MANAGEMENT SCIENCE
Vol. 58, No. 9, September 2012, pp. 1715–1731
ISSN 0025-1909 (print) | ISSN 1526-5501 (online)

http://dx.doi.org/10.1287/mnsc.1110.1513

Dynamic Pricing with Financial Milestones: Feedback-Form Policies

Omar Besbes, Costis Maglaras
Graduate School of Business, Columbia University, New York, New York 10025
{ob2105@columbia.edu, c.maglaras@gsb.columbia.edu}

We study a seller that starts with an initial inventory of goods, has a target horizon over which to sell the goods, and is subject to a set of financial *milestone* constraints on the revenues and sales that need to be achieved at different time points along the sales horizon. We characterize the revenue maximizing dynamic pricing policy for the seller and highlight the effect of revenue and sales milestones on its structure. The optimal policy can be written in feedback form, where the price at each point in time is selected so as to track the most stringent among all future milestones. Building on that observation, we propose a discrete-review policy that aims to dynamically track the appropriate milestone constraint and show that this simple and practical policy is near optimal in settings with large initial capacity and long sales horizons even in settings with no advance demand model information. One motivating application comes from the sales of new multiunit, residential real estate developments, where intermediate milestone constraints play an important role in their financing and construction.

Key words: revenue management; dynamic pricing; market uncertainty; estimation; feedback; real estate; asymptotic analysis
History: Received November 20, 2009; accepted December 16, 2011, by Gérard P. Cachon, stochastic models and simulation. Published online in *Articles in Advance* April 27, 2012.

1. Introduction

A popular practice in many industries is to make pricing decisions so as to "track" upcoming milestones or targets, essentially lowering prices when falling behind targets and raising prices if running ahead of the respective targets. This is often done without any consideration of the "underlying demand model" that may be unknown or may not be an integral part of the pricing decision itself even if known. Such behavior may be driven by habit, the incentive structure placed on the pricing managers, or by hard constraints on sales and revenue milestones across the horizon. Does this tracking practice make sense given the presence of milestones, and if so, what information about the underlying demand model is needed for such policies to perform well? This paper studies these questions, motivated first by the practical use of such feedback pricing policies, and second by the specific application of pricing new, multiunit, residential real estate projects where financial milestones arise naturally. We show that, yes, such pricing policies may be optimal in many settings, but that the pricing decision should not be selected to "myopically" track the upcoming milestone but rather to track the most stringent among all downstream milestones, in a way that is made precise in our analysis; myopic tracking of the next milestone may perform arbitrarily badly. Second, these policies can achieve good performance with little knowledge of the underlying demand model.

This paper studies the problem of a seller that starts with an initial inventory of goods to be sold over a finite horizon, operates in a setting with limited market information, and tries to select a dynamic pricing strategy to maximize her or his total revenues subject to a set of financial milestone targets, e.g., in terms of revenues or cumulative sales, that need to be met at various points along the sale horizon. This constrained revenue maximization problem can be viewed as a mathematical formalism of the practical approach adopted by some firms that change prices over time to track some implicit or explicit milestones imposed by the firm or by the pricing manager. A separate motivation for the above problem is from settings where financial milestones are inherently present because of the problem structure. One such application, which we will use as an example throughout this paper, is from the real estate industry, and specifically the problem of optimal dynamic pricing of new, multiunit residential developments. A typical timeline for new developments, at least in the residential market, is as follows. (1) The developer, perhaps together with some investors, acquires the land for the development, often involving a relatively short-term loan, e.g., for 24 months. (2) After the initial planning phase of the project (architectural drawings, construction plan, etc.) is completed so as to get approval from local authorities, the developer starts preselling[1] the development's units. At this point construction has not started, and presales continue until the developer reaches certain financial milestones set by the lender that are required to qualify for the construction loan. (3) After the developer secures the construction loan, which is also used to repay the land loan, the developer starts construction of the project and continues with the sale of the remaining inventory, while satisfying progress milestones set by the lender.

The financial milestones that need to be achieved in Step 2 above vary by geographic area, type of development, and the credentials of the developer. In the South Florida market, where significant new construction activity took place over the last decade, it was typical for the lender to require that presales would cover 100% of the loan amount, which itself should correspond to about 65%–70% of the development cost, inclusive of design marketing and sales. Finally, the anticipated sellout value of the project in the developer's proforma should result in 15%–18% profit margin. The latter imposes an implicit constraint on how should the presales target be met.[2] Finally, there may be additional sales and construction milestones while the project is in construction that the developer should meet to qualify for downstream installments of the construction loan.

The developer's goal is to select a pricing policy to maximize the lifetime profitability of the project, while satisfying the financial milestones that she or he may face, such as the cumulative revenue requirement needed to qualify for the construction loan as well as a projected total revenue target. Typically, the market model that captures the size of the various customer segments, their preferences, their willingness to pay (WtP), and the effect of relevant macroeconomic factors are generally unknown. The above formulation is more naturally stated in terms of revenue targets at intermediate points of the sales horizon. In other settings, it may also be appropriate to consider constraints on the number of units sold over different periods of time. A recent example comes from financial institutions and real estate investment trusts that have gained ownership of portfolios of distressed real estate assets that they wish to monetize over time in an uncertain market environment. It is common to manage this process by imposing milestones on the number of units to be sold by certain time points or, equivalently, on the inventory absorption speed. The policies described in this paper are particularly well suited for settings without accurate demand information. We note here that although revenue management tools have been introduced in the past few years in the real estate industry for tactical rental and leasing, as described in Broffman (2007), the sales process of multiunit real estate does not seem to have been addressed commercially or academically in a quantitative framework.[3]

1.1. Main Results and Contributions

The first contribution of this paper lies in the formulation of this problem, which in its many possible variants is of significant practical interest, but to the best of our knowledge has not been studied in the literature thus far.

From an analytical perspective, we first study the seller's problem of optimizing the total revenue assuming that the demand model is known and the demand process is continuous and deterministic. We characterize the optimal solution of this problem and show that it can be written in feedback form, i.e., the optimal price to apply only depends on the cumulative sales and revenues until the current time. In addition, at any point in time, the policy selects a price to track an appropriate envelope of the various milestone constraints, but not necessarily the upcoming constraint; this is an important distinction from the common practice that tends to myopically focus on the next constraint alone. These results extend the analysis of Gallego and van Ryzin (1994) that established the optimality of static pricing in the absence of milestone constraints, where the price to be selected is the maximum between the run-out price that would deplete all the inventory at the end of the horizon and the revenue maximizing price.

Second, we study the general version of the problem described above where demand is stochastic and the seller does not have access to the demand model.

[1] A presale of a unit corresponds to a contract between a buyer and the developer for a specific unit; the buyer pays a deposit (20% of the total sale price is typical) at the time of the signing, which goes in an escrow account, and pays the remainder balance only upon "closings" of the project, where the developer delivers the finished unit to the buyer. The latter needs to take place within a predefined time period, otherwise the buyer has the right to request her or his initial deposit and walk out of the contract.

[2] For example, if the loan amount is for $100 million, then the total pre-sale revenue counting the full price of the units (not just the initial deposits) should equal that amount, the anticipated cost of the development should be $143 million to $154 million; and, at a 15% profit margin, this would correspond to a target sellout value of $164 million to $177 million.

[3] In most settings, pricing of real estate developments is done in an ad hoc manner. Prices tend to be increased over time according to an agreed upon schedule that is not data driven. There tends to be an emphasis at each point in time on the upcoming milestone at a potential detriment of downstream milestones and the overall profitability of the project, and, finally, the joint adjustment of the price matrix of the different types of units of a project tends to also be ad hoc rather than data driven.

Building on the optimal solution of the full information deterministic problem, we propose a discrete review policy that updates prices at discrete points in time, and at each such point adjusts the price upward or downward according to a feedback signal that depends on how well the observed revenue and sales trajectories are tracking their milestones. We specify the appropriate length of discrete review periods in comparison with the potential time horizon such that the seller can estimate the instantaneous demand and revenue rate sufficiently accurately so as to compute the price feedback signal. We then prove that the proposed policy is asymptotically optimal in settings with large initial inventories and long sales horizons, and characterize the magnitude of the revenue loss compared to the performance of an optimal policy that would have full demand information. The good performance of the discrete-review heuristic even in settings without market information is due to its feedback structure. The mode of analysis is quite novel and could be of independent interest for the study of sampled stochastic dynamical systems.

From a practical viewpoint, the policy we propose performs a moderate number of price changes, and their magnitude is intuitive because it is based on a comparison of the estimated sales rate and the distance from the upcoming milestones. This, in conjunction with the simple review structure of the policy, lends it significant appeal. For example, in the real estate application described earlier, such a policy would review status and update its pricing every one to two weeks after each milestone, and less frequently thereafter; significant milestones occur 6 to 12 months apart. The policy would monitor whether the units were being sold too fast, in which case the prices would be raised by some appropriate amount; similarly, if the sales rate was too slow, then prices would be reduced.

1.2. Literature Review

Talluri and van Ryzin (2005) and Phillips (2005) review the area of dynamic pricing and revenue management. Our paper builds on the dynamic pricing model outlined in Gallego and van Ryzin (1994), with two important modifications: (i) the seller faces intermediate constraints, and (ii) the seller need not have access to an exact description of market characteristics.

Our study relates to recent work on pricing decisions when some demand characteristics are unknown. See, e.g., Lobo and Boyd (2003), Aviv and Pazgal (2005), Araman and Caldentey (2009), Besbes and Zeevi (2009), and Broder and Rusmevichientong (2009) for parametric approaches and Eren and Maglaras (2010) and Besbes and Zeevi (2009) for nonparametric approaches. The feedback pricing policy we propose does not rely on learning per se, but rather attempts to track a moving target at any point in time, which leads to a fundamentally different structure to what has appeared previously. When the demand characteristics are known, our model with targets bears some high-level connection to the work of Levin et al. (2007) that studies a dynamic pricing problem while imposing to reach a final revenue target with a given probability.

The policy we ultimately propose in the absence of demand model information can be written in feedback form, which is similar to the popular "resolving" revenue optimization heuristics often based on deterministic relaxations. The latter are tractable and lead to intuitive recommendations. The re-solving aspect allows the seller to correct for the stochastic variability in the sales process. This has been studied in the context of dynamic pricing problems in Maglaras and Meissner (2006) by recognizing that the fixed price heuristic proposed in Gallego and van Ryzin (1994) can be written in feedback form, and in the context of capacity allocation problems in Cooper (2002), Secomandi (2008), Reiman and Wang (2008), and Jasin and Kumar (2010). All of these papers assume full demand model information.

There is a broad literature on discrete- and continuous-review policies for control of stochastic networks; see, e.g., Tassiulas and Papavassiliou (1995), Harrison (1996), Gans and Van Ryzin (1997), and Maglaras (2000). Maglaras (2000) and Paschalidis et al. (2004) are examples of papers that describe tracking policies. Our policies share some of these characteristics, but unlike the above references they are applied here in a setting with uncertain model parameters (the demand model in our context). Discrete-review policies in systems with unknown model parameters were used in Bassambo et al. (2005, 2006) for the control of multiserver stochastic service systems with uncertain and fluctuating arrival rate, and a related revenue maximization problem was studied in Besbes and Maglaras (2009). In contrast to these two references, however, the time-scale separation that underlies their results is not in place in our setting, requiring a more detailed analysis to study the dynamics of the joint estimation and feedback process. Also, the seller's control, the price, affects sales and revenues in an uncertain manner, which adds a layer of complexity. Taken together, these elements suggest the use of review periods of varying length. From a methodological viewpoint, the policy analysis herein is qualitatively different than that of the above references, and of interest in its own right.

2. Model

We consider a seller wishing to sell C units of a homogeneous good over a sales horizon T, with the objective of maximizing the expected revenues accumulated over the horizon. Potential customers arrive according to a homogeneous Poisson process with instantaneous rate Λ, each with a WtP for one unit of the product, v, which is an independent draw from a cumulative distribution function $F(\cdot)$. A customer arriving at time t is quoted the prevailing price $p(t)$ and purchases one unit of the good if $v \geq p(t)$; otherwise, she or he leaves without purchasing a unit. The resulting instantaneous arrival rate at time t is

$$\lambda(p(t)) := \Lambda\,\mathbb{P}(v \geq p(t)) = \Lambda \bar{F}(p(t)), \qquad (1)$$

where $\bar{F}(\cdot)$ is the complementary cumulative distribution function, $\bar{F}(\cdot) := 1 - F(\cdot)$. For all $t \in [0, T]$, let $S(t)$ denote the number of units sold in $[0, t]$, and let $R(t)$ denote the revenues accumulated up to time t. We will use the convention that all time processes are right continuous with left limits.

2.1. Financial Constraints

The seller faces a sequence of financial milestones (constraints) that she or he needs to meet over the course of the sales horizon. We consider two types of constraints.

(1) *Sales targets*. There is a sequence of times $0 < t_1 < t_2 < \cdots < t_k \leq T$ at which sales targets $\xi_1 \leq \xi_2 \leq \cdots \leq \xi_k \leq C$ are imposed.

(2) *Revenue targets*. There is a sequence of times $0 < t'_1 < t'_2 < \cdots < t'_{k'} \leq T$ at which revenue targets $\zeta_1 \leq \zeta_2 \leq \cdots \leq \zeta_{k'}$ are imposed.

