

Revenue Management of a Make-to-Stock Queue

René Caldentey

Stern School of Business, New York University, New York, New York 10012,
rcaldent@stern.nyu.edu

Lawrence M. Wein

Graduate School of Business, Stanford University, Stanford, California 94305,
lwein@stanford.edu

Motivated by recent electronic marketplaces, we consider a single-product make-to-stock manufacturing system that uses two alternative selling channels: long-term contracts and a spot market of electronic orders. At time 0, the risk-averse manufacturer selects the long-term contract price, at which point buyers choose one of the two channels. The resulting long-term contract demand is a deterministic fluid, while the spot-market demand is modeled as a stochastic renewal process. An exponential reflected random walk model is used to model the spot-market price, which is correlated with the spot-market demand process. The manufacturer accepts or rejects each electronic order, and long-term contracts and accepted electronic orders are backordered if necessary. The manufacturer's control problem is to select the optimal long-term contract price as well as the optimal production (i.e., busy/idle) and electronic-order admission policies to maximize revenue minus inventory holding and backorder costs. Under heavy-traffic conditions, the problem is approximated by a diffusion-control problem, and analytical approximations are used to derive a policy that is simple, and reasonably accurate and robust.

Subject classifications: inventory/production: stochastic, approximations/heuristics; queues: diffusion models.

Area of review: Manufacturing, Service, and Supply Chain Operations.

History: Received November 2003; revisions received February 2004, April 2005, July 2005; accepted August 2005.

1. Introduction

Electronic marketplaces have changed the way many industries do business. Roughly speaking, an e-marketplace is an Internet trading community that includes buyers, suppliers, manufacturers, their channels of distribution, and support services that deliver finished goods and other value-added services. At the core of an e-marketplace is a secure communication channel that allows participants to interchange real-time market information. Depending on the stage of development of an e-marketplace, the type of information can vary from purely descriptive (such as technical specifications about products and services) to concrete commercial transactions. In terms of applications, business-to-business e-marketplaces have been developed for a wide variety of products such as electronic components (www.brokerforum.com), chemical and plastic products (www.skchem.com), steel (www.e-steel.com), jewelry (www.polygon.net), and industrial equipment (www.truckpartslocator.com), among many others.

As a concrete example of an e-marketplace, we briefly describe The Broker Forum (www.brokerforum.com), a provider of a Web-based platform that enables online transactions for electronic components on a 24/7 basis. Buyers or sellers can post items they wish to buy or sell at their chosen prices and can even target which sellers

or buyers have access to these offers. A buyer or seller, upon observing a posted item that is deemed attractive, can complete the transaction online in a neutral, secure environment. Buyers and sellers can post as many items as they want in a dynamic fashion; e.g., if a buyer placed an item with a low price and had no takers, he could raise the price in the hopes of attracting a seller. Hence, The Broker Forum has essentially created a spot market for a manufactured, storable good. Most buyers and sellers perform only a fraction of their business at an e-marketplace, and continue to buy and sell items using traditional long-term contracts. Further discussion and examples of e-marketplaces can be found in Kafka et al. (2000), Kaplan and Sawhney (2000), Geunes et al. (2002, part I), Keskinocak et al. (2001), and references therein.

In this paper, we develop an idealized model, which is formulated in §2, for how a single manufacturer uses an e-marketplace to sell a portion of his manufactured, storable product (such as electronic components). In addition to participating in an electronic spot market, the manufacturer in our model also employs long-term contracts to sell its product. We assume that the manufacturer, which is modeled here as a single-server queue operating in a make-to-stock mode of production, initially chooses a fixed price at which buyers can purchase units of the product via a long-term contract. The manufacturer knows the

deterministic flow of demand arising from the long-term contracts after choosing the contract price.

Customers who do not sign the long-term contract use the electronic spot market. A key assumption in our model is that the manufacturer acts as a price-taker in this e-marketplace: Rather than posting items for sale at specific prices, he observes requests to buy items at stated prices and decides which of these requests to pursue. More specifically, customer requests arrive to the spot market according to a renewal process. Each of these requests, which we refer to as “e-orders,” is for a single unit of product for a specific price. The manufacturer, in this spot market, acts as a price-taker and decides whether to accept or reject each arriving e-order at the time of its arrival. Although a rejected job could presumably be accepted by a competing manufacturer, this aspect of the situation is not part of our model; i.e., we do not attempt to model the competition among manufacturers.

To capture in an analytically tractable way the serial correlation of e-order prices, which is driven by exogenous factors affecting product demand and supply, such as changes in economic indicators or in prices of product inputs, substitutes or complements, the spot-price stochastic process (i.e., the price of an e-order that arrives at a specific time) is modeled as an exponential reflected random walk. This stochastic process is a bounded and piecewise-linear variant of geometric Brownian motion, which is the prototypical process for modeling dynamic prices in mathematical finance (e.g., Merton 1990). Rather than assume a behavioral model that expresses the arrival rate of e-orders as a function of the spot price and congestion, which has been an invaluable approach to studying queueing systems in which the system manager posts prices (Mendelson 1985), we simply assume that the e-order arrival process and the spot-price process are correlated. This can be viewed as a “data-driven” approach, where the correlation between the two stochastic processes (i.e., the sequence of interarrival times and prices) can be measured from historical data.

Our model idealizes from the real situation in certain industries (a detailed description from the paper industry can be found in Keskinocak et al. 2001), where arriving customers might also specify the size of an order, certain product characteristics (e.g., quality, size, color), and the time at which the order must be received. Rather than having customers state a desired deadline (e.g., Markowitz and Wein 2001, Plambeck et al. 2001), we assume that the manufacturer incurs a backorder cost (per unit of inventory per unit time) for accepted jobs that need to wait and also incurs a cost for holding inventory. Hence, the manufacturer must simultaneously decide on the traditional busy/idle policy for the machine in addition to an accept/reject policy for e-orders, with the aim of maximizing his revenue minus the inventory holding and backorder costs. Because this problem is difficult (two control processes and two states, inventory level and spot price), we employ a heavy

traffic approximation in §3 in an attempt at simplification. We cannot solve the approximating diffusion-control problem exactly, but approximate analytical approaches are used in §4, which are shown to coincide well with the numerically computed optimal solution. The optimal long-term contract price is investigated in §5 and concluding remarks are offered in §6.

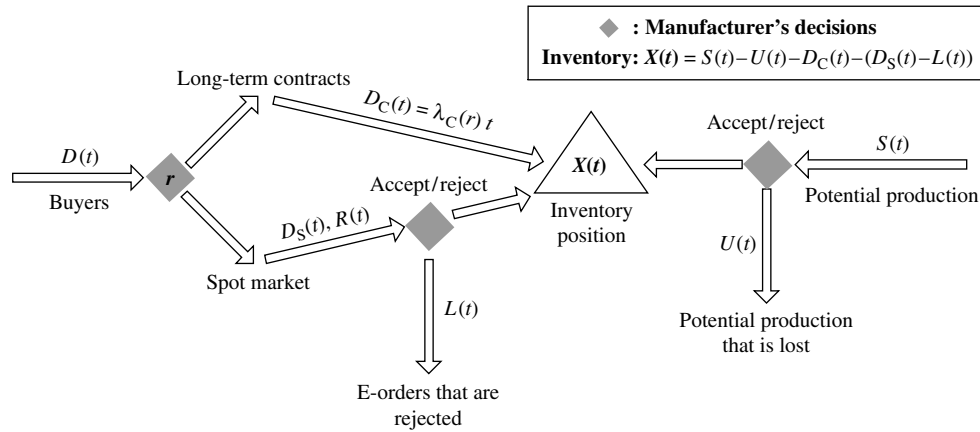
We conclude the introduction by attempting to position our paper within—but not reviewing—the vast literature on pricing and admission control in operations research. Arguably the two biggest application areas are the admission control of loss (queueing) networks for telecommunications (Kelly 1991) and revenue, or yield, management in the airline and related service industries (Talluri and van Ryzin 2004). Our problem differs from both of these problems because our manufactured product can be stored, but is closer in spirit to the former problem because of the continual operation of manufacturing and telecommunications systems (i.e., the airline problem terminates when the flight departs). In addition, there exists an extensive body of queueing theory, some of it with an economic orientation, devoted to admission control (going back to Naor 1969; see Stidham 1985 for a survey of some early work) and pricing (starting with Dolan 1978; see Van Mieghem 2000 for recent references). Moreover, dynamic pricing in stochastic inventory systems has received some recent attention (e.g., Federgruen and Heching 1999, Chen and Simchi-Levi 2004). Finally, make-to-stock queues, which combine aspects of queueing and inventory models, have also been studied in the context of pricing and admission control. Li (1992) allows the customer arrival rate to depend on the congestion and exogenous price, and derives optimal busy/idle policies for competing manufacturers. A closely related paper is Carr and Duenyas (2000), who consider a two-class queue with a make-to-stock class and a make-to-order class, where the controller sequences the jobs and makes accept/reject decisions on the make-to-stock class; however, prices are fixed in this model.

The distinguishing feature of our work is its explicit admission control in the case where the price of incoming orders is a stochastic process. To our knowledge, this has only been considered in a storable setting by Helm and Waldmann (1984), who derive structural results for admission control of a multiserver queue where the price (and holding cost rate and production capacity) evolves in a Markov-modulated manner.

2. The Problem Formulation

We consider a single-server manufacturing facility that produces and stocks a single type of product. This risk-averse manufacturer uses two alternative channels to sell its product: long-term contracts and a spot market of e-orders. The manufacturer has a two-step optimization problem: Determine the long-term contract price at time $t = 0$, and then dynamically control production and e-order admission over

Figure 1. A production-inventory system serving long-term contracts and spot-market e-orders.



an infinite time horizon. In practice, we expect the manufacturer to periodically (e.g., once a year) revise and renew the long-term contracts while using the spot market to accommodate unsold capacity. In this respect, our model is an approximation that focuses on a single cycle of the business. Figure 1 summarizes the different components of the model and notation that we describe in detail in the remainder of this section.

2.1. Long-Term Contracts

To model the long-term contract, we assume that the manufacturer quotes a price per unit, r , at time $t = 0$, and some of the potential buyers decide to sign the contract. Each buyer has a private reservation price ν , which is the maximum price that he is willing to pay for the product. The manufacturer correctly believes that $\Pr(\nu \geq r) = 1 - \Phi(r)$ for an arbitrary buyer. A fraction $\psi(r)$ of the buyers are *regular* buyers and sign the long-term contract if their reservation price is larger than r , while the remaining buyers are *speculators* that prefer to gamble on the spot market even if their reservation price is greater than r . We expect $\psi(r)$ to be decreasing in r to reflect the fact that an increase in r enhances buyers' propensity to reject the long-term contract and speculate on the spot market. The cumulative long-term contract demand $D_C(t)$ is modeled as a deterministic fluid with rate $\lambda_C(r) \triangleq \Lambda(1 - \Phi(r))\psi(r)$, where Λ is the total demand rate of all potential buyers. Consequently, revenue from the long-term contracts is accrued at a constant rate $r\lambda_C(r)$.

2.2. Spot Market

Buyers who do not select the long-term contract option procure the product on the spot market. This spot market is modeled by a pair of stochastic processes $(D_S(t), R(t))$, where $D_S(t)$ is the cumulative spot-market demand up to time t and $R(t)$ is the price posted by an e-order arriving at time t .

We define $\{x_n; n \geq 1\}$ to be the independent and identically distributed (i.i.d.) sequence of interarrival times of

e-orders with mean $E[x_1] = (\Lambda - \lambda_C(r))^{-1}$ and variance $\text{Var}(x_1) = \sigma_x^2$, which guarantees that the average arrival rate is equal to the average demand rate that is not satisfied by the long-term contracts. If we denote the corresponding partial sum process by $M_x(t) \triangleq \sum_{n=1}^{\lfloor t \rfloor} x_n$, then

$$D_S(t) = \sup_{n \geq 1} \{M_x(n) \leq t\}, \quad t \geq 0. \quad (1)$$

We model the spot-price process by

$$R(t) = R_1(r) + R_0 e^{\hat{\delta} Y(t)}, \quad (2)$$

where $R_1(r)$ is a deterministic function that captures the impact of the long-term contract price r , and $R_0 e^{\hat{\delta} Y(t)}$, where R_0 and $\hat{\delta}$ are positive constants and $Y(t)$ is a stochastic process that captures the dynamic nature of the price process. Both terms on the right side of (2) require a detailed development, and we begin with the latter term.

Our goal in constructing a price process is to capture the strong serial correlation that is present in spot markets, which may be due to seasonal factors affecting product supply and demand as well as other exogenous factors (e.g., changes in economic indicators or prices of product inputs). It is natural to consider a price process that is closely related to geometric Brownian motion (GBM), which is the standard model used to represent price processes in mathematical finance (Merton 1990). To this end, we define the i.i.d. sequence $\{y_n\}$ and the corresponding partial sum process $M_y(t) \triangleq \sum_{n=1}^{\lfloor t \rfloor} y_n$, where $E[y_1] = 0$ and $\text{Var}(y_1) = \sigma_y^2$. Because spot-price changes are observable only at the arrival epoch of e-orders, we model the spot-price process as a pure-jump process with jumps that coincide with those of the e-order arrival process. Hence, a natural candidate for $Y(t)$ in (2) is $M_y(D_S(t))$ because the exponential random walk $R_0 e^{M_y(D_S(t))}$ is the piecewise-constant version of GBM; here, the rate of price change depends on time only through the number of e-orders received in each time interval. However, this candidate is not suited to our setting because, in general, its expected

value does not remain bounded throughout the infinite time horizon. Although the operational decisions in our queueing control model (admit or reject customers, use or idle the machine) are typically considered (as they are here) in an infinite-horizon setting, the actual time frame over which these decisions are employed is finite, e.g., one year. Over the course of a year, it is hard to imagine realistic situations where the underlying spot price can be arbitrarily large because it will have an upper bound representing the price of a higher quality substitute product, which is bounded over a one-year time horizon. Consequently, we assume that the price process has a bounded support and define $Y(t)$ by

$$Y(t) = M_y(D_S(t)) + L_R(t) - U_R(t), \quad (3)$$

$$L_R(t) = \sup_{0 \leq s \leq t} [y_{\min} - M_y(D_S(t)) + U_R(s)]^+, \quad (4)$$

$$U_R(t) = \sup_{0 \leq s \leq t} [M_y(D_S(t)) + L_R(s) - y_{\max}]^+. \quad (5)$$

Here, $Y(t)$ is a two-sided regulated random walk, where the regulators $L_R(t)$ and $U_R(t)$ guarantee, with minimal effort, that $Y(t) \in [y_{\min}, y_{\max}]$ for all $t > 0$ (see §14.8 in Whitt 2002).

Although we could have constructed a price process by performing the bounding after taking the exponent—i.e., a reflected exponential random walk rather than an exponential reflected random walk—we chose the process in (3)–(5) for purposes of analytical tractability. Also, mean-reverting processes have been proposed as alternatives to GBM for modeling commodity prices (e.g., Dixit and Pindyck 1994, Schwartz 1997), and the question of whether a mean-reverting process or our exponential reflected random walk in (2)–(5) is a better model requires collecting transactional data and testing the fit of each model, which is beyond the scope of this research.

Turning to the term $R_1(r)$ on the right side of (2), we note that because a single population of buyers is split into two buying channels, the spot price needs to depend on the long-term contract price r . The mean reservation price of spot-market customers is

$$\begin{aligned} E[\nu_S] &= \left(\frac{1 - \psi(r)}{1 - \psi(r)(1 - \Phi(r))} \right) E[\nu] \\ &+ \left(\frac{\psi(r)\Phi(r)}{1 - \psi(r)(1 - \Phi(r))} \right) E[\nu \mid \nu < r], \end{aligned} \quad (6)$$

where the first term on the right side of (6) corresponds to the speculators, who operate exclusively in the spot market and have a reservation price distributed according to $\Phi(r)$, and the second term corresponds to the regular customers that use the spot market, whose reservation price is the conditional distribution of ν given $\nu < r$.

From the manufacturer's perspective, the time-average value of the spot price $R(t)$ —a quantity we denote by $E[R]$ —should be close to $E[\nu_S]$. However, spot-market buyers might behave strategically by quoting prices that differ

from their true reservation values. This line of reasoning would lead us to the analysis of a noncooperative equilibrium model for the spot market, which is beyond the scope of this work. In addition, the manufacturer's perception of future reservation prices may not be accurate at the time the long-term contract price is specified. To incorporate in a simple but reasonable way the manufacturer's imperfect information about the true value of $E[R]$ at $t = 0$, we assume that

$$E[R] = E[\nu_S] + \epsilon_R \quad \text{at } t = 0, \quad (7)$$

where ϵ_R is a normal random variable with mean zero and variance σ_ϵ^2 . Hence, by (2) and (7), we set

$$R_1(r) = E[\nu_S] - R_0 E[\exp(\delta Y)] + \epsilon_R, \quad (8)$$

where $Y = \lim_{t \rightarrow \infty} Y(t)$.

Although the value of ϵ_R is unknown to the manufacturer at time $t = 0$ when he decides the price of the long-term contract, for $t > 0$ the manufacturer is able to collect data from the spot market and improve his price forecast. Because the manufacturer is able to rapidly estimate the true value of ϵ_R in our continuous-time infinite-horizon framework, we assume that ϵ_R is unknown for the purpose of selecting the long-term price r but is known for $t > 0$ when the production and admission decisions are made.

We conclude this subsection by defining the relationship between the spot-market demand $D_S(t)$ and the spot price $R(t)$. Basic economic theory suggests that these two processes should be negatively correlated. We model this negative dependence by assuming that the pair (x_n, y_n) is positively correlated with coefficient of correlation $\hat{\rho}$. Note that, for the sake of analytical tractability, we are assuming that the interarrival time x_n is affected by the increment of the price process y_n rather than the actual price $\sum_{k=1}^n y_k$.

2.3. Manufacturing Process

The production system is driven by the i.i.d. sequence of service times $\{s_n \geq 0: n \geq 1\}$, which is independent of the spot-market demand and price processes. We also define the partial sum process $M_s(t) \triangleq \sum_{n=1}^{\lfloor t \rfloor} s_n$, so that

$$S(t) \triangleq \sup_{n \geq 1} \{M_s(n) \leq t\}, \quad t \geq 0, \quad (9)$$

is a nondecreasing stochastic process that represents the potential cumulative production of the system, i.e., the total number of units produced if the server had been producing continuously during $[0, t]$. This renewal process is characterized by the mean service time $E[s_1] = \mu^{-1}$ and variance $\text{Var}(s_1) = \sigma_s^2$. As a convenient normalization, we set $\mu = 1$ so that demand rates are expressed in units of production capacity, i.e., $\Lambda \equiv \rho$ is the system traffic intensity. In addition, under this normalization the inventory and workload formulations coincide.

The dynamics of the inventory process $X(t)$, which is negative whenever demand is backordered, are determined by the demand processes $D_C(t)$ and $D_S(t)$, the production process $S(t)$, and the manufacturer's decisions. The manufacturer has control over the admission and production processes. As is customary in the heavy-traffic literature, we model the admission and production controls using two nondecreasing and nonanticipating processes $A(t)$ and $T(t)$, respectively, satisfying

$$A(s) - A(t) \leq t - s \quad \text{and} \quad T(s) - T(t) \leq t - s$$

for $0 \leq s \leq t$.

