

MIT Open Access Articles

Dynamic Pricing with a Prior on Market Response

The MIT Faculty has made this article openly available. **Please share** how this access benefits you. Your story matters.

Citation: Farias, Vivek F. and Benjamin Van Roy. "Dynamic Pricing with a Prior on Market Response." (2010) Operations Research 58.1.

As Published: <http://dx.doi.org/10.1287/opre.1090.0729>

Publisher: Institute for Operations Research and the Management Sciences

Persistent URL: <http://hdl.handle.net/1721.1/64802>

Version: Author's final manuscript: final author's manuscript post peer review, without publisher's formatting or copy editing

Terms of use: Creative Commons Attribution-Noncommercial-Share Alike 3.0



Dynamic Pricing with a Prior on Market Response

Vivek F. Farias * Benjamin Van Roy †

January 20, 2009

Abstract

We study a problem of dynamic pricing faced by a vendor with limited inventory, uncertain about demand, aiming to maximize expected discounted revenue over an infinite time horizon. The vendor learns from purchase data, so his strategy must take into account the impact of price on both revenue and future observations. We focus on a model in which customers arrive according to a Poisson process of uncertain rate, each with an independent, identically distributed reservation price. Upon arrival, a customer purchases a unit of inventory if and only if his reservation price equals or exceeds the vendor's prevailing price.

We propose a simple heuristic approach to pricing in this context, which we refer to as decay balancing. Computational results demonstrate that decay balancing offers significant revenue gains over recently studied certainty equivalent and greedy heuristics. We also establish that changes in inventory and uncertainty in the arrival rate bear appropriate directional impacts on decay balancing prices in contrast to these alternatives, and we derive worst-case bounds on performance loss. We extend the three aforementioned heuristics to address a model involving multiple customer segments and stores and provide experimental results demonstrating similar relative merits in this context.

Keywords: dynamic pricing; demand learning; revenue management.

1 Introduction

Consider a vendor of Winter apparel. New items are stocked in the Autumn and sold over several months. Because of significant manufacturing lead times and fixed costs, items are not restocked over this period. Evolving fashion trends generate great uncertainty in the number of customers who will consider purchasing these items. To optimize revenue, the vendor should adjust prices over

*Sloan School of Management, Massachusetts Institute of Technology, email :vivekf@mit.edu

†Electrical Engineering and Management Science and Engineering, Stanford University, email:bvr@stanford.edu

time. But how should these prices be set as time passes and units are sold? This is representative of problems faced by many vendors of seasonal, fashion, and perishable goods.

There is a substantial literature on pricing strategies for such a vendor (see Talluri and van Ryzin (2004) and references therein). Gallego and van Ryzin (1994), in particular, formulated an elegant model in which the vendor starts with a finite number of identical indivisible units of inventory. Customers arrive according to a Poisson process, with independent, identically distributed reservation prices. In the case of exponentially distributed reservation prices the optimal pricing strategy is easily derived. The analysis of Gallego and van Ryzin (1994) can be used to derive pricing strategies that optimize expected revenue over a finite horizon and is easily extended to the optimization of discounted expected revenue over an infinite horizon. Resulting strategies provide insight into how prices should depend on the arrival rate, expected reservation price, and the length of the horizon or discount rate.

Our focus is on an extension of this model in which the arrival rate is uncertain and the vendor learns from sales data. Incorporating such uncertainty is undoubtedly important in many industries that practice revenue management. For instance, in the Winter fashion apparel example, there may be great uncertainty in how the market will respond to the product at the beginning of a sales season; the vendor must take into account how price influences both revenue and future observations from which he can learn.

In this setting, it is important to understand how uncertainty should influence price. However, uncertainty in the arrival rate makes the analysis challenging. Optimal pricing strategies can be characterized by a Hamilton-Jacobi-Bellman (HJB) Equation, but this approach is typically not analytically tractable (at least for the models we consider). Further, for arrival rate distributions of interest, grid-based numerical methods require discommoding computational resources and generate strategies that are difficult to interpret. As such researchers have designed and analyzed heuristic approaches.

Aviv and Pazgal (2005a) studied a *certainty equivalent* heuristic for exponentially distributed reservation prices which at each point in time computes the conditional expectation of the arrival rate, conditioned on observed sales data, and prices as though the arrival rate is equal to this expectation. Araman and Caldentey (2005) recently proposed a more sophisticated heuristic that takes arrival rate uncertainty into account when pricing. The idea is to use a strategy that is *greedy* with respect to a particular approximate value function. In this paper, we propose and analyze *decay balancing*, a new heuristic approach which makes use of the same approximate value function as the greedy approach of Araman and Caldentey (2005).

Several idiosyncrasies distinguish the models studied in Aviv and Pazgal (2005a) and Araman and Caldentey (2005). The former models uncertainty in the arrival rate in terms of a Gamma distribution, whereas the latter uses a two-point distribution. The former considers maximization of expected revenue over a finite horizon, whereas the latter considers expected discounted revenue over an infinite horizon. To elucidate relationships among the three heuristic strategies, we study them in the context of a common model. In particular, we take the arrival rate to be distributed according to a finite mixture of Gamma distributions. This is a very general class of priors and can closely approximate any bounded continuous density. We take the objective to be maximization of expected discounted revenue over an infinite horizon. It is worth noting that in the case of exponentially distributed reservation prices such a model is equivalent to one without discounting but where expected reservation prices diminish exponentially over time. This may make it an appropriate model for certain seasonal, fashion, or perishable products. Our modeling choices were made to provide a simple, yet fairly general context for our study. We expect that our results can be extended to other classes of models such as those with finite time horizons, though this is left for future work.

The certainty equivalent heuristic is natural and simple to implement. It does not, however, take uncertainty in the arrival rate into account. While the greedy heuristic does attempt to do this, we demonstrate through computational experiments that the performance of this approach can degrade severely at high levels of uncertainty in arrival rate. While being only slightly more complex to implement than the certainty equivalent heuristic and typically less so than greedy pricing, we demonstrate that decay balancing offers significant performance gains over these heuristics especially in scenarios where the seller begins with a high degree of uncertainty in arrival rate. From a more qualitative perspective, uncertainty in the arrival rate and changes in inventory bear appropriate directional impacts on decay balancing prices: uncertainty in the arrival rate increases price, while a decrease in inventory increases price. In contrast, uncertainty in the arrival rate has no impact on certainty equivalent prices while greedy prices can increase or decrease with inventory.

In addition to our base model, we consider a generalization which involves a vendor with multiple branches that can offer different prices and tend to attract different classes of customers. The branches share and learn from each others' data and price to maximize aggregate expected discounted revenue. We extend the three heuristics to the context of this model and present computational results demonstrating relative merits analogous to our base case of a single branch and single customer class.

We establish bounds on performance loss incurred by decay balancing relative to an optimal

policy. These bounds indicate a certain degree of robustness. For instance, when customer reservation prices are exponentially distributed and their arrival rate is Gamma distributed we establish that decay balancing always garners at least one third of the maximum expected discounted revenue. Allowing for a dependence on arrival rate uncertainty and/or restricting attention to specific classes of reservation price distributions leads to substantially stronger bounds. It is worth noting that no performance loss bounds (uniform or otherwise) have been established for the certainty equivalent and greedy heuristics. Further, computational results demonstrate that our bounds are not satisfied by the greedy heuristic.

Aside from Aviv and Pazgal (2005a) and Araman and Caldentey (2005), there is a significant literature on dynamic pricing while learning about demand. A recent paper in this regard is Aviv and Pazgal (2005b) which considers, in a discrete time setting, a partially observable Markov modulated demand model. As we will discuss further in 6, a special case of the heuristic they develop is closely related to decay balancing. Lin (2007) considers a model identical to Aviv and Pazgal (2005a) and develops heuristics which are motivated by the behavior of a seller who knows the arrival rate and anticipates all arriving customers. Bertsimas and Perakis (2003) develop several algorithms for a discrete, finite time-horizon problem where demand is an unknown linear function of price plus Gaussian noise. This allows for least-squares based estimation. Lobo and Boyd (2003) study a model similar to Bertsimas and Perakis (2003) and propose a “price-dithering” heuristic that involves the solution of a semi-definite convex program. All of the aforementioned work is experimental; no performance guarantees are provided for the heuristics proposed. Cope (2006) studies a Bayesian approach to pricing where inventory levels are unimportant (this is motivated by sales of on-line services) and there is uncertainty in the distribution of reservation price. His work uses a very general prior distribution - a Dirichlet mixture - on reservation price. Modeling this type of uncertainty within a framework where inventory levels do matter represents an interesting direction for future work. In contrast with the the above work, Burnetas and Smith (1998) and Kleinberg and Leighton (2004) consider non-parametric approaches to pricing with uncertainty in demand. However, those models again do not account for inventory levels. Recently, Besbes and Zeevi (2006) presented a non-parametric algorithm for ‘blind’ pricing of multiple products that use multiple resources, similar to the model considered in Gallego and van Ryzin (1997). Their algorithm requires essentially no knowledge of the demand function and is optimal in a certain asymptotic regime. However, the algorithm requires trying each alternative among a multidimensional grid in the space of price vectors, and therefore, could take a long time to adequately learn about demand.

The remainder of this paper is organized as follows: In Section 2, we formulate our model and cast our pricing problem as one of stochastic optimal control. Section 3 develops the HJB equation for the optimal pricing problem in the contexts of known and unknown arrival rates. Section 4 first introduces existing heuristics for the problem and then introduces a new heuristic – decay balancing – which is the focus of this paper. This section also discusses structural properties of the decay balancing policy. Section 5 presents a computational study that compares decay balancing to certainty equivalent and greedy pricing heuristics as well as a clairvoyant algorithm. Section 6 is devoted to a theoretical performance analysis of the decay balancing heuristic. We provide worst-case performance guarantees that depend on initial uncertainty in market response. When the arrival rate is Gamma distributed and reservation prices are exponentially distributed, we prove a uniform performance guarantee for our heuristic. Section 7 discusses an extension of our heuristic to a multidimensional version of the problem which involves a vendor with multiple stores and several customer segments with different demand characteristics. The section presents computational results that are qualitatively similar to those in Section 5. Section 8 concludes.

2 Problem Formulation

We consider a problem faced by a vendor who begins with x_0 identical indivisible units of a product and dynamically adjusts price p_t over time $t \in [0, \infty)$. Customers arrive according to a Poisson process with rate λ . As a convention, we will assume that the arrival process is right continuous with left limits. Each customer’s reservation price is an independent random variable with cumulative distribution $F(\cdot)$. A customer purchases a unit of the product if it is available at the time of his arrival at a price no greater than his reservation price; otherwise, the customer permanently leaves the system.

For convenience, we introduce the notation $\bar{F}(p) = 1 - F(p)$ for the tail probability. We place the following restrictions on $F(\cdot)$:

Assumption 1.

1. $F(\cdot)$ has a differentiable density $f(\cdot)$ with support \mathbb{R}_+ .
2. F has a non-decreasing hazard rate. That is, $\rho(p) = f(p)/\bar{F}(p)$ is non-decreasing in p .

Assumption 1 is satisfied by many relevant distributions including the exponential, extreme value, logistic and Weibull, to name a few. We introduce this assumption to facilitate use of first order optimality conditions to characterize solutions of various optimization problems that will

arise in our discussion. It is possible to extend our results to reservation price distributions that do not satisfy this assumption, though that would require additional technical work.

Let t_k denote the time of the k th purchase and $n_t = |\{t_k : t_k \leq t\}|$ denote the number of purchases made by customers arriving at or before time t . The vendor's expected revenue, discounted at a rate of $\alpha > 0$, is given by

$$E \left[\int_{t=0}^{\infty} e^{-\alpha t} p_t dn_t \right].$$

Let $\tau_0 = \inf\{t : x_t = 0\}$ be the time at which the final unit of inventory is sold. For $t \leq \tau_0$, n_t follows a Poisson process with intensity $\lambda \bar{F}(p_t)$. Consequently, (see Theorem III.T9 in Bremaud (1981)), one may show that

$$E \left[\int_{t=0}^{\infty} e^{-\alpha t} p_t dn_t \right] = E \left[\int_{t=0}^{\tau_0} e^{-\alpha t} p_t \lambda \bar{F}(p_t) dt \right].$$

We now describe the vendor's optimization problem. We first consider the case where the vendor knows λ and later allow for arrival rate uncertainty. In the case with known arrival rate, we consider pricing policies π that are measurable real-valued functions of the inventory level. The price is irrelevant when there is no inventory, and as a convention, we will require that $\pi(0) = \infty$. We denote the set of policies by Π_λ . A vendor who employs pricing policy $\pi \in \Pi_\lambda$ sets price according to $p_t = \pi(x_t)$, where $x_t = x_0 - n_t$, and receives expected discounted revenue

$$J_\lambda^\pi(x) = E_{x,\pi} \left[\int_{t=0}^{\tau_0} e^{-\alpha t} p_t \lambda \bar{F}(p_t) dt \right],$$

where the subscripts of the expectation indicate that $x_0 = x$ and $p_t = \pi(x_t)$. The optimal discounted revenue is given by $J_\lambda^*(x) = \sup_{\pi \in \Pi_\lambda} J_\lambda^\pi(x)$, and a policy π is said to be optimal if $J_\lambda^* = J_\lambda^\pi$.

Suppose now that the arrival rate λ is not known, but rather, the vendor starts with a prior on λ that is a finite mixture of Gamma distributions. A K th order mixture of this type is parameterized by vectors $a_0, b_0 \in \mathbb{R}_+^K$ and a vector of K weights $w_0 \in \mathbb{R}_+^K$ that sum to unity. The density, g , for such a prior is given by:

$$g(\lambda) = \sum_k w_{0,k} \frac{b_{0,k}^{a_{0,k}} \lambda^{a_{0,k}-1} e^{-\lambda b_{0,k}}}{\Gamma(a_{0,k})},$$

where Γ denotes the Gamma-function: $\Gamma(x) = \int_{s=0}^{\infty} s^{x-1} e^{-s} ds$. The expectation and variance are $E[\lambda] = \sum_k w_{0,k} a_{0,k} / b_{0,k} \triangleq \mu_0$ and $\text{Var}[\lambda] = \sum_k w_{0,k} a_{0,k} (a_{0,k} + 1) / b_{0,k}^2 - \mu_0^2$. Any prior on λ with a continuous, bounded density can be approximated to an arbitrary accuracy within such a family

(see Dalal and Hall (1983)). Moreover, as we describe below, posteriors on λ continue to remain within this family rendering such a model parsimonious as well as relatively tractable.

