

Dynamic control of a queue with adjustable service rate

Jennifer M. George
Melbourne Business School
200 Leicester St, Carlton
VIC 3053, Australia
Phone: (613) 9349-8145
Fax: (613) 9349-8133
Email: j.george@mbs.unimelb.edu.au

J. Michael Harrison
Graduate School of Business
Stanford University, Stanford
CA 94305, USA
Phone: (650) 723-4727
Fax:
Email: harrison_michael@gsb.stanford.edu

Dynamic control of a queue with adjustable service rate

May 23, 2000

Abstract

We consider a single-server queue with Poisson arrivals, where holding costs are continuously incurred as a non-decreasing function of the queue length. The queue length evolves as a birth-and-death process with constant arrival rate $\lambda = 1$ and with state-dependent service rates μ_n that can be chosen from a fixed subset A of $[0, \infty)$. Finally, there is a non-decreasing cost-of-effort function $c(\cdot)$ on A , and service costs are incurred at rate $c(\mu_n)$ when the queue length is n . The objective is to minimize average cost per time unit over an infinite planning horizon. The standard optimality equation of average-cost dynamic programming allows one to write out the optimal service rates in terms of the minimum achievable average cost z^* . Here we present a method for computing z^* which is so fast and so transparent that it may reasonably be described as an explicit solution for the problem of service rate control. The optimal service rates are non-decreasing as a function of queue length, and are bounded if the holding cost function is bounded. From a managerial standpoint it is natural to compare z^* , the minimum average cost achievable with state-dependent service rates, against the minimum average cost achievable with a single fixed service rate. The difference between those two minima represents the economic value of a responsive service mechanism, and numerical examples are presented which show that it can be substantial.

1 Introduction and Summary

This paper is concerned with dynamic control of the service rate in a single-server queuing system with Poisson arrivals and exponentially distributed service times. The objective is to minimize average cost per time unit over an infinite planning horizon, where cost has two elements: holding cost (or congestion cost) that increases with queue length, and a cost of effort that increases with the service level chosen. Many papers have been written on different versions of this problem over the last thirty years, dealing with both characterization and computation of optimal policies. Here we develop a new method for computing optimal policies, which has a number of important virtues: (a) its applicability does not depend on any extraneous technical assumptions about the problem data; (b) it proceeds by solving a sequence of approximating problems that are natural and interesting in their own right, each involving a truncation of the holding cost function; (c) the optimal policies for the approximating problems converge monotonically to a policy that is optimal under the original cost structure; (d) at each stage in the computation one has an implementable policy and a bound on its performance, relative to an optimal policy, under the original cost structure; and (e) our method appears to be more efficient than any proposed in earlier work, although data are not available for direct comparisons. In fact, the computations are so fast and so transparent that our approach can reasonably be said to provide an explicit solution for the problem of service rate control.

Throughout this paper the discrete units that flow through the queueing system will be referred to as “jobs” rather than “customers”. As is customary, we use the term “queue

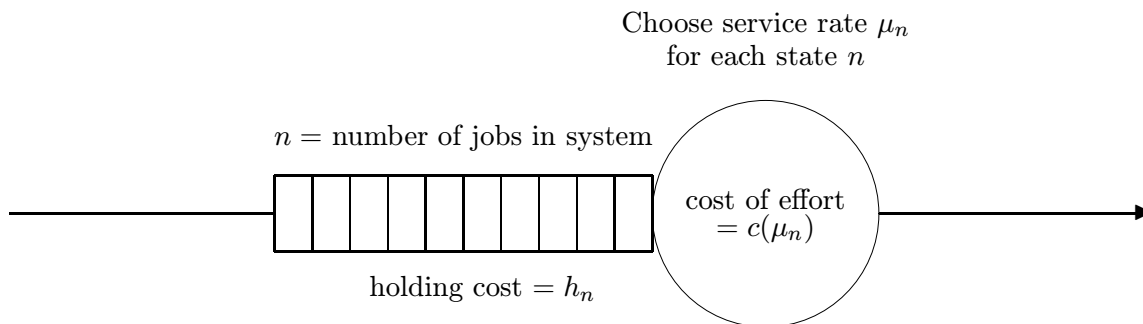


Figure 1: A Queueing System with Adjustable Service Rate

length” to mean the number of jobs in the system, including the job being served if there is one. The queue length evolves as a birth-and-death process with constant arrival rate $\lambda = 1$ and with state-dependent service rates μ_n that can be chosen from a closed subset A of $[0, \infty)$. It is assumed that A contains both the point $x = 0$ and some $x > 1$ (recall that $\lambda = 1$ by convention). Also given is a function $c(\cdot)$ on A , where $c(x)$ is a cost rate associated with service rate x . Imagining that a service rate x reflects or represents a level of effort by the server, we shall often refer to $c(x)$ as an effort cost. We assume that $c(\cdot)$ is non-decreasing and left-continuous with $c(0) = 0$. (The last of those assumptions is just a matter of convention, or a convenient normalization.) Also, if A is unbounded we further require that

$$\inf \{x^{-1}c(x) : x \in A, x \geq y\} \uparrow \infty \text{ as } y \uparrow \infty. \quad (1)$$

Despite its technical appearance, this assumption is substantive and indispensable. That assertion will be explained carefully in an appendix, but a quick summary of the argument is the following. If the infimum in (1) were to approach a finite limit α as $y \uparrow \infty$, then the

server would effectively be able to eject jobs from the system instantaneously, at a cost of α per job ejected. To avoid nonsensical conclusions in that case, one must adopt an alternative model formulation in which the ejection capability is explicitly recognized as a second mode of control.

As a final model element, we suppose that holding costs (or congestion costs) are continuously incurred at rate h_n when the queue length is n . The vector $h = (h_0, h_1, \dots)$ is assumed to be non-decreasing and to have less-than-geometric growth, as follows:

$$\sum_{n=0}^{\infty} h_n \theta^n < \infty \text{ for all } \theta \in [0, 1]. \quad (2)$$

This is a significant restriction from a mathematical standpoint, and it could be substantially weakened without changing any of our basic results (the arguments that would need to be modified come in section 7), but it makes for a clean theoretical development and is still quite mild from a practical standpoint. In particular, it is satisfied if one has a polynomial bound on h_n .

For us a *policy* is a vector $\mu = (\mu_1, \mu_2, \dots)$ with all components belonging to the set A . One interprets μ_n as the service rate to be used when the queue length is n . The problem is to choose a policy μ which minimizes the associated long-run average cost rate, including both holding costs and the cost of effort. Thus we are considering a Markov decision process with continuous time parameter, countable state space, time-invariant data and a long-run average cost criterion. In the language of dynamic programming, we are restricting attention to stationary Markov policies.

As stated earlier, we shall develop an explicit solution for this problem, imposing no as-

sumptions beyond the ones already set forth. The mathematical treatment is self-contained, making no use of general dynamic programming theory except to motivate the basic optimality equation (see section 3). The optimal policy that we obtain is monotone, meaning that the optimal service rate increases as a function of queue length, which is consistent with known results that will be reviewed shortly.

The paper is structured as follows. Section 2 gives a precise statement of the mathematical problem to be solved, and section 3 summarizes what is known about the problem from past work. The latter discussion includes a statement of the optimality equation (or Hamilton-Jacobi-Bellman equation) that is the starting point for our analysis, and in section 3 our approach to its solution is also described in broad outline. A reduced form of the optimality equation is derived in section 4, and the sufficiency of that equation for optimality of a given policy is rigorously proved. After some technical preliminaries have been dispensed with in section 5, we show in section 6 how to solve the optimality equation when holding costs are truncated, and then the method of approximation by means of successive truncations is developed in section 7, along with error bounds on the approximations and a rigorous proof of monotone convergence. Section 8 presents a family of numerical examples with quadratic cost of effort and holding costs of the form $h_n = s(n - M + 1)^+$, where s and M are positive constants. In discussing those examples we compare the minimum cost achievable with state-dependent service rates against the minimum achievable with a single service rate. The difference between those two minima represents the economic value of a responsive service mechanism, and our examples show that it can be substantial. Finally, we explain in the appendix how our problem formulation must be modified if one wants to

consider cost-of-effort functions $c(\cdot)$ for which assumption (1) fails.