We interpret $A(t)$ as the cumulative amount of time in $[0, t]$ during which the manufacturer is accepting e-orders. Similarly, $T(t)$ represents the cumulative amount of time in $[0, t]$ during which the manufacturing system is working.

Based on these control processes, we define $U(t) = S(t) - S(T(t))$ and $L(t) = D_S(t) - D_S(A(t))$ to be the cumulative lost output and cumulative rejected demand during $[0, t]$, respectively. It follows from these definitions that the inventory process $X(t)$ satisfies

$$X(t) = X(0) + S(t) - D(t) - U(t) + L(t),$$

where $X(0)$ is the initial inventory level. Finally, we define the *netput process*

$$Z(t) = X(0) + S(t) - D(t) \quad \text{for } t \geq 0$$

to be the inventory process if all spot-market demand is accepted and the server never idles, so that

$$X(t) = Z(t) - U(t) + L(t). \tag{10}$$

2.4. Optimization Problem

Our two-step optimization problem is solved by working backward in time. We first solve the infinite-horizon problem, assuming r and ϵ_R , and hence the average demand rate and price in the spot market are known. If we denote by $H(t, r, \epsilon_R)$ the profit rate that the manufacturer earns at time t , then this first optimization problem is

$$\mathcal{H}(r, \epsilon_R) = \max_{U, L} \lim_{T \rightarrow \infty} \frac{1}{T} \mathbb{E} \left[\int_0^T H(t, r, \epsilon_R) dt \right], \tag{11}$$

subject to the dynamics of X in (10). The solution to this problem, U^* and L^* , represent the optimal production and admission controls as a function of r and ϵ_R .

In the second optimization step, we choose the contract price r , where ϵ_R is unknown. Given the irreversible and strategic nature of this decision, we assume that the manufacturer adopts a risk-averse attitude and solves

$$\max_r \mathbb{E}[\Pi(\mathcal{H}(r, \epsilon_R))],$$

where Π is a concave utility function. In our computational examples, we assume an exponential utility function with constant measure of risk aversion β , that is, $\Pi(\mathcal{H}) = -\exp(-\beta\mathcal{H})/\beta$.

The remainder of this section formulates the infinite-horizon problem (11) in more detail. Revenue from the long-term contracts is accrued at constant rate $r\lambda_C(r)$, and the cumulative revenue earned from the spot market by time T is $\int_0^T R(t) dD_S(t) - \int_0^T R(t) dL(t)$. For later use, we note that (see Asmussen 2003, Chapter VI, Theorem 3.1)

$$\lim_{T \rightarrow \infty} \frac{1}{T} \mathbb{E} \left[\int_0^T R(t) dD_S(t) \right] = E[R]\lambda_S(r).$$

Once accepted, long-term contracts and e-orders are fully backlogged. Costs h and b are incurred for every unit held or backordered, respectively, in inventory per unit time. In this respect, we assume that long-term contracts and e-orders are indistinguishable. Hence, the cumulative inventory and backorder cost by time T is $\int_0^T c(X(t)) dt$, where $X(t)$ is the inventory level and $c(x) = hx^+ + bx^-$.

Taken together, we can express the manufacturer's optimization problem as

$$\max_r \mathbb{E}[\Pi(G(r, \epsilon_R) - C(r, \epsilon_R))], \tag{12}$$

subject to

$$G(r, \epsilon_R) = r\lambda_C(r) + E[R]\lambda_S(r), \tag{13}$$

$$C(r, \epsilon_R) = \min_{U, L} \lim_{T \rightarrow \infty} \frac{1}{T} \mathbb{E} \left[\int_0^T R(t) dL(t) + \int_0^T c(X(t)) dt \right], \tag{14}$$

$$X(t) = Z(t) - U(t) + L(t), \tag{15}$$

$$R(t) = R_1(r) + R_0 e^{\delta Y(t)}, \tag{16}$$

$$Y(t) = M_y(D_S(t)) + L_R(t) - U_R(t). \tag{17}$$

Note that the infinite-horizon subproblem, which is the subject of the next section, has been reduced to the cost-minimization problem (14)–(17). For general demand, spot-price and production processes, there is no clear method to tackle problem (14)–(17) unless we impose some additional Markovian structure. Instead, we use asymptotic analysis (§3) to approximate (14)–(17) by a diffusion-control problem for which an approximate solution is obtained (§4.2). The proposed policy that we derive is rather simple:

Production policy (U): Use a base-stock policy independent of the spot price to control production. That is, keep the machine producing as long as $X(t) \leq \xi$, where $\xi > 0$ is a fixed base-stock level.

Admission policy (L): Reject incoming e-orders when the inventory position is below a threshold level that is an affine function of the price. Specifically, reject an incoming e-order if $X(t) \leq c_0 + c_1 R(t)$ for two negative constants c_0 and c_1 .

The system is producing and accepting e-orders most of the time under this proposed policy. It is only when the inventory level reaches an upper or a lower bound that the decisions of stopping production or rejecting customers are, respectively, exercised. The production and admission thresholds in this policy are proportional to the cost of exercising these actions: the (opportunity) cost of rejecting a customer is given by the price this customer is willing to pay and so the admission policy is proportional to this price, and the cost of stopping production is independent of the price and so a fixed base-stock policy is employed. Although the form of the proposed policy is simple, the closed-form expressions for the constants ξ , c_0 , and c_1 are quite involved, and readers unfamiliar with heavy-traffic analysis might want to skip §3, §4.1, and §4.2.

We conclude this section with a brief discussion on how to extend the previous model to a multiproduct setting. Specifically, in many practical situations different customers require different order sizes. To handle this situation, similar to §2.1, we can introduce a vector of customer classes each with its own arrival process, reservation price distribution, and order size. As in the current model, within each class a fraction of buyers will sign the long-term agreement and the rest will select the spot market. To keep the state-space dimensionality of our control problem in (14)–(15) unchanged, we would need to pool together the multiclass spot-market demand into a single class. This condition boils down to assume that there is a unique per-unit spot-market price with dynamics given by Equation (2). That is, the time t spot price of an order of size i equals i times $R(t)$. To compute $R_1(r)$, we would need to repeat the analysis in §2.2 but replace the expected reservation price $E[\nu_s]$ by the average reservation price across all classes. In addition, the cumulative arrival rate of e-orders would be given by the weighted sum of the arrival rate of each class, where the weights are the order sizes of each class. Finally, on the manufacturing side, orders would join a single queue, and the same single-server production process described in §2.3 would process one unit at a time. Because our solution method is based on a heavy-traffic diffusion approximation, we can neglect the specification of holding/backordering costs at the order level and simply work with the same cost functional $c(X)$, introduced above, at the unit level.

3. The Diffusion-Control Problem

Because (14)–(17) is difficult to solve, we use heavy-traffic theory to approximate it by a diffusion-control problem. For a generic stochastic process $B(t)$ such that $\beta = \lim_{t \rightarrow \infty} B(t)/t$ with probability 1, and for a given positive number n (the heavy-traffic scaling parameter), define

$$\begin{aligned} \hat{B}_n(t) &= \frac{B(nt)}{\sqrt{n}}, & \tilde{B}_n(t) &= \frac{B(nt) - \beta nt}{\sqrt{n}}, & \text{and} \\ \bar{B}_n(t) &= \frac{B(nt)}{n}. \end{aligned} \tag{18}$$

Under technical conditions omitted here, Donsker’s theorem (Billingsley 1999, Whitt 2002) establishes that \hat{B}_n weakly converges to a driftless Brownian motion as $n \rightarrow \infty$. We use this functional central limit theorem (FCLT) to scale the original production, demand, and price processes to obtain a modified version of (14)–(17) that converges to a diffusion-control problem. Throughout this transformation, the main step is to ensure that the different processes that we construct do in fact converge as $n \rightarrow \infty$. In this respect, we use a heuristic derivation of the diffusion-control problem because we do not prove the convergence of the regulators \hat{L}_n and \hat{U}_n , which are the scaled processes used by the system manager to control the scaled netput process \hat{Z}_n .

The first step is to derive the joint convergence of the three primitive renewal processes $M_x(t)$, $M_y(t)$, and $M_s(t)$. Theorem 4.3.5 in Whitt (2000) implies that

$$((\tilde{M}_x)_n, (\tilde{M}_y)_n, (\tilde{M}_s)_n) \implies (\sigma_x \mathcal{W}_x, \sigma_y \mathcal{W}_y, \sigma_s \mathcal{W}_s), \tag{19}$$

where $(\mathcal{W}_x, \mathcal{W}_y, \mathcal{W}_s)$ is a $(0, \Sigma)$ -Brownian motion with covariance matrix

$$\Sigma \triangleq \begin{pmatrix} 1 & \hat{\rho} & 0 \\ \hat{\rho} & 1 & 0 \\ 0 & 0 & 1 \end{pmatrix}$$

We now use (19) to derive the convergence of our control problem. From (18) and the piecewise linearity of the cost function $c(x)$, we have that (14)–(17) is equivalent to

$$\begin{aligned} \mathcal{E}_n(r, \epsilon_R) &\triangleq \frac{C(r, \epsilon_R)}{\sqrt{n}} \\ &= \min_{\hat{U}_n, \hat{L}_n} \lim_{T \rightarrow \infty} \frac{1}{T} E \left[\int_0^T \bar{R}_n(t) d\hat{L}_n(t) + \int_0^T c(\hat{X}_n(t)) dt \right], \end{aligned} \tag{20}$$

$$\hat{X}_n(t) = \sqrt{n}(1 - \rho)t + \tilde{S}_n(t) - \tilde{D}_n(t) + \hat{L}_n(t) - \hat{U}_n(t), \tag{21}$$

$$\bar{R}_n(t) = \frac{R_1(r)}{n} + \frac{R_0}{n} \exp(\sqrt{n} \delta \hat{Y}_n(t)), \tag{22}$$

$$\hat{Y}_n(t) = \sum_{i=1}^{D_S(nt)} \frac{y_i}{\sqrt{n}} - (\hat{U}_R)_n(t) + (\hat{L}_R)_n(t). \tag{23}$$

Because the heavy-traffic scaling makes the original average cost $C(r, \epsilon_R)$ be order \sqrt{n} , we introduce in (20) the scaled version $\mathcal{E}_n(r, \epsilon_R)$ that remains bounded as n increases.

To guarantee convergence of \hat{X}_n , we invoke the *heavy-traffic condition*, which assumes the existence of a bounded real θ such that $\lim_{n \rightarrow \infty} \sqrt{n}(1 - \rho) = \theta$, where ρ is the traffic intensity. Corollary 13.8.1 in Whitt (2002) implies that

$$\tilde{S}_n(t) - \tilde{D}_n(t) \implies \lambda_S(r) \sigma_x \mathcal{W}_x(\lambda_S(r)t) - \sigma_s \mathcal{W}_s(t) \triangleq \mathcal{W}_Z(t),$$

where \mathcal{W}_Z is a $(0, \sigma)$ -Brownian motion with diffusion parameter $\sigma^2 \triangleq \lambda_S^3(r) \sigma_x^2 + \sigma_s^2$. We define \mathcal{U} and \mathcal{L} as the corresponding limits for \hat{U}_n and \hat{L}_n , respectively. If we

also define $\mathcal{X}(t) = \lim_{n \rightarrow \infty} \widehat{X}_n(t)$, then the limiting behavior of (21) is

$$\mathcal{X}(t) = \mathcal{Z}(t) + \mathcal{L}(t) - \mathcal{U}(t),$$

where $\mathcal{Z}(t) = \theta t + \mathcal{W}_Z(t)$ is a (θ, σ^2) Brownian motion.

From the renewal-reward FCLT for regenerative process (see Whitt 2002, Chapter 13), we have that the first summand on the right-hand side of the equality in (23) weakly converges to

$$\sum_{i=1}^{D_S(nt)} \frac{y_i}{\sqrt{n}} \implies \sigma_y \mathcal{W}_y(\lambda_S(r)t) \triangleq \mathcal{W}_R(t), \quad (24)$$

where $\mathcal{W}_R(t)$ is a $(0, \sigma_R)$ -Brownian motion with variance $\sigma_R^2 \triangleq \lambda_S(r)\sigma_y^2$. Therefore, to ensure convergence of the spot-price process in (22), we require $\sqrt{n}\widehat{\delta}$ to remain bounded as n increases, and assume the existence of a positive scalar $\widehat{\delta}$ such that

$$\lim_{n \rightarrow \infty} \sqrt{n}\widehat{\delta} = \widehat{\delta}. \quad (25)$$

In addition, to ensure a nontrivial limit for \bar{R}_n , we need the ratios $R_1(r)/n$ and R_0/n to remain positive for large n . We let $\mathcal{R}_1(r)$ and \mathcal{R}_0 be these limits, i.e.,

$$\lim_{n \rightarrow \infty} \frac{R_1(r)}{n} = \mathcal{R}_1(r) \quad \text{and} \quad \lim_{n \rightarrow \infty} \frac{R_0}{n} = \mathcal{R}_0. \quad (26)$$

Finally, let \mathcal{U}_R and \mathcal{L}_R be the corresponding limits of the two-sided regulators for $Y(t)$, and define $\mathcal{Y}(t) = \widehat{\delta}\mathcal{W}_R(t) + \mathcal{L}_R(t) - \mathcal{U}_R(t)$ to be the resulting regulated Brownian motion (RBM) on $[y_{\min}, y_{\max}]$. Hence, by (24)–(26) and the continuous mapping theorem (e.g., Whitt 2002), we conclude that

$$\mathcal{R}(t) \triangleq \lim_{n \rightarrow \infty} \bar{R}_n(t) = \mathcal{R}_1(r) + \mathcal{R}_0 e^{\mathcal{Y}(t)}. \quad (27)$$

To review the practical implications of the heavy-traffic approximation, the price has to be large enough (order n) to make the penalty of rejection have the same order of magnitude as the holding/backordering cost, and the diffusion parameter, $\widehat{\delta}$, has to be small (order $1/\sqrt{n}$) to keep the price variations bounded. For example, if $n = 100$ and the holding and backordering cost rates are order 1, then the minimum and initial prices (R_1, R_0) have to be order 100, the diffusion $\widehat{\delta}$ has to be order 0.1, and the idleness rate $(1 - \rho)$ has to be order 0.1.

Taking the limit as $n \rightarrow \infty$ in (20)–(22), we obtain

$$\mathcal{C}_\infty(r, \epsilon_R) = \min_{\mathcal{U}, \mathcal{L}} \lim_{T \rightarrow \infty} \frac{1}{T} \mathbb{E} \left[\int_0^T \mathcal{R}(t) d\mathcal{L}(t) + \int_0^T c(\mathcal{X}(t)) dt \right], \quad (28)$$

$$\mathcal{X}(t) = \mathcal{Z}(t) + \mathcal{L}(t) - \mathcal{U}(t), \quad (29)$$

$$\mathcal{R}(t) = \mathcal{R}_1(r) + \mathcal{R}_0 e^{\mathcal{Y}(t)}, \quad (30)$$

where \mathcal{Z} is a (θ, σ) -Brownian motion and \mathcal{Y} is a $(0, \widehat{\delta}\sigma_R)$ -RBM on $[y_{\min}, y_{\max}]$. To complete this diffusion formulation, we note that the covariance between \mathcal{Z} and \mathcal{Y} is given by $\widehat{\delta}\lambda_S^2(r)\sigma_x\sigma_y\widehat{\rho}$. Finally, for future references, we define

the diffusion and correlation coefficients

$$\delta \triangleq \widehat{\delta}\sigma_R \quad \text{and} \quad \rho \triangleq \frac{\widehat{\delta}\lambda_S^2(r)\sigma_x\sigma_y\widehat{\rho}}{\sigma\delta} = \frac{\lambda_S^2(r)\sigma_x\widehat{\rho}}{\sqrt{\lambda_S^4(r)\sigma_x^2 + \lambda_S(r)\sigma_S^2}}.$$

4. Analysis of the Diffusion-Control Problem

This section provides numerical and approximate analytical solutions to Problem (28)–(30). The optimality conditions are derived in §4.1, an approximate analytical solution is proposed in §4.2, a set of alternative approximate solutions is offered in §4.3 as a basis for comparison, and a performance comparison of the numerically computed policy and the policies in §§4.2 and 4.3 are compared in §4.4.

4.1. Optimality Conditions

Without loss of generality, we assume that $\mathcal{R}_0 = 1$, which is equivalent to normalizing the price coefficient $\mathcal{R}_1(r)$ and the holding and backordering cost rates, h and b , by \mathcal{R}_0 . This allows us to express the state space of our control problem as the pair $(\mathcal{X}(t), \mathcal{Y}(t))$ describing the inventory level and the natural logarithm of the price (minus $\mathcal{R}_1(r)$) at time t .

The optimality conditions are stated in terms of the minimal average cost, or gain, g , and the value function $V(x, y)$, which represents the additional cost incurred under the optimal policy when the initial state is (x, y) instead of an arbitrary reference state. Slightly abusing notation, we define the function $\mathcal{R}(y) = \mathcal{R}_1(r) + \mathcal{R}_0 e^y$, which represents the spot-market price associated with the particular value y of $\mathcal{Y}(t)$. The Hamilton-Jacobi-Bellman equation for (28)–(30) is

$$\begin{aligned} [-V_x(x, y)] \wedge [V_x(x, y) + \mathcal{R}(y)] \\ \wedge [\Gamma V(x, y) + c(x) - g] = 0, \end{aligned} \quad (31)$$

where $a \wedge b = \min\{a, b\}$, and

$$\Gamma = \theta \frac{\partial}{\partial x} + \frac{\sigma^2}{2} \frac{\partial^2}{\partial x^2} + \sigma\delta\rho \frac{\partial^2}{\partial xy} + \frac{\delta^2}{2} \frac{\partial^2}{\partial y^2}$$

is the infinitesimal generator of $(\mathcal{X}(t), \mathcal{Y}(t))$. The derivation of this equation is in §1 of the online companion at <http://or.pubs.informs.org/Pages/collect.html> and follows standard arguments (e.g., Taksar 1985) by deriving optimality conditions for the discounted problem (e.g., Harrison and Taksar 1983) and using a Tauberian argument as the discount rate goes to zero. This derivation is heuristic, however, because it assumes that the value function is twice continuously differentiable, an assumption known as the “principle of smooth fit” (Beneš et al. 1980).

Problem (28)–(30) is a two-dimensional singular control problem: Because the controls \mathcal{L} and \mathcal{U} can be exerted

instantaneously (i.e., in an unbounded manner), the resulting optimal control processes are continuous but singular (i.e., the set of time points at which they increase have measure zero). Consequently, the solution to our singular control problem is characterized by boundaries in $(\mathcal{X}, \mathcal{Y})$ space, and the controls are exerted (in a minimal fashion, similar to (4)–(5)) only when the two-dimensional diffusion process hits one of the boundaries.