The vendor revises his beliefs about λ as sales are observed. In particular, at time t , the vendor obtains a posterior that is a k th order mixture of Gamma distributions with parameters

$$a_{t,k} = a_{0,k} + n_t \quad \text{and} \quad b_{t,k} = b_{0,k} + \int_{\tau=0}^t \bar{F}(p_\tau) d\tau,$$

and weights that evolve according to:

$$dw_{t,k} = w_{t,k} \left(\frac{a_{t,k}/b_{t,k} - \sum_k w_{t,k}(a_{t,k}/b_{t,k})}{\sum_k w_{t,k}(a_{t,k}/b_{t,k})} dn_t + \left(a_{t,k}/b_{t,k} - \sum_k w_{t,k}(a_{t,k}/b_{t,k}) \right) dt. \right)$$

Note that the vendor does not observe all customer arrivals but only those that result in sales. Further, lowering price results in more frequent sales and therefore more accurate estimation of the demand rate.

We consider pricing policies π that are measurable real-valued functions of the inventory level and arrival rate distribution parameters. As a convention we require that $\pi(0, a, b, w) = \infty$ for all arrival rate distribution parameters a, b and w . We denote the domain by $\mathcal{S} = \mathbb{N} \times \mathbb{R}_+^k \times \mathbb{R}_+^k \times \mathbb{R}_+^k$ and the set of policies by Π . Let $z_t = (x_t, a_t, b_t, w_t)$. A vendor who employs pricing policy $\pi \in \Pi$ sets price according to $p_t = \pi(z_t)$ and receives expected discounted revenue

$$J^\pi(z) = E_{z,\pi} \left[\int_{t=0}^{\tau_0} e^{-\alpha t} p_t \lambda \bar{F}(p_t) dt \right],$$

where the subscripts of the expectation indicate that $z_0 = z$ and $p_t = \pi(z_t)$. Note that, unlike the case with known arrival rate, λ is a random variable in this expectation. The optimal discounted revenue is given by $J^*(z) = \sup_{\pi \in \Pi} J^\pi(z)$, and a policy π is said to be optimal if $J^* = J^\pi$. We will use the notation $J^{*,\alpha}$ for the optimal value function when we wish to emphasize the dependence on α .

3 Optimal Pricing

An optimal pricing policy can be derived from the value function J^* . The value function in turn solves the HJB equation. Unfortunately direct solution of the HJB equation, either analytically or computationally, does not appear to be a feasible task and one must resort to heuristic policies. With an end to deriving such heuristic policies we characterize optimal solutions to problems with

known and unknown arrival rates and discuss some of their properties.

3.1 The Case of a Known Arrival Rate

We begin with the case of a known arrival rate. For each $\lambda \geq 0$ and $\pi \in \Pi_\lambda$, define an operator $H_\lambda^\pi : \mathbb{R}^{x_0+1} \rightarrow \mathbb{R}^{x_0+1}$ by

$$(H_\lambda^\pi J)(x) = \lambda \bar{F}(\pi(x))(\pi(x) + J(x-1) - J(x)) - \alpha J(x).$$

Recall that $\pi(0) = \infty$. In this case, we interpret $\bar{F}(\pi(0))\pi(0)$ as a limit, and Assumption 1 (which ensures a finite, unique static revenue maximizing price) implies that $(H_\lambda^\pi J)(0) = -\alpha J(0)$. Further, we define the dynamic programming operator

$$(H_\lambda J)(x) = \sup_{\pi \in \Pi_\lambda} (H_\lambda^\pi J)(x).$$

It is easy to show that J_λ^* is the unique solution to the HJB Equation $H_\lambda J = 0$. The first order optimality condition for prices yields an optimal policy of the form

$$\pi_\lambda^*(x) = 1/\rho(\pi_\lambda^*(x)) + J_\lambda^*(x) - J_\lambda^*(x-1),$$

for $x > 0$. By Assumption 1 and the fact that $J_\lambda^*(x) \geq J_\lambda^*(x-1)$, the above equation always has a solution in \mathbb{R}_+ .

Given that J^* satisfies the HJB equation, we have

$$(1) \quad \alpha J_\lambda^*(x) = \begin{cases} \sup_{p \geq 0} \lambda \bar{F}(p)(p + J_\lambda^*(x-1) - J_\lambda^*(x)) & \text{if } x > 0 \\ 0 & \text{otherwise.} \end{cases}$$

Assumption 1 guarantees that $\sup_{p \geq 0} \bar{F}(p)(p - c)$ is a decreasing function of c on \mathbb{R}_+ . This allows one to compute $J_\lambda^*(x)$ given $J_\lambda^*(x-1)$ via bisection search. This offers an efficient algorithm that computes $J_\lambda^*(0), J_\lambda^*(1), \dots, J_\lambda^*(x)$ in x iterations. As a specific concrete example, consider the case where reservation prices are exponentially distributed with mean $r > 0$. We have

$$\alpha J_\lambda^*(x) = \begin{cases} \lambda r \exp\left(\frac{1}{r}(J_\lambda^*(x-1) - J_\lambda^*(x)) - 1\right) & \text{if } x > 0 \\ 0 & \text{otherwise.} \end{cases}$$

It follows that

$$(2) \quad J_\lambda^*(x) = rW\left(\left(e^{-1}\lambda/\alpha\right)\exp(J_\lambda^*(x-1)/r)\right)$$

for $x > 0$, where $W(\cdot)$ is the Lambert W-function (the inverse of xe^x).

We note that a derivation of the optimal policy for the case of a known arrival rate may also be found in Araman and Caldentey (2005), among other sources.

The value function for the case of a known arrival rate will be used in the design of our heuristic for the case with arrival rate uncertainty. We establish here properties of this value function J_λ^* and its associated optimal policy π_λ^* , which we will later use. We will make the following assumption to simplify our analysis.

Assumption 2. $J_\lambda^*(x)$ is a differentiable function of λ on \mathbb{R}_+ for all $x \in \mathbb{N}$.

Note that this assumption is satisfied for the case of exponential reservation prices. The following comparative statics for π_λ^* are proved in the appendix:

Lemma 1. $\pi_\lambda^*(x)$ is decreasing in x (on \mathbb{N}) and non-decreasing in λ (on \mathbb{R}_+).

For a fixed inventory level it is natural to expect decreasing returns to increases in the arrival rate λ ; the following Lemma, proved in the appendix, formalizes this intuition:

Lemma 2. For all $x \in \mathbb{N}$, $J_\lambda^*(x)$ is an increasing, concave function of λ on \mathbb{R}_+ .

3.2 The Case of an Unknown Arrival Rate

Let $S_{\tilde{x}, \tilde{a}, \tilde{b}} = \{(x, a, b, w) \in S : a + x = \tilde{a} + \tilde{x}, \tilde{b} \leq b, w \geq 0, 1'w = 1\}$ denote the set of states that might be visited starting at a state with $x_0 = \tilde{x}, a_0 = \tilde{a}, b_0 = \tilde{b}$. Let \mathcal{J} denote the set of functions $J : S \mapsto \mathbb{R}$ such that $\sup_{z \in S_{\tilde{x}, \tilde{a}, \tilde{b}}} |J(z)| < \infty$ for all \tilde{x} and $\tilde{b} > 0$ and that have bounded derivatives with respect to the third and fourth arguments. We define $\mu(z)$ to be the expectation for the prior on arrival rate in state z , so that $\mu(z) = \sum_k w_k a_k / b_k$.

For each policy $\pi \in \Pi$, we define an operator

$$(H^\pi J)(z) = \bar{F}(\pi(z)) \left(\mu(z) (\pi(z) + J(z') - J(z)) + (DJ)(z) \right) - \alpha J(z),$$

where $z \in S_{\tilde{x}, \tilde{a}, \tilde{b}}$, $z = (x, a, b, w)$ and $z' = (x - 1, a + 1, b, w')$. Here w' is defined according to

$w'_k = (w_k a_k)/(b_k \mu(z))$, and D is a differential operator given by:

$$(DJ)(z) = \sum_k w_k (\mu(z) - a_k/b_k) \frac{d}{dw_k} J(z) + \frac{d}{db_k} J(z).$$

We now define the dynamic programming operator H according to:

$$(3) \quad (HJ)(z) = \sup_{\pi} (H^{\pi} J)(z).$$

Using standard dynamic programming arguments, one can show the value function J^* solves the HJB equation:

$$(4) \quad (HJ)(z) = 0.$$

Further, a policy π is optimal if and only if $H^{\pi} J^* = 0$. This equation provides a prescription for efficient computation of an optimal policy given the value function. Unfortunately, there is no known analytical solution to the HJB Equation when the arrival rate is unknown, even for special cases such as a Gamma or two-point prior with exponential reservation prices. Further, grid-based numerical solution methods require discommoding computational resources and generate strategies that are difficult to interpret. As such, simple effective heuristics are desirable.

4 Heuristics

This section presents three heuristic pricing policies. The first two have been considered in prior literature and the third is one we propose and analyze in this paper.

4.1 Certainty Equivalent

Aviv and Pazgal (2005a) studied a *certainty equivalent* heuristic which at each point in time computes the conditional expectation of the arrival rate, conditioned on observed sales data, and prices as though the arrival rate is equal to this expectation. In our context, the price function for such a heuristic uniquely solves

$$\pi_{ce}(z) = \frac{1}{\rho(\pi_{ce}(z))} + J_{\mu(z)}^*(x) - J_{\mu(z)}^*(x-1),$$

for $x > 0$. The existence of a unique solution to this equation is guaranteed by Assumption 1. As derived in the preceding section, this is an optimal policy for the case where the arrival rate is known and equal to $\mu(z)$, which is the expectation of the arrival rate given a prior distribution with parameters a, b and w . The certainty equivalent policy is computationally attractive since J_λ^* is easily computed numerically (and in some cases, even analytically) as discussed in the previous section. As one would expect, prices generated by this heuristic increase as the inventory x decreases. However, arrival rate uncertainty bears no influence on price – the price only depends on the arrival rate distribution through its expectation $\mu(z)$. Hence, this pricing policy is unlikely to appropriately address information acquisition.

4.2 Greedy Pricing

We now present another heuristic which was recently proposed by Araman and Caldentey (2005) and does account for arrival rate uncertainty. To do so, we first introduce the notion of a *greedy policy*. A policy π is said to be *greedy* with respect to a function J if $H^\pi J = HJ$. The first-order necessary condition for optimality and Assumption 1 imply that the greedy price is given by the solution to

$$\pi(z) = \left(\frac{1}{\rho(\pi(z))} + J(z) - J(z') - \frac{1}{\mu(z)}(DJ)(z) \right)^+,$$

for $z = (x, a, b, w)$ with $x > 0$ and $z' = (x - 1, a + 1, b, w')$ with $w'_k = (w_k a_k)/(b_k \mu)$.

Perhaps the simplest approximation one might consider to $J^*(z)$ is $J_{\mu(z)}^*(x)$, the value for a problem with known arrival rate $\mu(z)$. One troubling aspect of this approximation is that it ignores the variance (as also higher moments) of the arrival rate. The alternative approximation proposed by Araman and Caldentey takes variance into account. In particular their heuristic employs a greedy policy with respect to the approximate value function \tilde{J} which takes the form

$$\tilde{J}(z) = E[J_\lambda^*(x)],$$

where the expectation is taken over the random variable λ , which is drawn from a Gamma mixture with parameters a, b and w . $\tilde{J}(z)$ can be thought of as the expected optimal value if λ is to be observed at the next time instant. It is interesting to note that this approximation is in the same spirit as the ‘full information’ approximation considered in Aviv and Pazgal (2005b). The greedy price, however, is distinct from the full information price considered there.

Since it can only help to know the value of λ , $J_\lambda^*(x) \geq E[J^*(z)|\lambda]$. Taking expectations of both

sides of this inequality, we see that \tilde{J} is an upper bound on J^* . The approximation $J_{\mu(z)}^*(x)$ is a looser upper bound on $J^*(z)$ (which follows from concavity of J_λ^* in λ). Consequently, we have the following result whose proof may be found in the appendix.

Lemma 3. *For all $z \in \mathcal{S}, \alpha > 0$*

$$J^*(z) \leq \tilde{J}(z) \leq J_{\mu(z)}^*(x) \leq \frac{\bar{F}(p^*)p^*\mu(z)}{\alpha}.$$

where p^* is the static revenue maximizing price.

The greedy price in state z is thus the solution to

$$\pi_{\text{gp}}(z) = \left(\frac{1}{\rho(\pi_{\text{gp}}(z))} + \tilde{J}(z) - \tilde{J}(z') - \frac{1}{\mu(z)}(D\tilde{J})(z) \right)^+,$$

for $z = (x, a, b, w)$ with $x > 0$ and $z' = (x - 1, a + 1, b, w')$ with $w'_k = (w_k a_k)/(b_k \mu(z))$.

We have observed through computational experiments (see Section 6) that when reservation prices are exponentially distributed and the vendor begins with a Gamma prior with scalar parameters a and b , greedy prices can increase or decrease with the inventory level x , keeping a and b fixed. This is clearly not optimal behavior.