2 Problem Formulation

Because our dynamic control problem has a countably infinite state space, an action set A which is only required to be closed, and potentially unbounded costs, it does not fit neatly within any standard theory of Markov decision processes. For example, no general result can be invoked to assure the existence of an optimal policy, cf. section 7.1. of [7], nor to assure that solutions of the Hamilton-Jacobi-Bellman equation (see section 3) correspond to optimal policies. Also, the standard technique of uniformization (cf. chapter 11 of [5] or chapter 5 of [1]) is not generally applicable to our model, because the set A of potential service rates may be unbounded and thus there is no positive lower bound on the expected time between state transitions, independent of the service rate chosen.

Thus we shall analyze the problem of service rate control “from first principles”. The following definitions, which rely on readers’ familiarity with the theory of birth and death processes, are intended to facilitate a streamlined treatment that is still mathematically rigorous.

A *policy* is simply defined as a vector $\mu = (\mu_1, \mu_2, \dots)$ with $\mu_n \in A$ for all n . A policy μ is said to be *ergodic* if there exists a probability distribution $p = (p_0, p_1, \dots)$ satisfying the balance equations (recall that $\lambda = 1$ by convention)

$$p_n = p_{n+1}\mu_{n+1} \quad \text{for all } n \geq 0. \quad (3)$$

The ergodic or stationary distribution p is known to be unique if it exists: if $\mu_n > 0$ for all

n then one has

$$p_n = p_0 \prod_{i=1}^n \mu_i^{-1} \quad \text{for } n \geq 0, \quad (4)$$

where p_0 is the appropriate normalization constant; if some service rates are zero but there exists a largest state N with $\mu_N = 0$, then one has $p_0 = \dots = p_{N-1} = 0$, and the remaining elements of p are given by the obvious modification of (4); finally, if $\mu_n = 0$ for infinitely many states n , then the policy μ cannot be ergodic. The stationary distribution p associated with an ergodic policy μ will hereafter be denoted by $p(\mu) = (p_0(\mu), p_1(\mu), \dots)$. The long-run average cost rate, or objective value, associated with an ergodic policy μ is

$$z(\mu) = \sum_{n=0}^{\infty} p_n(\mu) [h_n + c(\mu_n)]. \quad (5)$$

Because h is bounded below and $c(\cdot) \geq 0$, the quantity $z(\mu)$ is well defined but may be infinite. Assumption (2) guarantees the existence of an ergodic policy μ with $z(\mu) < \infty$, because one can simply take $\mu_n = x$ for all $n \geq 1$, where $x > 1$ and $x \in A$: the stationary distribution $p(\mu)$ is then geometric with parameter $\theta = x^{-1}$, so (2) implies that $z(\mu) < \infty$.

Now define

$$z^* = \inf z(\mu),$$

where the infimum is taken over ergodic policies (obviously, $z^* < \infty$). An ergodic policy μ is said to be *optimal* if $z(\mu) = z^*$.

In the case of bounded holding costs, where $h_n \uparrow h_\infty < \infty$, one encounters the following possibility. It may be that z^* , which is defined as an infimum over *ergodic* policies, is larger than h_∞ , which is achievable as the long-run average cost rate under the non-ergodic *do-nothing policy* $\mu = (0, 0, \dots)$. We shall say that our dynamic control problem is *degenerate*

if $z^* \geq h_\infty$, taking the view that further analysis of degenerate problems is uninteresting. That is, once a problem has been found to be degenerate, we implicitly declare the analysis to be complete, although there are other questions that could conceivably be investigated, some of which are quite subtle. To the best of our knowledge, the case with bounded holding costs has not been examined in any previous work, and thus the phenomenon of degeneracy has not previously been considered.

3 Literature Review

There are two streams of research that are relevant to the analysis undertaken in this paper, one concerned with characterization of optimal policies and the other with computation. To the best of our knowledge, the most general results in the former stream are those of Stidham and Weber [8], while the current state of knowledge with regard to computation of optimal policies is represented by the work of Wijngaard and Stidham [10], Jo [4], and Sennet [7]. In this section we shall summarize their results, simultaneously laying some groundwork for later analysis; readers interested in the historical development of the subject may consult the bibliographies of the works discussed here, especially [7].

With regard to characterization of optimal policies, there are several levels of analysis to be considered, and like the authors named above, we shall discuss these only in the context of the semi-Markov decision process (SMDP) that is obtained when one constrains the decision maker to maintain a constant service rate between times when the queue length changes. That is, we allow the decision maker to choose a new service rate whenever a new job arrives

or a service is completed, but not at any other time. Given the memoryless property of the exponential distributions that underlie our model, and the infinite planning horizon, it is more or less obvious that no advantage can be gained from more frequent changes of service rates.

In our context a “stationary” policy is one that chooses the same service rate μ_n whenever a transition to state n occurs. Under relatively general conditions, Stidham and Weber [8] prove that there exists a stationary policy that is optimal within the larger class of potentially non-stationary policies, making heavy use of their own results in an earlier paper [9]. In proving the existence of a stationary optimal policy for the problem of service rate control, Stidham and Weber [8] impose two assumptions that do not necessarily hold in our model: the set A must be bounded, so that there exists a largest available service rate $\bar{\mu}$, and the holding cost vector h must be unbounded. However, neither of those assumptions is essential, and presumably the Stidham-Weber analysis can be extended to justify our restriction to stationary policies, but we shall not do so.

The most important result obtained by Stidham and Weber [8] is their elegant proof that there exists a monotone optimal policy for the problem of service rate control; that is, they prove the existence of an optimal policy in which the service rate increases as a function of queue length. Several authors had proved this under stronger assumptions, beginning with Crabill [2, 3]. Although the conditions imposed by Stidham and Weber [8] are extremely weak by historical standards, they are still stronger than our assumptions, and the existence of a monotone optimal policy will be obtained as a by-product of our explicit computations.

In the approach developed by Wijngaard and Stidham [10] for the computation of optimal

policies, one begins with the standard optimality equation, or Hamilton-Jacobi-Bellman equation, for a semi-Markov decision process with average-cost criterion, cf. page 268 of Bertsekas [1] or page 554 of Puterman [5]. For the problem considered in this paper (recall that $\lambda = 1$ by convention), the optimality equation can be written as follows:

$$v_n = \inf_{x \in A} \left\{ \left(\frac{1}{1+x} \right) [c(x) + h_n - z] + \left(\frac{x}{1+x} \right) v_{n-1} + \left(\frac{1}{1+x} \right) v_{n+1} \right\} \quad \text{for } n \geq 1 \quad (6)$$

and

$$v_0 = (h_0 - z) + v_1. \quad (7)$$

Here one interprets z as a guess at the minimum average cost rate, or objective value, that was denoted by z^* in section 2. One interprets v_n as the minimum expected cost incurred until the next entry into an arbitrary reference state $m \geq 0$, starting in state n , under the z -revised cost structure that charges holding cost at rate $h_i - z$ in state $i \geq 0$. This interpretation depends on the fact that $1/(1+x)$ represents the expected time until the next state change when service rate x is chosen in any state $n \geq 1$; also, $x/(1+x)$ is the probability that the transition is to state $n-1$, and $1/(1+x)$ is the probability it is to state $n+1$.

The vector of unknowns $v = (v_0, v_1, \dots)$ is often called a *relative cost function* in average-cost dynamic programming, and one observes that even if z is treated as a known constant, the relative costs are only determined up to an additive constant by (6) and (7). Thus it is natural to define the relative cost differences

$$y_n = v_n - v_{n-1} \quad \text{for } n = 1, 2, \dots,$$

and then re-express (6)-(7) as follows:

$$0 = \sup_{x \in A} \left\{ \left(\frac{1}{1+x} \right) [z - h_n - c(x)] + \left(\frac{x}{1+x} \right) y_n - \left(\frac{1}{1+x} \right) y_{n+1} \right\} \quad \text{for } n \geq 1, \quad (8)$$

and

$$y_1 = z - h_1. \quad (9)$$

Proceeding as in Wijngaard and Stidham [10], one notes that (8) holds if and only if the quantity in braces is non-positive for all $x \in A$, with equality for some $x \in A$. Multiplying through by $(1+x)$ and rearranging terms, we can then re-express (8) as

$$y_{n+1} = \sup_{x \in A} \{ z - h_n - c(x) + y_n x \} \quad \text{for } n \geq 1 \quad (10)$$

Together, (9) and (10) constitute the analog for our problem of the transformed optimality equation (2) in [10]. It should be emphasized that the derivation sketched above serves only as motivation in our treatment of the service rate optimization problem; the only property of this optimality equation that we actually require (Proposition 1 in section 4) will be proved from first principles.