That is, ξ_i is the minimum number of units to be sold by t_i, and ζ_i is the cumulative revenue to be accrued up to t'_i. Without loss of generality, we will assume that both types of constraints are provided on a common grid: $\tau_0 = 0 < \tau_1 < \tau_2 < \cdots < \tau_m = T$. If there are no revenue or sales constraints at a time τ_i, we set the corresponding target to zero, and set $\xi_0 = 0$ and $\zeta_0 = 0$. Constraints at time $\tau_m = T$ are optional and may be used to impose terminal milestones, such as that the entire capacity is sold out by T or that a desired revenue target is achieved by T.

For the example discussed in the introduction, the seller would enforce two constraints that would impose that the preconstruction revenues in the first nine months exceed \$100 million and that the sellout value of the project after construction is completed in month 36 would exceed \$177 million.

2.1.1. Penalties and Feasibility.

In a stochastic setting the above constraints may be violated, in which case the seller will incur penalties. For $i = 1, \ldots, m$, $\beta_{1,i}$, and $\beta_{2,i}$ are positive penalty parameters such that the penalty for violating the sales target constraint $S(\tau_i) \geq \xi_i$ and the revenue target constraint $R(\tau_i) \geq \zeta_i$ are, respectively,

$$\beta_{1,i}(\xi_i - S(\tau_i))^+ \quad \text{and} \quad \beta_{2,i}(\zeta_i - R(\tau_i))^+. \qquad (2)$$

In other words, the penalty is proportional to the amount by which the target is missed.[4]

2.2. Objective

Suppose first that the market size, Λ, and WtP distribution, $F(\cdot)$, are known to the decision-maker. We say that $\{p(t)\colon 0 \leq t \leq T\}$ is nonanticipating if the value of $p(t)$ at each time $t \in [0, T]$ is only allowed to depend on past demand observations and price realizations. The seller's revenue maximization problem can be formulated as follows: select a nonanticipating pricing policy $\{p(t), t \in [0, T]\}$ to maximize

$$\mathbb{E}\bigg[\int_0^T p(t)\,dS(t) - \sum_{i=1}^{m}[\beta_{1,i}(\xi_i - S(\tau_i))^+ + \beta_{2,i}(\zeta_i - R(\tau_i))^+]\bigg], \qquad (3)$$

subject to the constraint that

$$S(T) \leq C \quad \text{almost surely.} \qquad (4)$$

We let J^* denote the supremum of (3) over all nonanticipating pricing policies satisfying (4). This paper focuses on a more general version of this problem where the seller does *not* know the market size and the buyers' WtP distribution apart from an upper bound on its support, and investigates the performance of adaptive policies that adjust pricing decisions based on sales observations. To tackle this problem, we first analyze an auxiliary deterministic problem in §3 whose optimal solution will guide the policy proposed for the general setting in §4.

We make the following assumptions throughout this paper.

Assumption 1. *The customers' WtP distribution has support in* $[0, \bar{p}]$, *where* $0 < \bar{p} < \infty$, *and* $\bar{F}(\cdot)$ *is continuously differentiable on* $[0, \bar{p}]$.

We let $\psi(\cdot)$ denote the inverse of $\bar{F}(\cdot)$ over $[0, 1]$ and define the revenue function $r(\lambda) := \lambda\psi(\lambda/\Lambda)$ for all $\lambda \in [0, \Lambda]$.

Assumption 2. (i) $r(\lambda)$ *is strictly concave for* $\lambda \in [0, \Lambda]$. (ii) *Let* $p^* = \arg\max_{p \in [0, \bar{p}]} p\bar{F}(p)$ *be the revenue maximizing price, and let* $\lambda^* = \lambda(p^*)$ *be the corresponding sales rate. Then, capacity is constrained in the sense that for some* $\sigma > 0$,

$$\lambda^* - \sigma \geq C/T, \quad \lambda^* - \sigma \geq \xi_i/\tau_i, \quad \text{and}$$
$$p^*\lambda^* - \sigma \geq \zeta_i/\tau_i, \quad i = 1, \ldots, m.$$

[4] Note that the revenue penalties are computed based on the actual revenues, excluding penalties; i.e., when penalties are incurred, they do not impact future revenue targets. The formulation can easily be modified to cover the case where penalties impact future revenue targets.

The first and second conditions in (ii) above say that the revenue maximizing sales rate will nominally deplete capacity faster than the sales rates that are implicitly defined through the sales constraints. If demand was deterministic, then the first two conditions in (ii) imply that under the revenue maximizing sales rate λ^* the seller would satisfy all of the sales constraints until it runs out of inventory, which occurs before the end of the sales horizon. The third condition in (ii) ensures that in a deterministic setting the revenue targets are also feasible. The conditions in (ii) also amount to assuming that the demand function $\lambda(p)$ is elastic in the relevant price range defined via the sales and revenue constraints. Specifically, let

$$\varphi(u) := \min\{\lambda \colon r(\lambda) = u\} \tag{5}$$

be the inverse of the revenue function with codomain $[0, \lambda^*]$ that defines the sales rate $\varphi(u)$ that would achieve an expected revenue rate of u. Let $p_0 = \psi(C/(T\Lambda))$ denote the price that would deplete the capacity exactly at time T, let $p_i^s = \psi(\xi_i/(\tau_i\Lambda))$ denote the price that would induce the target sales rate defined by the ith sales constraint, and let $p_i^r = \psi(\varphi(\zeta_i/\tau_i)/\Lambda)$ denote the price that would induce a revenue rate that would meet the ith revenue constraint. The conditions in (ii) imply that the demand function $\lambda(p)$ is elastic for all prices $p \geq \min(p_0, p_1^s, \ldots, p_m^s, p_1^r, \ldots, p_m^r)$. The presence of σ in the assumption ensures that the target prices above are separated from the revenue maximizing price, which, in the stochastic setting with limited information (§4), will ensure that the prescribed prices will stay in the elastic region with (appropriately) high probability.

We also make the following assumption on $F(\cdot)$ and the revenue function:

Assumption 3. *For some $\underline{K}, \bar{K}, \underline{\kappa} > 0$, we have that $\underline{K} \leq F'(p) \leq \bar{K}$ on $[p^*, \bar{p}]$ and $r'(\lambda^* - \sigma/2) \geq \underline{\kappa}$ for the value of σ used in Assumption* 2.

The second part of Assumption 3 is automatically satisfied under the other assumptions imposed on the distribution $F(\cdot)$. The upper bound on the derivative of F ensures that small price changes result in small changes in demand. The lower bound on the derivative ensures that the magnitude of a price change that is needed to induce a given demand change is bounded. The lower bound on the revenue derivative stems from similar considerations. These properties are key to prove that the policy we propose in §4 performs well; this policy tracks the sales and revenue targets by applying a price feedback, and the above bounds allow one to choose the feedback "gain" so as to guarantee that one can reliably approach the targets. In addition, the fact that the derivative of the revenue rate is bounded below ensures that it is never optimal to violate the constraints for sufficiently large penalty parameters, as laid out in §3.2 (see Proposition 2).

3. Analysis of a Related Deterministic Problem

This section analyzes a deterministic version of problem (3) that requires that the sales and revenue constraints are *always* met. The problem is as follows:

$$\max\left\{\int_0^T r(\lambda(s))\,ds \colon \int_0^{\tau_i} \lambda(s)\,ds \geq \xi_i,\ \int_0^{\tau_i} r(\lambda(s))\,ds \geq \zeta_i,\ i = 1, \ldots, m,\ \int_0^T \lambda(s)\,ds \leq C\right\}, \tag{6}$$

where (6) is obtained from (3) by replacing sales and revenue processes by their means and imposing that the decision maker never incurs any penalties. In addition, (6) is stated with the sales rate as a decision variable as opposed to the price, a classical transformation in revenue management problems (see, e.g., Talluri and van Ryzin 2005), with the implicit assumption that the decision maker infers the corresponding price using the demand relationship. This, of course, is not possible in settings where the demand model is not known, a problem that we will address in §4. When (6) is feasible, we let J^d denote the optimal value of the problem. In §3.1, we will offer a constructive criterion of feasibility for (6) and characterize its optimal policy; subsequently, in §3.2, we relate the optimal values J^d and J^*.

3.1. Optimal Pricing Policy

The formulation in (6) is a deterministic optimal control problem with interior constraints. Whereas the control literature (see, e.g., Bryson and Ho 1975 and related references) gives implicit conditions for feasibility and characterization of an optimal control, here, we will construct an explicit, feedback-form expression for the optimal policy. The first step is to consider at any time t the effect of all upcoming constraints at milestone times $\tau_i > t$. Specifically, for all $i = 1, \ldots, m$ such that $\tau_i > t$, for all cumulative sales $y \geq 0$ and cumulative revenues $z \geq 0$, define

$$\ell_i^s(y, t) := \frac{\xi_i - y}{\tau_i - t}, \tag{7}$$

$$r_i(z, t) := \frac{\zeta_i - z}{\tau_i - t}, \tag{8}$$

$$\ell_i^r(z, t) := \varphi\big(\min\{r(\lambda^*), \max\{r_i(z, t), 0\}\}\big), \tag{9}$$

where $\varphi(\cdot)$ was defined in (5). The quantity $\ell_i^s(y, t)$ in (7) defines the constant sales rate that is needed to ensure that starting from the current cumulative sales level y at time $t < \tau_i$ the seller will reach the target sales level of ξ_i at time τ_i (when $\xi_i \geq y$). The quantity $r_i(z, t)$ in (8) is the constant revenue rate that ensures that starting from the current cumulative revenue level of z at time $t < \tau_i$ the seller will reach the target revenue level of ζ_i at time τ_i (when

$\zeta_i \geq z$). The quantity $\ell_i^r(z, t)$ in (9) is the sales rate in $[0, \lambda^*]$ that will achieve the revenue rate $r_i(z, t)$ (when $r_i(z, t) \leq r(\lambda^*)$).

First, we construct an open loop control designed to track the various milestone constraints as follows. Let $S_0 = 0$ and $R_0 = 0$ and define, for $i = 0, \ldots, m-1$,

$$\tilde{\lambda}_i = \max_{i+1 \leq j \leq m} \{\ell_j^s(S_i, \tau_i), \ell_j^r(R_i, \tau_i), (C - S_i)/(T - \tau_i)\}, \quad (10)$$

where, for $i = 0, \ldots, m-1$,

$$S_i = \sum_{j=0}^{i-1} \tilde{\lambda}_j (\tau_{j+1} - \tau_j) \quad \text{and} \quad R_i = \sum_{j=0}^{i-1} r(\tilde{\lambda}_j)(\tau_{j+1} - \tau_j). \quad (11)$$

Proposition 1 (Feasibility). *Let Assumptions* 1 *and* 2 *hold. Then, problem* (6) *is feasible if and only if* $\sum_{i=0}^{m-1} \tilde{\lambda}_i (\tau_{i+1} - \tau_i) \leq C$.

The next result characterizes the optimal policy for problem (6).

Theorem 1 (Optimal Policy). *Let Assumptions* 1 *and* 2 *hold and suppose that* (6) *is feasible. Then, the following policy is optimal: for* $i = 0, \ldots, m-1$,

$$\tilde{\lambda}(t) = \tilde{\lambda}_i \quad \text{for } t \in [\tau_i, \tau_{i+1}). \quad (12)$$

The optimal control (12) *can equivalently be expressed as follows*:

$$\tilde{\lambda}(t) = \max\left\{\max_{i:\tau_i > t} \ell_i^s(S(t), t), \max_{i:\tau_i > t} \ell_i^r(R(t), t), (C - S(t))/(T - t)\right\}, \quad t \in [0, T), \quad (13)$$

where $\ell_i^s(\cdot, \cdot)$ *and* $\ell_i^r(\cdot, \cdot)$ *are as defined in* (7) *and* (9), *respectively.*

The optimal policy described above "tracks" the sales and revenue targets in the following way: at any point in time, it (i) computes the sales rates that one would need to apply to achieve each downstream sales and revenue target, (ii) selects the highest among these rates that corresponds to the most binding constraint at that time, and (iii) applies that control until the next milestone instance occurs; in other words, the optimal control is piecewise constant, which is clearly seen through the open loop version of (13), given in (12). The structure of the optimal policy and specifically the fact that one needs to track the most binding constraint is a consequence of the concavity of the revenue function (as a function of the sales rate) in conjunction with the elasticity conditions under which one operates (see Assumption 2(ii)).

The preceding analysis shows that milestone tracking is optimal in settings where demand is elastic in the relevant price region (see Assumption 2). It is important to note, however, that one should not necessarily myopically track the "next" milestone, but instead one should be tracking the most stringent among all downstream milestones, essentially creating a "least envelope" of all of the constraints that should guide the control decision.

In the absence of intermediate constraints, one recovers the result provided in Gallego and van Ryzin (1994, Proposition 2). Indeed, under Assumption 2, the latter states that it is optimal to apply a constant rate C/T throughout the horizon, which can be in feedback form as follows: $\lambda(t) = (C - S(t))/(T - t)$, $t \in [0, T)$. This is exactly what (13) reads when $\xi_i = \zeta_i = 0$ for $i = 0, \ldots, m$.

The optimal policy in this deterministic setting also satisfies the following structural property.