4.3 Decay Balancing

We now describe *decay balancing*, a new heuristic which will be the primary subject of the remainder of the paper. To motivate the heuristic, we start by deriving an alternative characterization of the optimal pricing policy. The HJB Equation (4) yields

$$\max_{p \geq 0} \bar{F}(p) (\mu(z) (p + J^*(z') - J^*(z)) + (DJ^*)(z)) = \alpha J^*(z),$$

for all $z = (x, a, b, w)$ and $z' = (x - 1, a + 1, b, w')$, with $x > 0$ and $w'_k = (w_k a_k)/(b_k \mu(z))$. This equation can be viewed as a balance condition. The right hand side represents the rate at which value decays over time; if the price were set to infinity so that no sales could take place for a time increment dt but an optimal policy is used thereafter, the current value would become $J^*(z) - \alpha J^*(z)dt$. The left hand side represents the rate at which value is generated from both sales and learning. The equation requires these two rates to balance so that the net value is conserved.

Note that the first order optimality condition implies that if $J(z') - J(z) + \frac{1}{\mu(z)}(DJ)(z) < 0$

(which must necessarily hold for $J = J^*$),

$$\frac{\bar{F}(p^*)}{\rho(p^*)}\mu(z) = \max_{p \geq 0} \bar{F}(p) (\mu(z) (p + J(z') - J(z)) + (DJ)(z)),$$

if p^* attains the maximum in the right hand side. Interestingly, the maximum depends on J only through p^* . Hence, the balance equation can alternatively be written in the following simpler form:

$$\frac{\bar{F}(\pi^*(z))}{\rho(\pi^*(z))}\mu(z) = \alpha J^*(z).$$

which implicitly characterizes π^* .

This alternative characterization of π^* makes obvious two properties of optimal prices. Note that $\bar{F}(p)/\rho(p)$ is decreasing in p . Consequently, holding a, b and w fixed, as x decreases, $J^*(z)$ decreases and therefore $\pi^*(z)$ increases. Further, since $J^*(z) \leq J_{\mu(z)}^*(x)$, we see that for a fixed inventory level x and expected arrival rate $\mu(z)$, the optimal price in the presence of uncertainty is higher than in the case where the arrival rate is known exactly.

Like greedy pricing, the decay balancing heuristic relies on an approximate value function. We will use the same approximation \tilde{J} . But instead of following a greedy policy with respect to \tilde{J} , the decay balancing approach chooses a policy π_{db} that satisfies the balance condition:

$$(5) \quad \frac{\bar{F}(\pi_{\text{db}}(z))}{\rho(\pi_{\text{db}}(z))}\mu(z) = \alpha \tilde{J}(z),$$

with the decay rate approximated using $\tilde{J}(z)$. The following Lemma guarantees that the above balance equation always has a unique solution so that our heuristic is well defined. The result is a straightforward consequence of Assumption 1 and the fact that $\frac{\bar{F}(p^*)}{\alpha \rho(p^*)}\mu(z) \geq \tilde{J}(z) \geq J^*(z) = \frac{\bar{F}(\pi^*(z))}{\alpha \rho(\pi^*(z))}\mu(z)$ where p^* is the static revenue maximizing price.

Lemma 4. *For all $z \in \mathcal{S}$, there is a unique $p \geq 0$ such that $\frac{\bar{F}(p)}{\rho(p)}\mu(z) = \alpha \tilde{J}(z)$.*

Unlike certainty equivalent and greedy pricing, uncertainty in the arrival rate and changes in inventory level have the correct directional impact on decay balancing prices. Holding a, b and w fixed, as x decreases, $\tilde{J}(z)$ decreases and therefore $\pi_{\text{db}}(z)$ increases. Holding x and the expected arrival rate $\mu(z)$ fixed, $\tilde{J}(z) \leq J_{\mu(z)}^*(x)$, so that the decay balance price with uncertainty in arrival rate is higher than when the arrival rate is known with certainty.

Table 1: Performance vs. a Clairvoyant algorithm

Inventory Level	$x_0 = 1$	$x_0 = 2$	$x_0 = 5$	$x_0 = 10$	$x_0 = 20$	$x_0 = 40$
Performance Gain (Gamma)	-13%	-10%	-6%	-3.7%	-2%	-0.5%
Performance Gain (Gamma Mixture)	-14.9%	-12.3%	-6.7%	-4.3%	-2.7%	-2.4%

5 Computational Study

This section will present computational results that highlight the performance of the decay balancing heuristic. We consider both Gamma as well as Gamma mixture priors and restrict attention to exponentially distributed reservation prices. Further, we will only consider problem instances where $\alpha = e^{-1}$ and $r = 1$; as we will discuss in Section 6, this is not restrictive (see Lemmas 5 and 10).

Performance relative to a Clairvoyant Algorithm: Consider a ‘clairvoyant’ algorithm that has access to the realization of λ at $t = 0$ and subsequently uses the pricing policy π_λ^* . The expected revenue garnered by such a pricing policy upon starting in state z is simply $E[J_\lambda^*(x)] = \tilde{J}(z)$ which, by Lemma 3, is an upper bound on $J^*(z)$. Our first experiment measures the average revenue earned using decay balancing with that earned using such a clairvoyant algorithm. The results are summarized in Table 1. We consider two cases: In the first, λ is drawn from a gamma distribution with shape parameter $a = 0.04$ and scale parameter $b = 0.001$ which corresponds to a mean of 40 and a coefficient of variation of 5. In the second, λ is drawn from a two point gamma mixture with parameters $a = [0.01023 \ 0.07161]'$, $b = [0.00102 \ 0.00102]'$, $w = [0.5 \ 0.5]'$ which correspond to a mean of 40 and a coefficient of variation of 5. These parameter values are representative of a high level of uncertainty in λ . As is seen in Table 2, the performance of the decay balancing heuristic is surprisingly close to that of the clairvoyant algorithm for both prior distributions.

Performance relative to Certainty Equivalent and Greedy Pricing Heuristics: We finally turn to studying the performance gains offered by decay balancing relative to the certainty equivalent and greedy policies.

1. *Dependence on initial inventory:* The gains offered by decay balancing relative to the certainty equivalent and greedy heuristic are pronounced at lower initial inventory levels i.e. in regimes where judiciously managing inventory is crucial. We consider relative performance gains for inventory levels between 1 and 10. We consider two initial priors for all heuristics: a gamma prior with

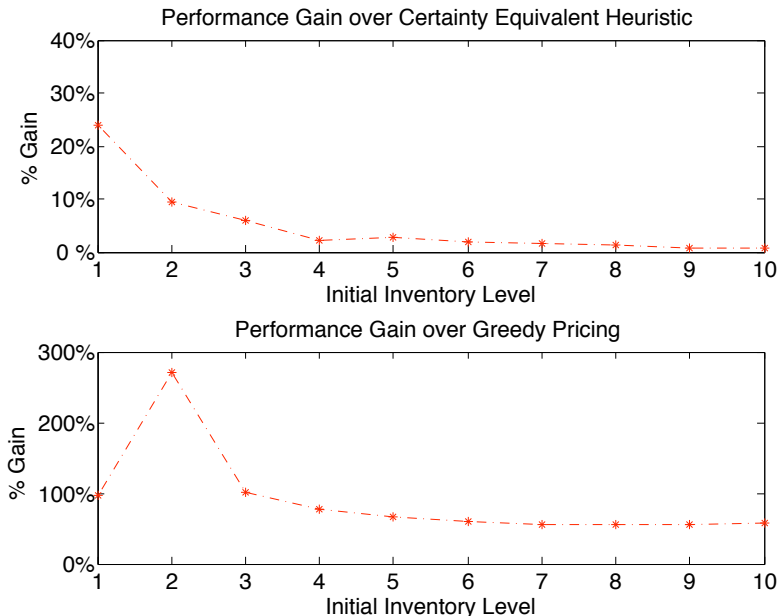


Figure 1: Performance gain over the Certainty Equivalent and Greedy Pricing heuristics at various inventory levels for Gamma prior

$a = 0.04$ and $b = 0.001$ (corresponding to a mean of 40 and a coefficient of variation of 5) and a gamma mixture prior with parameters $a = [0.01033 \ 0.09297]'$, $b = [0.00103 \ 0.00103]'$, $w = [0.5 \ 0.5]'$ (corresponding to a mean of 50 and a coefficient of variation of 4.47). Figures 1 and 2 indicate that we offer a substantial gain in performance over the certainty equivalent and greedy pricing heuristics, especially at lower initial inventory levels. The greedy policy performs particularly poorly. In addition, as discussed earlier, that policy exhibits qualitative behavior that is clearly suboptimal: for a problem with a gamma prior, mean reservation price 1 and discount factor e^{-1} , we compute $\pi_{gp}(1, 0.1, 0.1)(= 1.26) < \pi_{gp}(4, 0.1, 0.1)(= 1.61) > \pi_{gp}(10, 0.1, 0.1)(= 1.25)$ so that, all other factors remaining the same, prices may increase or decrease with an increase in inventory level. Our gain in performance falls at higher initial inventory levels. This is not surprising; intuitively, the control problem at hand is simpler there since we are essentially allowed to sacrifice a few units of inventory early on so as to learn quickly without incurring much of a penalty.

2. *Dependence on initial uncertainty:* The performance gains offered by the decay balancing heuristic are higher at higher initial levels of uncertainty. We present a lower bound on the maximal performance gain over the greedy and certainty equivalent heuristics for various coefficients of variation of an initial gamma prior on λ . See Figure 3 wherein the data point for each coefficient of variation c , corresponds to an experiment with $a = 1/c^2$, $b = 0.001$ (which corresponds to a mean of $1000/c^2$), and an inventory level of 1 and 2 for the certainty equivalent and greedy

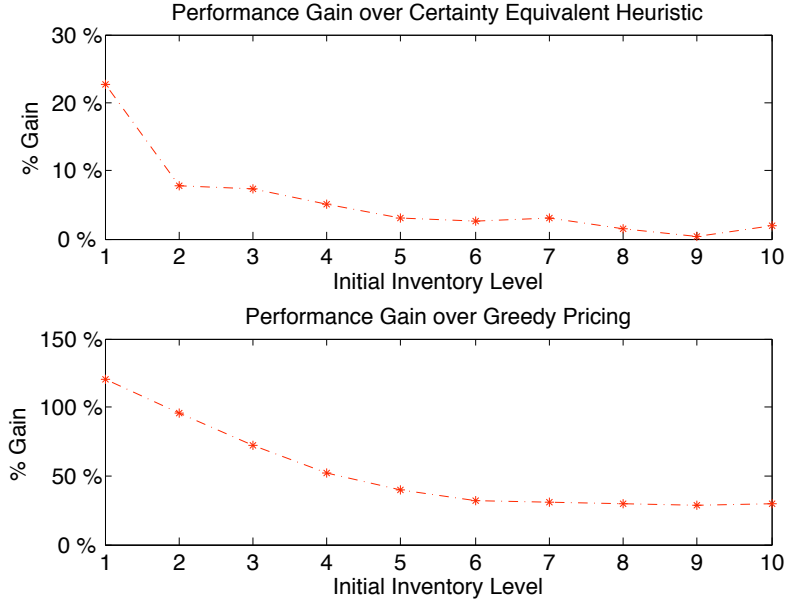


Figure 2: Performance gain over the Certainty Equivalent and Greedy Pricing heuristics at various inventory levels for Gamma mixture prior (right)

pricing heuristics respectively. These experiments indicate that the potential gain from using decay balancing increases with increasing uncertainty in λ , and that in fact the gain over certainty equivalence can be as much as a factor of 1.3, and that over the greedy policy can be as much as a factor of 4.

3. *Value of ‘active’ learning:* We examine the ‘learning gains’ offered by decay balancing relative to the certainty equivalent and greedy heuristics. In particular we do this by comparing the gains offered by all three pricing heuristics relative to a naive no-learning heuristic which, at every point in time, prices assuming that the arrival rate λ is equal to the initial prior mean. We observe that these gains are higher at lower initial levels of uncertainty. As before, we consider a gamma prior with mean 40 and co-efficient of variation 5 and measure performance for a range on starting inventory levels for each of the three heuristics as also the no learning heuristic. We define the learning gain relative to the certainty equivalent heuristic as $100 \times \frac{J^{\pi_{db}} - J^{\pi_{nl}}}{(J^{\pi_{ce}} - J^{\pi_{nl}})^+} - 100\%$ and that relative to the greedy pricing heuristic is defined similarly. Table 3 summarizes these relative learning gains. Since the greedy pricing policy consistently underperforms the no-learning policy, we only report learning gains relative to the certainty equivalent heuristic. In addition to noting that the decay balancing heuristic offers consistent gains relative to the no learning heuristic, the data reported suggests that the decay balancing heuristic likely captures a substantial portion of the gains to be had from using a learning scheme relative to a naive heuristic that does not learn.




Figure 3: Lower bound on maximal performance gain over the Certainty Equivalent Heuristic (left) and Greedy Heuristic (right) for various coefficients of variation

Table 2: Relative Learning Gains Offered By Decay Balancing

Inventory	$x_0 = 1$	$x_0 = 2$	$x_0 = 3$	$x_0 = 4$	$x_0 = 5$	$x_0 = 6$	$x_0 = 7$	$x_0 = 8$	$x_0 = 9$	$x_0 = 10$
vs. π_{ce}	∞	∞	109%	26.2%	27.7%	16.8%	13.1%	10.0%	5.7%	4.4%

To summarize our computational experience with Gamma and Gamma mixture priors, we observe that even at high levels of uncertainty in market response, the decay balancing heuristic offers a level of performance not far from that of a clairvoyant algorithm. Moreover, in all of our computational experiments, the decay balancing heuristic dominates both the certainty equivalent and the greedy pricing heuristics; this is especially so at high levels of uncertainty in the initial prior, and at high ‘load factors’, i.e. scenarios where judicious inventory management is important. One final issue is robustness - in its quest to intelligently leverage uncertainty in market response, the greedy heuristic experiences a drastic loss in performance. With this issue in mind, the following section provides a performance analysis that rules out the drastic performance decay experienced with greedy pricing and sheds some light on the critical determinants of performance for our pricing problem.