Given a value for z , the corresponding value of y_1 is obviously determined by (9), and then the values of y_2, y_3, \dots are recursively determined by (10). Thus one is led to the following question: what is the auxiliary condition that distinguishes the optimal objective value z^* ? Wijngaard and Stidham [10] assume that costs in their model satisfy a condition of less-than-geometric growth, which is precisely analogous to our assumption (2), and then observe that all trial values of z other than z^* cause the computed values of y_1, y_2, \dots to either decrease too quickly (if $z < z^*$) or else increase too quickly (if $z > z^*$). This leads them to a method

for successively refining an initial estimate of z^* , and they prove that the refined estimates do indeed converge to the optimal objective value z^* . However, that proof requires a number of extra assumptions that are not necessarily satisfied in our model, most notable being an assumption of “uniform tendency to the left” for large states.

Our computational method is based on the following key observation: if $h_n = h_{n+1} = \dots$ for some $n \geq 1$, then the required auxiliary condition for z is simply that $y_{n+1} = y_n$ (see section 5). There is only one value of z satisfying this auxiliary condition, and its computation is virtually trivial; the complete vector y of relative cost differences has $y_i = y_n$ for *all* $i \geq n$, and the optimal policy has $\mu_i = \mu_n$ for all $i \geq n$. By considering a sequence of truncated holding cost functions one generates a monotone sequence of approximations for both the optimal objective value and the optimal control policy. This method requires no extra assumptions for its justification, the computations are lightning fast, all intermediate quantities have ready interpretations, and the associated performance bounds are sharper than those obtained with the Wijngaard-Stidham method. On the other hand, our approach is tailored to a single problem, exploiting all of its special structure, while theirs is applicable to a large class of problems whose transition structure is skip-free-to-the-right.

The computational approach propounded by Jo [4] does not begin with the Hamilton-Jacobi-Bellman equation. Rather, he works directly with the algebraic expressions (3)-(5) that define the average cost $z(\mu)$ for a stationary policy μ , and then uses optimization theory to write out necessary conditions for the optimality of a given policy. He assumes both convex holding costs and existence of a largest service rate $\bar{\mu} \in A$, observing that in this case there exists a threshold level N such that the optimal service rate is $\mu_n = \bar{\mu}$ for all states $n \geq N$.

His computational method, which proceeds by generating a sequence of paired estimates for the optimal objective value z^* and the threshold level N , is close to ours in spirit, but it requires extra assumptions and is more complicated. Also, the justification of optimality is incomplete (the necessary conditions are never shown to be sufficient), and prior results are cited incorrectly. For example, Jo states on page 433 that “convexity of the holding costs is the weakest possible condition to achieve monotonicity of the optimal service rates;” Stidham and Weber [8] observe on page 611 that convexity *is* needed with a discounted cost criterion, but this assumption is superfluous in the average-cost case.

The recent book by Sennet [7] considers a wide variety of dynamic control problems associated with queueing models, and it develops a general computational method for such problems. Our problem of service rate control with average cost criterion is considered in section 10.4, assuming that A is a finite set, where Sennett illustrates the application of her general method in this particular context. The approach is to consider approximations of the original problem that have a finite state space, applying to each such problem a general solution technique such as value iteration. As one would expect given its broad applicability, this computational approach is not as efficient as our customized method, nor does it provide interesting characterizations of the optimal policy as a by-product.

4 The Optimality Equation

To further reduce our optimality equation (9)-(10), it is natural to define the function

$$\phi(y) = \sup_{x \in A} \{yx - c(x)\} \quad y \geq 0; \tag{11}$$

then (9)-(10) is equivalently expressed as

$$y_1 = z - h_0 \tag{12}$$

and

$$y_{n+1} = \phi(y_n) - h_n + z \quad \text{for } n \geq 1. \tag{13}$$

Given the assumptions on $c(\cdot)$ and A that were set forth in section 1, it is straightforward to prove the following: first, the supremum in (11) is finite for all $y \geq 0$, and second, there exists a smallest $x^* \in A$ that achieves the supremum. Hereafter that smallest maximizer x^* will be denoted by $\psi(y)$. Some important properties of the functions ϕ and ψ will be compiled in the next section.

We now provide a “verification lemma” which allows one to rigorously prove the optimality of a policy derived from a solution of (12)-(13), provided that the relative cost differences y_n are bounded. As we shall see later, one cannot expect bounded solutions of the optimality equation to exist in general, but the verification lemma can be used as part of a monotone approximation argument.

Proposition 1 *Let z and (y_1, y_2, \dots) be a solution of the optimality equation (12)-(13) such that y_1, y_2, \dots are bounded. Then $z \leq z(\mu)$ for every ergodic policy μ , and hence $z \leq z^*$. Moreover, if the policy μ^* defined by $\mu_n^* = \psi(y_n)$ for $n \geq 1$ is ergodic, then $z(\mu^*) = z = z^*$, implying that μ^* is optimal.*

Remark. The assumed monotonicity of h is not actually used in the following proof.

PROOF. From the definition (11) of $\phi(\cdot)$ we have that

$$xy_n - c(x) \leq \phi(y_n) \text{ for all } x \geq 0 \text{ and } n \geq 1. \quad (14)$$

Now let μ be an arbitrary ergodic policy. Setting $x = \mu_n$ in (14) and using (13) to substitute for $\phi(y_n)$, we have

$$\mu_n y_n - c(\mu_n) \leq y_{n+1} + h_n - z \text{ for } n \geq 1. \quad (15)$$

Multiplying both sides of (15) by $p_n(\mu)$, then making the substitution $\mu_n p_n(\mu) = p_{n-1}(\mu)$ by (3), we can rearrange terms to get

$$p_n(\mu)[h_n + c(\mu_n) - z] \geq p_{n-1}(\mu)y_n - p_n(\mu)y_{n+1} \text{ for } n \geq 1. \quad (16)$$

Because y is bounded, both terms on the right side of (16) are summable over $n \geq 1$ and the difference of those two sums is $p_0(\mu)y_1$. By definition, the first two terms on the left side of (16) sum to $z(\mu) - p_0(\mu)h_0$ and thus we have from (16) that

$$z(\mu) - p_0(\mu)h_0 - z[1 - p_0(\mu)] \geq p_0(\mu)y_1. \quad (17)$$

Now (12) says that $y_1 = z - h_0$, so (17) reduces to $z(\mu) \geq z$, as desired. To prove the last statement of the proposition, note that if we set $x = \psi(y_n)$ in (14), then (14) holds with equality for all $n \geq 1$ by definition of the maximizer $\psi(\cdot)$. Then (15) and (16) hold with equality as well when we take $\mu_n = \mu_n^* = \psi(y_n)$. Thus (17) holds with equality when μ^* is substituted for μ , implying that $z(\mu^*) = z$, because $y_1 = z - h_0$ by (12). \square

5 Properties of the Functions ϕ and ψ

This section constitutes something of a technical diversion, and to keep attention focused on the main flow of ideas, it may be advisable to just skim it on first reading. We consider the function $\phi(\cdot)$ defined by (11) and the associated maximizer $\psi(\cdot)$. First, because the cost-of-effort function $c(\cdot)$ is left-continuous by assumption, and because $\psi(y)$ was defined as the *smallest* $x \in A$ that achieves the supremum in (11), the function $\psi(\cdot)$ is itself left-continuous and non-decreasing. Next, fixing $y_0 \geq 0$, let $x_0 = \psi(y_0)$ so that

$$\phi(y_0) = y_0 x_0 - c(x_0). \quad (18)$$

For arbitrary $y \geq 0$ we then have

$$\phi(y) = \sup_{x \in A} \{yx - c(x)\} \geq yx_0 - c(x_0). \quad (19)$$

Combining (18) and (19) gives

$$\phi(y) \geq \phi(y_0) + x_0(y - y_0) \quad \text{for all } y \geq 0. \quad (20)$$

This state of affairs is portrayed graphically in Figure 2, and it follows easily from (20) that $\phi(\cdot)$ is convex on $[0, \infty)$. A finite-valued convex function is continuous and is differentiable almost everywhere, and (20) implies that the derivative $\phi'(\cdot)$ equals $\psi(\cdot)$ wherever the former exists. Thus we have

$$\phi(y) = \int_0^y \psi(u) \, du \quad \text{for all } y \geq 0, \quad (21)$$

with the integral understood in the ordinary Riemann sense (because ψ is left-continuous and non-decreasing, it is Riemann integrable). Finally, defining $a \geq 0$ via

$$a = \sup \{y \geq 0 : \phi(y) = 0\} = \inf \{x^{-1}c(x) : x \in A \text{ and } x > 0\} \quad (22)$$




Figure 2: The Function $\phi(\cdot)$ and Maximizer $\psi(\cdot)$

(see Figure 2), we observe that $\psi(\cdot)$ is *strictly* positive and non-decreasing on (a, ∞) .