The first two bracketed terms in (31) are the boundary conditions. Given the smoothness assumption, we can rewrite the second term, $V_x(x, y) + \mathcal{R}(y) = 0$, as

$$V(x, y) - V(x + \epsilon, y) = \mathcal{R}(y)\epsilon + o(\epsilon). \tag{32}$$

Given a price $\mathcal{R}(y)$, the left side of (32) is the marginal cost of accepting an order of size ϵ , and the right side is the marginal cost of rejecting an order of the same size. Therefore, this condition defines the set of states (x, y) for which the manager is indifferent between accepting or rejecting orders at the marginal price $\mathcal{R}(y)$. Similarly, $V_x(x, y) = 0$ characterizes those states for which the manager is indifferent between keeping the machine working or turning it off.

These observations allow us to re-express Equation (31) in terms of two boundary curves (see Figure 2), the rejection boundary $x = \eta(y)$ and the idleness boundary $x = \xi(y)$, which are defined implicitly by

$$V_x(\eta(y), y) + \mathcal{R}(y) = 0 \quad \text{and} \quad V_x(\xi(y), y) = 0. \tag{33}$$

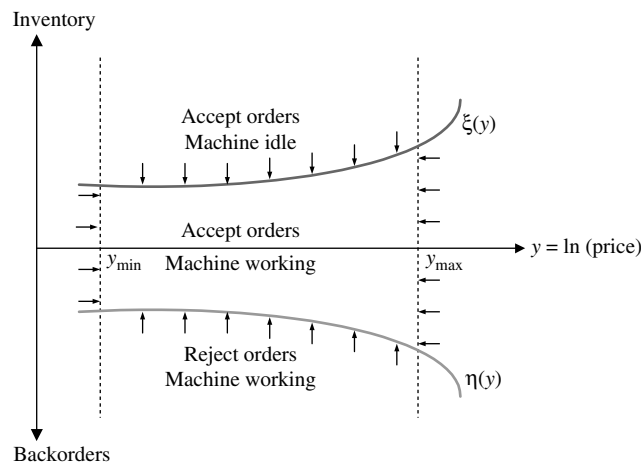
However, some additional notation is needed. For a given pair $(\eta(y), \xi(y))$, we define the region

$$\Omega = \{(x, y) : y_{\min} < y < y_{\max} \text{ and } \eta(y) < x < \xi(y)\}, \tag{34}$$

and divide the boundary $\partial\Omega$ of Ω into four components:

$$\begin{aligned} \partial\Omega^\eta &= \{(x, y) \in \partial\Omega : x = \eta(y)\}, \\ \partial\Omega^{\min} &= \{(x, y) \in \partial\Omega : y = y_{\min}\}, \\ \partial\Omega^\xi &= \{(x, y) \in \partial\Omega : x = \xi(y)\}, \\ \partial\Omega^{\max} &= \{(x, y) \in \partial\Omega : y = y_{\max}\}. \end{aligned} \tag{35}$$

Figure 2. Depiction of the optimal solution. (Arrows indicate the reflection field \vec{v} on the boundary.)



We also introduce the real-valued boundary function $F: \partial\Omega \rightarrow \Re$ and the reflection field $\vec{v}: \partial\Omega \rightarrow \Re^2$:

$$F(x, y) = \begin{cases} \mathcal{R}(y) & \text{if } (x, y) \in \partial\Omega^\eta, \\ 0 & \text{otherwise,} \end{cases} \quad \text{and}$$

$$\vec{v}(x, y) = \begin{cases} (1, 0) & \text{if } (x, y) \in \partial\Omega^\eta, \\ (-1, 0) & \text{if } (x, y) \in \partial\Omega^\xi, \\ (0, 1) & \text{if } (x, y) \in \partial\Omega^{\min}, \\ (0, -1) & \text{if } (x, y) \in \partial\Omega^{\max}. \end{cases} \tag{36}$$

The optimality equation (31) can now be expressed as

$$\Gamma V(x, y) + c(x) - g = 0 \quad \text{for all } (x, y) \in \Omega, \tag{37}$$

$$-\mathcal{R}(y) \leq V_x(x, y) \leq 0 \quad \text{for all } (x, y) \in \Omega, \tag{38}$$

$$\vec{v}(x, y) \cdot \nabla V(x, y) = -F(x, y) \quad \text{for all } (x, y) \in \partial\Omega, \tag{39}$$

$$V_{xx}(\eta(y), y) = V_{xx}(\xi(y), y) = 0 \quad \text{for all } y_{\min} \leq y \leq y_{\max}, \tag{40}$$

where ∇ is the gradient operator. This is a two-dimensional elliptic partial differential equation (PDE) problem with free boundary conditions. As depicted in Figure 2, the physical intuition behind this system is the following. By (37)–(38), if the inventory level x satisfies $\eta(y) < x < \xi(y)$, then the system manager should be accepting e-orders and producing. Equation (39) specifies the boundary conditions for the curves $x = \eta(y)$ and $x = \xi(y)$. If $x = \eta(y)$, then incoming e-orders should be rejected; and if $x = \xi(y)$, then production should stop. Finally, condition (40) imposes the smoothness of the value function $V(x, y)$ at the boundaries.

4.2. Proposed Policy

Closed-form solutions to multidimensional free boundary problems are rare, and we have been unable to solve (37)–(40) analytically. Therefore, we resort to analytical approximations in this subsection. Our approximate analysis of (37)–(40) modifies the angles of reflection at the boundaries, which reduces this PDE system to a simpler ordinary differential equation (ODE) system for which explicit solutions are presented. Suppose for the moment that the switching curves $\eta(y)$ and $\xi(y)$ are fixed and let $g(\eta, \xi)$ be the corresponding average cost associated with these boundaries. The optimality conditions (37) and (39) become

$$\Gamma V(x, y) = g(\eta, \xi) - c(x) \quad \text{for all } (x, y) \in \Omega, \tag{41}$$

$$\vec{v}(x, y) \cdot \nabla V(x, y) = -F(x, y) \quad \text{for all } (x, y) \in \partial\Omega. \tag{42}$$

Problem (41)–(42) is a two-dimensional elliptic PDE problem with oblique reflection \vec{v} . For this problem, let us

denote by $\pi_\Omega = \{\pi_\Omega(x, y) : (x, y) \in \Omega\}$ the steady-state distribution of $(\mathcal{X}, \mathcal{Y})$ in Ω . Combining Green's identity and the divergence theorem, it has been shown (Harrison and Williams 1987) that the stationary distribution π satisfies the following *basic adjoint relation* (BAR):

$$\int_\Omega \Gamma f \pi_\Omega \, ds + \frac{1}{2} \int_{\partial\Omega} \vec{v} \cdot \nabla f \pi_\Omega \, dl = 0$$

for all test functions f twice
 continuous and bounded, (43)

where ds denotes integration with respect to Lebesgue measure on Ω and dl denotes integration with respect to surface measure on $\partial\Omega$. Suppose we consider the test function $f = V(x, y)$ given by (41)–(42). Then, the BAR condition implies

$$g(\eta, \xi) = \int_\Omega c(x) \pi_\Omega(x, y) \, ds + \frac{1}{2} \int_{\partial\Omega} F(x, y) \pi_\Omega(x, y) \, dl. \quad (44)$$

Intuitively, (44) relates the average cost $g(\eta, \xi)$ to the average holding/backordering cost and the average cost of rejecting orders. Because our objective is to minimize this average cost, we formulate the following calculus-of-variations problem:

$$\min_{\eta, \xi} g(\eta, \xi) = \int_\Omega c(x) \pi_\Omega(x, y) \, ds + \frac{1}{2} \int_{\partial\Omega} F(x, y) \pi_\Omega(x, y) \, dl, \quad (45)$$

subject to π_Ω being the steady-state distribution of $(\mathcal{X}, \mathcal{Y})$ in Ω . (46)

Note that the objective function (45) depends on η and ξ through Ω and its boundary $\partial\Omega$. Interestingly, despite its complicated structure, the control problem (45)–(46) satisfies the following simple property (see §2 of the online companion for a proof).

PROPOSITION 1. *At optimality, the time-average fraction of the time that the system holds inventory equals the critical fractile $b/(h+b)$.*

This familiar newsvendor type of result is known to hold in a wide variety of make-to-stock queueing models with constant prices (e.g., Veatch and Wein 1996, Rubio and Wein 1996). Unfortunately, Proposition 1 alone does not fully characterize an optimal solution. Although our proposed policy does not satisfy Proposition 1, in §4.3 we compare this policy to another that does.

The main obstacle to solving this problem is constraint (46): Characterizing the stationary distribution π_Ω for an arbitrary reflection field \vec{v} is as demanding as solving the original PDE (41)–(42). However, a major simplification arises if we modify the reflection field \vec{v} by replacing (42) with the standard Neumann condition (John 1982),

i.e., substituting the inward unit normal vector field \vec{n} on $\partial\Omega$ (*normal derivative*) for \vec{v} . This substitution allows us to approximate the steady-state distribution of $(\mathcal{X}, \mathcal{Y})$ by an exponential. Note that we are modifying only the dynamics of the process on the boundary, where the controlled process spends little time (i.e., a set of measure zero).

We summarize the implications of this approximation in the following proposition, whose statement requires some additional notation and whose proof is in §3 of the online companion. Let Σ be the covariance matrix of $(\mathcal{X}, \mathcal{Y})$. Let V be the rotation matrix whose rows are the orthonormal eigenvectors of Σ and let E be the corresponding diagonal matrix of eigenvalues such that $\Sigma = V'EV$. We define the linear operator $T = E^{-1/2}V$ and $\Omega^* = T(\Omega)$.

PROPOSITION 2. *Suppose that $T\vec{v}$ is normal to $\partial\Omega^*$. Then, the steady-state distribution of $(\mathcal{X}, \mathcal{Y})$ has the exponential form*

$$\pi_\Omega(x, y) = K_\Omega \exp(m_x x + m_y y),$$

where $m_x = \frac{2\theta}{\sigma^2(1-\rho^2)}$, $m_y = \frac{-2\rho\theta}{\sigma\delta(1-\rho^2)}$, (47)

and K_Ω is a positive constant such that π_Ω integrates to one in Ω . In addition, from the definitions of Ω , $\partial\Omega$, and $F(x, y)$, we have that (45)–(46) is equivalent to

$$\min_{\eta, \xi} g(\eta, \xi) = \int_{y_{\min}}^{y_{\max}} \left[\underbrace{\frac{\sigma\sqrt{1-\rho^2}}{2} \mathcal{R}(y) \pi_\Omega(\eta(y), y)}_{M(\eta, \xi)} + \underbrace{\sqrt{\sigma^2 - 2\sigma\delta\rho\eta_y(y) + \delta^2\eta_y^2(y)} + \int_{\eta(y)}^{\xi(y)} c(x) \pi_\Omega(x, y) \, dx}_{M(\eta, \xi)} \right] dy \quad (48)$$

$$\text{subject to } \int_{y_{\min}}^{y_{\max}} \underbrace{\int_{\eta(y)}^{\xi(y)} \pi_\Omega(x, y) \, dx}_{N(\eta, \xi)} dy = 1. \quad (49)$$

We note that the normal reflection assumption not only allows us to replace the steady-state distribution of $(\mathcal{X}, \mathcal{Y})$ by an exponential but also provides a simple approximation for the surface measure of $(\mathcal{X}, \mathcal{Y})$ in Ω^η , which is needed to compute the average rejection cost (i.e., first summand inside the integral in (48)).

To solve (48)–(49), we define the Hamiltonian $H(\eta, \xi) = M(\eta, \xi) - \gamma N(\eta, \xi)$, where γ is the Lagrangian multiplier for (49) and $M(\eta, \xi)$ and $N(\eta, \xi)$ are defined in (48)–(49). The Euler-Lagrange necessary conditions for optimality are (e.g., Gelfand and Fomin 1963)

$$\frac{\partial H}{\partial \xi} - \frac{d}{dy} \left(\frac{\partial H}{\partial \xi_y} \right) = 0 \quad \text{and} \quad \frac{\partial H}{\partial \eta} - \frac{d}{dy} \left(\frac{\partial H}{\partial \eta_y} \right) = 0,$$

which in this case are equivalent to

$$c(\xi(y)) - \gamma = 0 \quad \text{and} \quad (50)$$

$$-c(\eta(y)) + \gamma = \frac{\sigma\sqrt{1-\rho^2}}{2} \left[\mathcal{R}(y)A(\eta_y) + \frac{d}{dy}(\mathcal{R}(y)B(\eta_y)) \right], \quad (51)$$

where

$$A(\eta_y) = \frac{-2\theta}{\sqrt{\sigma^2 - 2\sigma\delta\rho\eta_y(y) + \delta^2\eta_y^2(y)}} \quad \text{and} \quad (52)$$

$$B(\eta_y) = \frac{\delta^2\eta_y - \sigma\delta\rho}{\sqrt{\sigma^2 - 2\sigma\delta\rho\eta_y(y) + \delta^2\eta_y^2(y)}}.$$

Hence, the assumption that the reflection field is normal to the boundary allows us to replace the original PDE system by an ODE system that has two independent equations for η and ξ . In this case, (50) is not even a differential equation and yields

$$\xi(y) \triangleq \bar{\xi} = \frac{\gamma}{h}, \quad (53)$$

which implies that the idling boundary, $\xi(y)$, is characterized by a base-stock policy, $\bar{\xi}$, that is independent of the price.

Although we have been unable to solve the nonlinear ODE (51) in closed form, we have been able to construct an approximate solution by studying its asymptotic behavior as $y_{\min} \rightarrow -\infty$ and $y_{\max} \rightarrow +\infty$. In particular, we derive an admission policy that is *asymptotically consistent* with (51) in the following sense. Suppose we rewrite the differential equation as $\eta(y) = F(\mathcal{R}(y), \eta_y, \eta_{yy})$ for some suitable function F . Then, our proposed solution $\eta(y)$ satisfies

$$\lim_{y \rightarrow \pm\infty} \frac{\eta(y)}{\mathcal{R}(y)} = \lim_{y \rightarrow \pm\infty} \frac{F(\mathcal{R}(y), \eta_y, \eta_{yy})}{\mathcal{R}(y)}.$$

Section 4 of the online companion details the steps of this asymptotic analysis, which results in the following proposed policy.

Proposed policy:

$$\xi^P(y) = \bar{\xi} \quad \text{and} \quad (54)$$

$$\eta^P(y) = \bar{\eta} - \frac{\sigma\delta\sqrt{1-\rho^2}}{2b}(\mathcal{R}(y) - \mathcal{R}_1(r)).$$

Depending on the value of θ , the scalars $\bar{\xi}$ and $\bar{\eta}$ satisfy

$$\bar{\xi} = \sqrt{\frac{\sigma^2\sqrt{1-\rho^2}b\mathcal{R}_1(r)}{h(h+b)}} \quad \text{and} \quad (55)$$

$$\bar{\eta} = -\sqrt{\frac{\sigma^2\sqrt{1-\rho^2}h\mathcal{R}_1(r)}{b(h+b)}} \quad \text{if } \theta = 0, \text{ or}$$

$$h\bar{\xi} + b\bar{\eta} + \frac{\theta}{\sqrt{1-\rho^2}}\mathcal{R}_1(r) = 0 \quad \text{and} \quad (56)$$

$$\left(\frac{h}{h+b}\right)\exp(m_x\bar{\xi}) + \left(\frac{b}{h+b}\right)\exp(m_x\bar{\eta}) = 1 \quad \text{if } \theta \neq 0.$$

A few comments about our proposed policy are now in order. For general price processes, Equation (54) implies that η is an affine function of the price, i.e., $\eta = c_0 + c_1\mathcal{R}$ for appropriate constants c_0 and c_1 . Interestingly, the drift of the inventory process, θ , affects only the intercepts $\bar{\xi}$ and $\bar{\eta}$ but not the slope of the admission or production policies. Only the backorder cost b and diffusion parameters σ , δ , and ρ modulate the slope of the admission/rejection policy. As a general rule, an increase in the variance of the $(\mathcal{X}, \mathcal{Y})$ process (measured by $\sigma\delta\sqrt{1-\rho^2}$) results in an increase in the price sensitivity of η^P . As expected, in the proposed policy the base-stock level increases with the backorder cost rate b and decreases with the holding cost h . The opposite conclusion holds for the accept-reject threshold.

4.3. Alternative Policies

We consider three alternative policies for purposes of comparison. To assess the importance of using a price-dependent policy, we first consider a static policy, where admission and production decisions are independent of the price, i.e., $\eta(y)$ and $\xi(y)$ are constant. Let (ξ^S, η^S) be the optimal static solution. Again, closed-form solutions are only available if $\theta = 0$ or $h = b$. As in (55)–(56), this static solution solves

$$h\xi^S + b\eta^S = 0 \quad \text{and}$$

$$2h\xi^S(\xi^S - \eta^S) - \sigma^2\sqrt{1-\rho^2}\mathcal{R}^S - h(\xi^S)^2 - b(\eta^S)^2 = 0 \quad \text{if } \theta = 0,$$

$$h\xi^S + b\eta^S + \frac{\theta}{\sqrt{1-\rho^2}}\mathcal{R}^S = 0 \quad \text{and}$$

$$\left(\frac{h}{h+b}\right)\exp(m_x\xi^S) + \left(\frac{b}{h+b}\right)\exp(m_x\eta^S) = 1 \quad \text{if } \theta \neq 0,$$

where

$$\mathcal{R}^S = \frac{\int \mathcal{R}(y)\exp(m_y y) dy}{\int \exp(m_y y) dy}$$

is the average steady-state spot price. We note that for this static policy, the fact that $\eta(y)$ and $\xi(y)$ are constant implies that the normal reflection field assumption holds exactly. Therefore, the solution (ξ^S, η^S) is optimal within the family of static price-independent production and admission control policies.

Next, to assess the importance of capturing the dynamic and continuous nature of the spot price, we consider a policy that assumes i.i.d. spot prices. In this context, the manufacturer’s state space reduces to \mathcal{X} , the inventory position, and both production and admission decisions are based on this single state variable. To keep this model as close as possible to our original formulation, we assume that the distribution of the i.i.d. spot prices coincides with the steady-state distribution of our continuous price process.

Specifically, the limiting price process is given by $\mathcal{R}_1(r) + \mathcal{R}_0 \exp(\mathcal{Y})$, where \mathcal{Y} is uniformly distributed in $[y_{\min}, y_{\max}]$ (see §5.5 in Harrison 1985).

Computing the optimal production and admission policies when prices are i.i.d. is an open problem for general arrival and service processes. For this reason, we have restricted the analysis of this policy to a semi-MDP setting with Poisson arrivals and exponential service times. Section 5 in the online companion (<http://or.pubs.informs.org/Pages/collect.html>) is devoted to the formulation and computation of this i.i.d. price policy, which is based on the work by George and Harrison (2001) and Ata and Shneorson (2006).

The third policy uses the following myopic price-dependent admission/rejection policy: If the state of the system is (x, y) , then admit an incoming e-order only if the immediate reward $\mathcal{R}(y)$ exceeds the expected backordering cost incurred by this order, which is bx^- . Under this simple rule, the resulting admission/rejection boundary is given by

$$\eta^M(y) = -\frac{\mathcal{R}(y)}{b}. \quad (57)$$

This policy models the behavior of a myopic manager who balances revenues and immediate backorder costs on a case-by-case basis without taking into account the full effects that these decisions have on future congestion and costs. Similar to our proposed policy, this myopic admission policy is a linear function of the price. However, in contrast to our proposed policy, the myopic solution is insensitive to changes in the drift and diffusion parameters of the inventory and price processes. Furthermore, it can be shown that this myopic policy can lead to arbitrarily poor performance with respect to our proposed solution when the average spot price goes to infinity. Nevertheless, condition (57) captures in a parsimonious way the trade-off between admission and rejection costs that makes it appealing from a practical standpoint.