Corollary 1 (Monotonicity). *Let Assumptions* 1 *and* 2 *hold. Suppose that problem* (6) *is feasible; then the optimal pricing policy* (13) *is monotonically nondecreasing over time.*

Focusing on the feedback form of the optimal control in (13), it is important to note that the information needed to compute the optimal sales rate at any point in time is the future sales and revenue targets and the current accumulated sales and revenue. In particular, the market size and WtP distribution are not needed for the above calculation and are only (implicitly) used in converting the optimal sales rate to the corresponding price. This structural property of the optimal policy will prove very useful when we consider the limited demand model information setting in §4. To that end we present an alternative characterization of the feedback policy. Let $\tilde{p}(\cdot)$, $\tilde{S}(\cdot)$, and $\tilde{R}(\cdot)$ denote the price, cumulative sales, and cumulative revenue paths corresponding to the optimal policy (13), respectively. Under the optimal policy, $\tilde{p}(\cdot)$, $\tilde{S}(\cdot)$, and $\tilde{R}(\cdot)$ satisfy

$$\min\{\Delta_1(t, \tilde{p}(t), \tilde{S}(t)), \Delta_2(t, \tilde{p}(t), \tilde{R}(t))\} = 0 \quad \text{for all } t \in [0, T), \quad (14)$$

where for any $t, S, R \geq 0$ the feedback signals Δ_1 and Δ_2 are defined as

$$\Delta_1(t, p, S) := \lambda(p) - \max\left\{\max_{i:\tau_i > t}\left\{\frac{\xi_i - S}{\tau_i - t}\right\}, \frac{C - S}{T - t}\right\}, \quad (15)$$

$$\Delta_2(t, p, R) := r(\lambda(p)) - \max_{i:\tau_i > t}\left\{\frac{\zeta_i - R}{\tau_i - t}\right\}; \quad (16)$$

that is, Δ_1 measures how far is the current sales rate from the rate needed to achieve the most binding sales constraint, and Δ_2 measures how far is the current revenue rate from the rate needed to achieve the most stringent revenue target.

3.1.1. Numerical Illustration. We first depict below the structure of the optimal policy (13) for the following instance. The seller has $C = 200$ units to sell over a time horizon of $T = 1{,}000$ days. There are $m = 4$ milestones. After the first $\tau_1 = 300$ days, the seller needs to have generated $\zeta_1 = 1{,}000$ units of revenue. In addition, the seller needs to have sold $\xi_1 = 100$ units by τ_1, $\xi_2 = 160$ units by $\tau_2 = 500$ days, and $\xi_3 = 180$ units by $\tau_3 = 700$ days. We assume that the true underlying mean demand rate (mean number of sales per day) when the price is p is given by $\lambda(p) = 2 - 0.2 * p$ for $p \in [0, 10]$, which corresponds to $\Lambda = 2$ and F uniformly distributed in $[0, 10]$.

Figure 1 illustrates the structure of the optimal policy by depicting the rates $\max\{\ell_i^s(t, S(t)), 0\}$ for $i = 1, \ldots, 3$, $\max\{\ell_1^r(t, R(t)), 0\}$ and $(C - S(t))/(T - t)$ for $t \in [0, T)$, which are the only quantities needed to compute the optimal control. At any point in time, the optimal policy selects the maximum among these rates. At time $t = 0$, the next most binding constraint is the revenue constraint at time $\tau_1 = 300$; hence, the optimal policy selects $\ell_1^r(0)$ initially. By time $t = 300$, the next most binding constraint is the sales constraint at time $\tau_2 = 500$, and the corresponding rate is selected by the optimal policy, which continues in this fashion.

Also, Table 1 gives a numerical illustration of the potential consequences of myopic milestone tracking through a simple example. We consider a case with the same demand function, time horizon, and capacity as above but in which now there is a single revenue constraint at time τ_1 given by $\zeta_1 = 400$ and no sales constraints. Table 1 reports the ratio between the performance of the myopic policy that picks a price to always track the immediately upcoming constraint and that of the optimal policy for different values of τ_1. When $\tau_1 = 200$, the revenue constraint is the most stringent target, and both the optimal and the myopic policies coincide. When the milestone time τ_1 increases, the first revenue target is not anymore binding, and the optimal policy should price so as to achieve the run-out rate C/T. However, the myopic policy will still track the first revenue target until τ_1 and only after that switch to the run-out rate. This behavior is suboptimal and, depending on the model parameters, can be arbitrarily bad.

Figure 1 **Construction of the Optimal Policy in the Deterministic Setting**

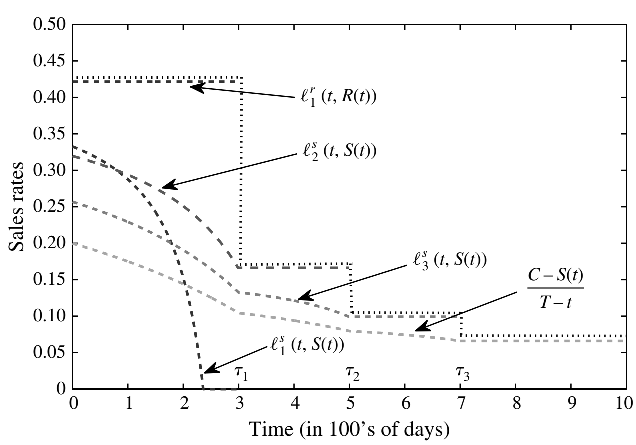

Note. This figure depicts the rates from which the optimal sales rates are derived; at any point in time, the maximum of these (dotted line) corresponds to the optimal sales rate.

Table 1 **Myopic Milestone Tracking**

	Milestone time τ_1			
	200	400	600	800
Performance ratio	1	0.98	0.93	0.75

Note. This table reports the ratio of the performance of the policy that myopically tracks the upcoming target relative to an optimal policy.

3.2. An Upper Bound on the Value of the Stochastic Problem

Proposition 2. *Let Assumptions* 1–3 *hold. Suppose that problem* (6) *is feasible. Then, if* $\min_{i=1,\ldots,m}\{\min\{\beta_{1,i}, \beta_{2,i}\}\} > \bar{p}\max\{1, \underline{\kappa}^{-1}\}$, *problem* (6) *and the problem*

$$\max\left\{\int_0^T r(\lambda(t))\,dt - \sum_{i=1}^m \beta_{1,i}(\xi_i - S(\tau_i))^+ - \sum_{i=1}^m \beta_{2,i}(\zeta_i - R(\tau_i))^+ \colon \int_0^T \lambda(t)\,dt \le C\right\} \tag{17}$$

have the same optimal solution.

This is an intuitive result that states that it is never optimal to violate the sales and revenue targets when the penalties are sufficiently high as long as avoiding to do so is feasible. Now, noting that the objective above is concave in $\{\lambda(t)\colon 0 \le t \le T\}$, one can use a similar argument as Gallego and van Ryzin (1994) to establish that the value of the deterministic relaxation (17) serves as an upper bound to the value of the original stochastic problem (3). We deduce that when (6) is feasible and the penalties are sufficiently large, then, by Proposition 2, the value of problem (6) serves as an upper bound to the value of (3), i.e.,

$$J^* \le J^d. \tag{18}$$

In all that follows, we assume that $\beta_{1,i}$, $\beta_{2,i}$ are indeed sufficiently large. Note that if this is not the case, when the demand model is known, the deterministic relaxation with the penalty terms is still a concave maximization problem that is tractable, and its solution can provide a good heuristic for the underlying stochastic problem.

4. The Limited Demand Model Information Setting

This section studies the stochastic problem (3) in a setting where the market size Λ and the WtP distribution $F(\cdot)$ are initially unknown to the decision-maker. It introduces a class of policies for this problem that are motivated by the structure identified in the solution of the full information deterministic setting, and establishes that these policies achieve "good" performance.

The key observation is that the optimal policy under full information (in the deterministic setting of the previous section) is to select a price so as to track an appropriately selected dynamic target. This policy takes the form of a feedback rule, where the feedback signals that drive the direction of the price adjustment can be approximated from directly observable quantities without any demand model information. Specifically, the proposed solution attempts to "mimic" the structure of the deterministic policy as described in (14). While imposing that $\min\{\Delta_1(t, p(t), S(t)), \Delta_2(t, p(t), R(t))\} = 0$ for all $t \in [0, T)$ is not possible when $\lambda(\cdot)$ is unknown, one can attempt to systematically drive this quantity to zero by estimating the demand and revenue rates at the current price. This is the essence of the feedback pricing policy we consider, which will (1) review the system status every δ time units, (2) estimate the running sales and revenue rates, (3) estimate the feedback signals, and (4) apply a price adjustment according to

$$p(t+\delta) = p(t) + \alpha \min\big\{\hat{\Delta}_1(t, p(t), S(t)), \hat{\Delta}_2(t, p(t), R(t))\big\}, \quad (19)$$

where α is a positive constant, the "gain" of the feedback control, and $\hat{\Delta}_1(t, p(t), S(t))$ and $\hat{\Delta}_2(t, p(t), R(t))$ are noisy estimates of $\Delta_1(t, p(t), R(t))$ and $\Delta_2(t, p(t), R(t))$, respectively (see (15) and (16)). The price is adjusted so as to drive $\min\{\hat{\Delta}_1, \hat{\Delta}_2\}$ to zero, and thereby follow the path that was shown to be optimal in the deterministic full information problem in §3.1. If $\min\{\hat{\Delta}_1, \hat{\Delta}_2\} > 0$, then the current sales rate is too high relative to what is needed to achieve the next most binding target, and as a result the price is increased. If $\min\{\hat{\Delta}_1, \hat{\Delta}_2\} < 0$, the current sales rate is too low, and to avoid missing the next most binding target, the price is decreased. If $\min\{\hat{\Delta}_1, \hat{\Delta}_2\} = 0$, the price is unchanged. The length of the review period δ is sufficiently long so that the seller can form accurate estimates of the sales and revenue rates at the current price, but short enough so that the policy "tracks" the desired trajectory with minimal revenue loss.

4.1. The Feedback Pricing Policy

The proposed policy is parametrized by a positive constant, α, and a sequence $\{\delta_j\colon 0 \le j \le k, \sum_{j=0}^{k} \delta_j = 1\}$. The constant α defines the strength of the feedback and $\{\delta_j\}$ defines how to split each interval $[\tau_i, \tau_{i+1})$ into successive review periods. Specifically, each time interval between successive milestones of the form $[\tau_i, \tau_{i+1})$ will be split in $k+1$ review periods of non-decreasing length equal to $\delta_j(\tau_{i+1} - \tau_i)$ for $j = 0, \ldots, k$. More precisely, the sequence of review times is as follows:

$$t_0^i = \tau_i,$$
$$t_{j+1}^i = t_j^i + (\tau_{i+1} - \tau_i)\delta_j, \quad j = 0, \ldots, k,$$

where the superscript denotes the nearest milestone τ_i and the subscript denotes the review period counter within the interval $[\tau_i, \tau_{i+1})$. We also let $\delta_j^i = t_{j+1}^i - t_j^i$. We describe below the policy in algorithmic form.

Feedback Pricing Policy: $\pi(\alpha, \delta)$
Select some initial price[5] $p(0)$ such that $\Lambda \bar{F}(p(0)) \le \lambda^* - \sigma$
`For` $i = 0, \ldots, m-1$
  If $i \ge 1$, set $p(t_0^i) = p(t_k^{i-1})$.
  `For` $j = 0, \ldots, k$
  1. **Price/Information Collection:**
     Set $p(t) = p(t_j^i)$ for $t \in (t_j^i, t_{j+1}^i]$
  2. **Estimation:** Compute

$$\hat{\lambda}(t_j^i) = \frac{S(t_{j+1}^i) - S(t_j^i)}{t_{j+1}^i - t_j^i} \quad \text{[estimated demand rate]} \quad (20)$$

$$\hat{r}(t_j^i) = p(t_j^i)\hat{\lambda}(t_j^i) \quad \text{[estimated revenue rate]} \quad (21)$$

  3. **Target Rate and Feedback Update:** Compute

$$\hat{\Delta}_1 = \hat{\lambda}(t_j^i) - \max\left\{\max_{i\colon \tau_i > t_{j+1}^i}\left\{\frac{\xi_i - S(t_{j+1}^i)}{\tau_i - t_{j+1}^i}\right\}, \frac{C - S(t)}{T - t}\right\} \quad (22)$$

$$\hat{\Delta}_2 = \hat{r}(t_j^i) - \max_{i\colon \tau_i > t_{j+1}^i}\left\{\frac{\zeta_i - R(t_{j+1}^i)}{\tau_i - t_{j+1}^i}\right\} \quad (23)$$

$$\hat{\phi}(t_{j+1}^i) = \min\{\hat{\Delta}_1, \hat{\Delta}_2\} \quad \text{[feedback direction]} \quad (24)$$

  4. **Price Update:**

$$\text{Set } p(t_{j+1}^i) = p(t_j^i) + \alpha\hat{\phi}(t_{j+1}^i) \quad \text{[apply feedback]} \quad (25)$$

  `End`
`End`

In the above description, we assume that as soon as the seller runs out of inventory, one changes the price to $\bar{p}$ until T. After each review period, one estimates

[5] In the limited demand model information setting, it is always possible to do so based on the current assumptions; for example, it possible to show that $p(0) = \bar{p} - \bar{K}^{-1}C/T$ satisfies this condition.

the demand rate at the current price $\hat{\lambda}(t_j^i)$ and the corresponding revenue rate $\hat{r}(t_j^i)$ (Step 2). Then, in Step 3, one estimates how far off one is from the target rate, the rate that an optimal policy in the deterministic setting (6) would apply given the sales and revenues accumulated up to the current time. In Step 4, one updates the price based on the sign and the magnitude of the difference between the current and actual target rates, as well as the feedback parameter α. The parameters defining the policy allow one to tune the strength of the feedback through α, as well as to control for the estimation error through the lengths δ_j that characterize the time windows over which the price is held constant.