6 Bounds on Performance Loss

For the decay balancing price to be a good approximation to the optimal price at a particular state, one requires only a good approximation to the value function at that state (and *not* its derivatives).

This section characterizes the quality of our approximation to J^* and uses such a characterization to ultimately bound the performance loss incurred by decay balancing relative to optimal pricing. Our analysis will focus primarily on the case of a Gamma prior and exponential reservation prices (although we will also provide performance guarantees for other types of reservation price distributions). A rigorous proof of the existence and uniqueness of a solution to the HJB equation for this case may be found in the appendix. We will show that in this case, decay balancing captures at least 33.3% of the expected revenue earned by the optimal algorithm for all choices of $x_0 > 1, a_0 > 0, b_0 > 0, \alpha > 0$ and $r > 0$ when reservation prices are exponentially distributed with mean $r > 0$. Such a bound is an indicator of robustness across all parameter regimes. Decay balancing is the first heuristic for problems of this type for which a uniform performance guarantee is available.

Before we launch into the proof of our performance bound, we present an overview of the analysis. Since our analysis will focus on a gamma prior we will suppress the state variable w in our notation, and a and b will be understood to be scalars. Without loss of generality, we will restrict attention to problems with $\alpha = e^{-1}$; in particular, the value function exhibits the following invariance where the notation $J^{\pi, \alpha}$ makes the dependence of the value function on α explicit:

Lemma 5. *Let $\pi : \mathcal{S} \rightarrow \mathbb{R}_+$ be an arbitrary policy and let $\pi' : \mathcal{S} \rightarrow \mathbb{R}_+$ be defined according to $\pi'(x, a, b) = \pi(x, a, b/\alpha)$. Then, for all $z \in \mathcal{S}, \alpha > 0$, $J^{\pi, \alpha}(z) = J^{\pi', 1}(x, a, \alpha b)$, and, in particular, $J^{*, \alpha}(z) = J^{*, 1}(x, a, \alpha b)$.*

Via the above Lemma, we see that any performance bound established for a heuristic assuming a discount factor e^{-1} applies to other discount factors $\alpha > 0$ as well. In particular, given a heuristic policy $\tilde{\pi}$ designed for discount factor e^{-1} and satisfying $J^{\tilde{\pi}, e^{-1}}(z) \leq \beta J^{*, e^{-1}}(z)$ for some $z \in \mathcal{S}$, the above lemma tells us that the policy $\tilde{\pi}'$ defined according to $\tilde{\pi}'(x, a, b) = \tilde{\pi}(x, a, \alpha b/e^{-1})$ satisfies $J^{\tilde{\pi}', \alpha}(x, a, e^{-1}b/\alpha) \leq \beta J^{*, \alpha}(x, a, e^{-1}b/\alpha)$.

As a natural first step, we attempt to find upper and lower bounds on $\pi_{\text{db}}(z)/\pi^*(z)$, the ratio of the decay balancing price in a particular state to the optimal price in that state. We are able to show that $1 \geq J^*(z)/\tilde{J}(z) \geq 1/\kappa(a)$ where $\kappa(\cdot)$ is a certain decreasing function. By then specializing attention to specific reservation price distributions, this suffices to establish that $1/f(\kappa(a)) \leq \pi_{\text{db}}(z)/\pi^*(z) \leq 1$ where f is some increasing, non-negative function dependent upon the reservation price distribution under consideration.¹

By considering a certain system under which revenue is higher than the optimal revenue, we

¹For instance, in the case of exponential reservation price distributions, $f(x) = 1 + \log x$.

then use the bound above and a dynamic programming argument to show that $1/f(\kappa(a)) \leq J^{\pi_{\text{db}}}(z)/J^*(z) \leq 1$ where $J^{\pi_{\text{db}}}(z)$ denotes the expected revenue earned by the decay balancing heuristic starting in state z . If z is a state reached after i sales then $a = a_0 + i > i$, so that the above bound guarantees that the decay balancing heuristic will demonstrate performance that is within a factor of $f(\kappa(i))$ of optimal moving forward after i sales.

Our general performance bound can be strengthened to a uniform bound in the special case of exponential reservation prices. In particular, a coupling argument that uses a refinement of the general bound above along with an analysis of the maximal loss in revenue up to the first sale for exponential reservation prices, establishes the uniform bound $1/3 \leq J^{\pi_{\text{db}}}(z)/J^*(z) \leq 1$.

We begin our proof with a simple dynamic programming result that we will have several opportunities to use. The proof is a consequence of Dynkin's formula and can be found in the appendix.

Lemma 6. *Let $J \in \mathcal{J}$ satisfy $J(0, a, b) = 0$. Let $\tau = \inf\{t : J(z_t) = 0\}$. Let $z_0 \in \mathcal{S}_{\bar{x}, \bar{a}, \bar{b}}$. Then,*

$$E \left[\int_0^\tau e^{-\alpha t} H^\pi J(z_t) dt \right] = J^\pi(z_0) - J(z_0)$$

Let $J : \mathbb{N} \rightarrow \mathbb{R}$ be bounded and satisfy $J(0) = 0$. Let $\tau = \inf\{t : J(x_t) = 0\}$. Let $x_0 \in \mathbb{N}$. Then,

$$E \left[\int_0^\tau e^{-\alpha t} H_\lambda^\pi J(x_t) dt \right] = J_\lambda^\pi(x_0) - J(x_0)$$

6.1 Decay Balancing Versus Optimal Prices

As discussed in the preceding outline, we will establish a lower bound on $J^*(z)/\tilde{J}(z)$ in order to establish a lower bound on $\pi_{\text{db}}(z)/\pi^*(z)$. Let $J^{nl}(z)$ be the expected revenue garnered by a pricing scheme that does not learn, upon starting in state z . Delaying a precise description of this scheme for just a moment, we will have $J^{nl}(z) \leq J^*(z) \leq \tilde{J}(z) \leq J_{a/b}^*(x)$. It follows that $J^{nl}(z)/J_{a/b}^*(x) \leq J^*(z)/\tilde{J}(z)$, so that a lower bound on $J^{nl}(z)/J_{a/b}^*(x)$ is also a lower bound on $J^*(z)/\tilde{J}(z)$. We will focus on developing a lower bound on $J^{nl}(z)/J_{a/b}^*(x)$.

Upon starting in state z , the 'no-learning' scheme assumes that $\lambda = a/b = \mu$ and does not update this estimate over time. Assuming we begin with a prior of mean μ , such a scheme would use a pricing policy given implicitly by:

$$(6) \quad \pi^{nl}(z) = \pi_\mu^*(x) = 1/\rho(\pi_\mu^*(x)) + J_\mu^*(x) - J_\mu^*(x-1).$$

Using the definition of $H_\lambda^{\pi^{nl}}$ and the fact that $H_\mu J_\mu^* = 0$, some simplification yields

$$H_\lambda^{\pi^{nl}} J_\mu^*(x) = (\lambda/\mu - 1)\alpha J_\mu^*(x).$$

The following two results are then essentially immediate consequences of Lemma 6; proofs can be found in the appendix.

Lemma 7. *If $\lambda < \mu$, $J_\lambda^{\pi^{nl}}(x) \geq (\lambda/\mu)J_\mu^*(x)$ for all $x \in \mathbb{N}$.*

Lemma 8. *If $\lambda \geq \mu$, $J_\lambda^{\pi^{nl}}(x) \geq J_\mu^*(x)$ for all $x \in \mathbb{N}$.*

Armed with these two results we can establish a lower bound on $J^{nl}(z)/J_{a/b}^*(x)$:

Theorem 1. *For all $z \in \mathcal{S}$,*

$$\frac{J^{nl}(z)}{J_{a/b}^*(x)} \geq \frac{\Gamma(a+1) - \Gamma(a+1, a) + a\Gamma(a, a)}{a\Gamma(a)} \equiv 1/\kappa(a)$$

Proof: Setting $\mu = a/b$, we have:

$$\begin{aligned} J^{nl}(z) &= E_\lambda \left[J_\lambda^{\pi^{nl}}(x) \right] \\ &\geq E_\lambda \left[1_{\lambda < \mu} \lambda/\mu J_\mu^*(x) + 1_{\lambda \geq \mu} J_\mu^*(x) \right] \\ &= \frac{\Gamma(a+1) - \Gamma(a+1, a) + a\Gamma(a, a)}{a\Gamma(a)} J_\mu^*(x) \end{aligned}$$

where the inequality follows from the two preceding Lemmas and the equality by direct integration of the Gamma(a, b) density. $\Gamma(\cdot, \cdot)$ is the incomplete Gamma function and is given by $\Gamma(x, y) = \int_y^\infty s^{x-1} e^{-s} ds$. \square

The decay balance equation allows one to use the above bound on the quality of our approximation \tilde{J} to compute bounds on the decay balance price relative to the optimal price at a given state. In particular, Corollary 1 establishes such bounds for exponential and logit reservation price distributions; see the appendix for bounds one may establish for general classes of reservation price distributions.

Corollary 1. *For all $z \in \mathcal{S}$, and exponential reservation price distributions with parameter r :*

$$\frac{1}{1 + \log \kappa(a)} \leq \frac{\pi_{\text{db}}(z)}{\pi^*(z)} \leq 1$$

For all $z \in \mathcal{S}$, and logit reservation price distributions with parameter r :

$$\frac{1.27}{1.27 + \log \kappa(a)} \leq \frac{\pi_{\text{db}}(z)}{\pi^*(z)} \leq 1$$

6.2 An Upper Bound on Performance Loss

We next establish a lower bound on $J^{\pi_{\text{db}}}(z)/J^*(z)$ that will depend on the coefficient of variation of the prior on λ , $1/\sqrt{a}$.

Let

$$R^{\text{db}}(z) = \sum_{k:t_k \leq \tau} e^{-e^{-1}t_k} \pi_{\text{db}}(z_{t_k}^-)$$

be the revenue under the decay balancing policy for a particular sample path of the sales process, starting in state z and define

$$R^{\text{ub}}(z) = \sum_{k:t_k \leq \tau} e^{-e^{-1}t_k} \pi^*(z_{t_k}^-).$$

This describes a system whose state evolution is identical to that under the decay balancing policy but whose revenues on a sale correspond to those that would be earned if the price set prior to the sale was that of the optimal pricing algorithm.

Of course, $J^{\pi_{\text{db}}}(z) = E_z[R^{\text{db}}(z)]$. Define $J^{\text{ub}}(z) = E_z[R^{\text{ub}}(z)]$, where the expectation is over $\{t_k\}$ and assumes that an arriving consumer at time t_k makes a purchase with probability $\bar{F}(\pi_{\text{db}}(z_{t_k}))$. That is, the expectation is understood to be according to the dynamics of the system controlled by π_{db} . The following result should be intuitive given our construction of the upper-bounding system and the fact that since $\pi_{\text{db}}(z) \leq \pi^*(z)$, the probability that a customer arriving in state z chooses to purchase is higher in the system controlled by the decay balancing policy. The proof uses a dynamic programming argument and may be found in the appendix.

Lemma 9. *For all $z \in \mathcal{S}$, and reservation price distributions satisfying Assumptions 1 and 2,*

$$J^{\text{ub}}(z) \geq J^*(z)$$

Now observe that since $\kappa(a)$ is decreasing in a , we have from Corollary 1 that

$$\frac{1}{f(\kappa(a))} \leq \frac{R^{\text{db}}(z)}{R^{\text{ub}}(z)} \leq 1$$

where $f(x) = 1 + \log(x)$ for exponential reservation price distributions and $f(x) = 1 + \log(x)/1.27$ for Logit distributions. Taking expectations, and employing Lemma 9, we then immediately have:

Theorem 2. For all $z \in \mathcal{S}$, and exponential reservation price distributions with parameter r

$$\frac{1}{1 + \log \kappa(a)} \leq \frac{J^{\pi_{\text{db}}}(z)}{J^*(z)} \leq 1,$$

while for logit reservation price distributions with parameter r

$$\frac{1.27}{1.27 + \log \kappa(a)} \leq \frac{J^{\pi_{\text{db}}}(z)}{J^*(z)} \leq 1.$$

See Figure 4 for an illustration of these bounds.




Figure 4: Lower bound on Decay Balancing performance

It is worth pausing to reflect on the bound we have just established. Our performance bound does *not* depend on x , b , α or the parameters of the reservation price distribution. Since the coefficient of variation of a Gamma prior with parameters a and b is given by $1/\sqrt{a}$, the bound illustrates how decay balancing performance approaches optimal performance as the coefficient of variation of the initial prior is decreased; below co-efficients of variation of 0.5, the bound guarantees performance levels that are at least within 80% of optimal. Nonetheless, the bound can be arbitrarily poor at high co-efficients of variation. Next, we further specialize our attention to exponential reservation price distributions and present a uniform performance guarantee for that case.