Finally, the corresponding base stock, $\xi^M(y)$, for the myopic policy is determined by solving $\min_{\xi} \{g(\eta^M(y), \xi)\}$, where $\eta^M(y)$ is given in (57). We solve this minimization numerically because we do not have an analytical representation of the steady-state distribution of the inventory-price process under this type of policy. However, the optimization is over a single parameter, ξ , which reduces the complexity of the computations.

4.4. Computational Experiments

In this subsection, we compare the policies in §§4.2 and 4.3 to the numerically computed optimal policy, which is denoted by η^N and ξ^N . We employ the Markov chain approximation technique (Kushner and Dupuis 2001) developed by Kushner (1977) to approximate the diffusion-control problem by a control problem for a finite-state Markov chain. However, the linear program that formulates the Markov chain control problem is cumbersome to solve,

and we instead combine the Markov chain approximation with a recent iterative method (Kumar and Muthuraman 2004) for updating the control boundary. Details of this algorithm, and a description of its convergence performance in our base-case example, can be found in §6 of the online companion at <http://or.pubs.informs.org/Pages/collect.html>. Although we do not attempt a proof of convergence of this procedure, in all computational experiments that we have performed the method converges to a solution to (37)–(40) in several iterations.

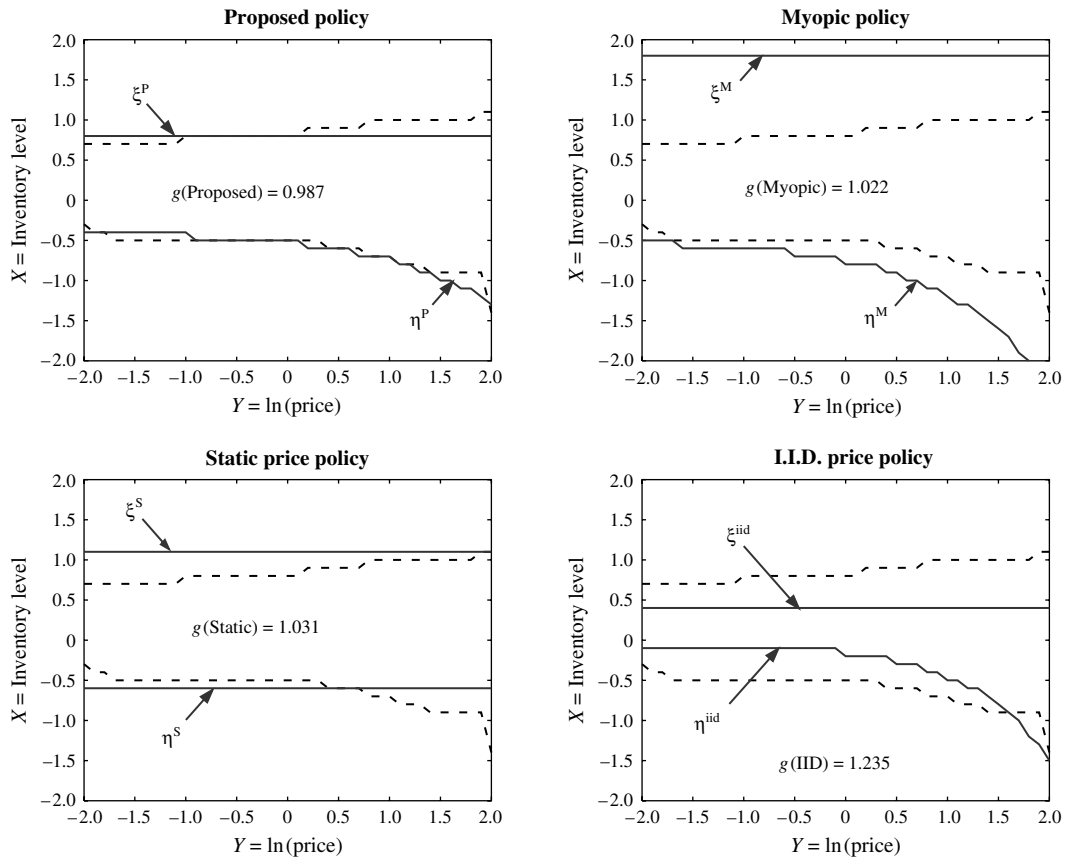
The data for our base case are $\theta = 0$, $\varrho = 0.2$, $\sigma^2 = \delta^2 = 1$, $b = 2h = 2$, $\mathcal{R}_1(r) = 1$, $\mathcal{R}_0 = 0.5$, $N_x = 100$, $-y_{\min} = y_{\max} = 2$, and $\bar{h} = 0.1$. Figure 3 shows the proposed, myopic, static, and i.i.d. price policies, along with the numerically computed optimal solution (dashed line). The proposed policy captures much of the horizontal and exponential shape of ξ and η , respectively, whereas the static policy is unable to describe the price dependency of η . On the other hand, the myopic policy captures this price dependency but rejects too few orders. The i.i.d. price policy exhibits the opposite behavior rejecting too many orders. Also, the proposed policy coincides with the numerical solution for small values of y . The numerically computed idling boundary ξ^N increases slightly with price, unlike the proposed policy, suggesting that ξ^P can be viewed as a lower bound for the optimal curve. In contrast, the base-stock level suggested by the static and myopic policies overestimate ξ^N . On the other hand, the i.i.d. price policy understocks. All of the above observations hold for other problem instances we encountered. Finally, let us define policy \mathcal{P} 's *suboptimality* by

$$\frac{g(\mathcal{P}) - g(\text{numerical})}{g(\text{numerical})} \times 100\%,$$

where $g(\mathcal{P})$ and $g(\text{numerical})$ are the expected average cost under policy \mathcal{P} and the numerical solution, respectively. The suboptimality of the proposed policy is 1.1%, while the static policy is 5.6% suboptimal, primarily because it rejects too many high-priced e-orders when the system is moderately backordered. The myopic policy is 4.7% suboptimal due to its high base-stock level. The i.i.d. price policy is 25.6% suboptimal because it rejects too many orders and holds insufficient inventory.

To assess the robustness of the four policies, we look at the effects on suboptimality of the backorder cost b , the price coefficient $\mathcal{R}_1(r)$, and the process variance-covariance data σ^2 , δ^2 , and ϱ . In each of these cases, we use the base-case data and change the value of the corresponding parameter. Table 1 presents policy suboptimality as a function of the backorder cost b and price level $\mathcal{R}_1(r)$. The proposed policy performs systematically better than the other strategies in Table 1. On average, as b varies, the suboptimality for the proposed policy is 1.63%, compared to 5.26% for the static policy, 6.44% for the myopic policy, and 31.91% for the i.i.d. price policy. Similarly, the proposed policy dominates the other three solutions as $\mathcal{R}_1(r)$ varies. Interestingly, the myopic policy is

Figure 3. Proposed (ξ^P, η^P), myopic (ξ^M, η^M), static (ξ^S, η^S), and i.i.d. price (ξ^{iid}, η^{iid}) policies.



Note. The dashed line corresponds to the numerically computed optimal solution, which has a steady-state average cost $g(\text{numerical}) = 0.976$.

much more sensitive to this parameter than the other policies and deteriorates rapidly as the fixed price component $\mathcal{R}_1(r)$ increases.

Tables 2 and 3 present the policy suboptimality as a function of the netput variance σ^2 , the price variance δ^2 , and the correlation coefficient ρ . Once again, the proposed

policy outperforms the static, myopic, and i.i.d. price policies, with average suboptimalities of 1.83%, 6.54%, 4.99%, and 57.34%, respectively.

The poor performance of the i.i.d. price policy with respect to the other three policies highlights the importance of taking into account the continuous-path nature of the

Table 1. Policy suboptimality as a function of the backorder cost (b) and the price level ($\mathcal{R}_1(r)$).

b	Prop. (%)	Static (%)	Myopic (%)	I.I.D. (%)
0.5	0.12	10.22	17.49	41.99
1	2.54	9.34	9.98	38.50
2	1.13	5.64	4.73	25.62
4	1.61	2.94	2.97	31.00
8	1.92	2.38	2.01	23.42
16	2.44	1.04	1.45	30.90
$\mathcal{R}_1(r)$	Prop. (%)	Static (%)	Myopic (%)	I.I.D. (%)
1	1.13	5.64	4.73	25.62
2	1.67	3.92	12.54	33.31
4	1.64	3.26	29.39	33.72
6	1.86	3.19	44.89	32.30
8	1.92	2.86	58.82	31.94
10	3.55	4.51	74.66	35.59

Table 2. Policy suboptimality as a function of the netput variance (σ^2) and the price variance (δ^2).

σ^2	Prop. (%)	Static (%)	Myopic (%)	I.I.D. (%)
0.5	0.75	3.91	14.45	6.82
1	1.13	5.64	4.73	25.62
2	2.10	5.53	1.41	70.21
4	1.77	6.14	1.71	108.71
8	2.66	6.38	5.29	183.15
16	2.67	6.25	9.87	291.69
δ^2	Prop. (%)	Static (%)	Myopic (%)	I.I.D. (%)
0.5	2.09	5.69	6.69	26.41
1	1.13	5.64	4.73	25.62
2	1.33	6.55	3.35	29.04
4	2.07	9.36	3.78	33.24
8	2.99	11.94	3.79	36.12
16	2.13	13.56	2.89	36.85

Table 3. Performance suboptimality as a function of the correlation coefficient (ρ).

ρ	Prop. (%)	Static (%)	Myopic (%)	I.I.D. (%)
0	3.13	5.69	6.69	26.41
0.1	2.72	5.92	5.89	26.81
0.2	1.13	5.64	4.73	25.62
0.3	0.76	4.88	3.27	25.83
0.4	0.45	4.02	2.71	25.98
0.5	1.88	5.03	3.76	27.98

price process. As we can see from this set of computational experiments, assuming i.i.d. prices (as most of the research on admission control does) when actual prices are highly correlated can have a deleterious impact on overall performance.

5. Optimal Long-Term Contract Price

In this section, we use the results in §4 to approximate the optimal value of $\mathcal{C}(r, \epsilon_R)$ in (14), and then solve Problem (12)–(14) for the optimal long-term contract price.

The first step is to reverse the heavy-traffic scaling to express (48) and (54) in terms of the original unscaled parameters. From §3, we recall that

$$\theta = \sqrt{n}(1 - \rho), \quad \mathcal{R}(y) = \frac{R(y)}{n}, \quad \mathcal{C}(r, \epsilon_R) = \frac{C(r, \epsilon_R)}{\sqrt{n}},$$

$$\delta = \sqrt{n}\tilde{\delta} = \sqrt{n}\hat{\delta}\sqrt{\lambda_S(r)\sigma_y}, \quad \sigma = \sqrt{\lambda_S^3(r)\sigma_x^2 + \sigma_s^2} \quad \text{and}$$

$$\rho = \frac{\lambda_S^2(r)\sigma_x\hat{\rho}}{\sqrt{\lambda_S^4(r)\sigma_x^2 + \lambda_S(r)\sigma_s^2}}.$$

In addition, the heavy-traffic transformation scales the inventory process $X(t)$ by \sqrt{n} , i.e., $\mathcal{X}(t) = X(nt)/\sqrt{n}$ (see §3 for details). Therefore, the unscaled switching boundaries $\eta^u(y)$ and $\xi^u(y)$ (the superscript u stands for *unscaled*) are equal to

$$\eta^u(y) = \sqrt{n}\eta(y) \quad \text{and} \quad \xi^u(y) = \sqrt{n}\xi(y).$$

On the other hand, the range $[y_{\min}, y_{\max}]$ of the regulated Brownian motion that modulates the spot price remains unchanged under the heavy-traffic transformation because $\mathcal{Y}(t) = Y(nt)$.

Under these transformations, and after some straightforward manipulations, we can rewrite the cost-rate function as follows:

$$C(r, \epsilon_R) = \int_{y_{\min}}^{y_{\max}} \left[\frac{\sigma\sqrt{1-\rho^2}}{2} R(y)\pi_{\Omega}^u(\eta^u(y), y) \cdot \sqrt{\sigma^2 - 2\sigma\tilde{\delta}\rho\eta^u(y) + \hat{\delta}^2(\eta^u(y))^2} + \int_{\eta^u(y)}^{\xi^u(y)} c(x)\pi_{\Omega}^u(x, y) dx \right] dy, \quad (58)$$

where

$$\xi^u(y) = \bar{\xi}^u \quad \text{and} \quad \eta^u(y) = \bar{\eta}^u - \frac{\sigma\hat{\delta}\sqrt{1-\rho^2}}{2b}(R(y) - R_1(r)). \quad (59)$$

The values of π_{Ω}^u , $\bar{\xi}^u$, and $\bar{\eta}^u$ are computed from (47), (55), and (56) respectively, using the unscaled values $\theta \leftarrow 1 - \rho$, $\delta \leftarrow \hat{\delta}$, and $\mathcal{R}_1(r) \leftarrow R_1(r)$, yielding an approximation for the original problem that is independent of the heavy-traffic scaling parameter n . Given the exponential utility function of the manufacturer and the fact that $E[R] = E[\nu_S(r)] + \epsilon_R$, we can write the manufacturer's problem as

$$\min_r E \left[-\exp(-\beta(r\lambda_C(r) + (E[\nu_S(r)] + \epsilon_R)\lambda_S(r)) - C(r, \epsilon_R)) \right] \cdot \beta^{-1}. \quad (60)$$


Although all the pieces are in place for solving this optimization problem, a closed-form solution is not available because of the functional form of $C(r, \epsilon_R)$. However, the computations can be simplified by the empirical observation (see Figure 4) that $C(r, \epsilon_R)$ is approximately linear in ϵ_R , with a slope that is independent of r , i.e., there is a function $C_0(r)$ and a constant C_{ϵ} such that

$$C(r, \epsilon_R) \approx C_0(r) + C_{\epsilon}\epsilon_R. \quad (61)$$

Using (61), we can approximate (60) by

$$\begin{aligned} \min_r \exp(-\beta(r\lambda_C(r) + E[\nu_S(r)]\lambda_S(r) - C_0(r))) \\ \cdot E[\exp(-\beta(\lambda_S(r) - C_{\epsilon})\epsilon_R)] \\ = \min_r \exp(-\beta(r\lambda_C(r) + E[\nu_S(r)]\lambda_S(r) - C_0(r)) \\ + \beta^2(\lambda_S(r) - C_{\epsilon})^2\sigma_{\epsilon}^2/2), \end{aligned} \quad (62)$$

Figure 4. Effect of ϵ_R on the average cost \mathcal{C}_{ϵ_R} for four different values of r .



where σ_ϵ^2 is the variance of ϵ_R . Note that if $\sigma_\epsilon \rightarrow 0^+$, then the manufacturer solves the risk-neutral version of the problem,

$$\max_r \{r\lambda_C(r) + E[\nu_S(r)]\lambda_S(r) - C_0(r)\},$$

and if the manufacturer is poorly equipped to predict the spot price so that $\sigma_\epsilon \gg 0$, then the optimal price r^* solves

$$\min_r (\lambda_S(r) - C_\epsilon)^2.$$

We now employ (62) to computationally investigate the behavior of three quantities: the long-term contract price r , the manufacturer’s expected utility Π , and the fraction of buyers that use the spot market f_{spot} . We report two normalized quantities,

$$\hat{r} = \frac{r}{E[\nu]} \quad \text{and} \quad \hat{\mathcal{H}} = \frac{\ln(\Pi)}{\Lambda E[\nu]},$$

where \hat{r} measures the relative value of the long-term contract price with respect to the average valuation of the product, and $\hat{\mathcal{H}}$ is the approximate average profit $\ln(\Pi)$ (if \mathcal{P} is the random profit, then $\ln(\Pi) \approx \ln(-E[1 - \mathcal{P}]) = \ln(E[\mathcal{P}] - 1) \approx E[\mathcal{P}]$) divided by the average revenue of a manufacturer that can do perfect price discrimination, $\Lambda E[\nu]$. The independent variables in these experiments are the traffic intensity ρ , the variability of the average spot-market price σ_ϵ , and the ratio of backorder cost to holding cost b/h .

The base-case data are the same as in §4.4, except that $\delta^2 = 2R_0 = 2$ and $\rho = 1.0$ ($\Lambda = 1.0, \mu = 1.0$). In addition, σ^2 cannot be viewed as a constant as in §4.4, and we set $\sigma_x = \sigma_y = \sigma_s = 1$. We also set the risk-aversion parameter $\beta = 1$, the noise parameter $\sigma_\epsilon = 4$, and let the reservation price have a Weibull distribution, $\Phi(r) = 1 - \exp(-ar^\alpha)$, with $a = 0.005$ and $\alpha = 2.0$, so that the reservation price has mean 12.53 and variance 22.76. Finally, we define

$$\psi(r) = \bar{\psi} + (1 - \bar{\psi})e^{-r},$$

where $\bar{\psi} = 0.4$, implying that 40% of customers are regular regardless of the price, and the remaining 60% are potential speculators that speculate with probability $1 - e^{-r}$.

Table 4 shows that the contract price and the fraction of buyers that participate on the spot market are U-shaped as a function of the traffic intensity with a minimum at $\rho = 1$. In contrast, the manufacturer’s expected utility decreases with ρ . Our intuition for this result is that as ρ increases, there is a tendency to accumulate costly backorders. This table also implies that as the manufacturer’s uncertainty about the true mean of the spot-market price increases, the optimal long-term contract price decreases, fewer customers engage in the spot market, and the manufacturer’s expected utility decreases. Table 5 looks at the effects of the backordering-to-holding cost ratio b/h and the correlation

Table 4. Performance measures as a function of the traffic intensity ρ (for fixed $\mu = 1$) and spot-price variability σ_ϵ .

ρ	\hat{r}	$\hat{\mathcal{H}}$	f_{spot}
0.7	0.861	0.425	0.776
0.8	0.819	0.386	0.764
0.9	0.788	0.352	0.754
1.0	0.788	0.321	0.754
1.1	0.792	0.300	0.756
1.2	0.810	0.279	0.761
1.3	0.842	0.260	0.771
σ_ϵ	\hat{r}	$\hat{\mathcal{H}}$	f_{spot}
0.0	1.631	0.590	0.950
0.5	1.612	0.585	0.948
1.0	1.522	0.568	0.935
1.5	1.397	0.542	0.914
2.0	1.265	0.507	0.886
2.5	1.127	0.466	0.853
3.0	0.990	0.421	0.815

coefficient ϱ on the optimal long-term price and manufacturer’s expected revenue. We note that the optimal long-term price and the fraction of buyers participating on the spot market increases with b/h . On the other hand, the manufacturer’s expected revenue decreases with b/h as backorder costs become more significant. Table 5 also suggests that the correlation between the demand and the spot price has a limited impact on the three performance measures.