An interesting and important aspect of the function $\phi(\cdot)$ and its associated maximizer $\psi(\cdot)$ is that neither one is affected if we replace the original cost-of-effort function $c(\cdot)$ by its convex hull $\hat{c}(\cdot)$. Following Rockafellar [6], one can define the convex hull $\hat{c}(\cdot)$ as the largest convex function f on $[0, \infty)$ such that $c(x) \geq f(x)$ for all $x \in A$; in the case where A is bounded and thus has a largest element $b < \infty$, this means that $\hat{c}(\cdot)$ is a finite-valued convex function on $[0, b]$ which is extended to all of $[0, \infty)$ by setting $\hat{c}(x) = \infty$ for $x > b$. Let us define A^* as the set of all points $x \in A$ such that $c(x) = \hat{c}(x)$. Then for each $y_0 \geq 0$, the maximizer $x_0 = \psi(y_0)$ lies in A^* , as follows: the definitive relationship (18) gives $c(x) \geq c(x_0) + y_0(x - x_0)$ for all $x \in A$; that is, our cost function $c(x)$ is minorized by the affine (and hence convex) function $f(x) = c(x_0) + y_0(x - x_0)$, with $c(x_0) = f(x_0)$, so one concludes that $\hat{c}(x_0) = c(x_0)$. Thus service rates x that are *not* on the convex hull can be excluded from the problem without affecting its optimal solution. This fact has long been




Figure 3: The Function $\phi(\cdot)$ when A is a Finite Set

recognized, and Jo [4] calls it *Crabill's exclusion principle*.

For example, suppose that A contains just the five points $0 = x_0 < x_1 < \dots < x_4$ pictured in the left panel of Figure 3, and that the associated cost rates $c(x_i)$ are as shown in that figure. The convex hull $\hat{c}(\cdot)$ is shown by a solid line in the left panel, and here one finds that $A^* = \{0, x_1, x_2, x_4\}$. Because $c(x_3) > \hat{c}(x_3)$, there is no value of $y \geq 0$ such that $\psi(y) = x_3$, and thus the service rate x_3 cannot appear in the optimal control policy μ^* to be derived below. In this case $\phi(\cdot)$ is the piecewise-linear, convex curve shown in the right panel of Figure 3: its break points are $a = y_1 = c(x_1)/x_1$, $y_2 = (c(x_2) - c(x_1))/(x_2 - x_1)$, and $y_4 = (c(x_4) - c(x_2))/(x_4 - x_2)$; the four line segments that make up the graph of $\phi(\cdot)$ have slopes zero, x_1 , x_2 and x_4 respectively as one proceeds from left to right.

Continuing with this same example, suppose now that all service rates $x \in [0, x_4]$ are available with associated cost rates $\hat{c}(x)$, so that the cost-of-effort function is piecewise linear and convex. One finds that this expanded capability is actually irrelevant because

the maximizer $\psi(y)$ always comes from the finite set A^* identified above, and hence the optimal policy μ^* derived below has $\mu_n^* \in A^*$ for all $n \geq 1$.

In certain ways, the following “standard case” makes for the simplest analysis. First suppose that all non-negative service rates x are available to the system controller, meaning that $A = [0, \infty)$. Also, assume that the cost-of-effort function $c(x)$ is strictly convex, strictly increasing and continuously differentiable on $[0, \infty)$ with $c(0) = 0$. Finally, to satisfy (1) we require that $c'(x) \uparrow \infty$ as $x \uparrow \infty$. Defining $a = c'(0)$, which is consistent with (22), the maximizer $\psi(\cdot)$ is then the (continuous) inverse of $c'(\cdot)$ on $[a, \infty)$. That is, one has

$$\phi(y) = 0 \quad \text{and} \quad \psi(y) = 0 \quad \text{for} \quad 0 \leq y \leq a, \quad (23)$$

$$\psi(y) = \{x \geq 0 : c'(x) = y\} \quad \text{for} \quad y > a \quad (24)$$

and

$$\phi(y) = y\psi(y) - c(\psi(y)) \quad \text{for} \quad y > 0. \quad (25)$$

6 Truncated Holding Costs

In this section we show how to solve a problem with “truncated holding costs”, where h_i is set equal to h_n for all $i \geq n$. (We take $n \geq 1$ to be fixed for purposes of this initial discussion, but it will be allowed to vary in the mathematical development that follows.)

The optimality equation (13) reduces to $y_{i+1} = \phi(y_i) - h_n + z$ for all $i \geq n$, so if we can find a value of z such that $y_{n+1} = y_n$ it follows that $y_{n+j} = y_n$ for all $j \geq 1$. In this section we show that there is indeed a unique value of z giving $y_{n+1} = y_n$. Proposition 1 will then be

used to verify the optimality of a policy μ that is derived from z in the obvious way (it has $\mu_n = \mu_{n+1} = \dots$). All the quantities of interest will prove to be monotone non-decreasing in n , and that monotonicity will be used heavily in the next section, both to characterize optimal policies for the general case and to show how nearly optimal policies can be computed.

To simplify the development that follows, let us extend $\phi(\cdot)$ from $[0, \infty)$ to all of \mathbb{R} by setting

$$\phi(y) = 0 \quad \text{for all } y < 0. \tag{26}$$

Then $\phi(\cdot)$ is convex (hence continuous) and non-decreasing on \mathbb{R} . Next, for each $z \in \mathbb{R}$ let us define $y_1(z) = z - h_0$ as in (12), then define $y_2(z), y_3(z), \dots$ recursively by means of (13). This defines a sequence of functions $y_n : \mathbb{R} \rightarrow \mathbb{R}$ indexed by $n \geq 1$, and using the properties of $\phi(\cdot)$ stated immediately above, one easily obtains the following by induction.

Proposition 2 *For each $n \geq 1$ the function $y_n(\cdot)$ is continuous, strictly increasing and unbounded both above and below.*

Recalling the definition (22) of a , let us now define $y_0(z) = a$ for all $z \in \mathbb{R}$ (this simplifies the recursive formulas that follow), then set

$$\Delta_n(z) = y_n(z) - y_{n-1}(z) \quad \text{for } z \in \mathbb{R} \quad \text{and } n \geq 1. \tag{27}$$

In particular, then, $\Delta_1(z) = y_1(z) - y_0(z) = z - h_0 - a$. Further setting

$$\delta_n = h_n - h_{n-1} \quad \text{for } n \geq 1, \tag{28}$$

observe that $\delta_n \geq 0$ for all $n \geq 1$, because h is assumed to be non-decreasing. It is immediate from (13) that

$$\Delta_{n+1}(z) = [\phi(y_n(z)) - \phi(y_{n-1}(z))] - \delta_n \text{ for } x \in \mathbb{R} \text{ and } n \geq 1. \quad (29)$$

We now identify values z_n ($n = 1, 2, \dots$) such that $\Delta_n(z_n) = 0$ for all $n \geq 1$. These will shortly be shown to equal the minimum average costs achievable with certain truncated holding costs.

Proposition 3 *There exists a unique monotone sequence $h_0 + a = z_1 \leq z_2 \leq \dots$ such that $\Delta_n(z_n) = 0$ for all $n \geq 1$. Moreover,*

$$a = y_0(z_n) \leq \dots \leq y_{n-1}(z_n) = y_n(z_n) \text{ for all } n \geq 1. \quad (30)$$

Remark The following proof shows how a simple one-dimensional search can be used to compute each successive value z_1, z_2, \dots . Given the simplicity of this computational task, we shall treat z_1, z_2, \dots as “known constants” hereafter.

PROOF. As noted immediately above, $\Delta_1(z) = z - h_0 - a$, so the only value of z_1 giving $\Delta_1(z_1) = 0$ is $z_1 = h_0 + a$, implying that $y_1(z_1) = a$. Moreover, $\Delta_1(\cdot)$ is continuous and strictly increasing on $[z_1, \infty)$.