6.3 A Uniform Performance Bound for Exponential Reservation Prices

We now consider the case of exponentially distributed prices. In particular, we have $\bar{F}(p) = \exp(-p/r)$ where $r > 0$. In light of the following Lemma (where the notation $J^{\pi,\alpha,r}$ makes the dependence of the value function on α and r explicit), we can assume without loss that the mean reservation price, r , is 1:

Lemma 10. *Let $\pi : \mathcal{S} \rightarrow \mathbb{R}_+$ be an arbitrary policy and let $\pi' : \mathcal{S} \rightarrow \mathbb{R}_+$ be defined according to $\pi'(z) = (1/r)\pi(z)$. Then, for all $z \in \mathcal{S}, \alpha > 0, r > 0$, $J^{\pi,\alpha,r}(z) = rJ^{\pi',\alpha,1}(z)$ and, in particular, $J^{*,\alpha,r}(z) = rJ^{*,\alpha,1}(z)$.*

Our proof of a uniform performance bound will use Theorem 2 along with a coupling argument to bound performance loss up to the time of the first sale. Begin by considering the following coupling (A superscript ‘db’ on a variable indicates that the variable is relevant to a system controlled by the π_{db} policy): For an arbitrary policy $\pi(\cdot) \in \Pi$, the sales processes n_t^{db} and n_t^π are coupled in the following sense: Denote by $\{t_k\}$ the points of the Poisson process corresponding to customer arrivals (not sales) to both systems. Assume $\pi_{\text{db}t_k^-} \leq \pi_{t_k^-}$. Then a jump in n^π at time t_k can occur if and only if a jump occurs in n^{db} at time t_k . Further, conditioned on a jump in n^{db} at t_k , the jump in n^π occurs with probability $\exp(-(\pi_{t_k^-} - \pi_{\text{db}t_k^-}))$. The situation is reversed if $\pi_{\text{db}t_k^-} > \pi_{t_k^-}$. Let τ denote the time of the first sale for the π_{db} system i.e. $\tau = \inf\{t : n_t^{\text{db}} = 1\}$. In the context of this coupling consider the optimal (i.e. π^*) and π_{db} controlled systems. Denoting

$$J^*(z|\tau) = E \left[\int_{t=0}^{\infty} e^{-\alpha t} \pi^*(z_t) \lambda \bar{F}(\pi^*(z_t)) dt \mid \tau, z_0 = z \right],$$

we then have:

Lemma 11. *For all $z \in \mathcal{S}$,*

$$J^*(z|\tau) \leq e^{-e^{-1}\tau} \left(e^{-(\pi^* - \pi_{\text{db}})} \left[\pi^* + J^*(x-1, a+1, b_\tau^{\text{db}}) \right] + (1 - e^{-(\pi^* - \pi_{\text{db}})}) J^*(x, a+1, b_\tau^{\text{db}}) \right)$$

where $\pi^* = \pi^*(x, a, b_\tau^*)$ and $\pi_{\text{db}} = \pi_{\text{db}}(x, a, b_\tau^{\text{db}})$.

The result above is essentially a consequence of the fact that it is never the case that the π^* controlled system sells its first item before the π_{db} system, and moreover, that conditioning on τ , and the information available in both systems up to τ^- yields a posterior with shape parameter $a+1$ and scale parameter b_τ^{db} . We will also find the following technical Lemma useful:

Lemma 12. For $x > 1, a > 1, b > 0$, $J^*(x, a, b) \leq 2.05J^*(x - 1, a, b)$.

The result above is intuitive; it would follow, for example, from decreasing returns to an additional unit of inventory. Unfortunately, we aren't able to show such a 'decreasing returns' property directly and a certain coupling argument is used to prove the Lemma (see the appendix). We are now poised to prove a uniform (over $x > 1$) performance bound for our pricing scheme:

Theorem 3. For all $z \in \mathcal{S}$ with $x > 1$,

$$\frac{J^{\pi_{\text{db}}}(z)}{J^*(z)} \geq 1/3.$$

Proof: In Lemma 11 we showed:

$$J^*(z) \leq E \left[e^{-e^{-1}\tau} \left(e^{-(\pi^* - \pi_{\text{db}})} \left[\pi^* + J^*(x - 1, a + 1, b_\tau^{\text{db}}) \right] + (1 - e^{-(\pi^* - \pi_{\text{db}})}) J^*(x, a + 1, b_\tau^{\text{db}}) \right) \right]$$

Now,

$$\begin{aligned} & e^{-e^{-1}\tau} \left(e^{-(\pi^* - \pi_{\text{db}})} \left[\pi^* + J^*(x - 1, a + 1, b_\tau^{\text{db}}) \right] + (1 - e^{-(\pi^* - \pi_{\text{db}})}) J^*(x, a + 1, b_\tau^{\text{db}}) \right) \\ & \leq e^{-e^{-1}\tau} \left(e^{-(\pi^* - \pi_{\text{db}})} \pi^* + J^*(x, a + 1, b_\tau^{\text{db}}) \right) \\ & \leq e^{-e^{-1}\tau} \left(e^{-(\pi^* - \pi_{\text{db}})} \pi^* + 2.05 J^*(x - 1, a + 1, b_\tau^{\text{db}}) \right) \\ & \leq e^{-e^{-1}\tau} \left(\pi_{\text{db}} + 2.05(1 + \log \kappa(a + 1)) J^{\pi_{\text{db}}}(x - 1, a + 1, b_\tau^{\text{db}}) \right) \\ & \leq e^{-e^{-1}\tau} 2.05(1 + \log \kappa(a + 1)) \left(\pi_{\text{db}} + J^{\pi_{\text{db}}}(x - 1, a + 1, b_\tau^{\text{db}}) \right) \end{aligned}$$

where the first inequality is because J^* is non-decreasing in x . The second inequality follows from Lemma 12. The third inequality follows from the fact that $\pi^* \geq \pi_{\text{db}} \geq 1$ so that $\pi^* e^{-\pi^*} \leq \pi_{\text{db}} e^{-\pi_{\text{db}}}$ and from Theorem 2. Finally, taking expectations of both sides we get:

$$\begin{aligned} & E \left[e^{-e^{-1}\tau} \left(e^{-(\pi^* - \pi_{\text{db}})} \left[\pi^* + J^*(x - 1, a + 1, b_\tau^{\text{db}}) \right] + (1 - e^{-(\pi^* - \pi_{\text{db}})}) J^*(x, a + 1, b_\tau^{\text{db}}) \right) \right] \\ & \leq 2.05(1 + \log \kappa(a + 1)) E \left[e^{-e^{-1}\tau} \left(\pi_{\text{db}} + J^{\pi_{\text{db}}}(x - 1, a + 1, b_\tau^{\text{db}}) \right) \right] \\ & \leq 2.05(1 + \log \kappa(1)) J^{\pi_{\text{db}}}(z) \end{aligned}$$

Thus,

$$\frac{J^{\pi_{\text{db}}}(z)}{J^*(z)} \geq \frac{1}{2.05(1 + \log \kappa(1))} \geq 1/3.$$

□

Perhaps the most crucial point borne out in the performance analysis we have just presented is that the decay balancing heuristic is *robust*; our analysis precludes the drastic performance breakdown observed for the greedy pricing heuristic in our computational experiments.

7 Multiple Stores and Consumer Segments

We now explore extensions of decay balancing to a model with multiple stores and consumer segments. We do not attempt to extend our performance analysis to this more general model but instead present numerical experiments, the goal being to show that decay balancing demonstrates the same qualitative behavior as in the one store, one customer segment case we have studied to this point.

More formally, we consider a model with N stores and M consumer segments. Each store is endowed with an initial inventory $x_{0,i}$ for $i \in \{1, \dots, N\}$. Customers from class j , for $j \in \{1, \dots, M\}$ arrive according to a Poisson process of rate λ_j where λ_j is a Gamma distributed random variable with shape parameter $a_{0,j}$ and scale parameter $b_{0,j}$. An arriving segment j customer considers visiting a single store and will consider store i with probability α_{ij} . Consequently, each store i sees a Poisson stream of customers having rate $\sum_j \alpha_{ij} \lambda_j$. We assume without loss of generality that $\sum_i \alpha_{ij} = 1$. We assume that customers in each segment have exponential reservation price distributions with mean r and moreover that upon a purchase the store has a mechanism in place to identify what segment the purchasing customer belongs to.

Let $p_t \in \mathbb{R}^N, t \in [0, \infty)$ represent the process of prices charged at the stores over time. Let $n_{t,i}^j$ represent the total number of type j customers served at store i up to time t and let $n_t^j = \sum_i n_{t,i}^j$. The parameter vectors a and b are then updated according to:

$$a_{t,j} = a_{0,j} + n_{t,j} \quad \text{and} \quad b_{t,j} = b_{0,j} + \int_{\tau=0}^t \sum_i e^{-p_{\tau,i}/r} d\tau$$

Our state at time t is now $z_t = (x_t, a_t, b_t)$. As before, we will consider prices generated by policies π that are measurable, non-negative vector-valued functions of state, so that $p_t = \pi(z_t) \geq 0$. Letting Π denote the set of all such policies, our objective will be to identify a policy $\pi^* \in \Pi$ that maximizes

$$\hat{J}^\pi(z) = E_{z,\pi} \left[\sum_i \int_0^{\tau^i} \hat{\rho}_{t,i} e^{-p_{t,i}/r} dt \right]$$

where $\tau^i = \inf\{t : \sum_j n_{t,i}^j = x_{0,i}\}$ and $\hat{\rho}_i = \sum_j \alpha_{i,j}(a_j/b_j)$. We define the operator

$$(\tilde{H}^\pi J)(z) = \sum_i \left[\hat{\rho}_i e^{-\pi(z)_i/r} \left(\pi(z)_i + \sum_j \frac{\alpha_{i,j}(a_j/b_j)}{\hat{\rho}_i} J(x - e_i, a + e_j, b) - J(z) \right) + \sum_j e^{-\pi(z)_i/r} \frac{d}{db_j} J(z) \right] - \alpha J(z).$$

where e_k is the vector that is 1 in the k th coordinate and 0 in other coordinates. One may show that $\hat{J}^* = \hat{J}^{\pi^*}$ is the unique solution to

$$\sup_{\pi \in \Pi} (\tilde{H}^\pi J)(z) = 0 \quad \forall z$$

satisfying $\hat{J}^*(0, a, b) = 0$, and that the corresponding optimal policy for $x_i > 0$ is given by

$$(7) \quad (\pi^*(z))_i = r + \hat{J}^*(z) - \sum_j \frac{\alpha_{i,j}(a_j/b_j)}{\hat{\rho}_i} \hat{J}^*(x - e_i, a + e_j, b) - \frac{1}{\hat{\rho}_i} \sum_j \frac{d}{db_j} \hat{J}^*(z).$$

Now, assuming that the λ_j 's are known perfectly a-priori, it is easy to see that the control problem decomposes across stores. In particular, the optimal strategy simply involves store i using as it's pricing policy $p_{t,i} = \pi_{\rho_i}^*(x_{t,i})$, where $\rho_i = \sum_j \alpha_{i,j} \lambda_j$. Consequently, a certainty equivalent policy would use the pricing policy $(\pi_{CE}(z))_i = \pi_{\rho_i}^*(x_i)$.

We can also consider as an approximation to \hat{J}^* , the following upper bound (which is in the spirit of the upper bound we derived in Section 4):

$$\bar{J}(z) = E \left[\sum_i J_{\rho_i}^*(x_i) \right].$$

The analogous greedy pricing policy π_{gp} is then given by (7) upon substituting $\bar{J}(\cdot)$ for $\hat{J}^*(\cdot)$ in that expression.

Motivated by the decay balancing policy derived for the single store case we consider using the following pricing policy at each store:

$$(\pi_{db}(z))_i = r \log \left(\frac{r \hat{\rho}_i}{\alpha E[J_{\rho_i}^*(x_i)]} \right).$$

The above pricing equation assumes that moving forward each store will operate as a separate entity. Nonetheless, the heuristic incorporates joint learning across stores and continues to account for the level of uncertainty in market size in the pricing process. Further, the structural properties




Figure 5: Performance relative to a clairvoyant algorithm

discussed in the single store case are retained. The joint learning under this heuristic does not however account for inventory levels across stores.

We now present computational results for the three heuristics. Our experiments will use the following model parameters. We take $N = 2, M = 2$ and assume $\alpha_{i,j} = 1/2$ for all i, j , and further that we begin with prior parameters $a_1 = a_2 = 0.04$, and $b_1 = b_2 = 0.001$ (which corresponds to a mean of 40 and a coefficient of variation of 5). As usual, $\alpha = e^{-1}, r = 1$. Our first set of results (Figure 5) compares the decay balancing heuristic’s performance against that of a clairvoyant algorithm which as in Section 5 has perfect a-priori knowledge of λ . As in the $N = 1, M = 1$ case, our performance is quite close to that of the clairvoyant algorithm. Figure 6 compares decay balancing performance to the certainty equivalent heuristic and the greedy heuristic. Figure 6 is indicative of performance that is qualitatively similar to that observed for the $N = 1, M = 1$ case; there is a significant gain over certainty equivalence at lower inventory levels, but this gain shrinks as inventory level increases. The performance of the greedy heuristic is particularly dismal, one explanation for which is that $\sum_j \frac{d}{db_j} \bar{J}(z)$ is a potentially poor approximation to $\sum_j \frac{d}{db_j} \tilde{J}^*(z)$.

8 Discussion and Conclusions

The dynamic pricing model proposed by Gallego and van Ryzin (1994) is central to a large body of the revenue management literature. This work considered an important extension to that model. In particular, we considered incorporating uncertainty in the customer arrival rate or ‘market response’ which is without doubt important in many industries that practice revenue management.

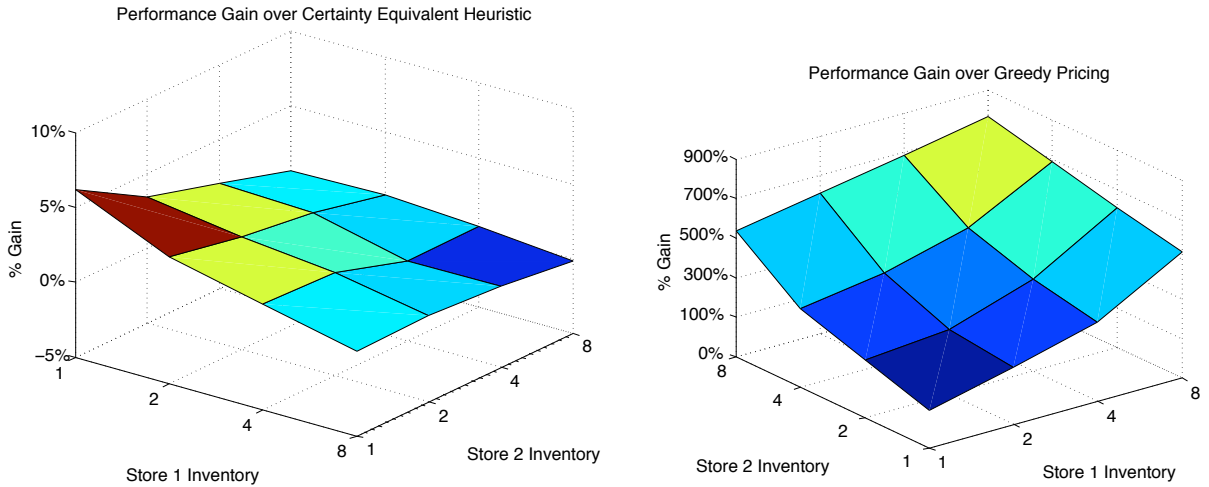


Figure 6: Performance relative to the Certainty Equivalent (left) heuristic and Greedy Pricing heuristic (right)

We proposed and analyzed a simple new heuristic for this problem: decay balancing. Decay balancing is computationally efficient and leverages the solution to problems with no uncertainty in market response. Our computational experiments (which focused on gamma priors and exponentially distributed reservation prices) suggest that decay balancing achieves near-optimal performance even on problems with high levels of uncertainty in market response. Pricing policies generated by decay balancing have the appealing structural property that, all other factors remaining the same, the price in the presence of uncertainty in market response is higher than the corresponding price with no uncertainty. This is reasonable from an operational perspective and is in fact a property possessed by the optimal policy. Our analysis also demonstrated a uniform performance guarantee for decay balancing when reservation prices are exponentially distributed, which is an indicator of robustness.