6. Concluding Remarks

The novelty of our problem formulation stems from the incorporation of two aspects of mathematical finance into a

Table 5. Performance measures as a function of the backordering-to-holding cost ratio b/h (for fixed h) and the coefficient of variation ϱ .

b/h	\hat{r}	$\hat{\mathcal{H}}$	f_{spot}
0.5	0.752	0.378	0.743
1.0	0.766	0.349	0.748
2.0	0.788	0.321	0.754
3.0	0.792	0.308	0.756
4.0	0.777	0.300	0.751
5.0	0.805	0.295	0.758
6.0	0.800	0.292	0.760
ϱ	\hat{r}	$\hat{\mathcal{H}}$	f_{spot}
0.0	0.784	0.314	0.753
0.1	0.792	0.317	0.755
0.2	0.788	0.321	0.754
0.3	0.784	0.325	0.753
0.4	0.796	0.329	0.757
0.5	0.796	0.334	0.757
0.6	0.791	0.339	0.755

traditional operations management model. The first aspect considers a spot market for e-orders whose prices behave as a variant of geometric Brownian motion. The second aspect borrows the portfolio optimization framework as we consider a risk-averse manufacturer that chooses a long-term contract price to optimally mix low-mean, riskless contracts with a high-mean high-variability spot market. Studies in other operations settings have incorporated continuous-time spot-price models (e.g., Huchzermeier and Cohen 1996, Semret and Lazar 1999, Kamrad and Ernst 2001, Caldentey and Haugh 2005) and portfolio models (e.g., Van Mieghem 2003), and we believe that the incorporation of finance models into operations is a fruitful area of research, akin to incorporating marketing models into standard operations management models (e.g., Kurawarwala and Matsuo 1996, van Ryzin and Mahajan 1999).

The embedded diffusion-control problem in §4 is two-dimensional (price and inventory level), and is much more difficult to analyze than the corresponding fixed-price problem (e.g., Wein 1992). Hence, we resort to numerical and approximate analytical methods. We use a numerical algorithm that has its basis in the Markov chain approximation algorithm (Kushner and Dupuis 2001), but uses a recent iterative technique that exploits the structure of singular control problems (Kumar and Muthuraman 2004). While we attempt no proof of convergence, this method performed reliably in our numerical study. We also develop an approximate analytical approach to the diffusion-control problem by altering the reflection field on the free boundary, so as to reduce the PDE system (37)–(40) to the simpler ODE system (50)–(52). The proposed policy derived from (50)–(52) is reasonably accurate and robust (suboptimality ranging between 1% and 4%).

This analytically-derived proposed policy is also quite simple: The machine busy/idle policy is characterized by a base-stock level that is independent of price, and an e-order with price p is accepted if the current inventory level is greater than $c_0 + c_1 p$ for constants c_0 and c_1 . Figure 3 shows that our proposed policy captures the salient features of the numerically computed optimal policy. While the latter policy's base-stock level increases with price, this dependence is quite weak because the price varies on a fast time scale. In contrast, if prices (or demands) were modeled by a Markov-modulated process with long holding times (e.g., Choudhury et al. 1997), then in the heavy-traffic limit one would likely observe stronger price dependence in the busy/idle policy. From a practical point of view, our analysis suggests that price variability impacts the admissions process, not the safety-stock level; in other words, it is not optimal to hold more inventory when the price is high; it is preferable to accept more e-orders.

The proposed policy performs better than a benchmark static policy (5% average suboptimality), where the busy/idle and accept/reject policies are independent of price. This static policy falters by failing to accept enough

high-priced e-orders when the system is moderately back-ordered (see Figure 3). To highlight the implications of our continuous-path price process, we compared our proposed policy to a policy that assumes i.i.d. prices. The average suboptimality of this policy is 40%, which reveals the strong impact that price correlation has on overall performance. We also compared the proposed policy to a myopic policy that accepts an incoming order if the revenue from this order exceeds the expected backorder cost incurred by the order. The optimal base-stock level for this myopic policy is computed numerically. This myopic policy incurs a 6% average suboptimality. The myopic policy's downfall is that by focusing on the immediate balance between backorder costs and revenues, it ignores the longer-term effects of increased congestion on the system. This shortsightedness leads to too many accepted orders, and consequently too much safety stock (Figure 3). In the loss network setting, Kelly (1991) has warned of the dangers of ignoring these “knock-on” congestion effects when making admission-rejection decisions.

Although our formulation has assumed that the system manager reacts to exogenous market prices by admitting or rejecting e-orders, we can also use our proposed policy to quote prices. The rejection boundary $x = \eta(y)$ represents the maximum level of backorders that we are willing to have at a price $\mathcal{R}(y)$. Let $y = \eta^{-1}(x)$ be the inverse of the rejection boundary $\eta(y)$. Then, at a given level x of backorders, the minimum price that we are willing to accept for a new incoming order is equal to $\tilde{\mathcal{R}}(x) \triangleq \mathcal{R}(\eta^{-1}(x))$. Depending on market conditions and the distribution of the bargaining power among the different players, the manufacturer should quote a price that is at least equal to $\tilde{\mathcal{R}}(x)$.

An approximate performance analysis of our proposed policy allows us to reduce the determination of an effective long-term contract price to a succinct optimization problem (Equation (62)), which is computed numerically. As in the portfolio optimization literature, we find that a poorly informed manufacturer (i.e., high value of σ_e) shies away from the spot market and allocates more production capacity to the riskless long-term contracts. A less obvious observation is that the contract price, and hence the fraction of buyers participating in the spot market, decreases as the traffic intensity ρ (potential demand divided by service rate) approaches one. That is, the optimal long-term contract price is minimized in those situations where demand and supply are relatively balanced. Also, as the backorder-to-holding cost ratio increases, the contract price increases and the manufacturer's expected revenue decreases. Finally, our numerical experiments suggest that the correlation coefficient between demand and spot price has a limited effect on the contract price and the manufacturer's expected utility.

In our view, there are several aspects of this problem that deserve further consideration. Our formulation assumes that there are two types of buyers: regular and speculators. However, the issue of how speculators act on the spot market was not properly addressed. It would be interesting to

extend our model to incorporate the strategic behavior of these speculators; initial progress on this difficult issue has been made in other queueing models (e.g., Parlakturk and Kumar 2004, Armony and Maglaras 2004). More generally, an empirical analysis of actual spot-market prices would no doubt shed considerable light on appropriate model formulations. In this regard, we note that our optimality conditions (50)–(51) are valid for general positive diffusion processes, including a mean-reverting process such as the Ornstein-Uhlenbeck process, which may be more realistic than a GRBM in some instances.

Finally, the promising numerical performance of our proposed policy suggests that our approach of tweaking the reflection field to obtain a more tractable problem could potentially provide useful engineering solutions to a variety of two-dimensional singular control problems that occur in queueing network control and finance. However, it could be that our particular problem is well suited to this approach because two of the boundaries are exogenously normal to the boundary (because of the bounded price process), and one of the two free boundaries is nearly normal (because of the near optimality of the price-independent base-stock policy). Hence, theoretical (e.g., deriving error bounds) and further numerical studies are warranted.

Acknowledgments

This research was partially supported by a fellowship from MIT's Sloan School of Management and by the Singapore-MIT Alliance (SMA) program. The first author is grateful to the participants at the seminars held at Carnegie Mellon, Yale, MIT, Northwestern, NYU, Michigan, Columbia, Duke, INSEAD, and Purdue for their comments on a preliminary version of this work and to Martin Haugh for helpful discussions. The second author thanks Mitchell Burman, whose ideas influenced the authors' problem formulation. The authors also thank an anonymous referee for suggesting the i.i.d. price policy and the associate editor for conjecturing Proposition 1 and proposing a comparison to the myopic accept-reject policy in (57).

References

- Armony, M., C. Maglaras. 2004. Contact centers with a call-back option and real-time delay information. *Oper. Res.* **52** 527–545.
- Asmussen, S. 2003. *Applied Probability and Queues*. Springer-Verlag, New York.
- Ata, B., S. Shneorson. 2006. Dynamic control of an M/M/1 service system with adjustable arrival and service rates. *Management Sci.* Forthcoming.
- Beneš, V. E., L. A. Shepp, H. S. Witsenhausen. 1980. Some solvable stochastic control problems. *Stochastics* **4** 39–83.
- Billingsley, P. 1999. *Convergence of Probability Measures*, 2nd ed. Wiley, New York.
- Caldentey, R., M. Haugh. 2005. Optimal control and hedging of operations in the presence of financial markets. *Math. Oper. Res.* Forthcoming.
- Carr, S., I. Duenyas. 2000. Optimal admission control and sequencing in a make-to-stock/make-to-order production system. *Oper. Res.* **48** 709–720.
- Chen, X., D. Simchi-Levi. 2004. Coordinating inventory control and pricing strategies with random demand and fixed ordering cost: The finite horizon case. *Oper. Res.* **52** 887–896.
- Choudhury, G. L., A. Mandelbaum, M. I. Reiman, W. Whitt. 1997. Fluid and diffusion limits for queues in slowly changing environments. *Comm. Statist.-Stochastic Models* **13** 121–146.
- Dixit, A., R. Pindyck. 1994. *Investment Under Uncertainty*. Princeton University Press, Princeton, NJ.
- Dolan, R. J. 1978. Incentive mechanisms for priority queueing problems. *Bell J. Econom.* **9** 421–436.
- Federgruen, A., A. Heching. 1999. Combined pricing and inventory control under uncertainty. *Oper. Res.* **47** 454–475.
- Gelfand, I. M., S. V. Fomin. 1963. *Calculus of Variations*. Prentice-Hall, Englewood Cliffs, NJ.
- George, J. M., J. M. Harrison. 2001. Dynamic control of a queue with adjustable service rate. *Oper. Res.* **49** 720–731.
- Geunes, J., P. M. Pardalos, H. E. Romeijn. 2002. *Supply Chain Management: Models, Applications, and Research Directions*. Kluwer Academic Publishers, Dordrecht, The Netherlands.
- Harrison, J. M. 1985. *Brownian Motion and Stochastic Flow Systems*. John Wiley and Sons, New York.
- Harrison, J. M., M. I. Taksar. 1983. Instantaneous control of Brownian motion. *Math. Oper. Res.* **8** 439–453.
- Harrison, M., R. Williams. 1987. Multidimensional reflected Brownian motions having exponential stationary distributions. *Ann. Probab.* **15** 115–137.
- Helm, W., K. Waldmann. 1984. Optimal control of arrivals to multiserver queues in a random environment. *J. Appl. Probab.* **21** 602–615.
- Huchzermeier, A., M. A. Cohen. 1996. Valuing operational flexibility under exchange rate risk. *Oper. Res.* **44** 100–113.
- John, F. 1982. *Partial Differential Equations*, 4th ed. Springer, New York.
- Kafka, S. J., B. D. Temkin, M. R. Sanders, J. Sharrard, T. O. Brown. 2000. eMarketplaces boost B2B trade. The Forrester Report, Forrester Research, Inc., Cambridge, MA.
- Kamrad, B., R. Ernst. 2001. An economic model for evaluating mining and manufacturing ventures with output yield uncertainty. *Oper. Res.* **49** 690–699.
- Kaplan, S., M. Sawhney. 2000. E-Hubs: The new B2B marketplaces. *Harvard Bus. Rev.* (May–June) 97–104.
- Karlin, S., H. M. Taylor. 1975. *A First Course in Stochastic Processes*, 2nd ed. Academic Press, New York.
- Kelly, F. P. 1991. Loss networks. *Ann. Appl. Probab.* **1** 319–378.
- Keskinocak, P., R. Goodwin, F. Wu, R. Akkiraju, S. Murthy. 2001. Decision support for managing an electronic supply chain. *Electronic Commerce Res. J.* **1** 15–31.
- Kumar, S., M. Muthuraman. 2004. A numerical method for solving singular Brownian control problems. *Oper. Res.* **52** 563–582.
- Kurawarwala, A. A., H. Matsuo. 1996. Forecasting and inventory management of short life-cycle products. *Oper. Res.* **44** 131–150.
- Kushner, H. J. 1977. *Probability Methods for Approximations in Stochastic Control and for Elliptic Equations*. Academic Press, New York.
- Kushner, H. J., P. G. Dupuis. 2001. *Numerical Methods for Stochastic Control Problems in Continuous Time*. Springer Verlag, New York.
- Li, L. 1992. The role of inventory in delivery-time competition. *Management Sci.* **38** 182–197.
- Markowitz, D. M., L. M. Wein. 2001. Heavy traffic analysis of dynamic cyclic policies: A unified treatment of the single machine scheduling problem. *Oper. Res.* **49** 246–270.
- Mendelson, H. 1985. Pricing computer services: Queueing effects. *Comm. ACM* **28** 312–321.
- Merton, R. 1990. *Continuous-Time Finance*. Basil Blackwell, Oxford and Cambridge, UK.

- Naor, P. 1969. The regulation of queue size by levying tolls. *Econometrica* **37** 15–24.
- Parlakturk, A., S. Kumar. 2004. Self interested routing in queueing networks. *Management Sci.* **50**(7) 949–967.
- Plambeck, E. L., S. Kumar, J. M. Harrison. 2001. A multiclass queue in heavy traffic with throughput time constraints: Asymptotically optimal dynamic controls. *Queueing System* **39** 23–54.
- Rubio, R., L. M. Wein. 1996. Setting base stock levels using product-form queueing networks. *Management Sci.* **42** 259–268.
- Schwartz, E. 1997. The stochastic behavior of commodity prices: Implications for valuation and hedging. *J. Finance* **52** 923–973.
- Semret, M., A. A. Lazar. 1999. Spot and derivative markets in admission control. *Proc. 16th Internat. Teletraffic Congress*. Edinburgh, UK, 925–941.
- Stidham, S. 1985. Optimal control of admission to a queueing system. *IEEE Trans. Automatic Control* **AC-30** 705–713.
- Taksar, M. I. 1985. Average optimal singular control and a related stopping problem. *Math. Oper. Res.* **10** 63–81.
- Talluri, K., G. J. van Ryzin. 2004. *The Theory and Practice of Revenue Management*. Kluwer Academic Publishers, Boston/Dordrecht/London.
- Van Mieghem, J. A. 2000. Price and service discrimination in queueing systems: Incentive compatibility of $Gc\mu$ scheduling. *Management Sci.* **46** 1249–1267.
- Van Mieghem, J. A. 2003. Capacity management, investment, and hedging: Review and recent developments. *Manufacturing Service Oper. Management* **5** 269–302.
- van Ryzin, G., S. Mahajan. 1999. On the relationship between inventory costs and variety benefits in retail assortments. *Management Sci.* **45** 1496–1509.
- Veatch, M. H., L. M. Wein. 1996. Scheduling a make-to-stock queue: Index policies and hedging points. *Oper. Res.* **44** 634–647.
- Wein, L. M. 1992. Dynamic scheduling of a multiclass make-to-stock queue. *Oper. Res.* **40** 724–735.
- Whitt, W. 2002. *Stochastic-Process Limits*. Springer-Verlag, New York.

Online Companion

Revenue Management of a Make-to-Stock Queue

René Caldentey

Stern School of Business, New York University, New York, NY 10012, 212-998-0298 Fax: 212-995-4227, rcaldent@stern.nyu.edu

Lawrence M. Wein

Graduate School of Business, Stanford University, Stanford, CA 94306, 650-326-0692, Fax: 650-725-0468, lwein@stanford.edu

1. Derivation of the Hamilton-Jacobi-Bellman Equation

Following Taksar (1985), we first consider the discounted version of the problem,

$$(P^\gamma) \quad \min_{\mathcal{L}, \mathcal{U}} \quad E \left[\int_0^\infty e^{-\gamma t} \mathcal{R}(t) d\mathcal{L}(t) + \int_0^\infty e^{-\gamma t} c(\mathcal{X}(t)) dt \right] \quad (\text{A1})$$

$$\text{subject to} \quad \mathcal{X}(t) = \mathcal{Z}(t) + \mathcal{L}(t) - \mathcal{U}(t), \quad (\text{A2})$$

where $\gamma > 0$ is the discount rate. After deriving the optimality equations for problem (P^γ) , we let $\gamma \rightarrow 0^+$ to derive the optimality equations for the average control problem by mean of a Tauberian argument. Our approach for solving (A1)-(A2) mimics page 442 of Harrison and Taksar (1983), and assumes the existence of a twice-continuously-differentiable function $f^\gamma(x, y)$ representing the optimal value in (A1) starting at (x, y) with $y \in (y_{\min}, y_{\max})$. This assumption is known as *the principle of smooth fit*, and allows us to write down the Hamilton-Jacobi-Bellman equation for problem (P^γ) :

$$[-f_x^\gamma(x, y)] \wedge [f_x^\gamma(x, y) + \mathcal{R}(y)] \wedge [\Gamma f^\gamma(x, y) + c(x) - \gamma f^\gamma(x, y)] = 0, \quad (\text{A3})$$

where $a \wedge b \triangleq \min\{a, b\}$ and

$$\Gamma \triangleq \theta \frac{\partial}{\partial x} + \frac{\sigma^2}{2} \frac{\partial^2}{\partial x^2} + \sigma \delta \varrho \frac{\partial^2}{\partial xy} + \frac{\delta^2}{2} \frac{\partial^2}{\partial y^2}$$

is the infinitesimal generator of $(\mathcal{X}(t), \mathcal{Y}(t))$.