4.2. Performance Analysis

A head-on analysis of the limited demand model information problem for the seller as well as an exact characterization of the performance of the proposed policy seems intractable. In this section, we focus on the performance of the policy specified above in settings where the seller starts with a large initial inventory that needs to be sold over a long time horizon, which, for example, is often relevant in residential real estate setting discussed in the introduction. Our main result is that for such problems the policy is asymptotically optimal in a manner that is made precise below.

Specifically, we focus on a sequence of problems, indexed by a positive integer n, where

$$\tau_{i,n} = n\tau_i, \quad C_n = nC, \quad \xi_{i,n} = n\xi_i, \quad \zeta_{i,n} = n\zeta_i, \quad i = 0, \ldots, m, \tag{26}$$

i.e., the problems get larger in terms of the seller's inventory and respective revenue and sales targets, and the horizon over which the sales are to take place grows proportionally large. The market size, WtP distribution, and penalty parameters $(\beta_{1,i}, \beta_{2,i})$, $i = 1, \ldots, m$, are unchanged. Because the targets are scaled proportionally with the inventory and the time horizon, the milestones and the capacity constraint continue to play a critical role as n grows. All quantities in a system with scale n will be denoted with a subscript n. For example, we let J_n^* denote the optimal value of the problem in the system with scale n. We also introduce the following notation; we let

$$Z_n = \sum_{i=1}^{m} \big[\beta_{1,i}(\xi_{i,n} - S_n(\tau_{i,n}))^+ + \beta_{2,i}(\zeta_{i,n} - R_n(\tau_{i,n}))^+\big] \tag{27}$$

denote the total penalties paid over the course of the horizon. The performance of any admissible policy can then be written as $\mathbb{E}[R_n(T_n) - Z_n]$ and is always upper bounded by J_n^*.

A prescription for the length of the review period lengths would need to balance (a) the estimation error of demand that decreases with the length of the period with (b) the tracking error of the targets that increases with the window length.

4.2.1. Intuition and Motivation for the Selection of Review Period Lengths. The optimal pricing policy for the deterministic problem (6) described in §3 is piecewise constant, with price changes occurring only at milestone times. Without knowledge of the demand model, in the stochastic system, after each milestone (or before the start of the season), the initial price is likely to be away from the price that would be prescribed under the optimal solution to the deterministic relaxation. The feedback policy we propose estimates the sales and revenue rates over discrete review periods of length δ_j^i units of time and attempts to track the target rate prescribed by (13). Assume for a moment that once the seller realizes that the current price level is wrong and identifies the correct sales rate target, she is able to change her price to a level such that the current estimate of the sales rate, $\hat{\lambda}(p)$, is exactly at the target sales rate. (Of course, the seller may not do so in a single step, but we will argue later on that she may do so in a "small" number of steps.) The estimate of the sales rate at a given price after a review period of length δ will be noisy. Because demand arrivals follow a Poisson process, the difference between the realized number of arrivals over a review period and its expected value is of order $\sqrt{\delta}$ with high probability; this noise translates into a possible error in the estimation of the sales rate of order $1/\sqrt{\delta}$, and roughly speaking, this defines a lower limit in performance in that one cannot approach the target with a better resolution than $1/\sqrt{\delta}$. Given the above discussion, the error accumulated over $[t_{j+1}^i, t_{j+2}^i]$ is of order $(\delta_j^i)^{-1/2}\delta_{j+1}^i$, i.e., the (pricing error due to the estimation noise)×(duration of the interval). Summing over all review periods between two successive milestones $\tau_{i,n}$ and $\tau_{i+1,n}$, and recognizing that over the first review period the pricing error may be of order 1, thus resulting into a revenue loss of order δ_0^i, the total potential revenue loss is of order

$$\delta_0^i + \sum_{j=1}^{k} (\delta_{j-1}^i)^{-1/2} \delta_j^i. \tag{28}$$

To reduce the estimation error, one would like to use longer review periods, but this would prolong the time spent at the wrong price levels and therefore incur less revenues. To optimize this trade-off, the seller should adjust the price over a sequence of review periods of increasing length that allow the seller to learn fast in the beginning and then refine her accuracy through longer review periods. Broadly, the problem of selecting the appropriate lengths of review periods is one of selecting a sequence $\{\delta_j^i\colon j = 0, \ldots, k\}$ that minimizes the potential revenue loss described above subject to the constraint that $\sum_{j=0}^{k}\delta_j^i = \tau_{i+1,n} - \tau_{i,n}$, which says that the sum of the review periods must be of order n because $\tau_{i+1,n} - \tau_{i,n}$ is of order n. One can show that no feasible sequence of reviews could guarantee losses with a lower order of magnitude than $O(\sqrt{n})$.[6] One way to select the δ_j^i's and (almost) achieve the lower bound of $O(\sqrt{n})$ on revenue loss is to pick the δ_j^i's to balance all terms in the sum in (28), while being of order $\sqrt{n}$; this leads to the increasing sequence

$$\delta_0^i = \sqrt{n}, \quad \delta_j^i = \sqrt{n}(\delta_{j-1}^i)^{1/2}, \quad j \geq 1,$$

which may be rewritten as $\delta_j^i = n^{1-1/2^{j+1}}$ for $j \geq 0$. With such a selection, the sum in (28) is of order $(k+1)\sqrt{n}$. One is left to select an appropriate number of reviews k such that the sequence is indeed feasible. It is easy to see that for

$$\bar{k} = (\log 2)^{-1} \text{ log-log } n, \tag{29}$$

$\delta_{\bar{k}}^i = n\exp\{-1/2\}$, and that $\delta_j^i = n\exp\{-(\bar{k}+1-j)/2\}$, $j = 0, \ldots, \bar{k}-1$. Hence, it suffices to have $\bar{k}$ reviews to ensure that $\sum_{j=0}^{\bar{k}}\delta_j^i$ is of order n, i.e., to cover the time interval between milestones.

The above discussion assumed that once the seller identified a target sales rate, she could pick the right price that will achieve that rate. When the demand model is unknown, the seller can only make directionally correct price adjustments, i.e., increasing the price when the sales rate needs to be reduced, and decreasing the price otherwise. To achieve the desired sales rate, she will have to apply several such directional feedback adjustments. To explain why a small number of such adjustments may be sufficient, consider a simpler recursion than (25), in which one attempts to reach a static as opposed to a dynamic target. For concreteness, suppose that one attempts to reach a sales rate $\lambda(\bar{q})$ for some target price $\bar{q}$ and uses the feedback rule $p(t_{j+1}^i) = p(t_j^i) + \alpha(\hat{\lambda}(p(t_j^i)) - \lambda(\bar{q}))$. In such a case, $p(t_{j+1}^i) - \bar{q} = p(t_j^i) - \bar{q} + \alpha(\lambda(p(t_j^i)) - \lambda(\bar{q})) + \alpha(\hat{\lambda}(p(t_j^i)) - \lambda(p(t_j^i)))$, and if $\alpha < (\Lambda\underline{K})^{-1}$, one may show that, under Assumption 3, $|p(t_{j+1}^i) - \bar{q}| \leq (1-\alpha\Lambda\underline{K})|p(t_j^i) - \bar{q}| + \alpha|\hat{\lambda}(p(t_j^i)) - \lambda(p(t_j^i))|$. Applying this inequality in a recursive fashion yields that

$$\left|p(t_{j+1}^i) - \bar{q}\right| \leq (1-\alpha\Lambda\underline{K})^{j+1}\left|p(t_0^i) - \bar{q}\right| + \sum_{k=0}^{j}(1-\alpha\Lambda\underline{K})^{j-k}\alpha\left|\hat{\lambda}(p(t_k^i)) - \lambda(p(t_k^i))\right|.$$

The first term on the right-hand side above shrinks to zero geometrically fast when $|1-\alpha\Lambda\underline{K}| < 1$; each term of the form $|\hat{\lambda}(p(t_k^i)) - \lambda(p(t_k^i))|$ in the sum represents fluctuations of the sample mean arrival rate around the true one of a Poisson process over δ_j^i units of time and is of order $O((\delta_j^i)^{-1/2})$ with high probability. As a result, if $j \approx \mathscr{C}\log(n)$ steps, one would have, with high probability, $|p(t_{j+1}^i) - \bar{q}| \leq n^{\mathscr{C}\log(1-\alpha\Lambda\underline{K})}|p(t_0^i) - \bar{q}| + O((\delta_j^i)^{-1/2}\log n)$. With a proper choice of $\mathscr{C}$, the first term can be arbitrarily small, and the price will be within order $((\delta_j^i)^{-1/2}\log n)$ of the target price, almost as assumed above in the discussion that led to (28). This motivates the definition of j^* and the construction below.

4.2.2. Definition of the Review Periods. Recall the definition of $\bar{k}$ in (29), and let

$$j^* = C_{j^*}\log n.$$

where C_{j^*} is a positive constant further specified in the proof of Theorem 2. For each time interval $[\tau_i, \tau_{i+1})$, the policy will make $\bar{k}j^*$ reviews and price adjustments, where $\bar{k}$ is defined in (29). The first j^* review period lengths are of order $n^{1-1/2}$, the next j^* reviews will be of length $n^{1-1/2^2}$, and so on, until the last review period, which is length $n^{1-1/2^{\bar{k}}}$. Define, for $k = 1, \ldots, \bar{k}$,

$$\delta_{(k-1)j^*+j} = C_\delta C_\nu n^{1-1/2^k}, \quad j = 0, \ldots, j^*-1, \quad \text{and} \tag{30}$$

$$\delta_{\bar{k}j^*} = C_\delta n^{1-1/2^{\bar{k}}}, \tag{31}$$

where C_ν is a positive constant in $(0,1]$ further specified in the proof of Theorem 2, and C_δ is a normalization constant selected so that $\sum_{j=0}^{j^*\bar{k}}\delta_j = n$.

Qualitatively, after each milestone, the proposed policy uses shorter review periods to quickly adjust the price near the appropriate level. Subsequently, the policy reviews its price less frequently to refine its estimate of the demand at the posted price and adjust more carefully its price. For example, when selling a few hundred units, over a period of approximately three years (1,000 days), the policy would prescribe to update prices every week after each milestone and slowly reduce the frequency of price changes so that one updates the price only every month or so when one approaches the next milestone. The nonuniform length of the review periods is intuitively appealing and serves to minimize the potential revenue loss. The next result characterizes the performance of the feedback pricing policy in the context of the asymptotic regime (26).

Theorem 2. *Let Assumptions* 1–3 *hold. Suppose that problem* (6) *is feasible. Then, for a sequence* δ *selected according to* (30)–(31) *with suitable choices of* C_{j^*} *and* C_ν

[6] Indeed, arguing by contradiction, suppose for a moment that there exists a sequence of review periods such that $\delta_0^i + \sum_{j=1}^{k}(\delta_{j-1}^i)^{-1/2}\delta_j^i \leq \frac{1}{2}\sqrt{n}$ and $\sum_{j=0}^{k}\delta_j^i = n$. Multiplying both sides of the first equation by $\sqrt{n}$ and lower bounding $\sqrt{n}(\delta_{j-1}^i)^{-1/2}$ by 1, one obtains $\sqrt{n}\delta_0^i + \sum_{i=1}^{k}\delta_j^i \leq n/2$. Given that $n \geq 1$, this implies that $\sum_{i=0}^{k}\delta_j^i \leq n/2$, which is a contradiction because $\sum_{i=0}^{k}\delta_j^i = n$. This heuristically provides an estimate of a lower bound for the losses than one may incur for this class of feedback policies.

