore, β can also be computed by solving a similar system of equations with different initial conditions. We remark that this result can also be obtained (with a fair amount of effort) using the transform methods developed by Puri (1972).

We now need some additional notation: Let \mathbf{u} denote the *leadtime demand process*; before L , \mathbf{u} behaves just like \mathbf{D} , but after L , \mathbf{u} stays constant at $u(t) = u(L) = D(L)$. Also, let Λ denote the diagonal matrix with diagonal entries $\lambda(x)$. We assume familiarity with the Kronecker matrix operations \otimes and \oplus . In this notation the generator of the joint process (\mathbf{x}, \mathbf{j}) , restricted to the transient states of \mathbf{j} , can be written $Q \oplus A = Q \otimes I + I \otimes A$.

Define the probability density function

$$P_{xjyk}(u, t) = \Pr\{x(t) = y, j(t) = k, u(t) \in [u, u + du] \mid x(0) = x, j(0) = j, u(0) = 0\}.$$

Also, define the matrix $P(u, t) = [p_{xjyk}(u, t)]$. Here, x and y range over all of X , but j and k range only over the transient states of \mathbf{j} . Then P satisfies the matrix partial differential equation

$$(5.3) \quad \begin{aligned} P(0, 0) &= I \otimes I, \\ P(u, 0) &= 0 \otimes 0, \quad u > 0, \\ P_t &= P(Q \oplus A) - P_u(\Lambda \otimes I), \quad t > 0. \end{aligned}$$

We shall also need the matrix $M(u, t)$ of the same dimensions as P , defined dynamically by

$$(5.4) \quad \begin{aligned} M(u, 0) &= 0 \otimes 0, \\ M_t &= -P(I \otimes A), \quad t \geq 0. \end{aligned}$$

In addition define the matrices

$$M(u) = \lim_{t \rightarrow \infty} M(u, t), \quad N(u) = \int_u^\infty M(v) dv.$$

We argue next that

$$(5.5) \quad F(u) = (\pi \otimes \alpha)N(u)(e \otimes e).$$

To see this, consider the matrix $M(u, t)(I \otimes e)$. From (5.4) we have

$$M_t(I \otimes e) = P(I \otimes -Ae).$$

From the definition of A we have $-Ae \geq 0$, so the right-hand side here is nonnegative, hence so is $M(u, t)(I \otimes e)$. Indeed, the k th element of the vector $-Ae$ is the rate at which \mathbf{j} jumps to its absorbing state, given $j(t) = k$. Thus, $M(u, t)(I \otimes e)$ is the accumulated joint probability density of $[x(L), u(L)]$ and the event $L \leq t$, given the starting conditions. Taking the limit as $t \rightarrow \infty$, $M(u)(I \otimes e)$ is just the density of $[x(L), u(L)]$, given the starting conditions. Premultiplying this matrix by $(\pi \otimes \alpha)$ simply weights the starting states by their actual probabilities, and postmultiplying by $(e \otimes I)$ sums over the final values of $x(L)$, so $(\pi \otimes \alpha)M(u)(e \otimes e)$ gives the density of $u(L)$, which is precisely the leadtime demand. Finally, integrating this last quantity over v yields (5.5).

In view of (5.5), our remaining task is to compute $N(u)$. Define the matrix

$$H = (Q \oplus A)(\Lambda^{-1} \otimes I).$$

We claim that $N(u)$ solves the ordinary matrix differential equation

$$(5.6) \quad \begin{aligned} N(0) &= (Q \oplus A)^{-1}(I \otimes A), \\ N_u &= NH. \end{aligned}$$

Thus, we may write

$$(5.7) \quad N(u) = (Q \oplus A)^{-1}(I \otimes A) \exp(Hu).$$

To verify (5.6), first notice that

$$M(u) = \int_0^\infty M_t(u, t) dt = -P(u)(I \otimes A), \quad u \geq 0,$$

where $P(u) = \int_0^\infty P(u, t) dt$. Also, for $u > 0$,

$$0 = \int_0^\infty P_t(u, t) dt = P(u)(Q \oplus A) - P_u(u)(\Lambda \otimes I).$$