To prove (A3), let us consider a small neighborhood $[x - \epsilon, x + \epsilon]$, and the stopping time $T_\epsilon = \inf\{t \geq 0 : |\mathcal{X}(t) - \mathcal{X}(0)| = \epsilon\}$. A controller facing the state (x, y) has three options: Increase \mathcal{U} moving *instantaneously* to $(x - \epsilon, y)$, increase \mathcal{L} moving to $(x + \epsilon, y)$, or do nothing until T_ϵ and re-evaluate what to do at that time. In these three cases, the objective function is given by, respectively,

$$f(x - \epsilon, y) = f(x, y) - f_x(x, y)\epsilon + o(\epsilon), \quad (\text{A4})$$

$$\mathcal{R}(y)\epsilon + f(x + \epsilon, y) = f(x, y) + [f_x(x, y) + \mathcal{R}(y)]\epsilon + o(\epsilon), \quad (\text{A5})$$

$$E_{(x,y)} \left[\int_0^{T_\epsilon} e^{-\gamma t} c(\mathcal{X}(t)) dt + e^{-\gamma T_\epsilon} f(\mathcal{X}(T_\epsilon), \mathcal{Y}(T_\epsilon)) \right]. \quad (\text{A6})$$

To compute (A6), we note that for any continuous function $A(\cdot)$,

$$\lim_{\epsilon \downarrow 0} \frac{\mathbb{E} \left[\int_0^{T_\epsilon} A(\mathcal{X}(t)) dt \right]}{\mathbb{E}[T_\epsilon]} = A(\mathcal{X}(0)) \quad \text{and} \quad \lim_{\epsilon \downarrow 0} \frac{\sigma^2 \mathbb{E}[T_\epsilon]}{\epsilon^2} = 1, \quad (\text{A7})$$

where σ^2 is the variance of \mathcal{X} . We can rewrite (A6) when $\epsilon \downarrow 0$ as follows

$$f(x, y) + \left(c(x) + \frac{\mathbb{E}_{(x,y)} [e^{-\gamma T_\epsilon} f(\mathcal{X}(T_\epsilon), \mathcal{Y}(T_\epsilon)) - f(x, y)]}{\mathbb{E}[T_\epsilon]} \right) \left(\frac{\epsilon}{\sigma} \right)^2 + o(\mathbb{E}_{(x,y)}[T_\epsilon]). \quad (\text{A8})$$

Since f is twice continuously differentiable, $f(\mathcal{X}, \mathcal{Y})$ is an Itô process. Integration by parts (e.g., Harrison 1985, §4.9, Proposition (2)) gives

$$e^{-\gamma t} f(\mathcal{X}(t), \mathcal{Y}(t)) = f(x, y) + \int_0^t e^{-\gamma s} df(\mathcal{X}(s), \mathcal{Y}(s)) - \gamma \int_0^t e^{-\gamma s} f(\mathcal{X}(s), \mathcal{Y}(s)) ds.$$

In addition, Itô's lemma implies that

$$df(\mathcal{X}, \mathcal{Y}) = f_x d\mathcal{X} + f_y d\mathcal{Y} + \frac{1}{2} (f_{xx} d\mathcal{X} + f_{xy} d\mathcal{Y}) d\mathcal{X} + \frac{1}{2} (f_{yx} d\mathcal{X} + f_{yy} d\mathcal{Y}) d\mathcal{Y}. \quad (\text{A9})$$

Moreover, we know that $d\mathcal{X}(t) = \theta dt + \sigma d\mathcal{W}_x(t)$ and $d\mathcal{Y}(t) = \delta d\mathcal{W}_y(t)$, where the processes $\mathcal{W}_x(t)$ and $\mathcal{W}_y(t)$ are two Wiener processes associated with \mathcal{X} and \mathcal{Y} . Therefore, replacing $d\mathcal{X}$ and $d\mathcal{Y}$ in (A9) we get

$$\begin{aligned} df(\mathcal{X}, \mathcal{Y}) &= \left[\theta f_x + \frac{1}{2} \sigma^2 f_{xx} + \frac{1}{2} \sigma \delta \varrho f_{xy} + \frac{1}{2} \sigma \delta \varrho f_{yx} + \frac{1}{2} \delta^2 f_{yy} \right] dt + \sigma f_x d\mathcal{W}_x + \delta f_y d\mathcal{W}_y \\ &= \Gamma f(\mathcal{X}, \mathcal{Y}) dt + \sigma f_x d\mathcal{W}_x + \delta f_y d\mathcal{W}_y, \end{aligned}$$

which implies

$$\begin{aligned} e^{-\gamma t} f(\mathcal{X}(t), \mathcal{Y}(t)) &= f(x, y) + \int_0^t e^{-\gamma s} (\Gamma f - \gamma f)(\mathcal{X}(s), \mathcal{Y}(s)) ds \\ &\quad + \underbrace{\int_0^t e^{-\gamma s} \sigma f_x d\mathcal{W}_x(s) + \int_0^t e^{-\gamma s} \delta f_y d\mathcal{W}_y(s)}_{M(t)}. \end{aligned} \quad (\text{A10})$$

Taking expectations in (A10) and noticing that the integrands in the two last stochastic integrals are bounded, we have that $\mathbb{E}[M(t)] = 0$, implying

$$\frac{\mathbb{E}_{(x,y)} [e^{-\gamma T_\epsilon} f(\mathcal{X}(T_\epsilon), \mathcal{Y}(T_\epsilon)) - f(x, y)]}{\mathbb{E}[T_\epsilon]} = \frac{\int_0^{T_\epsilon} e^{-\gamma s} (\Gamma f - \gamma f)(\mathcal{X}(s), \mathcal{Y}(s)) ds}{\mathbb{E}[T_\epsilon]}. \quad (\text{A11})$$

Substituting (A11) into (A8) allows us to replace (A6) by

$$f(x, y) + (c(x) + \Gamma f(x, y) - \gamma f(x, y)) \left(\frac{\epsilon}{\sigma} \right)^2 + o(\mathbb{E}_{(x,y)}[T_\epsilon]). \quad (\text{A12})$$

The optimality of $f(x, y)$ together with (A4), (A5), and (A12) imply that f satisfies

$$\begin{aligned} f(x, y) &= \min \{ f(x, y) - f_x(x, y)\epsilon + o(\epsilon), \\ &\quad f(x, y) + [f_x(x, y) + \mathcal{R}(y)]\epsilon + o(\epsilon), \\ &\quad f(x, y) + (c(x) + \Gamma f(x, y) - \gamma f(x, y)) \left(\frac{\epsilon}{\sigma} \right)^2 + o(\mathbb{E}_{(x,y)}[T_\epsilon]). \end{aligned}$$

Subtracting $f(x, y)$ from both sides, letting $\epsilon \downarrow 0$ and using (A7) yields (A3).

Turning to the average cost case, we define the value function $V^\gamma(x, y) = f^\gamma(x, y) - f^\gamma(\hat{x}, \hat{y})$, which represents the relative discounted cost of being in state (x, y) instead of being in an arbitrary fixed state (\hat{x}, \hat{y}) . We then let $\gamma \downarrow 0$ and define the average cost $g = \lim_{\gamma \rightarrow 0} \gamma f^\gamma(\hat{x}, \hat{y})$, whose existence is guaranteed by the boundedness of the market price and holding/backorder costs. Passing to the limit in (A3) yields the optimality equation (31) in the main text.

2. Proof of Proposition 1

Let us define $\pi(x, y; \eta, \xi)$ as the stationary probability measure of the inventory-price process (X, Y) in the domain $\Omega(\eta, \xi) \triangleq \{(x, y) : \eta(y) \leq x \leq \xi(y) \text{ and } y_{\min} \leq y \leq y_{\max}\}$. The corresponding stationary average cost per unit time is given by

$$\mathbb{E}_{\eta, \xi}[X^+] + b \mathbb{E}_{\eta, \xi}[X^-] + \tilde{\mathbb{E}}_{\eta, \xi}(R(y) \mathbb{1}(X = \eta(y))),$$

where $\mathbb{E}_{\eta, \xi}$ is the expected value operator associated to $\pi(x, y; \eta, \xi)$ and $\tilde{\mathbb{E}}_{\eta, \xi}$ is an associated measure of the rejection cost. We note that this representation of $g(\eta, \xi)$ is just a convenient way of rewriting the BAR condition (44) in the main text for the case $c(x) = h x^+ + b x^-$ and $F(x, y) = R(y) \mathbb{1}(x = \eta(y))$.

Because the infinitesimal generator of process (X, Y) is time and space homogeneous, its stationary distribution $\pi(x, y; \eta, \xi)$ satisfies the following translation invariance property:

$$\pi(x, y; \eta, \xi) = \pi(x + \alpha, y; \eta + \alpha, \xi + \alpha) \quad \text{for all } \alpha \in \mathbb{R}.$$

In other words, shifting both boundaries $\eta(y)$ and $\xi(y)$ by the same constant amount (independent of y) does not affect the steady-state probability distribution of the process. This result together with the fact that the boundary cost $R(y)$ does not depend on x imply that the boundary measure $\tilde{\mathbb{E}}_{\eta, \xi}$ is also invariant to translation in the following sense:

$$\tilde{\mathbb{E}}_{\eta, \xi}(R(y) \mathbb{1}(X = \eta(y))) = \tilde{\mathbb{E}}_{\eta + \alpha, \xi + \alpha}(R(y) \mathbb{1}(X = \eta(y) + \alpha)) \quad \text{for all } \alpha \in \mathbb{R}.$$

Therefore, for a given fixed pair of boundaries $(\eta(y), \xi(y))$, it follows that the one-dimensional optimization problem $\min_{\alpha} g(\eta + \alpha, \xi + \alpha)$ has the same optimal solution as

$$\min_{\alpha} \{h \mathbb{E}_{\eta + \alpha, \xi + \alpha}[X^+] + b \mathbb{E}_{\eta + \alpha, \xi + \alpha}[X^-]\} \iff \min_{\alpha} \{(h + b) \mathbb{E}_{\eta + \alpha, \xi + \alpha}[X^+] - b \mathbb{E}_{\eta + \alpha, \xi + \alpha}[X]\}, \quad (\text{A13})$$

where the equivalence follows from the identity $X^- = X^+ - X$. Because

$$\frac{d}{d\alpha} \mathbb{E}_{\eta + \alpha, \xi + \alpha}[X] = 1 \quad \text{and} \quad \frac{d}{d\alpha} \mathbb{E}_{\eta + \alpha, \xi + \alpha}[X^+] = \mathbb{P}(X \geq 0),$$

the first-order optimality condition of problem (A13) yields the familiar fractile condition

$$\mathbb{P}(X \geq 0) = \frac{b}{h + b}. \quad (\text{A14})$$

Since this result holds for any pair $(\eta(y), \xi(y))$, we conclude that it must also hold at the optimal solution.

3. Proof of Proposition 2

Recall that $(\mathcal{X}, \mathcal{Y})$ is a (\vec{d}, Σ) -Brownian motion where the drift $\vec{d} = (\theta, 0)$ and the covariance matrix is given by

$$\Sigma = \begin{bmatrix} \sigma^2 & \sigma \delta \varrho \\ \sigma \delta \varrho & \delta^2 \end{bmatrix}.$$

Let V be the rotation matrix whose rows are the orthonormal eigenvectors of Σ and let E be the corresponding diagonal matrix of eigenvalues such that $\Sigma = V' E V$. We define the linear operator $T = E^{-1/2} V$ and apply it to $(\mathcal{X}, \mathcal{Y})$. The resulting process $(\mathcal{X}^*, \mathcal{Y}^*) = T \cdot (\mathcal{X}, \mathcal{Y})$ is a two-dimensional (\vec{d}^*, I) -Brownian motion with drift $\vec{d}^* = T \cdot \vec{d}$ and covariance matrix I , the identity matrix. Given this transformation, system (41)-(42) in the main text can be written in *normal form* in the $(\mathcal{X}^*, \mathcal{Y}^*)$ coordinates as follows.

$$\frac{1}{2} \Delta V(x^*, y^*) + \vec{d}^* \cdot \nabla V(x^*, y^*) = g(\eta, \xi) - c^*(x^*, y^*) \quad \text{for all } (x^*, y^*) \in \Omega^* \quad (\text{A15})$$

$$\vec{\nu}^* \cdot \nabla V(x^*, y^*) = -\|\vec{\nu}^*\|^{-1} F^*(x^*, y^*) \quad \text{for all } (x^*, y^*) \in \partial\Omega^*, \quad (\text{A16})$$

where $\vec{\nu}^* = T \vec{\nu}$, $c^*(x^*, y^*) = c(x)$, $F^*(x^*, y^*) = F(x, y)$, and $\|\vec{x}\|$ is the norm of \vec{x} . Let us now impose the assumption that $\vec{\nu}^*$ is normal to $\partial\Omega^*$. That is, the tangential vector field $\vec{q}^* = \vec{\nu}^* - \vec{n}^* = 0$, where \vec{n}^* is the inward unit vector field to $\partial\Omega^*$. Then the skew-symmetry condition in Theorem 2.1 in Harrison and Williams (1987) holds and we have that the stationary distribution of $(\mathcal{X}^*, \mathcal{Y}^*)$ has a density of the exponential form

$$\pi_\Omega^*(x, y) = K_\Omega \exp(\vec{m}^* \cdot (x, y)') = K_\Omega \exp(m_x^* x + m_y^* y) \quad \text{for all } (x, y) \in \Omega,$$

where $\vec{m}^* = 2 \vec{d}^*$. Thus, the original process $(\mathcal{X}, \mathcal{Y})$ has also a steady-state distribution $\pi_\Omega(x, y)$ of the exponential form $\pi_\Omega(x, y) = K_\Omega \exp((x, y)' \cdot \vec{m})$ with $\vec{m} = T' \cdot \vec{m}^* = 2 \Sigma^{-1} \cdot \vec{d}$. After some algebra we get

$$m_x = \frac{2\theta}{\sigma^2(1 - \varrho^2)}, \quad m_y = \frac{-2\varrho\theta}{\sigma\delta(1 - \varrho^2)}, \quad \text{and} \quad \pi_\Omega(x, y) = K_\Omega \exp(m_x x + m_y y). \quad (\text{A17})$$

The value of K_Ω is chosen such that $\pi_\Omega(x, y)$ integrates to one inside Ω . We now apply the BAR condition ((43) in the main text) to (A15)-(A16) to get

$$0 = \int_{\Omega^*} (c^*(x^*, y^*) - g(\eta, \xi)) \pi_\Omega^*(x, y) ds^* + \frac{1}{2} \int_{\partial\Omega^*} \|\vec{\nu}^*\|^{-1} F^*(x^*, y^*) \pi_\Omega^*(x^*, y^*) dl^*, \quad (\text{A18})$$

where ds^* denotes integration with respect to Lebesgue measure on Ω^* and dl^* denotes integration with respect to surface measure on $\partial\Omega^*$. Using the linear transformation $(x^*, y^*) = T(x, y)$ we can rewrite the BAR condition in terms of the original variables (x, y) . In particular, letting $\|T\|$ be the determinant of T , we have

$$ds^* = \|T\| ds, \quad dl^* = \|T\| \left(\sigma^2 dx^2 - 2\sigma\delta\varrho dx dy + \delta^2 dy^2 \right)^{1/2}, \quad \|\vec{\nu}^*\| = \|T\| \delta \quad \text{and} \quad \|T\| = \frac{1}{\sigma\delta\sqrt{1 - \varrho^2}}.$$

Finally, substituting these values into (A18) and using the definition of the boundary function $F(x, y) = \mathcal{R}(y)$ in $\partial\Omega^\eta$ and $F(x, y) = 0$ elsewhere, we get (48) in the main text.

4. Derivation of the Proposed Policy

In this section, we derive the proposed policy in (54) by approximating the solution to the ODE in (51) and computing the value of the Lagrange multiplier, γ , for constraint (49) in the main text. Equation (51) is a second-order nonlinear ODE in the interval $[y_{\min}, y_{\max}]$ that we have not been able to solve in closed form. Instead, we derive an approximate solution by considering two extreme cases in which $y_{\min} \rightarrow -\infty$ and $y_{\max} \rightarrow +\infty$. We note that the numerical computations that we report in §4.4 in the main text support our methodology to derive our proposed policy.

Our heuristic procedure has four steps. First, we introduce the notion of an “asymptotically-consistent” solution in the sense of condition (A21) below. We note that by definition every solution to the ODE is an asymptotically-consistent solution, and so asymptotic consistency is a weaker solution concept for the ODE. Second, we rely on two intuitive properties that we think a solution should possess at $y \rightarrow -\infty$ and $y \rightarrow \infty$ to propose a candidate policy. This policy turns out to be an affine function of the price, that is, $\eta(y) = \alpha_1 + \alpha_2 \mathcal{R}(y)$ for two constants α_1 and α_2 . Third, we show (in Proposition A1) that there is a unique pair (α_1, α_2) , with $\alpha_2 < 0$, that is asymptotically consistent with the ODE. The final step is to compute γ , the Lagrange multiplier for constraint (49) in the main text. We tackle this problem by requiring our proposed solution to coincide with the optimal solution if the price was constant at $\mathcal{R}(y) = \mathcal{R}_1(r)$ (*i.e.*, $\mathcal{R}_0 = 0$). Because we have a closed-form solution for the optimal policy in this particular case (see Proposition A2), we are able to fully characterize the value of γ and, hence, the value of the constants α_1 and α_2 .

Turning to the first step, we rewrite the ODE for $\eta(y)$ as

$$\eta(y) = \frac{-\gamma}{b} + \frac{\sigma \sqrt{1 - \varrho^2}}{2b} \mathcal{R}(y) L(\mathcal{R}(y), \eta_y, \eta_{yy}), \quad y \in [y_{\min}, y_{\max}], \quad (\text{A19})$$

where the auxiliary function L satisfies

$$L(\mathcal{R}(y), \eta_y, \eta_{yy}) \triangleq \frac{(1 - \varrho^2)\sigma^2\delta^2 \eta_{yy} + [(1 - \mathcal{R}_1(r)\mathcal{R}^{-1}(y))(\delta^2\eta_y - \sigma\delta\varrho) - 2\theta] [\sigma^2 - 2\sigma\delta\varrho\eta_y(y) + \delta^2\eta_y^2(y)]}{[\sigma^2 - 2\sigma\delta\varrho\eta_y(y) + \delta^2\eta_y^2(y)]^{\frac{3}{2}}}. \quad (\text{A20})$$

Based on the ODE (A19), we look for an approximate solution $\eta(y)$ that is *asymptotically consistent* in the sense that

$$\lim_{y \rightarrow \pm\infty} \frac{\eta(y)}{\mathcal{R}(y)} = \lim_{y \rightarrow \pm\infty} \frac{1}{\mathcal{R}(y)} \left(\frac{-\gamma}{b} + \frac{\sigma \sqrt{1 - \varrho^2}}{2b} \mathcal{R}(y) L(\mathcal{R}(y), \eta_y, \eta_{yy}) \right). \quad (\text{A21})$$

In the second step, we specify the form of our proposed policy by assessing its behavior as $y \rightarrow \pm\infty$. First, as $y \rightarrow -\infty$ the price $\mathcal{R}(y)$ is approximately constant and equal to $\mathcal{R}_1(r)$. Note that if $\mathcal{R}(y)$ is independent of y in (A19), then the solution $\eta(y)$ is a constant. Hence, we assume that $\eta(y)$ converges to a constant as $y \rightarrow -\infty$. Second, as $y \rightarrow \infty$, we expect the rejection of e-orders to decrease with the spot market price, that is, we expect $\eta(y)$ to increase in y . Because the price increases exponentially with y , we assume that $\eta(y)$ also has exponential growth as $y \rightarrow +\infty$; although we do not prove that this needs to be so, we have been able to show (details not given) that if $\eta(y)$ is polynomial in y , then it cannot be asymptotically consistent as $y \rightarrow \infty$. Based on these two assumptions, and the fact that $\eta(y)$ depends on y only through the price function $\mathcal{R}(y)$, we propose the following affine (in the price) policy

$$\eta(y) = \alpha_1 + \alpha_2 \mathcal{R}(y),$$

where α_1 and α_2 are two constants. Furthermore, we also expect the number of rejections to decrease with the spot price, that is, $\alpha_2 \leq 0$.

The third step consists of the following result, which establishes that there is a unique pair (α_1, α_2) , with $\alpha_2 < 0$, for which the solution $\eta(y)$ above is asymptotically consistent in the sense of (A21).