Arguing by induction, let us now fix $n \geq 1$, assume there exist unique values $z_1 \leq \dots \leq z_n$ satisfying $\Delta_1(z_1) = \dots = \Delta_n(z_n) = 0$, and further assume that $\Delta_i(\cdot)$ is continuous and strictly increasing and unbounded on $[z_i, \infty)$ for all $i = 1, \dots, n$. Because $\Delta_n(z_n) = y_n(z_n) -$

$y_{n-1}(z_n) = 0$, we have from (29) that $\Delta_{n+1}(z_n) = -\delta_n \leq 0$. Now for any $z \geq z_n$ we know that $\Delta_n(z) \geq 0$ so (21) can be used to rewrite (29) as

$$\Delta_{n+1}(z) = \int_{y_{n-1}(z)}^{y_{n-1}(z)+\Delta_n(z)} \psi(u) du - \delta_n \quad \text{for } z \geq z_n. \quad (31)$$

From this, Proposition 2, the induction hypotheses concerning $\Delta_n(\cdot)$, and the fact that $\psi(\cdot)$ is strictly positive and non-decreasing on (a, ∞) , we have that $\Delta_{n+1}(\cdot)$ is continuous, strictly increasing and unbounded on $[z_n, \infty)$. Thus there exists a unique $z_{n+1} \geq z_n$ such that $\Delta_n(z_{n+1}) = 0$, which extends the induction hypothesis to $n + 1$ and completes the proof of the first statement. Property (30) follows from the monotonicity of each function $\Delta_i(\cdot)$ that was proved as a byproduct. That is, one knows that $z_n \geq z_i$, and hence that $\Delta_i(z_n) = y_i(z_n) - y_{i-1}(z_n) \geq 0$, for each $i = 1, \dots, n - 1$. \square

Proposition 4 *Fix $n \geq 1$ and consider the modified control problem with holding cost vector $(h_0, h_1, \dots, h_{n-1}, h_n, h_n, h_n, \dots)$. The optimal objective value for that modified problem (that is, the infimum of average costs achievable with ergodic policies) is greater than or equal to z_{n+1} . If $z_{n+1} < h_n$ then that optimal objective value actually equals z_{n+1} , and the policy μ defined by*

$$\mu_i = \begin{cases} \psi(y_i(z_{n+1})) & \text{for } i = 1, \dots, n, \\ \psi(y_n(z_{n+1})) & \text{for } i > n \end{cases} \quad (32)$$

is ergodic with $z(\mu) = z_{n+1}$, and thus it is optimal. If $z_{n+1} \geq h_n$ then the modified control problem is degenerate (see section 2).

PROOF. Observe that $(y_1(z_{n+1}), \dots, y_n(z_{n+1}), y_n(z_{n+1}), \dots)$ is non-negative and bounded, and that together with z_{n+1} it satisfies the optimality equation (12)-(13) when the modified (truncated) holding cost vector is substituted for the original one. The desired conclusions are then immediate from Proposition 1 except for the following: it remains to show that if $z_{n+1} < h_n$ (which equals h_∞ for the problem under discussion) then the policy μ identified above is ergodic. That is, we need to show that $z_{n+1} < h_n$ implies $\psi(y_n(z_{n+1})) > 1$.

Accordingly, assume $z_{n+1} < h_n$, so that the affine function $l(y) = h_n - z_{n+1} + y$ (for $y \geq 0$) has slope 1 and $l(0) > 0$. (Readers may find it helpful to imagine this line added to Figure 2.) The defining property of z_{n+1} is that $y_n(z_{n+1}) = y_{n+1}(z_{n+1})$, but we also have $y_{n+1}(z_{n+1}) = \phi(y_n(z_{n+1})) - h_n + z_{n+1}$ by (13), so $y_n(z_{n+1})$ must be a solution of the equation $l(y) = \phi(y)$. Because $l(0) > 0$ and $\phi(0) = 0$, the line $l(\cdot)$ can only intersect the convex function $\phi(\cdot)$ at one point, and because $l(\cdot)$ has slope 1, the intersection must be at a point y where the slope $x = \psi(y)$ that supports $\phi(\cdot)$ is strictly greater than 1 (see Figure 2). That is, it must be that $\psi(y_n(z_{n+1})) > 1$. \square

7 Optimal and Nearly Optimal Policies

In this section we consider the n^{th} problem with truncated holding costs, which was solved in the previous section, and show that its optimal solution converges monotonically, as $n \rightarrow \infty$, to an optimal solution under our original cost structure. Moreover, the final inequality displayed in this section provides a performance bound for the policy derived from the n^{th} truncation, comparing its average cost under the original cost structure against the minimum

average cost z^* .

Recall from section 2 that z^* , defined as the infimum of average cost rates achievable with ergodic policies using our original holding cost vector h , is necessarily finite. Of course, the corresponding infimum with truncated holding costs is less than or equal to z^* , and the sequence z_1, z_2, \dots identified in Proposition 3 is non-decreasing, so from Proposition 4 we have

$$z_n \uparrow z_\infty \leq z^* \text{ as } n \uparrow \infty. \quad (33)$$

Next, Propositions 2 and 3 give the following:

$$y_i(z_n) \uparrow y_i^* = y_i(z_\infty) \text{ for } i \geq 1, \quad (34)$$

and

$$a < y_1^* \leq y_2^* \leq \dots. \quad (35)$$

That is, the sequence y_1^*, y_2^*, \dots derived from z_∞ by means of our optimality equation (12)-(13) is non-decreasing. From the proof of Proposition 3 it is clear that if $z < z_\infty$ (and hence $z < z_n$ for some $n \geq 1$), then the sequence $y_1(z), y_2(z), \dots$ cannot be non-decreasing, and thus we have the following.

Proposition 5 z_∞ is the infimum of those $z \in \mathbb{R}$ for which the sequence $y_1(z), y_2(z), \dots$ derived from z via (12)-(13) is non-decreasing.

Recall from section 4 that the maximizer $\psi(\cdot)$ is strictly positive, left-continuous and non-decreasing on (a, ∞) . Thus (34) and (35) imply that

$$\psi(y_i(z_n)) \uparrow \mu_i^* = \psi(y_i^*) \text{ as } n \uparrow \infty \text{ for each } i \geq 1, \quad (36)$$

and

$$0 < \mu_1^* \leq \mu_2^* \leq \dots . \quad (37)$$

In this section we confirm that the policy μ^* defined by (36) is indeed optimal, with associated average cost rate $z^* = z_\infty$, provided that our dynamic control problem is non-degenerate. Dispensing first with the degenerate case, the following is immediate from (33) and the definition of degeneracy (see section 2).

Proposition 6 *Suppose that $h_n \uparrow h_\infty < \infty$ as $n \uparrow \infty$, and that moreover $z_\infty \geq h_\infty$. Then our original dynamic control problem is degenerate.*

Proposition 7 *If $z_\infty < h_\infty$ (that is, excluding the degenerate case treated in Proposition 6) then $z_\infty = z^* = z(\mu^*)$, so μ^* is an optimal policy. Moreover, the policies μ^1, μ^2, \dots derived from z_1, z_2, \dots via (32) satisfy $z(\mu^n) \rightarrow z^*$ as $n \rightarrow \infty$.*

PROOF. Because the degenerate case is excluded, one has $z_n \leq z_\infty < h_\infty$ for all $n \geq 1$, which implies the following: there exists an integer $N \geq 1$ such that

$$z_n \leq z_\infty < h_n \quad \text{for all } n \geq N. \quad (38)$$

Given (38), one can now argue exactly as in the proof of Proposition 4 that

$$\psi(y_n(z_{n+1})) > 1 \quad \text{for all } n \geq N. \quad (39)$$

Hereafter, for each $n \geq N$ let us denote by $\mu^n = (\mu_1^n, \mu_2^n, \dots)$ the policy defined by (32). In particular, then, we have that

$$\mu_i^n = \psi(y_n(z_{n+1})) \quad \text{for all } i \geq n \geq N, \quad (40)$$

so (39) shows that μ^n is an ergodic policy for all $n \geq N$ (recall that $\lambda = 1$ by convention in our formulation). Moreover, denoting by $p(\mu^n)$ the stationary distribution under policy μ^n as in section 2, we have from (3) and (40) that

$$p_{n+j}(\mu^n) = p_n(\mu^n)(\mu_n^n)^{-j} \quad \text{for all } j \geq 0, \quad (41)$$

provided $n \geq N$, so we have from assumption (2) that

$$z(\mu^n) = \sum_{i=0}^{\infty} p_i(\mu^n)[h_i + c(\mu_i^n)] < \infty \quad (42)$$

for all $n \geq N$. Finally, defining truncated holding cost vectors h^n via

$$h_i^n = h_i \wedge h_n \quad \text{for } i \geq 0 \quad \text{and } n \geq N, \quad (43)$$

we have from Proposition 4 that

$$z_{n+1} = \sum_{i=0}^{\infty} p_i(\mu^n)[h_i^n + c(\mu_i^n)] \quad \text{for } n \geq N. \quad (44)$$