Since $(I \otimes A)$ commutes with both $(Q \oplus A)$ and $(\Lambda \otimes I)$, we can write this as

$$-P_u(u)(I \otimes A)(\Lambda \otimes I) = -P(u)(I \otimes A)(Q \oplus A)$$

or

$$M_u(u) = M(u)H.$$

Integrating both sides of this equation yields the second part of (5.6). To evaluate $N(0)$, define

$$K(t) = \int_0^\infty P(u, t) du.$$

By the definition of P , $K(t) = \exp[(Q \oplus A)t]$, so

$$N(0) = \int_0^\infty M(u) du = - \int_0^\infty K(t) dt (I \otimes A) = (Q \oplus A)^{-1} (I \otimes A).$$

Thus, (5.6) holds.

When we combine (5.5) and (5.7), we obtain a further simplification: We can write

$$N(0) = [(I \otimes I) - (Q \otimes - A^{-1})]^{-1} = \sum_{n=0}^{\infty} (Q \otimes - A^{-1}),$$

so, since $\pi Q = 0$, $(\pi \otimes \alpha)N(0) = (\pi \otimes \alpha)$. Finally, we obtain

$$F(u) = (\pi \otimes \alpha) \exp(Hu)(e \otimes e).$$

Thus, F has precisely the same form as (5.2); that is, the leadtime demand does have a phase-type distribution, as claimed. Specifically, we can compute F by solving the differential equations

$$(5.8) \quad \begin{aligned} f(0) &= \pi \otimes \alpha, \\ f_t &= fH, \end{aligned}$$

and setting $F(t) = f(t)(e \otimes e)$.

Also, from the definition of β we immediately obtain

$$\beta(u) = (\pi \otimes \alpha) H^{-2} \exp(Hu)(e \otimes e).$$

Thus, β can also be computed through an equation of the form of (5.8), using the alternative initial condition $f(0) = (\pi \otimes \alpha)H^{-2}$.

These methods thus allow us to evaluate any (r, q) policy relatively easily.

We remark that, while this result applies only to the special model considered here, similar methods lead to analogous results for other models. For example, suppose \mathbf{x} is a diffusion process and L is the time until absorption of some other diffusion. Then F and β can be computed by solving a linear *partial* differential equation; the form of this equation is analogous to those derived above.

REFERENCES

- ARNOLD, L. (1974). *Stochastic Differential Equations*. Wiley, New York.
- ATHREYA, K., McDONALD, D. and NEY, P. (1978). Limit theorems for semi-Markov processes and renewal theory for Markov chains. *Ann. Probab.* **6** 788–797.
- BATHER, J. (1966). A continuous time inventory model. *J. Appl. Probab.* **3** 538–549.
- CAPLIN, A. (1985). The variability of aggregate demand with (S, s) inventory policies. *Econometrica* **53** 1395–1409.