Proposition A1 *Within the family of affine functions $\{\alpha_1 + \alpha_2 \mathcal{R}(y) : (\alpha_1, \alpha_2) \in \mathbb{R} \times \mathbb{R}_-\}$ there is a unique member that is asymptotically consistent with (A19) in the sense of (A21). This function is given by*

$$\eta^p(y) = -\frac{\gamma}{b} - \frac{(2\theta - \sigma\delta)\sqrt{1 - \varrho^2}}{2b} \mathcal{R}_1(r) - \frac{\sigma\delta\sqrt{1 - \varrho^2}}{2b} \mathcal{R}(y). \quad (\text{A22})$$

Proof: Consider an arbitrary affine policy $\tilde{\eta}(y) = \alpha_1 + \alpha_2 \mathcal{R}(y)$ with $\alpha_2 < 0$. Using the definitions of L in (A20) and $\mathcal{R}(y) = \mathcal{R}_1(r) + \mathcal{R}_0 e^y$, it follows that

$$\lim_{y \rightarrow -\infty} \frac{\tilde{\eta}(y)}{\mathcal{R}(y)} = \frac{\alpha_1 + \alpha_2 \mathcal{R}_1(r)}{\mathcal{R}_1(r)} \quad (\text{A23})$$

and

$$\lim_{y \rightarrow -\infty} \frac{1}{\mathcal{R}(y)} \left(\frac{-\gamma}{b} + \frac{\sigma \sqrt{1 - \varrho^2}}{2b} \mathcal{R}(y) L(\mathcal{R}(y), \tilde{\eta}_y, \tilde{\eta}_{yy}) \right) = \frac{-1}{\mathcal{R}_1(r)} \left[\frac{\gamma}{b} + \frac{\theta \sqrt{1 - \varrho^2}}{b} \mathcal{R}_1(r) \right]. \quad (\text{A24})$$

Therefore, by definition (A21) and equations (A23)-(A24), $\tilde{\eta}(y)$ is asymptotically consistent as $y \rightarrow -\infty$ if and only if

$$\alpha_1 + \alpha_2 \mathcal{R}_1(r) = -\frac{\gamma}{b} - \frac{\theta \sqrt{1 - \varrho^2}}{b} \mathcal{R}_1(r). \quad (\text{A25})$$

As $y \rightarrow \infty$, we have that

$$\lim_{y \rightarrow \infty} \frac{\tilde{\eta}(y)}{\mathcal{R}(y)} = \alpha_2.$$

Letting $y \rightarrow \infty$ in (A20) gives

$$\lim_{y \rightarrow \infty} L(\mathcal{R}(y), \tilde{\eta}_y, \tilde{\eta}_{yy}) = \delta \text{sign}(\alpha_2) = -\delta,$$

so that

$$\lim_{y \rightarrow \infty} \frac{1}{\mathcal{R}(y)} \left(\frac{-\gamma}{b} + \frac{\sigma \sqrt{1 - \varrho^2}}{2b} \mathcal{R}(y) L(\mathcal{R}(y), \tilde{\eta}_y, \tilde{\eta}_{yy}) \right) = -\frac{\sigma \delta \sqrt{1 - \varrho^2}}{2b}.$$

Hence, $\tilde{\eta}(y)$ is asymptotically consistent as $y \rightarrow \infty$ if and only if

$$\alpha_2 = -\frac{\sigma \delta \sqrt{1 - \varrho^2}}{2b}. \quad (\text{A26})$$

The proposition follows from equations (A25) and (A26). \square

To complete the characterization of our proposed policy, we need to estimate the value of the Lagrange multiplier γ . To do this, we force our solution in (A22) to be optimal for the special case of a constant price, which is a case that is solved in closed form in the following proposition. The obvious choice for the constant price is $\mathcal{R}(y) = \mathcal{R}_1(r)$, which ensures consistency as $y \rightarrow -\infty$.

Proposition A2 *If $\mathcal{R}(y) = \mathcal{R}_1(r)$ then the optimal solution \mathcal{L} , \mathcal{U} to problem (28)-(30) in the main text is characterized by a pair of constants $(\bar{\eta}, \bar{\xi})$, where*

$$\bar{\xi} = \sqrt{\frac{\sigma^2 \sqrt{1 - \varrho^2} b \mathcal{R}_1(r)}{h(h+b)}} \quad \text{and} \quad \bar{\eta} = -\sqrt{\frac{\sigma^2 \sqrt{1 - \varrho^2} h \mathcal{R}_1(r)}{b(h+b)}} \quad \text{if } \theta = 0, \quad (\text{A27})$$

and

$$\bar{\xi} = \frac{\ln(1 + \text{sign}(\theta) \sqrt{1 - \chi})}{m_x} \quad \text{and} \quad \bar{\eta} = -\bar{\xi} - \frac{\theta \mathcal{R}_1(r)}{\sqrt{1 - \varrho^2} h} \quad \text{if } \theta \neq 0 \quad \text{and} \quad h = b, \quad (\text{A28})$$

where $\chi = \exp(-m_x \theta \mathcal{R}_1(r) (\sqrt{1 - \varrho^2} h)^{-1})$.

Proof: The fact that the optimal solution is characterized by a pair of constants is well known (e.g. Harrison 1985). We derive the thresholds separately for the cases $\theta = 0$ and $\theta \neq 0$. In the driftless case, the steady-state inventory distribution is uniform (e.g. Harrison 1985) so that π_Ω is constant and equal to $(\bar{\xi} - \bar{\eta})^{-1}$. Therefore, the optimization problem in (48)-(49) becomes

$$\min_{\bar{\eta}, \bar{\xi}} \frac{1}{\bar{\xi} - \bar{\eta}} \left(\frac{\sigma^2 \sqrt{1 - \varrho^2} \mathcal{R}_1(r)}{2} + \int_{\bar{\eta}}^{\bar{\xi}} c(x) dx \right).$$

The optimality conditions for this problem are

$$h \bar{\xi} + b \bar{\eta} = 0 \quad \text{and} \quad 2h \bar{\xi}(\bar{\xi} - \bar{\eta}) - \sigma^2 \sqrt{1 - \varrho^2} \mathcal{R}_1(r) - h \bar{\xi}^2 - b \bar{\eta}^2 = 0, \quad (\text{A29})$$

and its solution is given by (A27).

In the case where $\theta \neq 0$, the steady-state inventory distribution is a truncated exponential in the interval $[\bar{\eta}, \bar{\xi}]$ with rate $m_x = \frac{2\theta}{\sigma^2(1-\varrho^2)}$ (e.g. Harrison 1985). Hence, the optimization problem (48)-(49) is given by

$$\min_{\bar{\eta}, \bar{\xi}} \frac{1}{\exp(m_x \bar{\xi}) - \exp(m_x \bar{\eta})} \left(\frac{\sigma^2 \sqrt{1 - \varrho^2} \mathcal{R}_1(r) \exp(m_x \bar{\eta})}{2} + \int_{\bar{\eta}}^{\bar{\xi}} c(x) \exp(m_x x) dx \right).$$

The first-order optimality conditions for this problem are

$$h \bar{\xi} + b \bar{\eta} + \frac{\theta}{\sqrt{1 - \varrho^2}} \mathcal{R}_1(r) = 0 \quad \text{and} \quad \left(\frac{h}{h + b} \right) \exp(m_x \bar{\xi}) + \left(\frac{b}{h + b} \right) \exp(m_x \bar{\eta}) = 1. \quad (\text{A30})$$

A closed-form solution to (A30) is only available for the special case $h = b$, and is given in (A28). \square

We are now in a position to state our proposed policy. To derive our proposed admission control, $\eta^P(y)$, we first note that by (A22) our candidate policy satisfies

$$\eta^P(y) = \alpha - \frac{\sigma \delta \sqrt{1 - \varrho^2}}{2b} (\mathcal{R}(y) - \mathcal{R}_1(r)),$$

for a constant α that depends on γ . Furthermore, as $y \rightarrow -\infty$, the price $\mathcal{R}(y)$ converges to $\mathcal{R}_1(r)$ and so $\eta^P(y)$ converges to α . Hence, by Proposition A2 we choose α equal to $\bar{\eta}$, which solves (A29) or (A30).

To derive our proposed production policy, $\xi^P(y)$, we note that the optimality equation (53) in the main text suggests a fixed base-stock level. Proposition A2 and the fact that $\mathcal{R}(y)$ converges to $\mathcal{R}_1(r)$ as $y \rightarrow -\infty$ suggest that we choose this fixed base-stock level to be $\bar{\xi}$, which solves (A29) or (A30).

In summary, we characterize our proposed policy by

$$\mathbf{Proposed Policy:} \quad \xi^P(y) = \bar{\xi} \quad \text{and} \quad \eta^P(y) = \bar{\eta} - \frac{\sigma \delta \sqrt{1 - \varrho^2}}{2b} (\mathcal{R}(y) - \mathcal{R}_1(r)).$$

Depending on the value of θ , the pair $(\bar{\xi}, \bar{\eta})$ solves (A29) or (A30).

5. Derivation of the iid Price Policy

In this section, we derive the optimal admission/production policy for the case in which the sequence of spot prices is iid. Our model and solution techniques are based on George and Harrison (2001) and Ata and Shneorson (2004).

For tractability reasons (and consistent with existent literature), we consider the case in which the arrivals of e-orders is a Poisson process $N(t)$ with intensity Λ and the service time is exponentially distributed with mean μ^{-1} . Recall that without loss of generality we have set $\mu = 1$.

Let $R(n)$ be the spot price of the n^{th} e-order. The sequence $\{R(n) : n \geq 1\}$ is iid with probability distribution $G(R)$. As in our current model, we assume that arriving buyers truthfully declare their reservation prices $R(n)$. Also, to keep this iid price model as consistent as possible with our model, we assume that $G(\cdot)$ coincides with the steady-state distribution of our original price process. That is,

$$R(1) \stackrel{d}{=} R_1(r) + R_0 e^Y \quad (\text{A31})$$

with Y uniformly distributed in $[Y_{\min}, Y_{\max}]$ (see section §5.5 in Harrison 1985). Hence, $G(\cdot)$ is given by

$$G(r) = \frac{\ln(r - R_1(r)) - \ln(R_0) - Y_{\min}}{Y_{\max} - Y_{\min}}, \quad r \in [R_1(r) + R_0 e^{Y_{\min}}, R_1(r) + R_0 e^{Y_{\max}}].$$

Let τ_n be the n^{th} arrival epoch. At each of these epochs, the manufacturer observes the inventory position $X(\tau_n)$ and the incoming price $R(n)$ and decides whether to accept or reject the order. Because of the Markovian nature of the problem, both $X(\tau_n)$ and $R(n)$ are sufficient statistics for the admission decision. Moreover, by monotonicity it is reasonable to argue that for each inventory level x , there is a threshold price r_x such that the manufacturer accepts an incoming e-order if and only if the price exceeds r_x . Thus, the effective arrival rate when the inventory position is x equals

$$\lambda_x \triangleq \Lambda \mathbb{P}(R \geq r_x) = \Lambda (1 - G(r_x)).$$

For future references, we define the inverse function

$$r(\lambda) \triangleq G^{-1}\left(1 - \frac{\lambda}{\Lambda}\right), \quad \lambda \in [0, \Lambda].$$

On the production side, the manufacturer has the ability to stop production when the inventory level reaches a base-stock level ξ . Since the sequence of spot prices is iid, this threshold value ξ depends exclusively on the inventory position. Therefore, the production rate as a function of the inventory position is given by

$$\mu_x \triangleq \mathbb{1}(x < \xi). \quad (\text{A32})$$

As a matter of the fact, we could have considered more general production strategies of the form $\vec{\mu} = \{\mu_x \in [0, 1] : x \in \mathbb{Z}\}$. However, because controlling production is costless it turns that the optimal production strategies takes the bang-bang form (A32) for an appropriate ξ .

The manufacturer's choice of the \vec{r} and $\vec{\mu}$ is based on profits maximization. Revenues are accrued every time an incoming order is accepted. For an inventory level of x , the choice of a threshold price r_x implies that revenues are generated at a rate of

$$\mathcal{B}(\lambda_x) \triangleq \lambda_x \mathbb{E}[R | R \geq r(\lambda_x)].$$

It is not hard to see that this revenue function is given by

$$\mathcal{B}(\lambda) = \Lambda \int_{r(\lambda)}^{R^{\max}} R dG(R), \quad \lambda \in [0, \Lambda].$$

On the other hand, the holding/inventory cost rate is $c(x) = h x^+ + b x^-$. Thus, the net profit rate at state x is

$$\mathcal{P}(x) \triangleq \mathcal{B}(\lambda_x) - c(x).$$

Note that, as opposed to $c(x)$, the revenue term $\mathcal{B}(\lambda_x)$ is bounded above by $\Lambda \mathbb{E}[R]$. Therefore, as $|x| \rightarrow \infty$, the profit rate $\mathcal{P}(x) \rightarrow -\infty$ and so it is in the manufacturer interest to bound the domain of the inventory process $X(t)$ to a closed interval $[\eta, \xi]$.

In this setting, the manufacturer's problem is to select a pair of boundaries (η, ξ) and the sequence of arrival rates $\vec{\lambda} \triangleq \{\lambda_x \in [0, \Lambda] : \eta + 1 \leq x \leq \xi\}$ ¹ that maximize the expected long-term time-average profits. Interestingly, note that the structure of an optimal solution in this model coincides with the structure of our proposed solution in the paper. That is, a base-stock level ξ that regulates production decisions independent of the posted price and an admission threshold r_x for the posted price that is state dependent.

Similar to Ata and Shneorson (2004), we define the set of *admissible* controls \mathcal{A} to be the set of pairs $(\vec{\lambda}, \vec{\mu})$, for which the corresponding inventory process is a stable birth-death process with unique equilibrium distribution $\{\pi_x : x \in \mathbb{Z}\}$ satisfying the usual balance equations and summing up to unity. Using the ergodicity of the Markov process $X(t)$, the manufacturer's objective function can be written as

$$V(\vec{\lambda}, \vec{\mu}) = \sum_{x \in \mathbb{Z}} [\mathcal{B}(\lambda_x) - c(x)] \pi_x \quad (\text{A33})$$

and the optimization problem reduces to

$$\max_{(\vec{\lambda}, \vec{\mu}) \in \mathcal{A}} \sum_{x \in \mathbb{Z}} [\mathcal{B}(\lambda_x) - c(x)] \pi(x). \quad (\text{A34})$$

This iid price model is closely related to the one studied by George and Harrison (2001) and Ata and Shneorson (2004). However, there are a number of differences that prevent us from applying directly their results to problem (A34). The most notorious is the fact that we are dealing with a make-to-stock model with full backorders. In this case, the inventory level x can go negative and the function $c(x)$ ² is not monotonic on x . Similarly, the reward term $\mathcal{B}(\lambda)$ is not guaranteed to satisfy the conditions of the corresponding function $b(\cdot)$ in Ata and Shneorson (2004).

To solve problem (A34), let us first fix (η, ξ) and compute the optimal arrival rate vector $\vec{\lambda}^*$. Using the semi-MDP structure we can define the optimality conditions for our problem as follows³

$$V = y_\eta^+ - c(\eta) \quad (\text{A35})$$

$$V = \Phi(y_{x-1}) + y_x^+ - c(x), \quad \eta + 1 \leq x \leq \xi - 1 \quad (\text{A36})$$

$$V = \Phi(y_{\xi-1}) - c(\xi), \quad (\text{A37})$$

where V is a constant representing the optimal steady-state average profit, the vector y_x captures the relative value difference (see George and Harrison 2001 or Ata and Shneorson 2004 for details), and the auxiliary function Φ is defined by

$$\Phi(y) \triangleq \max_{\lambda \in [0, \Lambda]} \{\mathcal{B}(\lambda) - \lambda y\}.$$

Some useful properties of this function⁴, such as continuity, monotonicity (strictly decreasing), and concavity, can be found in section §5 of George and Harrison (2001) or section §3.2 in Ata and

¹Or equivalently a sequence of vector of prices $\vec{r} \triangleq \{r_x \in [R^{\min}, R^{\max}] : \eta + 1 \leq x \leq \xi\}$.

²Denoted by h_n in Ata and Shneorson (2004).

³These are the same as equations (12)-(13) in George and Harrison (2001) or equations (17)-(19) in Ata and Shneorson (2004).

⁴Also known as the Fenchel-Legendre transform of $\mathcal{B}(\lambda)$.

Shneorson (2004). Furthermore, because of the special form of the revenue component $\mathcal{B}(\lambda)$, we can show that

$$\Phi(y) = \Lambda \int_{r^*(y)}^{R^{\max}} [R - y] dG(R), \quad \text{where } r^*(y) \triangleq \min\{R^{\max}, \max\{y, R^{\min}\}\}.$$

The following *verification* result is similar to proposition 1 in George and Harrison (2001) or Ata and Shneorson (2004).

Proposition A3 *Suppose we can find constants V, ξ, η , and a vector $\vec{y} = (y_\eta, \dots, y_{\xi-1})$ such that*

1. V and \vec{y} jointly satisfy (A35)-(A37) and $\vec{y} \geq 0$,
2. $V + c(x) \geq R^{\max}$ for $x \leq \eta - 1$,
3. $V + c(x) \geq \Phi(0) = \Lambda E[R]$ for $x \geq \xi + 1$.

Define the associated admission and production policies by

$$\lambda_x^* = \begin{cases} \Lambda(1 - G(r^*(y_{x-1}))) & \text{if } x = \eta + 1, \dots, \xi \\ 0 & \text{otherwise} \end{cases} \quad \text{and} \quad \mu_x^* = \mathbb{1}(x < \xi), \quad (\text{A38})$$

respectively. Then, the resulting policy is optimal and V is the optimal steady-state average cost.

Proof: Suppose (V, \vec{y}) satisfies the conditions in the proposition. Let $(\vec{\lambda}, \vec{\mu})$ be an arbitrary admissible admission policy that induces a steady-state distribution π_x . Because (V, \vec{y}) satisfies (a) it also satisfies (A36) and it follows that

$$\begin{aligned} V \pi_x &= [\Phi(y_{x-1}) + y_x - c(x)] \pi_x \geq [\mathcal{B}(\lambda_x) - c(x)] \pi_x + y_x \pi_x - \lambda_x y_{x-1} \pi_x \\ &= [\mathcal{B}(\lambda_x) - c(x)] \pi_x + y_x \pi_x - y_{x-1} \pi_{x-1}, \end{aligned}$$

where the inequality follows from the definition of Φ and the last equality uses the balance condition $\lambda_x \pi_x = \pi_{x-1}$. Summing this condition over $x = \eta + 1, \dots, \xi - 1$ it follows that

$$V \sum_{x=\eta+1}^{\xi-1} \pi_x \geq \sum_{x=\eta+1}^{\xi-1} [\mathcal{B}(\lambda_x) - c(x)] \pi_x + y_{\xi-1} \pi_{\xi-1} - y_\eta \pi_\eta.$$

Also, conditions (A35) and (A37) it imply that

$$V \pi(\eta) = y_\eta \pi_\eta - c(\eta) \pi_\eta$$

and

$$V \pi_\xi = [\Phi(y_{\xi-1}) - c(\xi)] \pi_\xi \geq [\mathcal{B}(\lambda_\xi) - c(\xi)] \pi_\xi - \lambda_\xi y_{\xi-1} \pi_\xi = [\mathcal{B}(\lambda_\xi) - c(\xi)] \pi_\xi - y_{\xi-1} \pi_{\xi-1}.$$

Summing the previous three equations we get

$$V \sum_{x=\eta}^{\xi} \pi_x \geq \sum_{x=\eta}^{\xi} [\mathcal{B}(\lambda_x) - c(x)] \pi_x - \mathcal{B}(\lambda_\eta) \pi_\eta. \quad (\text{A39})$$

In addition, condition (b) implies

$$[V + c(x)] \pi(x) \geq R^{\max} \pi_x \geq R^{\max} \lambda_{x+1} \pi_{x+1} \geq \mathcal{B}(\lambda_{x+1}) \pi_{x+1}, \quad x \leq \eta - 1$$

and so

$$V \pi_x \geq [\mathcal{B}(\lambda_x) - c(x)] \pi_x + \mathcal{B}(\lambda_{x+1}) \pi_{x+1} - \mathcal{B}(\lambda_x) \pi_x, \quad x \leq \eta - 1.$$

Summing over $x \leq \eta - 1$ we obtain

$$V \sum_{x=-\infty}^{\eta-1} \pi_x \geq \sum_{x=-\infty}^{\eta-1} [\mathcal{B}(\lambda_x) - c(x)] \pi_x + \mathcal{B}(\lambda_\eta) \pi_\eta - \lim_{x \rightarrow -\infty} \mathcal{B}(\lambda_x) \pi_x. \quad (\text{A40})$$

Because $\mathcal{B}(\lambda)$ is bounded and the policy under consideration is admissible, it follows that the limit vanishes.