Two heuristics proposed for problems of this nature prior to our work were the certainty equivalent heuristic (by Aviv and Pazgal (2005a)) and the greedy pricing heuristic (by Araman and Caldentey (2005)). Our computational results suggest that decay balancing offers significant performance advantages over these heuristics. These advantages are especially clear at high levels of uncertainty in market response which is arguably the regime of greatest interest. Decay balancing relies only on a good approximation to the value of an optimal policy at a given state. This is in contrast with greedy pricing that requires not only a good approximation to value but further to derivatives of value with respect to the scale parameter. At the same time, uncertainty in arrival rate and changes in inventory levels bear the appropriate directional impact on decay balancing

prices: uncertainty in the arrival rate calls for higher prices than in corresponding situations with no uncertainty, while a decrease in inventory calls for an increase in prices. In contrast, uncertainty in the arrival rate has no impact on certainty equivalent prices while greedy prices can increase or decrease with decreasing inventory.

Our computational study and performance analysis were focused on exponentially distributed reservation prices and gamma priors, but we expect favorable performance for other distributions as well. In particular, the analysis of Theorem 1 can be extended to Gamma mixture priors yielding encouraging estimates on the quality of approximation provided by \tilde{J} . Since the decay balancing price at state z is likely to be a good approximation to the optimal price at z if $\tilde{J}(z)$ is a good approximation to $J^*(z)$, this suggests that decay balancing is likely to do a good job of approximating the optimal price for general reservation price distributions and priors on arrival rate, which in turn should lead to superior performance.

There is ample room for further work in the general area of pricing with uncertainty in market response and other factors that impact demand. One direction is considering more complex models. We considered applying the decay balancing heuristic to a problem with multiple stores and consumer segments; our computational results there are promising. There are other models one might hope to consider. For example, the multi product model proposed by Gallego and van Ryzin (1997). Another potential direction is exploring new approximations to the value function beyond the approximation considered here and applying such approximations with either the greedy pricing or decay balancing heuristics. Finally, it would be interesting to extend the approaches in this paper to problems with uncertainty in other factors that impact demand such as price elasticity.

Acknowledgments

The first author would like to thank Mike Harrison, Paat Rusmevichientong and Gabriel Weintraub for useful discussions, suggestions, and pointers to the literature. The first author also acknowledges support from the INFORMS MSOM society in the form of a best student paper award for a preliminary version of this work. This research was supported in part by the National Science Foundation through grant IIS-0428868.

References

- Araman, V., R. Caldentey. 2005. Dynamic pricing for non-perishable products with demand learning. *Submitted* .

- Aviv, Y., A. Pazgal. 2005a. A partially observable markov decision process for dynamic pricing. *Working paper* .
- Aviv, Y., A. Pazgal. 2005b. Pricing of short life-cycle products through active learning. *Management Science* **51**(9) 1400–1416.
- Bertsimas, D., G. Perakis. 2003. Dynamic pricing: A learning approach. *to appear in Models for Congestion Charging/Network Pricing* .
- Besbes, O., A. Zeevi. 2006. Blind nonparametric revenue management: Asymptotic optimality of a joint learning and pricing method. *Submitted* .
- Bremaud, P. 1981. *Point Processes and Queues: Martingale Dynamics*. 1st ed. Springer-Verlag.
- Burnetas, A. N., C. E. Smith. 1998. Adaptive ordering and pricing for perishable products. *Operations Research* **48**(3) 436–443.
- Cope, E. 2006. Bayesian strategies for dynamic pricing in e-commerce. *Naval Research Logistics* **54**(3) 265–281.
- Dalal, S. R., W. J. Hall. 1983. Approximating priors by mixtures of natural conjugate priors. *J. R. Statist. Soc. B.* **45**(2) 278–286.
- Gallego, G., G. van Ryzin. 1994. Optimal dynamic pricing of inventories with stochastic demand over finite horizons. *Management Science* **40**(8) 999–1020.
- Gallego, G., G. van Ryzin. 1997. A multiproduct dynamic pricing problem and its applications to network yield management. *Operations Research* **45**(1) 24–41.
- Kleinberg, R., F. T. Leighton. 2004. The value of knowing a demand curve: Bounds on regret for on-line posted-price auctions. *Proceedings of the 44th IEEE Symposium on Foundations of Computer Science (FOCS 2003)*. 594–605.
- Lin, K. Y. 2007. Dynamic pricing with real-time demand learning. *European Journal of Operational Research* (*forthcoming*) .
- Lobo, M. S., S. Boyd. 2003. Pricing and learning with uncertain demand. *Submitted* .
- Talluri, Kalyan T., Garrett J. van Ryzin. 2004. *The Theory and Practice of Revenue Management*. Springer Science+Business Media.

Appendix

Results in this appendix are numbered consistently with those in the main paper. Results that do not appear in the paper (auxiliary Lemmas or additional theorems omitted from the exposition in the main paper) are numbered using the convention ‘SectionLetter.Number’ (eg. Theorem E.1).

We recall the following assumptions in several proofs that follow and so find it convenient to repeat them here.

Assumption 1.

1. $F(\cdot)$ has a differentiable density $f(\cdot)$ with support \mathbb{R}^+ .
2. F has a non-decreasing hazard rate. That is, $\rho(p) = \frac{f(p)}{F(p)}$ is non-decreasing in p .

Assumption 2. $J_\lambda^*(x)$ is a differentiable function of λ on \mathbb{R}^+ for all $x \in \mathbb{N}$.

A Proofs for Section 3

Lemma 1. $\pi_\lambda^*(x)$ is decreasing in x (on \mathbb{N}) and non-decreasing in λ (on \mathbb{R}_+).

Proof: We find it convenient to prove the following sub-homogeneity property for $J_\lambda^*(x)$ viewed as function of λ : For $\lambda_2 \geq \lambda_1 > 0$, $J_{\lambda_2}^*(x) \leq \frac{\lambda_2}{\lambda_1} J_{\lambda_1}^*(x)$. To see this, consider a system beginning with x units of inventory facing arrivals at rate λ_2 . Every arrival to the system is marked as either ‘real’ or ‘fictitious’ with probability $\frac{\lambda_1}{\lambda_2}$ and $1 - \frac{\lambda_1}{\lambda_2}$ respectively, independent of all other arrivals. Consider using the pricing policy $\pi_{\lambda_2}^*(\cdot)$, and denote by $J_{\lambda_2}^{*,f}(x)$ and $J_{\lambda_2}^{*,r}(x)$ the expected revenues earned under this policy from sales to arrivals marked as fictitious and real respectively. By construction, we have $J_{\lambda_2}^*(x) = J_{\lambda_2}^{*,f}(x) + J_{\lambda_2}^{*,r}(x)$ and further, $J_{\lambda_2}^{*,r}(x) = \frac{\lambda_1}{\lambda_2} J_{\lambda_1}^*(x)$. But $J_{\lambda_2}^{*,r}(x)$ is the expected revenue earned under a randomized non-anticipatory policy for a system beginning with x units of inventory and arrival rate λ_1 , so that $J_{\lambda_2}^{*,r}(x) \leq J_{\lambda_1}^*(x)$. Thus $\frac{\lambda_1}{\lambda_2} J_{\lambda_2}^*(x) \leq J_{\lambda_1}^*(x)$ which is the inequality we require.

We now turn to the proof of the Lemma. We have from the HJB equation for the case of a known arrival rate and $x > 0$:

$$\frac{\alpha J_\lambda^*(x)}{\lambda} = \sup_p \bar{F}(p)(p + J_\lambda^*(x-1) - J_\lambda^*(x))$$

Now $\frac{\alpha J_\lambda^*(x)}{\lambda}$ is trivially increasing in x . $\frac{\alpha J_\lambda^*(x)}{\lambda}$ is non-increasing in λ by the inequality we have just shown (i.e. since $J_\lambda^*(x)$ is a sub-homogenous function of λ). Further, observe that $\sup_p \bar{F}(p)(p - c)$ is decreasing in c . It follows that $J_\lambda^*(x) - J_\lambda^*(x-1)$ is decreasing in x and non-decreasing in λ . But, $\pi_\lambda^*(x) - \frac{1}{\rho(\pi_\lambda^*(x))} = J_\lambda^*(x) - J_\lambda^*(x-1)$ and $p - 1/\rho(p)$ is an increasing function of p by Assumption 1. The claim follows. \square

Lemma 2. For all $x \in \mathbb{N}$, $J_\lambda^*(x)$ is an increasing, concave function of λ on \mathbb{R}_+ .

Proof: Consider two systems with $\lambda_1 < \lambda_2$. We will show that $\frac{d}{d\lambda} J_\lambda^*(x)|_{\lambda=\lambda_1} \geq \frac{d}{d\lambda} J_\lambda^*(x)|_{\lambda=\lambda_2}$. Delaying a proof until later in our argument, we have:

$$(1) \quad \frac{d}{d\lambda} J_\lambda^*(x) \Big|_{\lambda=\bar{\lambda}} = E \left[T_\alpha \pi_{\bar{\lambda}}^*(x_{T_\alpha}) \bar{F} \left(\pi_{\bar{\lambda}}^*(x_{T_\alpha}) \right) \right]$$

where T_α is exponentially distributed with mean $1/\alpha$. Now, the instantaneous rate at which a sale occurs in a system with arrival rate λ and x units of inventory on hand is given by $\lambda \bar{F}(\pi_\lambda^*(x)) = \lambda \frac{\alpha J_\lambda^*(x) \rho(\pi_\lambda^*(x))}{\lambda} = \alpha J_\lambda^*(x) \rho(\pi_\lambda^*(x))$, which is an increasing function of λ , since $\pi_\lambda^*(x)$ and $J_\lambda^*(x)$ are increasing functions of λ (see Lemma 1) and $\rho(\cdot)$ is a non-decreasing function by Assumption 1. Thus, letting $x_{T_\alpha}^{\lambda_i}$ be the inventory on hand at time T_α in the i th system (for $i = 1, 2$), we must have that $x_{T_\alpha}^{\lambda_1}$ stochastically dominates $x_{T_\alpha}^{\lambda_2}$. We consequently have:

$$\begin{aligned} \frac{d}{d\lambda} J_\lambda^*(x) \Big|_{\lambda=\lambda_2} &= E \left[T_\alpha \pi_{\lambda_2}^*(x_{T_\alpha}^{\lambda_2}) \bar{F} \left(\pi_{\lambda_2}^*(x_{T_\alpha}^{\lambda_2}) \right) \right] \\ &\leq E \left[T_\alpha \pi_{\lambda_2}^*(x_{T_\alpha}^{\lambda_1}) \bar{F} \left(\pi_{\lambda_2}^*(x_{T_\alpha}^{\lambda_1}) \right) \right] \\ &\leq E \left[T_\alpha \pi_{\lambda_1}^*(x_{T_\alpha}^{\lambda_1}) \bar{F} \left(\pi_{\lambda_1}^*(x_{T_\alpha}^{\lambda_1}) \right) \right] \\ &= \frac{d}{d\lambda} J_\lambda^*(x) \Big|_{\lambda=\lambda_1} \end{aligned}$$

The first inequality follows from the fact that $\pi_\lambda^*(x)$ is decreasing in x by Lemma 1 and since $p\bar{F}(p)$ is decreasing in p for $p \geq p^*$ (the static revenue maximizing price). The second inequality follows from the fact that $\pi_\lambda^*(x)$ is increasing in λ by Lemma 1 and since $p\bar{F}(p)$ is decreasing in p for $p \geq p^*$. That $p\bar{F}(p)$ is decreasing in p for $p \geq p^*$ follows from the fact that $\frac{d}{dp} p\bar{F}(p) = f(p)(1/\rho(p) - p)$ which by Assumption 1 is negative for $p > p^*$ and 0 at $p = p^*$.

That $J_\lambda^*(x)$ is increasing in λ follows from the positivity of the right hand side in (1).