Combining (41)-(44), one then has that

$$z(\mu^n) = z_{n+1} + p_n(\mu^n) \sum_{j=0}^{\infty} (h_{n+j} - h_n)(\mu_n^n)^{-j}. \quad (45)$$

Because $z(\mu^n)$ is the average cost rate achieved by a specific (and readily computable) policy μ^n , we know from (33) that

$$z_{n+1} \leq z^* \leq z(\mu^n). \quad (46)$$

Thus, defining

$$g_n(x) = \sum_{j=0}^{\infty} (h_{n+j} - h_n)x^{-j} \quad \text{for } x > 1 \quad \text{and } n \geq 1, \quad (47)$$

we have from (45) that

$$z_{n+1} \leq z^* \leq z_{n+1} + p_n(\mu^n)g_n(\mu^n) \quad \text{for } n \geq N. \quad (48)$$

Because μ_i^n is non-decreasing in both i and n , it is easy to show that $p_n(\mu^n) \rightarrow 0$ as $n \rightarrow \infty$, and $g_n(\mu^n) \rightarrow 0$ as $n \rightarrow \infty$ as well by assumption (2) (here one uses the fact that $0 \leq h_{n+j} - h_n \leq h_{n+j}$). Thus we have from (48) that $z_n \rightarrow z^*$ as $n \rightarrow \infty$, meaning that $z_\infty = z^*$.

Now (36) says that $\mu_i^n \uparrow \mu_i^*$ as $n \uparrow \infty$ for each $i \geq 1$, from which it follows that $p_i(\mu^n) \rightarrow p_i(\mu^*)$ for each $i \geq 0$ as $n \uparrow \infty$. Also, $c(\mu_i^n) \rightarrow c(\mu_i^*)$ for each $i \geq 1$ as $n \uparrow \infty$, because $c(\cdot)$ was assumed to be left-continuous on A , and these facts together imply that $z(\mu^n) \rightarrow z(\mu^*)$ as $n \uparrow \infty$. But $z(\mu^n) \rightarrow z_\infty$ by (45), and this completes the proof. \square

Of course, calculating the policy μ^n for any given n is a finite computational task, and one can use (45) to bound the difference between its average cost rate and the optimal average cost rate z^* , as follows:

$$0 \leq z(\mu^n) - z^* \leq p_n(\mu^n)g_n(\mu^n), \quad (49)$$

where $g_n(\cdot)$ is defined by (47). To make practical use of this bound, one must be able to compute $g_n(x)$ for given n and x , as is the case when h_n is defined by a polynomial formula for sufficiently large n (such an example is treated in the next section).

8 Numerical Examples

In this section we consider a family of numerical examples that fit the “standard case” identified at the end of section 5. To be specific, the cost-of-effort function $c(x)$ in all our

examples is $c(x) = 0.5x^2$ for $x \geq 0$, implying that $a = c'(0) = 0$, that $c'(x) = x$, and hence by (24) and (25) that

$$\psi(y) = y \quad \text{and} \quad \phi(y) = 0.5y^2 \quad \text{for all } y \geq 0.$$

With regard to holding costs, we assume that $h_n = s(n - M + 1)^+$ for all $n \geq 0$, where s is strictly positive and $M \geq 1$ is an integer: thus holding costs are zero until the queue length n reaches a minimum value M , after which h_n increases linearly with slope s as n increases. A holding cost vector of this form captures the notion that congestion costs are negligible until the queue length n exceeds some “acceptable level”, that being $M - 1$ in our notation. Of course, linear holding costs are represented by the case $M = 1$. In a sense, the holding cost vector is “most convex” when M has an intermediate value, because small values of M approximate the linear case, and as M gets large, holding costs become a negligible problem element (that is, h approaches the zero vector).

Tables 1 and 2 and Figures 4 through 7 give numerical results for 48 different cases, corresponding to eight different values of the slope s and six different values of M . The different values of M are easy to interpret, but to put in perspective the different s values, it is useful to note the following. If both arrivals and services were deterministic (that is, perfectly regular) then the system manager could implement a constant service rate $\mu = 1$, matching the arrival rate $\lambda = 1$, and still never have more than one job in the system. The average cost-of-effort per time unit would then be $c(1) = 0.5$, and using the convexity of $c(\cdot)$ and Jensen’s inequality, one sees that this is a lower bound on the long-run average effort cost per time unit under any ergodic policy. Thus, we shall refer to $c(1) = 0.5$ as the *baseline*

Table 1: Summary of Incremental Costs

s	$M = 1$						$M = 3$					
	Total Cost		Holding Cost		Effort Cost		Total Cost		Holding Cost		Effort Cost	
	Fixed	Controllable	Fixed	Controllable	Fixed	Controllable	Fixed	Controllable	Fixed	Controllable	Fixed	Controllable
	μ	μ	μ	μ	μ	μ	μ	μ	μ	μ	μ	μ
0.1	0.4472	0.3873	0.2236	0.2127	0.2236	0.1746	0.3276	0.2462	0.1274	0.1007	0.2002	0.1455
0.2	0.6325	0.5669	0.3162	0.3083	0.3162	0.2586	0.4262	0.3234	0.1579	0.1224	0.2682	0.2011
0.3	0.7746	0.7062	0.3873	0.3807	0.3873	0.3255	0.4942	0.3758	0.1781	0.1358	0.3161	0.2399
0.4	0.8944	0.8242	0.4472	0.4411	0.4472	0.3830	0.5477	0.4162	0.1936	0.1456	0.3541	0.2706
0.5	1.0000	0.9284	0.5000	0.4943	0.5000	0.4342	0.5923	0.4496	0.2063	0.1535	0.3860	0.2961
1	1.4142	1.3391	0.7071	0.7026	0.7071	0.6365	0.7500	0.5647	0.2500	0.1790	0.5000	0.3857
10	4.4721	4.3901	2.2361	2.2343	2.2361	2.1558	1.5436	1.0848	0.4581	0.2758	1.0854	0.8089
100	14.1421	14.0576	7.0711	7.0705	7.0711	6.9871	2.9779	1.8500	0.8218	0.3925	2.1561	1.4575

s	$M = 5$						$M = 10$					
	Total Cost		Holding Cost		Effort Cost		Total Cost		Holding Cost		Effort Cost	
	Fixed	Controllable	Fixed	Controllable	Fixed	Controllable	Fixed	Controllable	Fixed	Controllable	Fixed	Controllable
	μ	μ	μ	μ	μ	μ	μ	μ	μ	μ	μ	μ
0.1	0.2610	0.1661	0.0859	0.0529	0.1751	0.1133	0.1780	0.0782	0.0459	0.0160	0.1321	0.0622
0.2	0.3255	0.2049	0.1004	0.0588	0.2251	0.1462	0.2114	0.0893	0.0504	0.0160	0.1609	0.0733
0.3	0.3681	0.2294	0.1094	0.0618	0.2586	0.1676	0.2324	0.0958	0.0531	0.0159	0.1792	0.0799
0.4	0.4005	0.2474	0.1161	0.0638	0.2844	0.1836	0.2479	0.1004	0.0550	0.0157	0.1929	0.0846
0.5	0.4270	0.2619	0.1214	0.0653	0.3056	0.1965	0.2603	0.1039	0.0565	0.0156	0.2038	0.0883
1	0.5171	0.3086	0.1389	0.0696	0.3782	0.2391	0.3012	0.1145	0.0612	0.0151	0.2399	0.0994
10	0.9152	0.4814	0.2109	0.0794	0.7043	0.4020	0.4611	0.1469	0.0781	0.0129	0.3830	0.1340
100	1.5118	0.6705	0.3135	0.0843	1.1983	0.5862	0.6631	0.1740	0.0979	0.0107	0.5652	0.1633

s	$M = 15$						$M = 20$					
	Total Cost		Holding Cost		Effort Cost		Total Cost		Holding Cost		Effort Cost	
	Fixed	Controllable	Fixed	Controllable	Fixed	Controllable	Fixed	Controllable	Fixed	Controllable	Fixed	Controllable
	μ	μ	μ	μ	μ	μ	μ	μ	μ	μ	μ	μ
0.1	0.1379	0.0452	0.0309	0.0069	0.1071	0.0383	0.1139	0.0294	0.0232	0.0036	0.0908	0.0258
0.2	0.1601	0.0498	0.0331	0.0065	0.1270	0.0433	0.1304	0.0318	0.0245	0.0033	0.1059	0.0285
0.3	0.1738	0.0524	0.0344	0.0063	0.1394	0.0461	0.1405	0.0331	0.0253	0.0031	0.1153	0.0299
0.4	0.1838	0.0542	0.0353	0.0061	0.1485	0.0481	0.1479	0.0340	0.0258	0.0030	0.1221	0.0309
0.5	0.1918	0.0556	0.0360	0.0060	0.1558	0.0496	0.1537	0.0346	0.0262	0.0029	0.1275	0.0317
1	0.2175	0.0596	0.0382	0.0056	0.1793	0.0540	0.1723	0.0366	0.0275	0.0027	0.1448	0.0339
10	0.3139	0.0710	0.0456	0.0043	0.2683	0.0667	0.2404	0.0418	0.0317	0.0019	0.2088	0.0399
100	0.4278	0.0796	0.0535	0.0033	0.3743	0.0763	0.3182	0.0456	0.0359	0.0014	0.2822	0.0442