Figure 2 **Sample Path of the Price, Cumulative Sales, and Revenue Processes**

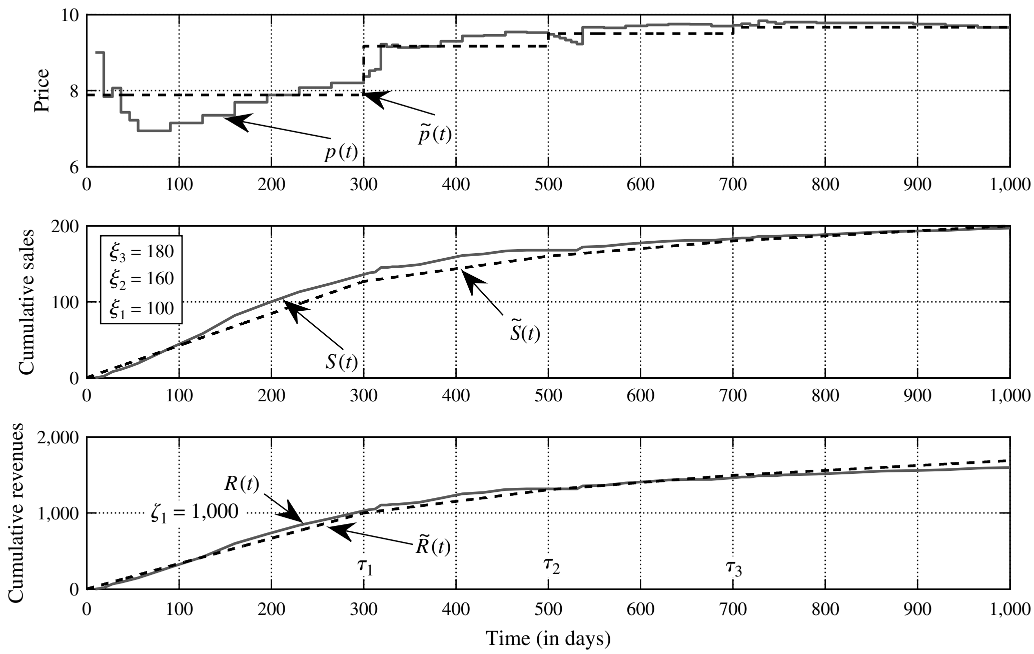

Note. The policy π_f uses a feedback parameter $\alpha = 0.75$ and starts at the price $p_0 = 9$.

and of a constant feedback parameter α, *the performance of the feedback pricing policy with parameters* (α, δ) *satisfies*

$$\frac{E[R(T_n) - Z_n]}{J_n^*} \geq 1 - \frac{\mathscr{C}(\log n)^2}{n^{1/2}}$$

for some positive constant $\mathscr{C}$.

Noting that J_n^* is of order n, Theorem 2 asserts that the proposed policy, which does not make any use of the market size or the WtP distribution, is able to track the relevant constraints appropriately and approximately achieve the performance of an optimal policy that would have access to the demand model.

To put this result into perspective, suppose that one were to restrict attention to fixed length review periods. In such a case, the length of a period, δ, should minimize $\delta + k\delta^{1/2} \approx \delta + (\tau_{i+1,n} - \tau_{i,n})\delta^{-1/2}$. An "optimal" choice would set $\delta \approx (\tau_{i+1,n} - \tau_{i,n})^{2/3}$, with $k \approx n^{1/3}$; this, in turn, would yield absolute losses of $O(n^{2/3})$. The optimized sequence of review periods with varying lengths, apart from the associated intuitive structure, results in a revenue loss of essentially $O(n^{1/2})$. This advantage is illustrated numerically in §5. One also notes that Theorem 2 applies to the case where there are no milestones. In such settings and when the demand model is fully known, the fixed price heuristic proposed in Gallego and van Ryzin (1994) achieves a performance that is within an order of $n^{1/2}$ of the optimal, which is close to that of our policy that does not make use of demand model information.

5. Numerical Illustrations

5.1. An Illustrative Example

We first consider an example with problem parameters as outlined in §3.1. The seller has $C = 200$ units to sell over a time horizon of $T = 1{,}000$ days. There are $m = 4$ milestones. After the first $\tau_1 = 300$ days, the seller needs to have generated $\zeta_1 = 1{,}000$ in units of revenues. In addition, the seller needs to have sold $\xi_1 = 100$ units by τ_1, $\xi_2 = 160$ units by $\tau_2 = 500$ days, and $\xi_3 = 180$ units by $\tau_3 = 700$ days. We consider an example where given a posted price p, the true underlying demand function (mean number of sales per day) is given by $\Lambda\bar{F}(p) = (2 - 0.2 * p)^+$, but the latter is unknown to the decision maker.

This problem can be seen as one with $m = 4$ milestones, and if we take $C = 200$ to be the proxy for the "scale" of the project, the proposed feedback pricing policy, denoted by π_f, is characterized by a sequence of review windows as specified in (30)–(31),[7] where we have taken $C_\nu = 1$. In what follows, we take the penalty parameters to be $\beta_{1,i} = \beta_{2,1} = 10$ for $i = 1, \ldots, m$. The feedback parameter α will vary throughout the experiments.

In Figure 2, we depict a sample path of the price, cumulative sales and revenues associated with π_f in conjunction with the (deterministic) paths associated with the optimal solution of the auxiliary problem (6): $\tilde{p}(\cdot)$, $\tilde{S}(\cdot)$, and $\tilde{R}(\cdot)$. As expected, we observe that

[7] This roughly corresponds to update prices initially every two weeks before decreasing the frequency to an update every month.

Table 2 Performance of the Proposed Policy

Feedback parameter α	Scale of the project: $n=100$			$n=200$			$n=500$		
	0.5	0.75	1	0.5	0.75	1	0.5	0.75	1
Initial price									
$p_0=6$	0.24	0.29	0.44	0.16	0.14	0.18	0.10	0.07	0.07
$p_0=7$	0.18	0.23	0.29	0.12	0.12	0.16	0.07	0.06	0.07
$p_0=8$	0.18	0.22	0.40	0.11	0.12	0.18	0.07	0.06	0.07
$p_0=9$	0.20	0.24	0.43	0.11	0.12	0.17	0.07	0.06	0.06

Note. Performance $1-J^{\pi}/J^{d}$ of the policy π_f for various values of the feedback parameter α and the initial price p_0 is shown.

the cumulative sales and revenues associated with π_f closely track $\tilde{S}(\cdot)$ and $\tilde{R}(\cdot)$. In addition, we see that although the price is initially off after each milestone, it roughly stabilizes after a few iterations to a price that "almost" tracks the sales and revenue targets. This illustrates the corrective feature of the feedback policy whenever the price is away from the appropriate target.

5.2. The Impact of the Scale of the System on Relative Performance

Table 2 studies the effect of problem size as captured by n on performance. The problem parameters are the ones given in Table 3, and $n=200$ corresponds to the instance analyzed above. The performance is measured relative to the value of the auxiliary deterministic problem J^d, i.e., we report an estimate of $1-J^{\pi}/J^{d}$, where J^{π} was evaluated through simulation by using 10^4 replications. The standard error associated with the estimate of $1-J^{\pi}/J^{d}$ was always below 10^{-3}. Note that J^d is an upper bound on the optimal performance J^*, and hence the performance estimates presented below are conservative. Although in theory it would be ideal to compare the performance to J^*, computing the latter is likely to be an intractable problem in general, as alluded to in Gallego and van Ryzin (1994) for the case where there are no milestones. As expected from Theorem 2, the relative performance improves with the scale of the system, and with scale of $n=500$, one is able to be within approximately 7% of the optimal performance for most of the initial prices and feedback parameters tested.

The tuning parameters of the policy play an important role in its performance. It is difficult, if not intractable, to provide an "optimal" selection of such parameters. Some guidance can be obtained by simulating the policy performance over a range of possible demand models and choosing parameters that lead to robust performance. Intuitively, higher values of α lead to faster tracking and better overall performance, but will also imply more jittery price recommendations. In practice, their selection may depend on the application setting, e.g., weekly to monthly updates and quantized price increments of \$5,000 or \$10,000 in the real estate setting, which will moderate price jitters.

In addition, the initial price is a key driver to performance, because it affects the time it takes to recover form the initial transient. The best performance is reached for an initial price of 8, which is not surprising because this is actually the closest price (among those tested) to the initial price associated with the optimal solution of the auxiliary deterministic problem that is $\tilde{p}(0)\approx 7.88$; this price allows one to reach exactly the revenue constraint at $\tau_1=300$ days, which is the most binding initial constraint in this problem.

Table 3 Problem Instance Parameters

Capacity	$C_n=n$			
Milestone times	$\tau_1^{(n)}=1.5n$	$\tau_2^{(n)}=2.5n$	$\tau_3^{(n)}=3.5n$	$\tau_4^{(n)}=5n$
Sales milestones	$\xi_1^{(n)}=0.5n$	$\xi_2^{(n)}=0.8n$	$\xi_3^{(n)}=0.9n$	$\xi_4^{(n)}=0$
Revenue milestones	$\zeta_1^{(n)}=5n$	$\zeta_2^{(n)}=0$	$\zeta_3^{(n)}=0$	$\zeta_4^{(n)}=0$

5.3. Benchmark Performance in the Absence of Milestone Constraints

In the remainder of this section, we report on experiments without milestones, which allows us to compare the performance of the proposed policy J^{π} against that of the fixed price heuristic presented in Gallego and van Ryzin (1994), J^{FP}, which uses the price p such that $\Lambda\bar{F}(p)=C/T$. It is worth noting that the latter is based on the market size and WtP distribution, whereas the former does not use this information.

We consider instances with no milestones and various values of initial inventory C and time horizon $T=5C$, where the capacity is constrained in the sense that $C/T<\lambda^*$. For each, we take the "scale" of the system to be C. We report in Table 4 the performance of the feedback pricing policy relative to that of the fixed price heuristic for two demand models and different values of α and p_0. Note that in the linear demand case, the fixed price heuristic (that corresponds to a demand rate of C/T) is $p=9$, whereas in the exponential case it is given by $p=7.675$. The feedback pricing policy that does not rely on demand model information performs almost as well as the fixed price heuristic when the initial price is close to the price used by the fixed price heuristic. When the initial price is far from the fixed price heuristic price, which is the case, for example, when demand is linear and $p_0=7$, then the time needed to "catch up" with the appropriate target introduces a loss compared to the fixed price heuristic.

5.4. The Value of a Nonuniform Grid

As discussed in §4, a nonuniform discrete-review grid allows the seller to initially update prices faster to

Table 4 Performance Comparison with the Fixed Price Heuristic

	Initial capacity C					
	100		200		500	
Feedback parameter α	0.5	0.75	0.5	0.75	0.5	0.75
Demand model: $\Lambda\bar{F}(p) = (2-0.2*p)^+$						
$p_0 = 7$	0.118	0.093	0.111	0.080	0.103	0.066
$p_0 = 8$	0.028	0.009	0.031	0.014	0.031	0.014
$p_0 = 9$	−0.004	−0.006	−0.005	−0.006	−0.001	−0.003
Demand model: $\Lambda\bar{F}(p) = 2*\exp\{-0.3p\}$						
$p_0 = 7$	0.031	0.030	0.038	0.035	0.045	0.038
$p_0 = 8$	0.046	0.033	0.039	0.025	0.038	0.028
$p_0 = 9$	0.160	0.146	0.157	0.133	0.151	0.128

Note. This table reports $1 - J^{\pi}/J^{FP}$, the performance of the proposed policy J^{π} relative to J^{FP}, the performance of the fixed price heuristic that relies on demand model information.

Table 5 Performance Comparison with Fixed Length Interval Updates

	Initial capacity C					
	100		200		500	
$(j^*, \bar{k})$	(10, 1)	(5, 2)	(12, 1)	(6, 2)	(14, 1)	(7, 2)
$p_0 = 7$	0.182	0.158	0.180	0.139	0.171	0.121
$p_0 = 8$	0.083	0.070	0.077	0.061	0.071	0.051
$p_0 = 9$	0.048	0.036	0.037	0.025	0.024	0.018

Notes. This table reports $1 - J^{\pi}/J^{d}$, the performance of the proposed policy π_f for various values of j^* and $\bar{k}$. The case $\bar{k} = 1$ corresponds to fixed length intervals. Here, $\alpha = 0.5$, and $\Lambda\bar{F}(p) = (2-0.2*p)^+$.

"zoom in" on the region of a near-optimal price, before decreasing the frequency of price changes to adjust in a more refined fashion toward a near-optimal price. To illustrate this point, we focus again on the simplest instance of the problem when there are no milestones. We analyze the performance of the proposed policy as a function of j^* and $\bar{k}$ for a fixed number of price changes $\bar{k}j^*$. Given that C_ν is taken equal to 1, when $\bar{k} = 1$, the decision maker updates prices at equal length intervals, and when $\bar{k} = 2$, the frequency of price changes is initially higher than when one approaches the end of the horizon. We take as a reference case $\bar{k} = 2$, and in this case, we let $j^* = \lceil \log C \rceil$. Table 5 reports the performance for the cases with $\bar{k} = 1$ and $\bar{k} = 2$. As expected from our results, the use of a nonuniform grid leads to 2%–5% performance improvements for the problem parameters of the examples. The magnitude of these numbers can be significant, but perhaps more importantly the nature of the nonuniform grid seems intuitively appealing for practical applications, allowing the seller to start with more frequent price adjustments.

Acknowledgments

The authors are grateful to Camilo Galvis (RE Optima) and Soulaymane Kachani (Columbia University) for many discussions on residential real estate, which motivated this work. They are also grateful to Gérard Cachon, an associate editor, and three referees for their feedback that helped improve this paper.