Using a similar argument, condition (c) implies

$$[V + c(x)] \pi_x \geq \Phi(0) \pi_x \geq \mathcal{B}(\lambda_x) \pi_x \quad \text{and so} \quad V \pi_x \geq [\mathcal{B}(\lambda_x) - c(x)] \pi_x, \quad x \geq \xi + 1.$$

Summing over $x \geq \xi + 1$ it follows that

$$V \sum_{x=\xi+1}^{\infty} \pi_x \geq \sum_{x=\xi+1}^{\infty} [\mathcal{B}(\lambda_x) - c(x)] \pi_x. \quad (\text{A41})$$

Combining (A39), (A40), and (A41) we conclude that

$$V \geq \sum_{x=-\infty}^{\infty} [\mathcal{B}(\lambda_x) - c(x)] \pi_x.$$

Since this inequality holds for an arbitrary admissible policy it follows that V is an upper bound for the optimal steady-state average profit.

To conclude, we need to show that the proposed policy in (A38) is optimal. For this, we simply repeat the steps above replacing $(\vec{\lambda}, \vec{\mu})$ by $(\vec{\lambda}^*, \vec{\mu}^*)$ and all inequalities by equalities. \square

The previous results provides a simple way to check whether a candidate solution is optimal or not. We use this result to solve system (A35)-(A36) using the following algorithm.

(η, ξ) -Algorithm

I) **Initialization:** Fix the boundaries (η, ξ) and choose an initial value for V in the range $[0, \Lambda E[R]]$.

II) **Iteration:** Use recursively (forward in x) conditions (A35) and (A36) to compute y_x for all $x = \eta, \eta + 2, \dots, \xi - 1$. This step can be done efficiently because of the triangular form of the system. In fact,

$$y_\eta = V + c(\eta) \quad \text{and} \quad y_x = V + c(x) - \Phi(y_{x-1}), \quad \eta + 1 \leq x \leq \xi - 1.$$

III) **Optimality Check:** If $V + c(\xi) = \Phi(y_{\xi-1})$ then stop; otherwise change the value of V and iterate at step II.

A few comments about this algorithm are now in order. First, the range for V in step (I) follows from two observations: (i) V is trivially nonnegative and (ii) an upper bound can be found assuming $c(x) = 0$ for all x . In this case, the optimal decision is to admit every order independently of its price and so $V = \Lambda E[R]$. Second, the optimality check in (III) is based on condition (A37). Also, to operationalize this method we need to specify in what sense we interpret the equality in (III). In particular, we can decide to stop if for some small tolerance $\epsilon > 0$, $|V + c(\xi) - \Phi(y_{\xi-1})| < \epsilon$.

The algorithm find the optimal prices for a given region (η, ξ) but it does not compute the optimal boundaries. This is the free-boundary part of the problem that complicates the solution. We solve for these boundaries through an exhaustive search in the region

$$\{(\eta, \xi) \text{ such that } -M_\eta \leq \eta \leq 0 \leq \xi \leq -M_\xi\},$$

where the constants $M_\eta \leq 0$ and $M_\xi \geq 0$ are chosen large enough to ensure that the optimal solution (η^*, ξ^*) belong to the set above. We can use the optimality conditions (b) and (c) in Proposition A3 to find these bounds. In fact, since $V \geq 0$, it follows that for all $x \leq 0$ such that $c(x) \geq R^{\max}$ (that is, for $x \leq -R^{\max}/b$) condition (b) is automatically satisfied. Similarly, for $x \geq \Lambda E[R]/h$ condition (c) always holds. Then, we can choose

$$M_\eta = - \left\lceil \frac{R^{\max}}{b} \right\rceil \quad \text{and} \quad M_\xi = \left\lceil \frac{\Lambda E[R]}{h} \right\rceil,$$

where $\lceil x \rceil$ is the smaller integer greater than or equal to x .

6. Numerical Solution to the Diffusion Control Problem

As with any numerical method for solving PDEs, we start by discretizing the state space. We approach this problem using the *Markov chain approximation* technique (Kushner and Dupuis 2001) first developed by Kushner (1977), which relies on the fact that a regulated diffusion process can be approximated by a finite-state Markov chain. Let us first define a bounded region Ω^0 on the plane $\{(x, y) \in \mathbb{R}^2\}$ where the process $(\mathcal{X}, \mathcal{Y})$ resides. We have to choose the set Ω^0 large enough to ensure that $(\mathcal{X}, \mathcal{Y})$ will rarely reach the boundaries given the optimal control policies $(\mathcal{L}^*, \mathcal{U}^*)$. Let \hbar be the *finite difference interval* which defines how finely state and time are discretized. The initial state space of the Markov chain associated to \hbar is given by the lattice

$$\Omega_\hbar^0 = \{(x, y) \in \Omega^0 : x = n_x \hbar, y = n_y \hbar, |n_x| \leq N_x, y_{\min} \leq y \leq y_{\max}, n_x, n_y \text{ integers}\},$$

where N_x is a positive integer that defines the dimension of Ω_\hbar^0 . Starting with Ω_\hbar^0 , we iteratively generate a sequence of regions $\Omega_\hbar^1, \Omega_\hbar^2, \dots$ such that $\Omega_\hbar^0 \supseteq \Omega_\hbar^1 \supseteq \Omega_\hbar^2 \supseteq \dots$ that approach the optimal region Ω_\hbar^* defined by the boundaries $\eta(y)$ and $\xi(y)$:

$$\Omega_\hbar^* = \{(x, y) \in \Omega : x = n_x \hbar, y = n_y \hbar, \eta(y) \leq x \leq \xi(y), y_{\min} \leq y \leq y_{\max}, n_x, n_y \text{ integers}\}.$$

Therefore, we require N_x to be large enough to ensure $\Omega_\hbar^* \subseteq \Omega_\hbar^0$. Suppose that we are at stage k and that the current region is Ω_\hbar^k . The remainder of our algorithm description provides the details of the iteration that generates Ω_\hbar^{k+1} . For notational convenience, we omit the dependence of Ω_\hbar^k on k . As in (35) in the main text, we divide the boundary $\partial\Omega_\hbar$ into the four components $\partial\Omega_\hbar^\eta$, $\partial\Omega_\hbar^\xi$, $\partial\Omega_\hbar^{\min}$, and $\partial\Omega_\hbar^{\max}$.

To compute the transition probabilities for the Markov chain associated with Ω_\hbar , we assume (as implied by (37)-(40)) that the controls are only exerted on the boundaries, and that the directions of control are dictated by the reflection field $\vec{\nu}$ (see Figure 2 in the main text). Let $P_\hbar((x, y), (x', y'))$ be the probability of a transition from state (x, y) to state (x', y') and define (see chapter 5 of Kushner and Dupuis 2001) $Q_\hbar = \hbar |\theta| + \sigma^2 + \delta^2 - |\sigma \delta \varrho|$. Then, $P_\hbar((x, y), (x', y'))$ is given by

- If (x, y) is in the interior, *i.e.*, $(x, y) \in \Omega_{\hbar} - \partial\Omega_{\hbar}$, then

$$\begin{aligned}
P_{\hbar}((x, y), (x \pm \hbar, y)) &= \frac{\sigma^2 - |\sigma\delta\rho| + 2\hbar\theta^{\pm}}{2Q_{\hbar}}, \\
P_{\hbar}((x, y), (x, y \pm \hbar)) &= \frac{\delta^2 - |\sigma\delta\rho|}{2Q_{\hbar}}, \\
P_{\hbar}((x, y), (x \pm \hbar, y \pm \hbar)) &= \frac{\sigma\delta\rho^+}{2Q_{\hbar}}, \\
P_{\hbar}((x, y), (x \pm \hbar, y \mp \hbar)) &= \frac{\sigma\delta\rho^-}{2Q_{\hbar}}, \\
P_{\hbar}((x, y), (x', y')) &= 0 \quad \text{otherwise,}
\end{aligned} \tag{A42}$$

where the superscripts $+$ and $-$ denote the positive and negative parts of the parameters, respectively.

- If $(x, y) \in \partial\Omega_{\hbar}^{\eta}$ then $P_{\hbar}((x, y), (x + \hbar, y)) = 1$.
- If $(x, y) \in \partial\Omega_{\hbar}^{\xi}$ then $P_{\hbar}((x, y), (x - \hbar, y)) = 1$.
- If $(x, y) \in \partial\Omega_{\hbar}^{\max}$ then $P_{\hbar}((x, y), (x, y - \hbar)) = 1$.
- If $(x, y) \in \partial\Omega_{\hbar}^{\min}$ then $P_{\hbar}((x, y), (x, y + \hbar)) = 1$.

The *interpolation interval* Δt^{\hbar} is defined by

$$\Delta t^{\hbar} = \frac{\hbar^2}{Q_{\hbar}},$$

which forces the first two moments of the Markov chain to coincide with those of the diffusion process.

This Markov chain is used to approximate the value function $V(x, y)$ by $V_{\hbar}(x, y)$ and the expected average cost g by g_{\hbar} . Let $\{\pi_{\hbar}(x, y) : (x, y) \in \Omega_{\hbar}\}$ be the steady-state probability distribution for the Markov chain. Given the singular nature of the control (see chapter 8 of Kushner and Dupuis 2001), we have that

$$g_{\hbar} = \left[\sum_{\partial\Omega_{\hbar}^{\eta}} \frac{\mathcal{R}(y) Q_{\hbar} \pi_{\hbar}(x, y)}{\hbar} + \sum_{\Omega_{\hbar} - \partial\Omega_{\hbar}} c(x) \pi_{\hbar}(x, y) \right] \left(\sum_{\Omega_{\hbar} - \partial\Omega_{\hbar}} \pi_{\hbar}(x, y) \right)^{-1}. \tag{A43}$$

On the other hand, the value function satisfies

$$V_{\hbar}(x, y) = k(x, y) + \sum_{(x', y') \in \Omega_{\hbar}} P_{\hbar}((x, y), (x', y')) V_{\hbar}(x', y'), \tag{A44}$$

where

$$k(x, y) = \begin{cases} (c(x) - g_{\hbar})\Delta t^{\hbar} & \text{if } (x, y) \in \Omega_{\hbar} - \partial\Omega_{\hbar}; \\ \mathcal{R}(y)\hbar & \text{if } (x, y) \in \partial\Omega_{\hbar}^{\eta}; \\ 0 & \text{otherwise.} \end{cases}$$

Because the solution in (A44) is unique up to an additive constant, we can arbitrarily set the value of $V(\hat{x}, \hat{y})$ for a given state $(\hat{x}, \hat{y}) \in \Omega_{\hbar}$. Given the sparse nature of the transition matrix P_{\hbar} , Gaussian elimination can efficiently solve (A44). However, the dimensions of the system increase as $\hbar \rightarrow 0$ and value iteration methods may also be considered.

The final step for updating Ω_{\hbar} is based on a numerical method proposed by Kumar and Muthuraman (2004). First, we observe that the Markov chain approximation asymptotically satisfies (37) as $\hbar \rightarrow 0$, leaving conditions (38)-(40) to be checked. For this purpose, we approximate $V_x(x, y)$ for any $(x, y) \in \Omega_{\hbar}$ by

$$V_x(x, y) \approx \frac{V_{\hbar}(x + \hbar, y) - V_{\hbar}(x, y)}{\hbar}. \quad (\text{A45})$$

Given our current solution Ω_{\hbar} , the boundaries $\partial\Omega_{\hbar}^{\eta}$ and $\partial\Omega_{\hbar}^{\xi}$ are the natural candidates for $\eta(y)$ and $\xi(y)$, respectively. Moreover, from (A44) we have

$$V_{\hbar}(x, y) = \mathcal{R}(y) \hbar + V_{\hbar}(x + \hbar, y), \quad \forall (x, y) \in \partial\Omega_{\hbar}^{\eta},$$

which is equivalent to (using (A45))

$$\frac{V_{\hbar}(x + \hbar, y) - V_{\hbar}(x, y)}{\hbar} + \mathcal{R}(y) = 0 \quad \text{or} \quad V_x(x, y) + \mathcal{R}(y) = 0, \quad \forall (x, y) \in \partial\Omega_{\hbar}^{\eta}.$$

Similarly, (A44)-(A45) imply

$$V_{\hbar}(x, y) = V_{\hbar}(x - \hbar, y), \quad \forall (x, y) \in \partial\Omega_{\hbar}^{\xi} \quad \text{or} \quad V_x(x, y) = 0, \quad \forall (x, y) \in \partial\Omega_{\hbar}^{\xi}.$$

Therefore, given the approximation for $V_x(x, y)$ in (A45), we have that (39) is satisfied with $\{(\eta(y), y)\} = \partial\Omega_{\hbar}^{\eta}$ and $\{(\xi(y), y)\} = \partial\Omega_{\hbar}^{\xi}$. It only remains to check (38) and (40) in the main text. Unless Ω_{\hbar} is optimal at least one of these conditions is violated and we need to update Ω_{\hbar} .




Figure A1: Boundary Update.

Figure A1 plots a prototypical example showing how conditions (38) and (40) in the main text are usually violated and how the update of Ω_{\hbar} is done. Given the two-dimensional nature of the process, Figure A1 shows the update of Ω_{\hbar} for a fixed value $y = \hat{y}$. The solid line represents the value of $V_x(x, \hat{y})$. The abscissas η and ξ are the current values of the boundaries $\partial\Omega_{\hbar}^{\eta}$ and $\partial\Omega_{\hbar}^{\xi}$ at the level \hat{y} . Since condition (38) in the main text is not satisfied in the example, we find the new values η^* and ξ^* by looking for the minimum and maximum values of $V_x(x, \hat{y})$ ($V_{xx}(x, \hat{y}) = 0$) in the range $[\eta, \xi]$. Repeating this procedure for all values of y provides the updated values of $\partial\Omega_{\hbar}^{\eta}$ and $\partial\Omega_{\hbar}^{\xi}$.

We conclude this subsection with a description of the algorithm’s convergence for the base-case data in the main text. To get an idea of the computational complexity of this instance, solving (A44) requires the inversion of a $16,000 \times 16,000$ matrix with 80,000 non-zero entries for the initial state space Ω_h^0 . Figure A2 plots the evolution of Ω_h^k . Initially, $k = 0$ and Ω_h^0 is a rectangle defined by N_x and h . As k increases, region Ω_h^k shrinks as predicted by our previous discussion. The algorithm stops at $k = 6$ when the boundaries remain unchanged from the previous iteration. Figure A2 also shows the values of the expected average cost (g) for each region Ω_h^k . We observe that the improvements obtained during early iterations are significantly larger than those obtained at the end.

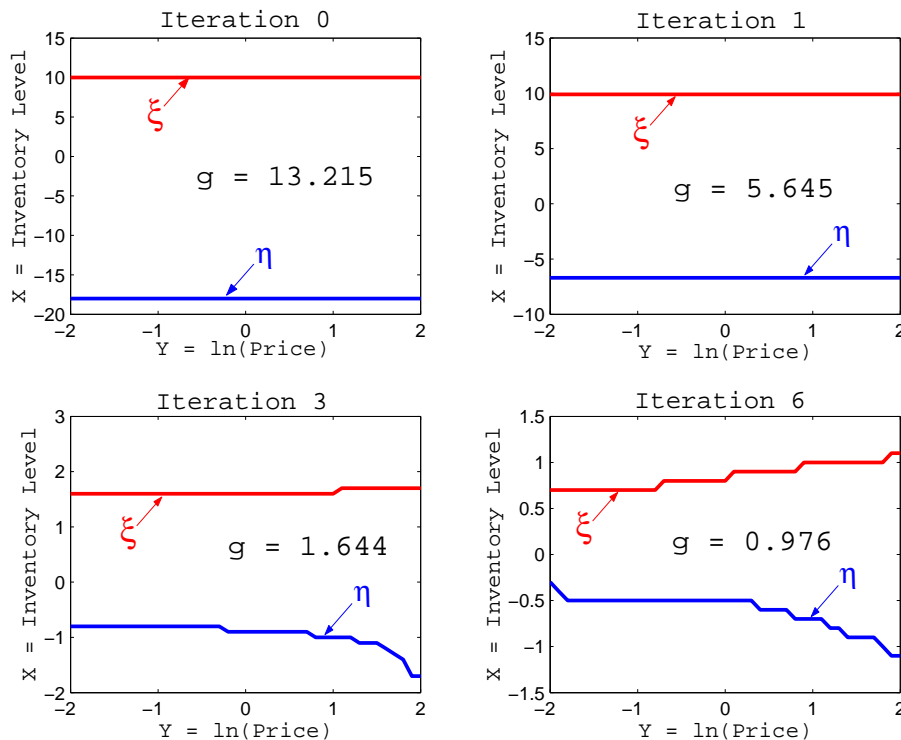


Figure A2: Numerical solution for the base case.

References

- Ata, B., S. Shneerson. 2004. Optimal dynamic control of a service facility. Working paper, Kellogg School of Management.
- George, J.M., J.M. Harrison. 2001. Dynamic control of a queue with adjustable service rate. *Oper. Res.* **49** 720–731.
- Harrison, J.M. 1985. *Brownian Motion and Stochastic Flow Systems*. John Wiley & Sons, New York.
- Harrison, J.M, M.I. Taksar. 1983. Instantaneous control of brownian motion. *Math. Oper. Res.* **8** 439–453.
- Harrison, M., R. Williams. 1987. Multidimensional reflected brownian motions having exponential stationary distributions. *Ann. Probab.* **15** 115–137.
- Kumar, S., M. Muthuraman. 2004. A numerical method for solving singular brownian control problems. *Oper. Res.* **52** 563–582.
- Kushner, H.J. 1977. *Probability Methods for Approximations in Stochastic Control and for Elliptic Equations*. Academic Press, New York.

- Kushner, H.J., P.G. Dupuis. 2001. *Numerical Methods for Stochastic Control Problems in Continuous Time*. Springer Verlag, New York.
- Taksar, M.I. 1985. Average optimal singular control and a related stopping problem. *Math. Oper. Res.* **10** 63–81.