Figure 4: Holding and Effort Costs as a Proportion of Total Costs ($s = 0.1$)




Figure 5: Value of Responsiveness (Controllable μ versus Fixed μ)

operating cost hereafter. The smallest value of s considered in our study is $s = 0.1$, in which case the baseline operating cost is equivalent to the holding cost for five jobs. At the other extreme we consider a largest slope of $s = 100$, which means that the holding cost for just one job is 200 times larger than the baseline operating cost.

For each of the 48 parameter combinations described above, we determine the minimum average cost rate z^* achievable through dynamic control of the service rate, calling this the *Controllable μ Solution*, and compare that against the lowest average cost rate achievable using a single service rate in every state n , called the *Fixed μ Solution*. With a fixed service rate the system operates as an ordinary M/M/1 queue, so one can develop a formula for the average total cost and then use calculus to optimize the service rate. To derive the Controllable μ Solution, we use the truncated-holding-cost approximations described in sections 5 and 6, increasing the truncation level until the bound (49) guarantees that the average cost under our “nearly optimal” policy is no more than one tenth of one percent above the true optimal value. (The required truncation level n typically exceeded M by only 4 or 5.) To avoid unnecessary verbiage, the policy derived by these means, and its associated costs, are referred to as “optimal” rather than “nearly optimal”.

All of the quantities reported in our tables and graphs are long-run average values, but that modifying phrase is deleted in the headings, and in the text that follows, to avoid tedious repetition. Thus, for example, we speak of the “holding cost” under a given policy rather than the “long-run average holding cost per time unit”. Also, all of the costs reported are *incremental costs*, by which we mean increases in average cost per time unit over the baseline operating cost of $c(1) = 0.5$ described above. That is, the baseline value of 0.5 has

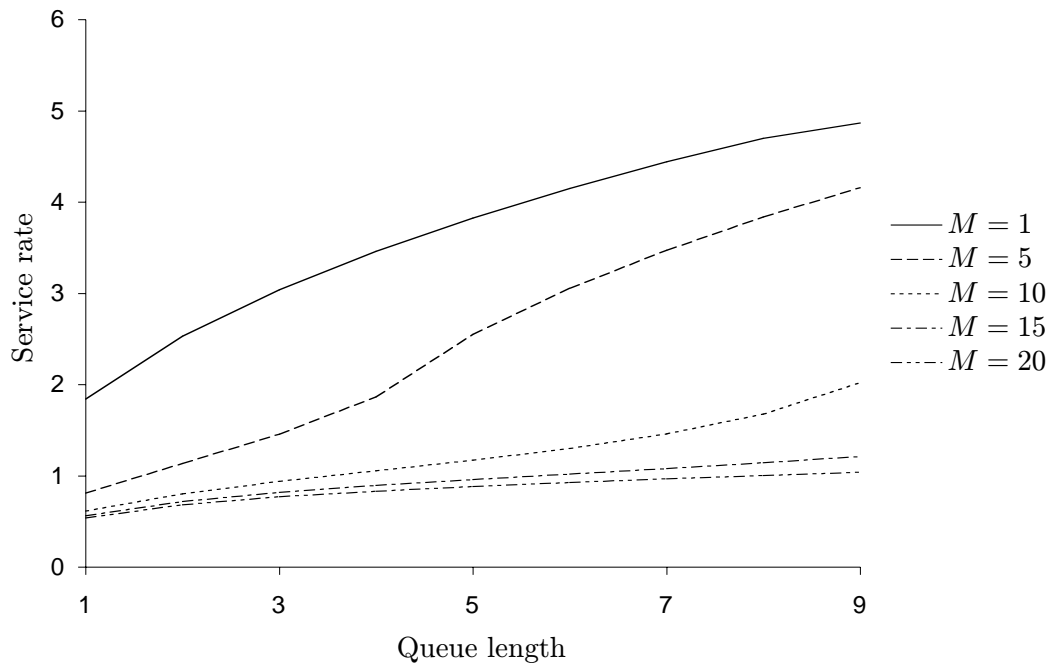


Figure 6: Controllable Service Rates ($s = 1$)

been subtracted from all of the effort costs, and hence all of the total costs, reported in our tables and graphs. All of the costs reported are rates, in units like dollars per hour.

Table 1 shows the total incremental cost for both our Controllable μ Solution and the Fixed μ Solution with different parameter combinations, further breaking that total cost into its holding cost and effort cost constituents. In all cases, the Controllable μ Solution yields both lower holding cost and lower effort cost. Figure 4 presents these cost comparisons in graphical form for the smallest of our eight holding cost slope parameters ($s = 0.1$). Notice that as M increases the effort costs are proportionally larger when compared to holding costs and that the flexibility in service rates allowed by the Controllable μ model therefore is able to drive both holding and effort costs down.

Figure 5 shows that responsiveness tends to be more valuable as M increases and for

$M = 1$

s	Fixed	Controllable μ								
	μ	1	2	3	4	5	6	7	8	9
0.1	1.4472	0.8864	1.1793	1.3818	1.5412	1.6741	1.7877	1.8843	1.9618	2.0107
1	2.4142	1.8391	2.5301	3.0399	3.4595	3.8232	4.1475	4.4400	4.6960	4.8651
10	5.4721	4.8901	6.8466	8.3279	9.5669	10.6527	11.6304	12.5235	13.3093	13.4586
100	15.1422	14.5576	20.5191	25.0737	28.9038	32.2725	35.3143	38.1081	40.6716	41.6471

$M = 5$

s	Fixed	Controllable μ								
	μ	1	2	3	4	5	6	7	8	9
0.1	1.3502	0.6660	0.8878	1.0602	1.2280	1.4201	1.5743	1.7053	1.8201	1.9223
1	1.7564	0.8086	1.1356	1.4534	1.8649	2.5476	3.0538	3.4713	3.8338	4.1575
10	2.4086	0.9814	1.4629	2.0514	3.0855	5.7415	7.4637	8.8346	10.0065	11.0461
100	3.3966	1.1705	1.8555	2.8919	5.3521	15.4930	21.1873	25.6206	29.3777	32.6965

$M = 10$

s	Fixed	Controllable μ								
	μ	1	2	3	4	5	6	7	8	9
0.1	1.2643	0.5782	0.7453	0.8559	0.9445	1.0242	1.1027	1.1861	1.2816	1.3995
1	1.4798	0.6145	0.8033	0.9372	1.0537	1.1696	1.2985	1.4576	1.6768	2.0204
10	1.7661	0.6469	0.8561	1.0133	1.1603	1.3200	1.5181	1.7992	2.2654	3.2129
100	2.1304	0.6740	0.9011	1.0800	1.2572	1.4643	1.7461	2.1985	3.0907	5.4502

$M = 15$

s	Fixed	Controllable μ								
	μ	1	2	3	4	5	6	7	8	9
0.1	1.2142	0.5452	0.6938	0.7858	0.8539	0.9098	0.9590	1.0050	1.0502	1.0966
1	1.3586	0.5596	0.7162	0.8161	0.8926	0.9580	1.0185	1.0783	1.1410	1.2105
10	1.5366	0.5710	0.7339	0.8403	0.9240	0.9978	1.0688	1.1421	1.2232	1.3190
100	1.7485	0.5796	0.7475	0.8590	0.9485	1.0294	1.1094	1.1950	1.2936	1.4163