Appendix A. Proofs

Proof of Proposition 1. Recall the definition of $\tilde{\lambda}_i$ in (10). If $\sum_{i=0}^{m-1} \tilde{\lambda}_i(\tau_{i+1} - \tau_i) \le C$, it is easy to see that the following control $\tilde{\lambda}(t) = \tilde{\lambda}_i$ for $t \in [\tau_{i-1}, \tau_i)$ is feasible for (6). Conversely, if (6) is feasible, then, by Theorem 1, the above control is feasible, and therefore the condition in the statement of the proposition must be satisfied. □

Proof of Theorem 1. Note that when optimizing (6), by concavity of $r(\cdot)$, one can restrict attention to policies that are constant over each interval $[\tau_i, \tau_{i+1})$, $i = 0, \ldots, m-1$. In addition, it is easy to see that under Assumption 2 one can restrict attention to policies with rates in $[0, \lambda^*]$. Let $\{\gamma(s)\colon 0 \le s \le T\}$ denote a feasible solution to (6), and let γ_i denote the value (rate) of the solution over $[\tau_i, \tau_{i+1})$, $i = 0, \ldots, m-1$. Note that $\tilde{\lambda}(s)$ is also constant over each such interval, taking value $\tilde{\lambda}_i$ over $[\tau_i, \tau_{i+1}]$, $i = 0, \ldots, m-1$. The proof establishes that any policy for which $\gamma_i \ne \tilde{\lambda}_i$ for some $i \in \{0, 1, \ldots, m-1\}$ cannot be optimal.

Suppose that for some $i \in \{0, 1, \ldots, m-1\}$, $\gamma_i \ne \tilde{\lambda}_i$, and let j denote the smallest such index. For any policy $\mu(\cdot)$, we let $S_\mu(t)$ and $R_\mu(t)$ denote the cumulative sales and revenues up to time t under this policy. Note that $S_\gamma(\tau_j) = S_{\tilde{\lambda}}(\tau_j)$ and $R_\gamma(\tau_j) = R_{\tilde{\lambda}}(\tau_j)$. Let

$$\begin{aligned} C' &= C - S_\gamma(\tau_j), \quad \xi_i' = \xi_i - S_\gamma(\tau_j), \\ \zeta_i' &= \zeta_i - R_\gamma(\tau_j), \quad i = j+1, \ldots, m, \\ k &= \underset{i:\, j+1 \le i \le m}{\arg\max} \left\{ \frac{\xi_i'}{\tau_i - \tau_j}, \varphi \frac{\zeta_i'}{\tau_i - \tau_j} \right\}, \end{aligned} \tag{A1}$$

where C' is the remaining capacity at τ_j under control γ; ξ_i'; and ζ_i' are the additional number of units to be sold and remaining revenue to be accrued in $(\tau_j, \tau_i]$, respectively, for the sales and revenue targets to be met at time $\tau_i > \tau_j$; and, k is the most stringent of the downstream sales and revenue constraints at time τ_j.

Case 1. $\gamma_j < \max\{\xi_k'/(\tau_k - \tau_j), \varphi(\zeta_k'/(\tau_k - \tau_j)), C'/(T - \tau_j)\}$. In this case, the control γ_j is less than the static sales rate that would eventually satisfy all of the constraints until τ_k, where k is the most stringent milestone starting at time τ_j. Note that because γ is feasible, necessarily $k \geq j+2$ (and hence $j \leq m-2$), and $\sum_{i=j}^{k-1}(\tau_{i+1}-\tau_i)r(\gamma_i) \geq \zeta'_k$. Define $\alpha_i^{(k)} = (\tau_i - \tau_{i-1})/(\tau_k - \tau_j)$ for $i = j, \ldots, k-1$ and $\bar{\lambda} = \sum_{i=j}^{k-1} \alpha_i^{(k)} \gamma_i$. Consider the policy θ that applies $\gamma(\cdot)$ on $[0, \tau_j)$, $\bar{\lambda}$ on $[\tau_j, \tau_k)$, and $\gamma(\cdot)$ on $[\tau_k, T]$. We check that θ is feasible. All milestone constraints $i = 1, \ldots, j$ are satisfied because γ is feasible. In addition, we have

$$\begin{aligned} R_\theta(\tau_k) - R_\theta(\tau_j) &= (\tau_k - \tau_j)r(\bar{\lambda}) \geq \sum_{i=j}^{k-1}(\tau_{i+1}-\tau_i)r(\gamma_i) \\ &= R_\gamma(\tau_k) - R_\gamma(\tau_j) \geq \zeta'_k, \end{aligned} \tag{A2}$$

where we have used the strict concavity of $r(\cdot)$ and the fact that $\gamma_j \neq \bar{\lambda}$. Similarly,

$$\begin{aligned} S_\theta(\tau_k) - S_\theta(\tau_j) &= (\tau_k - \tau_j)\bar{\lambda} = \sum_{i=j}^{k-1}(\tau_{i+1}-\tau_i)\gamma_i \\ &= S_\gamma(\tau_k) - S_\gamma(\tau_j) \geq \xi'_k. \end{aligned} \tag{A3}$$

Now, note that (A2) and (A3) imply that $\bar{\lambda} \geq \max\{\xi'_k/(\tau_k - \tau_j), \varphi(\zeta'_k/(\tau_k - \tau_j))\}$. We deduce that for $i = j+1, \ldots, k$,

$$\begin{aligned} R_\theta(\tau_i) - R_\theta(\tau_j) &= (\tau_i - \tau_j)r(\bar{\lambda}) \geq (\tau_i - \tau_j) \\ &\quad \cdot r\left(\max\left\{\frac{\xi'_k}{\tau_k - \tau_j}, \varphi\frac{\zeta'_k}{\tau_k - \tau_j}\right\}\right) \\ &\overset{(a)}{\geq} (\tau_i - \tau_j) r\left(\varphi\left(\frac{\xi'_i}{\tau_i - \tau_j}\right)\right) = \xi'_i, \end{aligned}$$

where we have used the fact that $r(\cdot)$ is increasing on $[0, \lambda^*]$ (note that $\bar{\lambda} \leq \lambda^*$ because we restrict attention to policies γ that use rates in $[0, \lambda^*]$) and the definition of k (see (A1)). It is straightforward to verify that the other milestone constraints as well as the final inventory constraint are satisfied under θ.

Suppose first that $\gamma_j < \max\{\xi'_k/(\tau_k - \tau_j), \varphi(\zeta'_k/(\tau_k - \tau_j))\}$, in which case, necessarily, $\gamma_j \neq \bar{\lambda}$, and the strict concavity of $r(\cdot)$ implies that the total revenues generated by the policy θ satisfy

$$R_\theta(T) - R_\gamma(T) = (\tau_k - \tau_j)r(\bar{\lambda}) - \sum_{i=j}^{k-1}(\tau_{i+1}-\tau_i)\gamma_i > 0,$$

and hence γ cannot be optimal.

Suppose now that $\gamma_j \geq \max\{\xi'_k/(\tau_k - \tau_j), \varphi(\zeta'_k/(\tau_k - \tau_j))\}$, in which case, necessarily, $\max\{\xi'_k/(\tau_k - \tau_j), \varphi(\zeta'_k/(\tau_k - \tau_j))\} < C'/(T - \tau_j)$. Noting that $\gamma_j < C'/(T-\tau_j) \leq \lambda^*$, we observe that the policy that applies $\gamma(\cdot)$ on $[0, \tau_j)$, and $C'/(T - \tau_j)$ on $[\tau_j, T]$ is feasible and achieves revenues strictly greater than γ (because applying the static rate $\min\{\lambda^*; C'/(T - \tau_j)\} = C'/(T - \tau_j)$ $[\tau_j, T]$ maximizes $R(T) - R(\tau_j)$ in the absence of milestone constraints). We deduce that in this case, again, γ cannot be optimal.

Case 2. $\gamma_j > \max\{\xi'_k/(\tau_k - \tau_j), \varphi(\zeta'_k/(\tau_k - \tau_j)), C'/(T - \tau_j)\}$. In this case, γ_j exceeds the sales rate that is needed to satisfy the most stringent downstream milestone that occurs at time τ_k. Consider the policy $\mu(\cdot)$ that applies $\gamma(\cdot)$ on $[0, \tau_j]$, $\max\{\xi'_k/(\tau_k - \tau_j), \varphi(\zeta'_k/(\tau_k - \tau_j)), C'/(T - \tau_j)\}$ on $[\tau_j, T')$, where $T' = \min\{T, C'/\max\{\xi'_k/(\tau_k - \tau_j), \varphi(\zeta'_k/(\tau_k - \tau_j)), C'/(T - \tau_j)\}\}$, and 0 on $[T', T]$.

Because $S_\gamma(\tau_j) = S_\mu(\tau_j)$ and

$$\gamma_j > \max\left\{\frac{\xi'_k}{\tau_k - \tau_j}, \varphi\left(\frac{\zeta'_k}{\tau_k - \tau_j}\right)\right\},$$

we necessarily have $S_\gamma(\tau_{j+1}) > S_\mu(\tau_{j+1})$. Noting that $S_\mu(T') = C \geq S_\gamma(T')$, there exists some $t \in [\tau_{j+1}, T']$ such that $S_\gamma(t) = S_\mu(t)$. Let t' denote the smallest such value, and let

$$\bar{\lambda} = \frac{1}{t' - \tau_j}\int_{\tau_j}^{t'} \gamma(s)\,ds.$$

Note that $\bar{\lambda} = \max\{\xi'_k/(\tau_k - \tau_j), \varphi(\zeta'_k/(\tau_k - \tau_j)), C'/(T - \tau_j)\} < \gamma_j$. Consider the policy ν that applies $\gamma(\cdot)$ on $[0, \tau_j]$, $\bar{\lambda}$ on $[\tau_j, t']$, and $\gamma(\cdot)$ on $(t', T]$. It is clearly feasible by concavity of $r(\cdot)$ and the fact that $\bar{\lambda} = \max\{\xi'_k/(\tau_k - \tau_j), \varphi(\zeta'_k/(\tau_k - \tau_j)), C'/(T - \tau_j)\} \geq \varphi(\xi'_i/(\tau_i - \tau_j))$ for $i = j+1, \ldots, m$. The revenues it achieves satisfy

$$\begin{aligned} &\int_0^T r(\nu(s))\,ds \\ &\quad = \int_0^{\tau_j} r(\gamma(s))\,ds + (t' - \tau_j)r(\bar{\lambda}) + \int_{t'}^T r(\gamma(s))\,ds \\ &\quad \overset{(a)}{>} \int_0^{\tau_j} r(\gamma(s))\,ds + \int_{\tau_j}^{t'} r(\gamma(s))\,ds + \int_{t'}^T r(\gamma(s))\,ds \\ &\quad = \int_0^T r(\gamma(s))\,ds, \end{aligned}$$

where (a) follows from the strict concavity of $r(\cdot)$ in conjunction with the fact that $\bar{\lambda} \neq \gamma_j$. We deduce that γ cannot be optimal.

We conclude from both cases that $\tilde{\lambda}(\cdot)$ is necessarily optimal. □

Proof of Proposition 2. We will establish that when $\min_{i=1,\ldots,m}\{\min\{\beta_{1,i}, \beta_{2,i}\}\} > \bar{p}\max\{1, \underline{\kappa}^{-1}\}$, and problem (6) is feasible, then the optimal solutions of (6) and (17) coincide.

Suppose that (6) is feasible and consider an optimal solution to (17), $\{\gamma(s): 0 \leq s \leq T\}$, which, without loss of generality, is assumed to be constant over each interval $[\tau_j, \tau_{j+1})$, $j = 0, \ldots, m-1$. We let γ_j denote the value (rate) of the optimal solution over $[\tau_j, \tau_{j+1})$, $j = 0, \ldots, m-1$. Suppose that γ is not feasible for (6), which implies that $\gamma_i \neq \tilde{\lambda}_i$ for some $j \in \{0, \ldots, m-1\}$. Let j be the smallest such index.

Case 1. $\gamma_j < \tilde{\lambda}_j$. Consider the following problem:

$$\begin{aligned} \max\bigg\{\int_0^T r(\lambda(t))\,dt &- \sum_{i=1}^m \beta_{1,i}(\xi_i - S(\tau_i))^+ \\ &- \sum_{i=1}^m \beta_{2,i}(\zeta_i - R(\tau_i))^+ - \bar{p}(S(T) - C)^+\bigg\}. \end{aligned} \tag{A4}$$

This is a problem where the seller has no resource constraints but has to pay a penalty of $\bar{p}$ for every unit sold after the first C units. Given that γ is optimal for (17), it is also optimal for (A4) because no units can ever be sold for more than $\bar{p}$. (A4) is a concave problem and the optimality of γ for (A4) implies that

$$\begin{aligned} r'(\gamma_j) &+ \sum_{i=j}^m [\beta_{1,i}\mathbf{1}\{\xi_i - S(\tau_i) > 0\} + \beta_{2,i}r'(\gamma_j)\mathbf{1}\{\zeta_i - R(\tau_i) > 0\}] \\ &- \bar{p}\mathbf{1}\{S(T) - C \geq 0\} \leq 0. \end{aligned} \tag{A5}$$

Note that since γ is not feasible for (6) and since $\gamma_i = \tilde{\lambda}_i$ for $i \le j-1$, there exists some $k \ge j+1$ such that $\xi_k - S(\tau_k) > 0$ or $\zeta_k - R(\tau_k) > 0$. Suppose first that $\xi_k - S(\tau_k) > 0$. In this case, noting that $r'(\gamma_j) \ge r'(\tilde{\lambda}_j) \ge r'(\lambda^* - \sigma) \ge \underline{\kappa} \ge 0$, (A5) would imply that

$$\beta_{1,i} - \bar{p}\mathbf{1}\{S(T) - C \ge 0\} \le 0,$$

which is a contradiction since $\beta_{1,i} > \bar{p}$. Similarly, if $\zeta_k - R(\tau_k) > 0$, then (A5) would imply that

$$\beta_{2,i}\underline{\kappa} - \bar{p}\mathbf{1}\{S(T) - C \ge 0\} \le \beta_{2,i} r'(\gamma_j) - \bar{p}\mathbf{1}\{S(T) - C \ge 0\} \le 0,$$

which is again a contradiction since $\beta_{2,i}\underline{\kappa} > \bar{p}$.