We now establish the equality (1). Consider a system with arrival rate λ . The expected revenue from this system is equal to the expected revenue from an un-discounted system, where after a random time $T_\alpha \sim \exp(1/\alpha)$, no revenues are recorded. This can be seen by simply noting that the HJB equations for the respective problems are identical and given by

$$(2) \quad \alpha J_\lambda^*(x) = \begin{cases} \sup_{p \geq 0} \lambda \bar{F}(p) (p + J_\lambda^*(x-1) - J_\lambda^*(x)) & \text{if } x > 0 \\ 0 & \text{otherwise.} \end{cases}$$

In particular,

$$(3) \quad J_\lambda^*(x) = E \left[\int_0^{T_\alpha} \pi_\lambda^*(x_t) \bar{F}(\pi_\lambda^*(x_t)) \lambda dt \Big| x_0 = x \right]$$

Next, we observe that increasing λ to $\lambda + \delta$ is equivalent to decreasing α to $\alpha(\frac{\lambda}{\lambda + \delta})$. That is,

$$J_{\lambda+\delta}^{*,\alpha}(x) = J_\lambda^{*,\alpha(\frac{\lambda}{\lambda+\delta})}(x)$$

which is immediate from the HJB equation for a known arrival rate. This in turn is equivalent to

increasing T_α on each sample path to $T_\alpha(1 + \delta/\lambda)$. In particular, we have:

$$\begin{aligned}
(4) \quad J_{\lambda+\delta}^{*,\alpha}(x) &= E \left[\int_0^\infty \pi_{\lambda+\delta}^*(x_t) \bar{F}(\pi_{\lambda+\delta}^*(x_t)) (\lambda + \delta) \exp(-\alpha t) dt \middle| x_0 = x \right] \\
&= E \left[\int_0^\infty \pi_{\lambda+\delta}^*(x_t) \bar{F}(\pi_{\lambda+\delta}^*(x_t)) \lambda \exp\left(-\alpha \frac{\lambda}{\lambda + \delta} t\right) dt \middle| x_0 = x \right] \\
&= E \left[\int_0^{T_\alpha(1+\delta/\lambda)} \pi_{\lambda+\delta}^*(x_t) \bar{F}(\pi_{\lambda+\delta}^*(x_t)) \lambda dt \middle| x_0 = x \right]
\end{aligned}$$

where the second equality follows by noting that the optimal policy for the system with arrival rate $\lambda + \delta$ and discount factor α is identical to that for the system with arrival rate λ and discount factor $\alpha \frac{\lambda}{\lambda + \delta}$ which in turn follows from the fact that the HJB equations for the two systems are identical. The third equality follows as in (3).

Now $\pi_\lambda^*(x)$ is a differentiable function of λ for all $x \in \mathbb{N}$. To see this we note that $\pi_\lambda^*(x)$ is given implicitly by

$$\pi_\lambda^*(x) = 1/\rho(\pi_\lambda^*(x)) + J_\lambda^*(x) - 1_{x>0} J_\lambda^*(x-1).$$

Since $p - 1/\rho(p)$ is increasing on $p \geq 0$ with \mathbb{R}^+ in its range (and therefore invertible on \mathbb{R}^+) and differentiable in p (all of which follows from Assumptions 1) and since $J_\lambda^*(x)$ was assumed differentiable in λ (Assumption 2) we may invoke the Inverse Function Theorem to conclude that $\pi_\lambda^*(x)$ is a differentiable function of λ on \mathbb{R}^+ .

Let x'_t denote the inventory on hand at time t in an optimally controlled system with arrival rate $\lambda + \delta$. Let us couple the sales processes in the systems with arrival rate λ and $\lambda + \delta$ as follows: assume the prevailing prices in the two systems are p and p' respectively. If $\lambda \bar{F}(p) \leq (\lambda + \delta) \bar{F}(p')$ then the system with arrival rate λ will witness its next sale no sooner than the system with arrival rate $\lambda + \delta$; the next sale to the system with arrival rate $\lambda + \delta$ will arrive at rate $(\lambda + \delta) \bar{F}(p')$ and will constitute a sale in the system with arrival rate λ with probability $\lambda \bar{F}(p) / ((\lambda + \delta) \bar{F}(p'))$. The situation is reversed if $(\lambda + \delta) \bar{F}(p') < \lambda \bar{F}(p)$. By the continuity of π_λ^* in λ , we have $x'_{T_\alpha} \rightarrow x_{T_\alpha}$ in probability under this coupling. Then, by the Cauchy-Schwarz inequality,

$$|E[T_\alpha \pi_{\lambda+\delta}^*(x_{T_\alpha}) \bar{F}(\pi_{\lambda+\delta}^*(x_{T_\alpha}))] - E[T_\alpha \pi_{\lambda+\delta}^*(x'_{T_\alpha}) \bar{F}(\pi_{\lambda+\delta}^*(x'_{T_\alpha}))]| \leq 2 \sup_p p \bar{F}(p) \sqrt{\Pr(x'_{T_\alpha} \neq x_{T_\alpha}) E[T_\alpha^2]},$$

where $\Pr(\cdot)$ is the joint measure induced by our coupling. Since $\sup_p p \bar{F}(p) < \infty$ by Assumption 1, we thus have:

$$E[T_\alpha \pi_{\lambda+\delta}^*(x_{T_\alpha}) \bar{F}(\pi_{\lambda+\delta}^*(x_{T_\alpha}))] - E[T_\alpha \pi_{\lambda+\delta}^*(x'_{T_\alpha}) \bar{F}(\pi_{\lambda+\delta}^*(x'_{T_\alpha}))] \rightarrow 0$$

Again, via the continuity of π_λ^* in λ , the dominated convergence theorem yields

$$E[T_\alpha \pi_\lambda^*(x_{T_\alpha}) \bar{F}(\pi_\lambda^*(x_{T_\alpha}))] - E[T_\alpha \pi_{\lambda+\delta}^*(x_{T_\alpha}) \bar{F}(\pi_{\lambda+\delta}^*(x_{T_\alpha}))] \rightarrow 0$$

by considering the dominating random variable $2T_\alpha \sup_p p \bar{F}(p)$. Together, the preceding two limits

let us conclude that

$$E[T_\alpha \pi_\lambda^*(x_{T_\alpha}) \bar{F}(\pi_\lambda^*(x_{T_\alpha}))] - E[T_\alpha \pi_{\lambda+\delta}^*(x'_{T_\alpha}) \bar{F}(\pi_{\lambda+\delta}^*(x'_{T_\alpha}))] \rightarrow 0$$

Together with (3) and (4) this yields:

$$\begin{aligned} (5) \quad & \left. \frac{d}{d\lambda} J_\lambda^*(x) \right|_{\lambda=\bar{\lambda}} \\ &= \lim_{\delta \rightarrow 0} \left(E \left[\int_0^{T_\alpha(1+\delta/\bar{\lambda})} \pi_{\lambda+\delta}^*(x_t) \bar{F}(\pi_{\lambda+\delta}^*(x_t)) \bar{\lambda} dt \middle| x_0 = x \right] - E \left[\int_0^{T_\alpha} \pi_{\bar{\lambda}}^*(x_t) \bar{F}(\pi_{\bar{\lambda}}^*(x_t)) \bar{\lambda} dt \middle| x_0 = x \right] \right) / \delta \\ &= \frac{d}{d\lambda} E \left[\int_0^{T_\alpha} \pi_{\bar{\lambda}}^*(x_t) \bar{F}(\pi_{\bar{\lambda}}^*(x_t)) \bar{\lambda} dt \middle| x_0 = x \right] \Big|_{\lambda=\bar{\lambda}} + \lim_{\delta \rightarrow 0} \left(E \left[T_\alpha \pi_{\bar{\lambda}+\delta}^*(x'_{T_\alpha}) \bar{F}(\pi_{\bar{\lambda}+\delta}^*(x'_{T_\alpha})) \right] + O(\delta) \right) \\ &= \frac{d}{d\lambda} E \left[\int_0^{T_\alpha} \pi_{\bar{\lambda}}^*(x_t) \bar{F}(\pi_{\bar{\lambda}}^*(x_t)) \bar{\lambda} dt \middle| x_0 = x \right] \Big|_{\lambda=\bar{\lambda}} + E \left[T_\alpha \pi_{\bar{\lambda}}^*(x_{T_\alpha}) \bar{F}(\pi_{\bar{\lambda}}^*(x_{T_\alpha})) \right] \end{aligned}$$

We note that

$$E \left[\int_0^{T_\alpha} \pi_\lambda(x_t) \bar{F}(\pi_\lambda(x_t)) \bar{\lambda} dt \middle| x_0 = x \right]$$

is differentiable with respect to $\pi_\lambda(\cdot)$. This follows from the differentiability of $E[\exp(-\alpha\tau)]$ with respect to η when τ is distributed as an exponential random variable with parameter η , and since \bar{F} is differentiable by Assumption 1.

Now,

$$\begin{aligned} & \left. \frac{d}{d\lambda} E \left[\int_0^{T_\alpha} \pi_\lambda^*(x_t) \bar{F}(\pi_\lambda^*(x_t)) \bar{\lambda} dt \middle| x_0 = x \right] \right|_{\lambda=\bar{\lambda}} \\ &= \sum_{X=0}^x \left(\frac{d}{d\lambda} \pi_\lambda^*(X) \right) \left(\frac{d}{d\pi_\lambda(X)} E \left[\int_0^{T_\alpha} \pi_\lambda(x_t) \bar{F}(\pi_\lambda(x_t)) \bar{\lambda} dt \middle| x_0 = x \right] \Big|_{\pi_\lambda(X)=\pi_\lambda^*(X)} \right) \Big|_{\lambda=\bar{\lambda}} \\ &= 0 \end{aligned}$$

where we use fact that since π_λ^* attains maximum revenue for an arrival rate $\lambda = \bar{\lambda}$,

$$\left. \frac{d}{d\pi_\lambda(X)} E \left[\int_0^{T_\alpha} \pi_\lambda(x_t) \bar{F}(\pi_\lambda(x_t)) \bar{\lambda} dt \middle| x_0 = x \right] \right|_{\pi_\lambda(X)=\pi_\lambda^*(X)} = 0$$

With (5), this yields equality (1) and the proof. \square

B Proofs for Section 4

Lemma 3. For all $z \in \mathcal{S}$, $\alpha > 0$

$$J^*(z) \leq \tilde{J}(z) \leq J_{\mu^*(z)}^*(x) \leq \frac{\bar{F}(p^*) p^* \mu(z)}{\alpha}.$$

where p^* is the static revenue maximizing price.

Proof: Since $J_\lambda^*(x)$ is concave in λ by Lemma 2, Jensen's inequality gives us that $J_{a/b}^*(x) = J_{E[\lambda]}^*(x) \geq E[J_\lambda^*(x)] = \tilde{J}(z)$. Note that $J_\lambda^*(x)$ is bounded above by the value of a system with customer arrival rate λ but without a finite capacity constraint. The optimal policy in such a system is simply to charge the static revenue maximizing price, p^* , garnering a value of $\frac{\bar{F}(p^*)p^*\lambda}{\alpha}$ yielding $J_\lambda^*(x) \leq \frac{\bar{F}(p^*)p^*\lambda}{\alpha}$. \square

Lemma 4. For all $z \in \mathcal{S}$, there is a unique $p \geq 0$ such that $\frac{\bar{F}(p)}{\rho(p)}\mu(z) = \alpha\tilde{J}(z)$.

Proof: Note that $\frac{\bar{F}(p)p\mu(z)}{\alpha}$ is a continuous, monotone decreasing function of p for $p \geq p^*$ under Assumption 1. But since $\frac{\bar{F}(\pi^*(z))\pi^*(z)\mu(z)}{\alpha} = J^*(z)$, the result is immediate from Lemma 3; in fact the unique solution to $\frac{\bar{F}(p)}{\rho(p)}\mu(z) = \alpha\tilde{J}(z)$ must be in $[p^*, \pi^*(z)]$. \square

C Proofs for Section 6

Lemma 5. Let $\pi : \mathcal{S} \rightarrow \mathbb{R}_+$ be an arbitrary policy and let $\pi' : \mathcal{S} \rightarrow \mathbb{R}_+$ be defined according to $\pi'(x, a, b) = \pi(x, a, b/\alpha)$. Then, for all $z \in \mathcal{S}, \alpha > 0$, $J^{\pi, \alpha}(z) = J^{\pi', 1}(x, a, \alpha b)$, and, in particular, $J^{*, \alpha}(z) = J^{*, 1}(x, a, \alpha b)$.

Proof: Let $\hat{z} \equiv (\hat{x}, \hat{a}, \hat{b}) \in \mathcal{S}$ be arbitrary. Restricting attention to the pricing policy π , we have that $J^{\pi, \alpha}$ is given by the unique solution to the HJB equation $H^\pi J = 0$. That is, $J^{\pi, \alpha}$ uniquely satisfies

$$(6) \quad \bar{F}(\pi(x, a, b)) \left(\frac{a}{b} (\pi(x, a, b) + J(x - a, a + 1, b) - J(x, a, b)) + \frac{d}{db} J(x, a, b) \right) - \alpha J(x, a, b) = 0,$$

for all $z \in \mathcal{S}_{\hat{x}, \hat{a}, \hat{b}}$ and similarly for $J^{\pi', 1}$. In particular,

$$\bar{F}(\pi(x, a, b)) \left(\frac{a}{b} (\pi(x, a, b) + J^{\pi, \alpha}(x - a, a + 1, b) - J^{\pi, \alpha}(x, a, b)) + \frac{d}{db} J^{\pi, \alpha}(x, a, b) \right) - \alpha J^{\pi, \alpha}(x, a, b) = 0,$$

for all $z \in \mathcal{S}_{\hat{x}, \hat{a}, \hat{b}}$ and

$$\bar{F}(\pi'(x, a, b)) \left(\frac{a}{b} (\pi'(x, a, b) + J^{\pi', 1}(x - a, a + 1, b) - J^{\pi', 1}(x, a, b)) + \frac{d}{db} J^{\pi', 1}(x, a, b) \right) - J^{\pi', 1}(x, a, b) = 0$$

for all $z \in \mathcal{S}_{\hat{x}, \hat{a}, \alpha \hat{b}}$.