$M = 20$

s	Fixed	Controllable μ								
	μ	1	2	3	4	5	6	7	8	9
0.1	1.1815	0.5294	0.6695	0.7535	0.8133	0.8601	0.8992	0.9337	0.9653	0.9953
1	1.2896	0.5366	0.6805	0.7681	0.8316	0.8823	0.9258	0.9651	1.0023	1.0388
10	1.4175	0.5418	0.6886	0.7789	0.8452	0.8990	0.9459	0.9892	1.0310	1.0733
100	1.5645	0.5456	0.6945	0.7868	0.8551	0.9112	0.9608	1.0072	1.0528	1.0998

Table 2: Service Rates

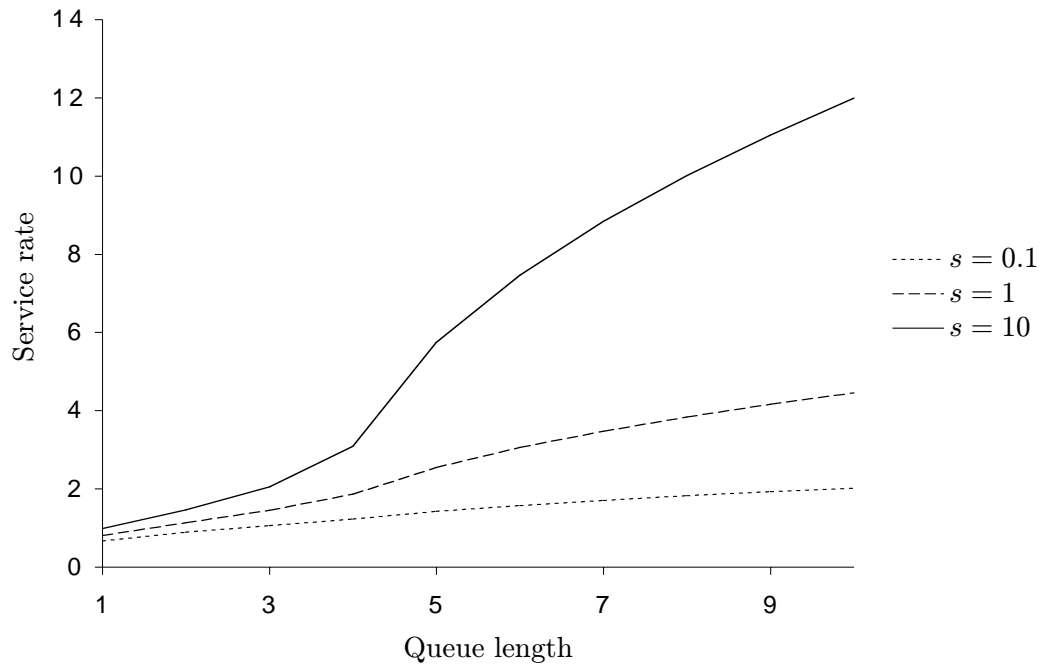


Figure 7: Controllable Service Rates ($M = 5$)

larger values of s . The “value of responsiveness” is the percentage decrease in Total Cost achieved by the Controllable μ policy, relative to the best Fixed μ policy. It is noteworthy that for $M = 1$ in Figure 5, the value of responsiveness decreases in s for straight linear costs but for more convex costs ($M > 1$) the value of responsiveness is increasing in s . This is an example where linear holding costs give a result that is not representative of the general case.

Service rates for the Controllable μ solution are unbounded as the queue length increases, because h is unbounded. Table 2 sets out service rates for a selection of M and s values. The optimal Fixed μ service rate is given in the first column and the first nine service rates (that is, service rates for queue lengths 1 through 9) for the Controllable μ solution are given. Figures 6 and 7 illustrate an interesting phenomenon: whereas service rates for $M = 1$ are

concave, the service rates for higher values of M are S-curves, growing rapidly just before M (where the holding costs become positive) and then looking similar to the $M = 1$ curve for higher values of the service rate.

Appendix: Instantaneous Ejection Capability

To understand what happens when assumption (1) fails, suppose that $A = [0, b]$ and $c(x) = \alpha x$ for $x \in A$, where $\alpha > 0$. Specializing results derived in sections 5 and 6, one has the following: either the problem is degenerate (this can only happen if h is bounded, of course) or else the optimal policy μ^* has $\mu_n^* = b$ for all $n \geq 1$. Now what happens if we let $b \uparrow \infty$ in this model? Strictly speaking, the limiting problem is one in which no optimal policy exists, but the real answer is that one must switch to a different formulation to capture the limiting scenario in a mathematically meaningful fashion, as follows. If the system begins in some state $n > 0$ and service rate $\mu_n = b$ is chosen, where b is large, then it is very likely one will see a service completion before the next arrival occurs. The expected time required for that service completion is b^{-1} , and the expected effort cost incurred before the service completion is $c(b)b^{-1} = (\alpha b)b^{-1} = \alpha$. Taking the limit as $b \uparrow \infty$, one sees that the original model becomes one in which the system manager can, at any desired time, effect an instantaneous downward transition for a fixed charge of α . That is, starting with a queue length $n \geq 1$, the system manager can immediately eject $i \leq n$ of those jobs at a cost of $i\alpha$, if that is deemed desirable. The natural “limiting model” is one in which both instantaneous ejection and service at a finite rate are available as control modes, but given the linear cost-of-effort

function that we have hypothesized, one ultimately finds that the latter option is dominated. To be more specific, one ultimately finds that just two possibilities exist: either the limiting problem is degenerate, or else it is optimal to instantaneously eject each new job at the moment of its arrival.

Extending this discussion to a more interesting and more general setting, imagine that $c(x)$ is piecewise linear and convex, like the solid line in the left panel of Figure 3, with its last linear segment having slope α and right endpoint b . The right way to formulate the limiting control problem (as $b \uparrow \infty$) is to retain the discrete control modes that correspond to the breakpoints of the piecewise linear function $c(\cdot)$, but add an ability to eject any job at any time for a charge of α . Finite service rates other than the breakpoints can be eliminated from the formulation, because the supremum in (11) is achieved either at a breakpoint or else at $x = \infty$ (corresponding to ejection in the proposed formulation).

Finally, extending these observations to an arbitrary problem where (1) fails to hold, let us suppose that the infimum in (1) increases to $\alpha < \infty$ as $y \uparrow \infty$, rather than increasing without bound. To avoid an outcome where no optimal policy exists, one can simply add an ability to eject any job at any desired time for a charge of α . This will “close” the model in the sense described above, or at least we presume it will. Problems of this hybrid type, where the system manager can either serve at a finite rate or effect instantaneous transitions, incorporate an “admission control” capability of the sort described in section 11.5.4 of Puterman [5]. It seems likely that all of the results developed here can be extended to such hybrid formulations, but we have not investigated that matter.

References

- [1] Bertsekas, D. P., *Dynamic Programming and Optimal Control*, Volume 2, Athena Scientific, 1995.
- [2] Crabill, T. B., Optimal Control of a Service Facility with Variable Exponential Service Times and Constant Arrival Rate, *Management Science*, 1972, 560-566.
- [3] Crabill, T. B., Optimal Control of a Maintenance System with Variable Service Rates, *Operations Research*, 22(1974), 736-745.
- [4] Jo, K.Y., A Lagrangian Algorithm for Computing the Optimal Service Rates in Jackson Queuing Networks, *Computers and Operations Research*, 16(1989), 431-440.
- [5] Puterman, M. L., *Markov Decision Processes*, Wiley-Interscience, New York, 1994.
- [6] Rockafellar, R. T., *Convex Analysis*, Princeton University Press, Princeton, New Jersey, 1970.
- [7] Sennet, L., *Stochastic Dynamic Programming and the Control of Queuing Systems*, John Wiley and Sons, New York, 1999.
- [8] Stidham, S., and Weber, R.R., Monotonic and Insensitive Optimal Policies for Control of Queues with Undiscounted Costs, *Operations Research*, 87(1989), 611-625.
- [9] Weber, R.R., and Stidham, S., Optimal Control of Service Rates in Networks of Queues, *Advances in Applied Probability*, 19(1987), 202-218.
- [10] Wijngaard, J., and Stidham, S., Forward Recursion for Markov Decision Processes with Skip-Free-To-the-Right Transitions, Part I: Theory and Algorithm, *Mathematics of Operations Research*, 11(1986), 295-308.