- FEDERGRUEN, A. and ZHENG, Y. (1988). A simple and efficient algorithm for computing optimal (r, Q) policies in continuous-review stochastic inventory systems. Working Paper, Graduate School of Business, Columbia Univ.
- FELLER, W. (1971). *An Introduction to Probability Theory and Its Applications 2*, 2nd ed. Wiley, New York.
- FINCH, P. (1961). Some probability theorems in inventory control. *Publ. Math. Debrecen* **8** 241–261.
- GALLIHER, H., MORSE, P. and SIMOND, M. (1959). Dynamics of two classes of continuous-review inventory systems. *Oper. Res.* **7** 362–384.
- HADLEY, G. and WHITIN, T. (1963). *Analysis of Inventory Systems*. Prentice-Hall, Englewood Cliffs, N.J.
- KARLIN, S. and TAYLOR, H. (1981). *A Second Course in Stochastic Processes*. Academic, New York.
- KINGMAN, J. (1963). Ergodic properties of continuous time Markov processes and their discrete skeletons. *Proc. London Math. Soc.* **13** 593–604.
- KINGMAN, J. (1972). *Regenerative Phenomena*. Wiley, New York.
- MAHAJAN, V. and WIND, Y. (1986). *Innovation Diffusion Models for New Product Acceptance*. Ballinger, Cambridge, Mass.
- MALLIARIS, A. and BROCK, W. (1982). *Stochastic Methods in Economics and Finance*. North-Holland, Amsterdam.
- MOSSER, P. (1986). Trade inventories and (S, s) . Working Paper, Dept. Economics, Columbia Univ.
- NEUTS, M. (1981). *Matrix-Geometric Solutions in Stochastic Models*. Johns Hopkins Univ. Press.
- NUMMELIN, E. (1984). *General Irreducible Markov Chains and Non-Negative Operators*. Cambridge Univ. Press.
- OREY, S. (1971). *Limit Theorems for Markov Chain Transition Probabilities*. Van Nostrand Reinhold, New York.
- PURI, P. (1972). A method for studying the integral functionals of stochastic processes with applications. In *Proc. Sixth Berkeley Symp. Math. Statist. Probab.* 481–500. Univ. California Press, Berkeley.
- PUTERMAN, M. (1975). A diffusion process model for a storage system. In *Studies in the Management Sciences. I. Logistics* (M. Geisler, ed.) 143–159. North-Holland, Amsterdam.
- ROSS, S. (1970). *Applied Probability Models with Optimization Applications*. Holden-Day, San Francisco.
- SAHIN, I. (1979). On the stationary analysis of continuous review (s, S) inventory systems with constant lead times. *Oper. Res.* **27** 717–729.
- SAHIN, I. (1982). On the objective function behavior in (s, S) inventory models. *Oper. Res.* **30** 709–724.
- SAHIN, I. (1983). On the continuous-review (s, S) inventory model under compound renewal demand and random lead times. *J. Appl. Probab.* **20** 213–219.
- SAHIN, I. (1990). *Regenerative Inventory Systems*. Springer, New York.
- SCHMIDT, V. (1986). On the asymptotic uniform distribution of stochastic clearing processes. *Math. Operationsforsch. Statist. Ser. Optim.* **17** 125–134.
- SERFOZO, R. and STIDHAM, S. (1978). Semi-stationary clearing processes. *Stochastic Process. Appl.* **6** 165–178.
- SETHI, S. and THOMPSON, G. (1981). *Optimal Control Theory: Applications to Management Science*. Martinus Nijhoff, Boston.
- SIVAZLIAN, B. (1974). A continuous-review (s, S) inventory system with arbitrary interarrival distribution between unit demand. *Oper. Res.* **22** 65–71.
- STIDHAM, S. (1974). Stochastic clearing systems. *Stochastic Process. Appl.* **2** 85–113.
- STIDHAM, S. (1977). Cost models for stochastic clearing systems. *Oper. Res.* **25** 100–127.
- WHITT, W. (1981). The stationary distribution of a stochastic clearing process. *Oper. Res.* **29** 294–308.
- WILLIAMS, D. (1979). *Diffusions, Markov Processes, and Martingales 1*. Wiley, New York.

- ZIPKIN, P. (1986a). Stochastic leadtimes in continuous-time inventory models. *Naval Res. Logist. Quart.* **33** 763–774.
- ZIPKIN, P. (1986b). Inventory service-level measures: Convexity and approximation. *Management Sci.* **32** 975–981.
- ZIPKIN, P. (1987). *Lecture Notes on Inventory Theory*, part 5, version 1. Graduate School of Business, Columbia Univ.

GRADUATE SCHOOL OF BUSINESS
COLUMBIA UNIVERSITY
NEW YORK, NEW YORK 10027