Case 2. $\gamma_j > \tilde{\lambda}_j$. In this case, a similar analysis as that conducted in Case 2 of the proof of Theorem 1 would yield that γ cannot be optimal, a contradiction.

We conclude that γ is necessarily feasible for (6), and the proof is complete. □

Appendix B. Proof of Theorem 2

Overview

The guiding idea of the proof lies in comparing the performance of the feedback pricing policy to the optimal value of the deterministic problem (6). The comparison to the deterministic problem is driven by the results of §3 and, in particular, by the fact that the optimal value of the deterministic problem, J_n^d, serves as an upper bound for J_n^*, the optimal performance in the stochastic system when full knowledge of the demand model is available (see §3.2). Hence, this comparison will ultimately provide an upper bound on the performance gap between an optimal policy in the stochastic system with knowledge of the demand model and the proposed policy.

To quantify the above performance gap, we first focus on a subset of sample paths $\mathscr{A}$ that has high probability and establish in Lemma 1 that for each intermilestone segment $[\tau_i, \tau_{i+1}]$, for all sample paths in $\mathscr{A}$, the following "physical" properties hold:

(i) After each milestone, the feedback force $|\min\{\hat{\Delta}_1, \hat{\Delta}_2\}|$ defined in (24) decreases in a "fast" fashion over time. This is spelled out in (C18) in the online companion (available at http://www.columbia.edu/~ob2105/OC_milestones.pdf).

(ii) The price stays in the region where demand is elastic, as detailed in (B11).

(iii) The cumulative sales follow closely the optimal path of cumulative sales associated with the optimal solution of the deterministic problem (6). (B12) is the formal statement of this property.

(iv) The cumulative revenues follow closely the optimal path of cumulative revenues associated with the the optimal solution of the deterministic problem (6) (see (B13)).

Roughly speaking, based on the analysis of §3 (see in particular (14)), $|\min\{\hat{\Delta}_1, \hat{\Delta}_2\}|$ can be seen as a proxy for performance losses per unit of time. Property (i) above says that one quickly reaches a point where $\min\{\hat{\Delta}_1, \hat{\Delta}_2\} \approx 0$, and hence performance losses are "small." Then the selection of the nonuniform grid of review points enables one to refine the estimate of the price that would allow to reach the next most binding target while mitigating performance losses (this is the result presented in Claim 2 in the online companion). Property (ii) is important to ensure, as the proposed feedback policy will only allow to adjust prices in the proper direction if one operates in the elastic region of the demand function. Finally, properties (iii) and (iv) are what one is really after because our ultimate aim is to compare the revenues accumulated and the penalties incurred (that depend on the sales and revenue paths) by the proposed policy and those by an optimal policy in the deterministic system.

The above analysis is conducted on a subset of sample paths $\mathscr{A}$. We conclude the proof by lower bounding the probability of the set $\mathscr{A}$ in Lemma 2 and providing an upper bound on the performance of the proposed policy, $\mathbb{E}[R_n(T_n) - Z_n]$, in terms of J_n^*.

Preliminaries

We denote by $\{t_j^i, i = 0, \ldots, m-1, j = 0, \ldots, \bar{k}j^* + 1\}$ the discrete review points in the system with scale n and for any $(i, j) \in \{0, \ldots, m-1\} \times \{0, \ldots, \bar{k}j^*\}$, $\delta_j^i = t_{j+1}^i - t_j^i$. Here we do not make the dependence on n explicit to avoid cluttering the notation. Let also $\ell_m^s(t_j^i)$ denote $(nC - S_n(t_j^i))/(\tau_{m,n} - t_j^i)$. As long as one does not run out of inventory, the evolution of $(p_n(\cdot), S_n(\cdot), R_n(\cdot))$ under the proposed feedback pricing policy can be written in a recursive fashion as follows: for $i = 0, \ldots, m-1$, for $j = 0, \ldots, \bar{k}j^*$,

$$p_n(t_{j+1}^i) = p_n(t_j^i) + \alpha_n\big[\phi\big(t_j^i, p(t_j^i), S_n(t_j^i), R_n(t_j^i)\big) + \varepsilon_p(t_j^i)\big], \quad \text{(B1)}$$

$$S_n(t_{j+1}^i) = S_n(t_j^i) + \Lambda\bar{F}(p(t_j^i))(t_{j+1}^i - t_j^i) + \varepsilon_S(t_j^i), \quad \text{(B2)}$$

$$R_n(t_{j+1}^i) = R_n(t_j^i) + r\big(\Lambda\bar{F}(p(t_j^i))\big)(t_{j+1}^i - t_j^i) + \varepsilon_R(t_j^i), \quad \text{(B3)}$$

where $(p(0), S(0), R(0)) = (p_0, 0, 0)$, and

$$\begin{aligned}
\phi(t, p, S, R) &= \min\big\{\Delta_1(t, p, S), \Delta_2(t, p, S)\big\},\\
\varepsilon_p(t_j^i) &= \big[\hat{\phi}(t_j^i) - \phi\big(t_j^i, p(t_j^i), S_n(t_j^i), R_n(t_j^i)\big)\big],\\
\varepsilon_S(t_j^i) &= S_n(t_{j+1}^i) - S_n(t_j^i) - \Lambda\bar{F}\big(p_n(t_j^i)\big)(t_{j+1}^i - t_j^i),\\
\varepsilon_R(t_j^i) &= R_n(t_{j+1}^i) - R_n(t_j^i) - r\big(\Lambda\bar{F}(p_n(t_j^i))\big)(t_{j+1}^i - t_j^i)\\
&= p_n(t_j^i)\varepsilon_S(t_j^i).
\end{aligned}$$

In the rest of the proof, we consider a system where the seller can sell more than C units, but for all the units sold after the first C units, there is a penalty to be paid of $\bar{p}$ per unit. It is clear that the revenues generated by the proposed policy in such a system are upper bounded by those generated in the original system, and hence any lower bound on revenues in this new system is also a lower bound in the original system.

For $(i, j) \in \{0, \ldots, m-1\} \times \{0, \ldots, \bar{k}j^*\}$, we define

$$\begin{aligned}
\gamma_j^i &= \mathscr{C}_\gamma(\delta_j^i)^{1/2}(\log n)^{1/2},\\
\mathscr{A} &= \big\{\omega\colon |\varepsilon_p(t_j^i)| \le \gamma_j^i/\delta_j^i,\ |\varepsilon_S(t_j^i)| \le \gamma_j^i,\ |\varepsilon_R(t_j^i)|\\
&\qquad \le \gamma_j^i\colon 0 \le i \le m-1,\ 0 \le j \le \bar{k}j^*\big\},
\end{aligned}$$

where $\mathscr{C}_\gamma$ is a positive constant that is suitably large (a condition on the value of C_γ is provided in the proof of

Lemma 2). We first conduct the analysis over sample paths in $\mathcal{A}$ and then evaluate the probability of the event $\mathcal{A}$.

Relation to the Deterministic Relaxation Paths
Let $(p_n(\cdot), S_n(\cdot), R_n(\cdot))$ be defined according to (B1)–(B3) with initial conditions $(p_0, 0, 0)$ and let $(\tilde{p}_n(\cdot), \tilde{S}_n(\cdot), \tilde{R}_n(\cdot))$ denote the optimal price, cumulative sales, and revenues paths respectively, for problem (6) with scale n. We will establish by induction on i that for all $\omega \in \mathcal{A}$, for $i = 0, \ldots, m$,

$$p(\tau_{i,n}^-) \geq \min\{p(0), \tilde{p}(0)\} - i(\bar{k}j^* + 1)\mathcal{D}_3(\log n)^{1/2} n^{-1/4}, \quad \text{(B4)}$$

$$\left|S_n(\tau_{i,n}) - \tilde{S}_n(\tau_{i,n})\right| \leq (\bar{k}j^* + 1)\rho_n\left(\sum_{l=1}^{i}\exp\{\mathcal{D}_2 l\}\right), \quad \text{(B5)}$$

$$\left|R_n(\tau_{i,n}) - \tilde{R}_n(\tau_{i,n})\right| \leq (\bar{k}j^* + 1)\rho_n\left(\sum_{l=1}^{i}\exp\{\mathcal{D}_2 l\}\right), \quad \text{(B6)}$$

where

$$\mathcal{D}_1 = \max\{1, \underline{\kappa}^{-1}\}, \quad \text{(B7)}$$

$$\mathcal{D}_2 = \mathcal{D}_1 \max\{\bar{\kappa}, 1\}\exp\{1/2\}, \quad \text{(B8)}$$

$$\mathcal{D}_3 = \alpha + \alpha \max\big\{n \geq 2\colon\ (\bar{k}j^* + 1)\rho_n \cdot \exp\{\mathcal{D}_2(m+1)\}(\delta^i_{\bar{k}j^*})^{-1}(\eta^i_0)^{-1}\big\}, \quad \text{(B9)}$$

$$\rho_n = C_\rho n^{1/2}(\log n)^{1/2}, \quad \text{(B10)}$$

and C_ρ is a suitably large positive constant (a condition on the value of C_ρ is provided in the online companion (C1)). Note that the maximum in the definition of $\mathcal{D}_3$ is well defined because $(\bar{k}j^* + 1)\rho_n(\delta^i_{\bar{k}j^*})^{-1}(\eta^i_0)^{-1}$ converges to 0 as $n \to \infty$.

The result is clearly true for $i = 0$. Suppose it is true for some $i \in \{0, \ldots, m-1\}$. The next result shows that the sales and revenue paths stay appropriately close to the ones associated with the deterministic relaxation optimal policy over $[\tau_i, \tau_{i+1}]$. The proof is provided in the online companion.

Lemma 1. *Let Assumptions* 1–3 *hold. Let* $t^0_{-1} = 0$, *and for all* $i = 1, \ldots, m$, *let* $t^i_{-1} = t^{i-1}_{\bar{k}j^*}$. *For all* $\omega \in \mathcal{A}$, *for all* $j = 0, \ldots, \bar{k}j^* + 1$,

$$p(t^i_{j-1}) \geq \min\{p(0), \tilde{p}(0)\} - \big(i(\bar{k}j^* + 1) + j\big)\mathcal{D}_3(\log n)^{1/2} n^{-1/4}, \quad \text{(B11)}$$

$$\left|S_n(t^i_j) - \tilde{S}_n(t^i_j)\right| \leq \rho_n\left[(\bar{k}j^* + 1)\sum_{l=1}^{i}\exp\left\{\mathcal{D}_2\left(l + \frac{t^i_j - \tau_{i,n}}{\tau_{i+1,n} - \tau_{i,n}}\right)\right\} + \sum_{l=1}^{j}\exp\left\{\mathcal{D}_2\frac{(t^i_j - t^i_l)}{\tau_{i+1,n} - \tau_{i,n}}\right\}\right], \quad \text{(B12)}$$

$$\left|R_n(t^i_j) - \tilde{R}_n(t^i_j)\right| \leq \rho_n\left[(\bar{k}j^* + 1)\sum_{l=1}^{i}\exp\left\{\mathcal{D}_2\left(l + \frac{t^i_j - \tau_{i,n}}{\tau_{i+1,n} - \tau_{i,n}}\right)\right\} + \sum_{l=1}^{j}\exp\left\{\mathcal{D}_2\frac{(t^i_j - t^i_l)}{\tau_{i+1,n} - \tau_{i,n}}\right\}\right]. \quad \text{(B13)}$$

Taking $j = \bar{k}j^* + 1$ and upper bounding all the terms $t^i_j - t^i_l$ by $\tau_{i+1,n} - \tau_{i,n}$ in the two inequalities above yields that (B4)–(B6) are satisfied for $i+1$. In other words, for all $\omega \in \mathcal{A}$, starting at the same level, the two paths are guaranteed to be "close" up to a factor of $(\bar{k}j^* + 1)\rho_n$, an error stemming from the random nature of arrivals and the absence of demand model information. In addition, let n_0 be such that

$$m\mathcal{D}_3(\bar{k}j^* + 1)(\log n)^{1/2} n^{-1/4} \leq \bar{F}^{-1}\left(\frac{\lambda^* - \sigma}{\Lambda} - \bar{F}^{-1}\frac{\lambda^* - \sigma/2}{\Lambda}\right).$$

Then for $n \geq n_0$, (B11) ensures that the price is always in the elastic region for all sample paths in $\mathcal{A}$ and in particular is such that $\Lambda\bar{F}(p) \leq \lambda^* - \sigma/2$.

Final Performance Analysis
First note that, from (B6), one has

$$\begin{aligned}&\mathbb{E}[\tilde{R}_n(T_n) - R_n(T_n)]\\ &\quad= \mathbb{E}\big[\tilde{R}_n(T_n) - R_n(T_n) \mid \mathcal{A}\big]\mathbb{P}\{\mathcal{A}\} + \mathbb{E}\big[\tilde{R}_n(T_n) - R_n(T_n) \mid \mathcal{A}^c\big]\mathbb{P}\{\mathcal{A}^c\}\\ &\quad\leq (\bar{k}j^* + 1)\rho_n \exp\{\mathcal{D}_2(m+1)\} + n\Lambda\bar{p}T\,\mathbb{P}\{\mathcal{A}^c\}.\end{aligned}$$

Turning to the penalties, recall that in the system analyzed, the seller may sell more than C units but incurs a penalty of $\bar{p}$ for any unit sold after the first C units. Hence the total penalties incurred are given by

$$\mathbb{E}\big[Z_n + \bar{p}(S_n(\tau_{i,n}) - C)^+\big] = \mathbb{E}\big[Z_n + \bar{p}(S_n(\tau_{m,n}) - C)^+ \mid \mathcal{A}\big]\mathbb{P}\{\mathcal{A}\} + \mathbb{E}\big[Z_n + \bar{p}(S_n(\tau_{m,n}) - C)^+ \mid \mathcal{A}^c\big]\mathbb{P}\{\mathcal{A}^c\}.$$

We next analyze each term on the right-hand side. The first conditional expectation can be upper bounded by noting that the the optimal solution of the deterministic problem incurs zero penalties, and hence using Lemma 1, the penalties incurred by the proposed policy $Z_n = \sum_{i=1}^m \beta_{1,i}(\xi_{i,n} - S_n(\tau_{i,n}))^+ + \beta_{2,i}(\zeta_{i,n} - R_n(\tau_{i,n}))^+$ will be upper bounded as follows for all $\omega \in \mathcal{A}$:

$$Z_n \leq \left[\sum_{i=1}^m \beta_{1,i}\exp\{\mathcal{D}_2(i+1)\} + \beta_{2,i}\exp\{\mathcal{D}_2(i+1)\}\right](\bar{k}j^* + 1)\rho_n,$$

where the inequality follows from (B5) and (B6). Noting that $\mathcal{D}_2 \geq 1$, $\exp\{\mathcal{D}_2\} - 1 \geq 1$ and

$$\sum_{l=1}^{i}\exp\{\mathcal{D}_2 l\} = \sum_{i=1}^{i}\big(\exp\{\mathcal{D}_2\}\big)^l = \exp\{\mathcal{D}_2\}\frac{\exp\{\mathcal{D}_2 i\} - 1}{\exp\{\mathcal{D}_2\} - 1} \leq \exp\{\mathcal{D}_2(i+1)\},$$

one deduces that

$$Z_n \leq 2\max_{i=1,\ldots,m}\{\max\{\beta_{1,i}, \beta_{2,i}\}\}\exp\{\mathcal{D}_2(m+2)\}(\bar{k}j^* + 1)\rho_n.$$

In addition, for all $\omega \in \mathcal{A}$, one has $\bar{p}(S_n(\tau_{m,n}) - C)^+ \leq \bar{p}(\bar{k}j^* + 1)\exp\{\mathcal{D}_2(m+1)\}\rho_n$.

Upper bounding the penalties by their maximum possible values, one obtains $\mathbb{E}[Z_n \mid \mathcal{A}^c] \leq n\sum_{i=1}^m(\beta_{1,i}\xi_i + \beta_{2,i}\zeta_i)$.

We now analyze $\mathbb{E}[\bar{p}(S_n(\tau_{m,n})-C)^+ \mid \mathcal{A}^c]$. First note that

$$\begin{aligned}
&\mathbb{E}\big[S_n(\tau_{m,n}) \mid \mathcal{A}^c\big]\\
&=\sum_{i=0}^{m-1}\sum_{j=0}^{\bar{k}j^*}\mathbb{E}\big[S_n(t^i_{j+1})-S_n(t^i_j) \mid \mathcal{A}^c\big]\\
&=\sum_{i=0}^{m-1}\sum_{j=0}^{\bar{k}j^*}\mathbb{E}\Big[\mathbb{E}\big[S_n(t^i_{j+1})-S_n(t^i_j) \mid \mathcal{A}^c, p(t^i_j), \varepsilon_S(t^i_j)\ge\gamma^i_j\big]\\
&\qquad\cdot\mathbb{P}\big\{\varepsilon_S(t^i_j)\ge\gamma^i_j \mid \mathcal{A}^c, p(t^i_j)\big\}\\
&\qquad+\mathbb{E}\big[S_n(t^i_{j+1})-S_n(t^i_j) \mid \mathcal{A}^c, p(t^i_j), \varepsilon_S(t^i_j)<\gamma^i_j\big]\\
&\qquad\qquad\cdot\mathbb{P}\big\{\varepsilon_S(t^i_j)<\gamma^i_j \mid \mathcal{A}^c, p(t^i_j)\big\}\Big]\\
&=\sum_{i=0}^{m-1}\sum_{j=0}^{\bar{k}j^*}\mathbb{E}\Big[\mathbb{E}\big[S_n(t^i_{j+1})-S_n(t^i_j) \mid p(t^i_j), \varepsilon_S(t^i_j)\ge\gamma^i_j\big]\\
&\qquad\cdot\mathbb{P}\big\{\varepsilon_S(t^i_j)\ge\gamma^i_j \mid \mathcal{A}^c, p(t^i_j)\big\}\\
&\qquad+\mathbb{E}\big[S_n(t^i_{j+1})-S_n(t^i_j) \mid p(t^i_j), \varepsilon_S(t^i_j)<\gamma^i_j\big]\\
&\qquad\qquad\cdot\mathbb{P}\big\{\varepsilon_S(t^i_j)<\gamma^i_j \mid \mathcal{A}^c, p(t^i_j)\big\}\Big]\\
&\overset{(a)}{\le}\sum_{i=0}^{m-1}\sum_{j=0}^{\bar{k}j^*}\mathbb{E}\Big[\big[\Lambda\bar{F}(p(t^i_j))\delta^i_j+\gamma^i_j+1\big]\mathbb{P}\big\{\varepsilon_S(t^i_j)\ge\gamma^i_j \mid \mathcal{A}^c, p(t^i_j)\big\}\\
&\qquad+\big[\Lambda\bar{F}(p(t^i_j))\delta^i_j+\gamma^i_j\big]\mathbb{P}\big\{\varepsilon_S(t^i_j)\le\gamma^i_j \mid \mathcal{A}^c, p(t^i_j)\big\}\Big],
\end{aligned}$$

where (a) follows from the fact that for a Poisson random variable Y, for any $a>0$, $\mathbb{E}[Y \mid Y\ge a]=\mathbb{E}[Y]+\mathbb{E}[Y]\mathbb{P}\{Y\ge\lceil a\rceil-1\}/\mathbb{P}\{Y\ge\lceil a\rceil\}\le\mathbb{E}[Y]+\lceil a\rceil$. This, in turn, implies that

$$\begin{aligned}
\mathbb{E}\big[(S_n(\tau_{m,n})-C)^+ \mid \mathcal{A}^c\big]&\le\mathbb{E}\big[S_n(\tau_{m,n}) \mid \mathcal{A}^c\big]\\
&\le\sum_{i=0}^{m-1}\sum_{j=0}^{\bar{k}j^*}\big[\Lambda\delta^i_j+\gamma^i_j\big]\\
&=\Lambda\tau_{m,n}+m(\bar{k}j^*+1)(\gamma^i_j+1)\le\mathcal{D}_5 n,
\end{aligned}$$

where $\mathcal{D}_5$ is an appropriate positive constant. To conclude the proof, the next result, whose proof is presented in the online companion, provides a lower bound on the probability of $\mathcal{A}$.

Lemma 2. *For $\alpha_n=\alpha$ and δ specified as in* (30)–(31), $\mathbb{P}\{\mathcal{A}\}\ge 1-\mathcal{D}_4/n$, *for some* $\mathcal{D}_4>0$.

From this, one obtains

$$\begin{aligned}
&\mathbb{E}[\tilde{R}_n(T_n)-R_n(T_n)]+\mathbb{E}[Z_n+\bar{p}(S_n(\tau_{m,n})-C)^+]\\
&\le(\bar{k}j^*+1)\rho_n\exp\{\mathcal{D}_2(m+1)\}+n\Lambda\bar{p}T\mathcal{D}_4\\
&\quad+\Big(2\max_{i=1,\dots,m}\{\max\{\beta_{1,i},\beta_{2,i}\}\}+\bar{p}\Big)\exp\{\mathcal{D}_2(m+2)\}(\bar{k}j^*+1)\rho_n\\
&\quad+\Big[\sum_{i=1}^{m}(\beta_{1,i}\xi_i+\beta_{2,i}\zeta_i)+\bar{p}\mathcal{D}_5\Big]\mathcal{D}_4.
\end{aligned}$$

Noting that $\tilde{R}_n(T)=J^d_n=nJ^d$, $J^*_n\le J^d_n$ and that log-log $n\le(\log n)^{1/2}$, we deduce that for some positive constant $\mathcal{D}_6$,

$$\begin{aligned}
&\mathbb{E}[R_n(T_n)-Z_n-\bar{p}(S_n(\tau_{m,n})-C)^+]\\
&\quad\ge J^d_n-\mathcal{D}_6(\log n)^2 n^{1/2}\ge J^*_n-\mathcal{D}_5(\log n)^2 n^{1/2}.
\end{aligned}$$

This completes the proof. □

References

Araman, V. F., R. A. Caldentey. 2009. Dynamic pricing for nonperishable products with demand learning. *Oper. Res.* **57**(5) 1169–1188.

Aviv, Y., A. Pazgal. 2005. Pricing of short life-cycle products through active learning. Working paper, Washington University in St. Louis, St. Louis.

Bassambo, A., J. M. Harrison, A. Zeevi. 2005. Dynamic routing and admission control in high-volume service systems: Asymptotic analysis via multi-scale fluid limits. *Queueing Systems* **51**(3–4) 249–285.

Bassambo, A., J. M. Harrison, A. Zeevi. 2006. Design and control of a large call center: Asymptotic analysis of an LP-based method. *Oper. Res.* **54**(3) 419–435.

Besbes, O., C. Maglaras. 2009. Revenue optimization of a make-to-order queue in an uncertain market environment. *Oper. Res.* **57**(6) 1438–1450.

Besbes, O., A. Zeevi. 2009. Dynamic pricing without knowing the demand function: Risk bounds and near-optimal algorithms. *Oper. Res.* **57**(6) 1407–1420.

Broder, J., P. Rusmevichientong. 2009. Dynamic pricing under a general parametric choice model. Working paper, Cornell University, Ithaca, NY.

Broffman, W. 2007. The bottom line on revenue management. *Multihousing Professional* (March), http://www.multihousingpro.com/article.php?AID=230.

Bryson, A. E., Y. C. Ho. 1975. *Applied Optimal Control*. Hemisphere Publishing, New York.

Cooper, W. 2002. Asymptotic behavior of an allocation policy for revenue management. *Oper. Res.* **50**(4) 720–727.

Eren, S., C. Maglaras. 2010. Monopoly pricing with limited demand information. *J. Revenue Pricing Management* **9**(1–2) 23–48.

Gallego, G., G. van Ryzin. 1994. Optimal dynamic pricing of inventories with stochastic demand over finite horizons. *Management Sci.* **50**(8) 999–1020.

Gans, N., G. Van Ryzin. 1997. Optimal control of a multiclass, flexible queueing system. *Oper. Res.* **45**(5) 677–693.

Harrison, J. M. 1996. The BIGSTEP approach to flow management in stochastic processing networks. F. Kelly, S. Zachary, I. Ziedins, eds. *Stochastic Networks: Theory and Applications*. Oxford University Press, Oxford, UK, 57–90.

Jasin, S., S. Kumar. 2010. A re-solving heuristic with bounded revenue loss for network revenue management with customer choice. Working paper, Stanford University, Stanford, CA.

Levin, Y., J. McGill, M. Nediak. 2007. Risk in revenue management and dynamic pricing. *Oper. Res.* **56**(2) 326–343.

Lobo, M. S., S. Boyd. 2003. Pricing and learning with uncertain demand. Working paper, Duke University, Durham, NC.

Maglaras, C. 2000. Discrete-review policies for scheduling stochastic networks: Trajectory tracking and fluid-scale asymptotic optimality. *Ann. Appl. Prob.* **10**(3) 897–929.

Maglaras, C., J. Meissner. 2006. Dynamic pricing strategies for multi-product revenue management problems. *Manufacturing and Service Oper. Management* **8**(2) 136–148.

Paschalidis, I., C. Su, M. Caramanis. 2004. Target-pursuing scheduling and routing policies for multiclass queueing networks. *IEEE Trans. Automatic Control* **49**(10) 1709–1722.

Phillips, R. 2005. *Pricing and Revenue Optimization*. Stanford University Press, Palo Alto, CA.

Reiman, M. I., Q. Wang. 2008. An asymptotically optimal policy for a quantity-based network revenue management problem. *Math. Oper. Res.* **33**(2) 257–282.

Secomandi, N. 2008. An analysis of the control-algorithm re-solving issue in inventory and revenue management. *Manufacturing Service Oper. Management* **10**(3) 468–483.

Talluri, K. T., G. J. van Ryzin. 2005. *Theory and Practice of Revenue Management*. Springer-Verlag, New York.

Tassiulas, L., S. Papavassiliou. 1995. Optimal anticipative scheduling with asynchronous transmission opportunities. *IEEE Trans. Automatic Control* **40**(12) 2052–2062.