Now, in order to prove our claim it will suffice to show that $\bar{J}(z)$ defined according to $\bar{J}(x, a, b) =$

$J^{\pi',1}(x, a, \alpha b)$ satisfies (6). But, identifying the change of variables $b' = \alpha b$, we have:

$$\begin{aligned}
& \bar{F}(\pi(x, a, b)) \left(\frac{a}{b} \left(\pi(x, a, b) + J^{\pi',1}(x - a, a + 1, \alpha b) - J^{\pi',1}(x, a, \alpha b) \right) + \frac{d}{db} J^{\pi',1}(x, a, \alpha b) \right) \\
& - \alpha J^{\pi',1}(x, a, \alpha b) \\
& = \bar{F}(\pi(x, a, b'/\alpha)) \left(\frac{a\alpha}{b'} \left(\pi(x, a, b'/\alpha) + J^{\pi',1}(x - a, a + 1, b') - J^{\pi',1}(x, a, b') \right) + \frac{d}{db} J^{\pi',1}(x, a, b') \right) \\
& - \alpha J^{\pi',1}(x, a, b') \\
& = \alpha \left(\bar{F}(\pi'(x, a, b')) \left(\frac{a}{b'} \left(\pi'(x, a, b') + J^{\pi',1}(x - a, a + 1, b') - J^{\pi',1}(x, a, \alpha b) \right) + \frac{d}{db'} J^{\pi',1}(x, a, b') \right) \right) \\
& - \alpha J^{\pi',1}(x, a, b') \\
& = 0.
\end{aligned}$$

This suffices for the proof. \square

Lemma 6. Let $J \in \mathcal{J}$ satisfy $J(0, a, b) = 0$. Let $\tau = \inf\{t : J(z_t) = 0\}$. Let $z_0 \in \mathcal{S}_{\bar{x}, \bar{a}, \bar{b}}$. Then,

$$E \left[\int_0^\tau e^{-\alpha t} H^\pi J(z_t) dt \right] = J^\pi(z_0) - J(z_0)$$

Let $J : \mathbb{N} \rightarrow \mathbb{R}$ be bounded and satisfy $J(0) = 0$. Let $\tau = \inf\{t : J(x_t) = 0\}$. Let $x_0 \in \mathbb{N}$. Then,

$$E \left[\int_0^\tau e^{-\alpha t} H_\lambda^\pi J(x_t) dt \right] = J_\lambda^\pi(x_0) - J(x_0)$$

Proof: Define for $J \in \mathcal{J}$, and $\pi \in \Pi$,

$$\mathcal{A}_{\pi,z} J(z) = \lim_{t>0, t \rightarrow 0} \frac{e^{-\alpha t} E_{z,\pi}[J(z(t))] - J(z)}{t}.$$

Further, define

$$H^\pi J(z) = \bar{F}(\pi(z)) \frac{a}{b} \pi(z) + \mathcal{A}_{\pi,z} J(z)$$

Lemma E.5 verifies that this definition is in agreement with our previous definition provided $J \in \mathcal{J}$.

Let τ be a stopping time of the filtration $\sigma(z^t)$ (where $z^t = \{z_{t'} : 0 \leq t' \leq t\}$). We then have:

$$\begin{aligned}
E \left[\int_0^\tau e^{-\alpha t} H^\pi J(z_t) dt \right] &= E \left[\int_0^\tau e^{-\alpha t} \left(\bar{F}(\pi(z_t)) \frac{a_t}{b_t} \pi(z_t) + \mathcal{A}_{\pi,z} J(z_t) \right) dt \right] \\
&= J^\pi(z_0) + E_{z_0} [e^{-\alpha \tau} J(z_\tau)] - J(z_0) \\
&= J^\pi(z_0) - J(z_0)
\end{aligned}$$

The second equality follows from the fact that

$$E \left[\int_0^\tau e^{-\alpha t} \mathcal{A}_{\pi,z} J(z_t) dt \right] = E_{z_0} [e^{-\alpha \tau} J(z_\tau)] - J(z_0)$$

which is Dynkin's formula for Markov processes (see III.10 in Rogers and Williams (2000)). The third equality follows by the definition of τ and the assumption that $J(0, a, b) = 0$. The proof of

the second assertion is identical. \square

Lemma 7. *If $\lambda < \mu$, $J_\lambda^{\pi^{nl}}(x) \geq (\lambda/\mu)J_\mu^*(x)$ for all $x \in \mathbb{N}$.*

Proof: Letting $\tau = \inf\{t : n_t = x_0\}$ as usual, we have

$$\begin{aligned} -E \left[\int_0^\tau e^{-\alpha t} H_\lambda^{\pi^{nl}} J_\mu^*(x_t) dt \right] &= E \left[\int_0^\tau e^{-\alpha t} (1 - \lambda/\mu) \alpha J_\mu^*(x_t) dt \right] \\ &\leq E \left[\int_0^\tau e^{-\alpha t} (1 - \lambda/\mu) \alpha J_\mu^*(x_0) dt \right] \\ &\leq (1 - \lambda/\mu) J_\mu^*(x_0) \end{aligned}$$

where the inequality follows from the fact that $J_\mu^*(x)$ is decreasing in x and since $\lambda < \mu$ here. So, from Lemma 6, we immediately have:

$$J_\mu^*(x_0) - J_\lambda^{\pi^{nl}}(x_0) \leq (1 - \lambda/\mu) J_\mu^*(x_0)$$

which is the result. \square

Lemma 8. *If $\lambda \geq \mu$, $J_\lambda^{\pi^{nl}}(x) \geq J_\mu^*(x)$ for all $x \in \mathbb{N}$.*

Proof: Here,

$$-E \left[\int_0^\tau e^{-\alpha t} H_\lambda^{\pi^{nl}} J_\mu^*(x(t)) dt \right] \leq 0$$

so the result follows immediately from Lemma 6. \square

Corollary 1. *For all $z \in \mathcal{S}$, and exponential reservation price distributions with parameter r :*

$$\frac{1}{1 + \log \kappa(a)} \leq \frac{\pi_{\text{db}}(z)}{\pi^*(z)} \leq 1$$

For all $z \in \mathcal{S}$, and logit reservation price distributions with parameter r :

$$\frac{1.27}{1.27 + \log \kappa(a)} \leq \frac{\pi_{\text{db}}(z)}{\pi^*(z)} \leq 1$$

Proof: The decay balancing equation for exponential reservation prices yields:

$$\begin{aligned} \frac{\pi_{\text{db}}(z)}{\pi^*(z)} &= \frac{r \log \frac{ra}{be^{-1}\tilde{J}(z)}}{r \log \frac{ra}{be^{-1}J^*(z)}} \\ &\geq \frac{\log \frac{ra}{be^{-1}\tilde{J}(z)}}{\log \frac{ra\kappa(a)}{be^{-1}\tilde{J}(z)}} \\ &= \frac{\log \frac{ra}{be^{-1}\tilde{J}(z)}}{\log \frac{ra}{be^{-1}\tilde{J}(z)} + \log \kappa(a)} \\ &\geq \frac{1}{1 + \log \kappa(a)} \end{aligned}$$

where the first inequality follows from Theorem 1 and the second inequality follows from the fact that by Lemma 3, $\tilde{J}(z) \leq \frac{a}{b}r$. That $\pi_{\text{db}}(z) \leq \pi^*(z)$ is immediate from the decay balance equation

and the fact that $\tilde{J}(z) \geq J^*(z)$. The proof of the bound for logit reservation prices is identical; we employ the fact that for logit reservation prices, $\bar{F}(p^*)p^* = e^{-1.27}r$, so that $\tilde{J}(z) \leq \frac{a}{b}re^{-0.27}$. \square

Lemma 9. *For all $z \in \mathcal{S}$, and reservation price distributions satisfying Assumptions 1 and 2,*

$$J^{\text{ub}}(z) \geq J^*(z)$$

Proof: Define the operator:

$$(H^{\text{ub}}J)(z) = \bar{F}(\pi_{\text{db}}(z)) \left(\frac{a}{b} (\pi^*(z) + J(z') - J(z)) + \frac{d}{db} J(z) \right) - e^{-1}J(z).$$

Analogous to the proof of Theorem E.1, one may verify that J^{ub} is the unique bounded solution to $(H^{\text{ub}}J)(z) = 0$ for all $z \in \mathcal{S}_{\tilde{x}, \tilde{a}, \tilde{b}}$ satisfying $J^{\text{ub}}(0, a, b) = 0$. Identically to the proof of Lemma 6, we can then show for $J \in \mathcal{J}$ satisfying $J(0, a, b) = 0$, and $z_0 \in \mathcal{S}_{\tilde{x}, \tilde{a}, \tilde{b}}$ that

$$(7) \quad E \left[\int_0^\tau e^{-\alpha t} H^{\text{ub}} J(z_t) dt \right] = J^{\text{ub}}(z_0) - J(z_0)$$

Now, observe that for $x > 0$,

$$\begin{aligned} & (H^{\text{ub}}J^*)(z) \\ &= \bar{F}(\pi_{\text{db}}(z)) \left(\frac{a}{b} (\pi^*(z) + J^*(z') - J^*(z)) + \frac{d}{db} J^*(z) \right) - e^{-1}J^*(z) \\ &\geq \bar{F}(\pi^*(z)) \left(\frac{a}{b} (\pi^*(z) + J^*(z') - J^*(z)) + \frac{d}{db} J^*(z) \right) - e^{-1}J^*(z) \\ &= 0 \end{aligned}$$

where for the inequality, we use the fact that

$$\pi^*(z) + J^*(z') - J^*(z) + \frac{b}{a} \frac{d}{db} J^*(z) = 1/\rho(\pi^*(z)) \geq 0$$

and that $\pi_{\text{db}}(z) \leq \pi^*(z)$ from Corollary 1. The equality is simply the HJB equation. We consequently have

$$H^{\text{ub}}J^*(z) \geq 0 \quad \forall z \in \mathcal{S}_{\tilde{x}, \tilde{a}, \tilde{b}}$$

so that (7) applied to J^* immediately gives:

$$J^{\text{ub}}(x, a, b) \geq J^*(x, a, b)$$

\square

Lemma 10. *Let $\pi : \mathcal{S} \rightarrow \mathbb{R}_+$ be an arbitrary policy and let $\pi' : \mathcal{S} \rightarrow \mathbb{R}_+$ be defined according to $\pi'(z) = (1/r)\pi(z)$. Then, for all $z \in \mathcal{S}, \alpha > 0, r > 0$, $J^{\pi, \alpha, r}(z) = rJ^{\pi', \alpha, 1}(z)$ and, in particular, $J^{*, \alpha, r}(z) = rJ^{*, \alpha, 1}(z)$.*

Proof: Consider the following coupling of the r system starting at state $z = (x, a, b)$, and of the 1 system starting at state z . The first system is controlled by the price function $\pi(\cdot)$ while the second

is controlled by the price function $\pi'(\cdot) = (1/r)\pi(\cdot)$. Consider the evolution of both systems under a sample path with arrivals at $\{t_k\}$ and a corresponding binary valued sequence $\{\psi_k\}$ indicating whether or not the consumer chose to make a purchase. Let $E[\cdot]$ be a joint expectation over $\{t_k, \psi_k; k \leq x\}$ assuming $\{t_k\}$ are the points of a Poisson(λ) process where $\lambda \sim \Gamma(a, b)$, and ψ_k is a Bernoulli random variable with parameter $\exp(-\pi(t_k^-)/r) = \exp(-\pi'(t_k^-))$. We then have:

$$\begin{aligned} J^{\pi, \alpha, r}(z) &= E \left[\sum_{k=1}^x \psi_k \pi(t_k^-) \exp(-\alpha(t_k)) \right] \\ &= r E \left[\sum_{k=1}^x \psi_k \pi'(t_k^-) \exp(-\alpha(t_k)) \right] \\ &= r J^{\pi', \alpha, 1}(z) \end{aligned}$$

The result follows. □

Lemma 11. For all $z \in \mathcal{S}$,

$$J^*(z|\tau) \leq e^{-e^{-1}\tau} \left(e^{-(\pi^* - \pi_{\text{db}})} \left[\pi^* + J^*(x-1, a+1, b_\tau^{\text{db}}) \right] + (1 - e^{-(\pi^* - \pi_{\text{db}})}) J^*(x, a+1, b_\tau^{\text{db}}) \right)$$

where $\pi^* = \pi^*(x, a, b_\tau^*)$ and $\pi_{\text{db}} = \pi_{\text{db}}(x, a, b_\tau^{\text{db}})$.

Proof: Since $\pi^*(\cdot) \geq \pi_{\text{db}}(\cdot)$, and further since $\pi_{\text{db}}(\cdot)$ is decreasing in b ¹, we must have that $\pi_t^* \geq \pi_{\text{db}t}$ on $t < \tau$. Thus, by our coupling we must have that $n_t^* \leq n_t^{\text{db}}$ on $t \leq \tau$; $n_\tau^* = 1$ with probability $e^{-(\pi^* - \pi_{\text{db}})}$ and $n_\tau^* = 0$ with the remaining probability. Moreover, conditioned on τ and n_τ^* , λ is distributed as a Gamma random variable with shape parameter $a+1$ and scale parameter b_τ^{db} .

We thus have

$$\begin{aligned} J^*(z|\tau) &= E \left[\int_{t=0}^{\infty} e^{-e^{-1}t} \pi^*(z_t^*) \lambda \bar{F}(\pi^*(z_t^*)) dt \middle| \tau, z_0^* = z \right] \\ &= e^{-e^{-1}\tau} e^{-(\pi^* - \pi_{\text{db}})} \pi^* + e^{-(\pi^* - \pi_{\text{db}})} E \left[\int_{t=\tau}^{\infty} e^{-e^{-1}t} \pi^*(z_t^*) \lambda \bar{F}(\pi^*(z_t^*)) dt \middle| \tau, x_\tau^* = x-1, z_0^* = z \right] \\ &\quad + (1 - e^{-(\pi^* - \pi_{\text{db}})}) E \left[\int_{t=\tau}^{\infty} e^{-e^{-1}t} \pi^*(z_t^*) \lambda \bar{F}(\pi^*(z_t^*)) dt \middle| \tau, x_\tau^* = x, z_0^* = z \right] \end{aligned}$$

But by our observation on the posterior statistics of λ given τ and n_τ^* ,

$$\begin{aligned} &E \left[\int_{t=\tau}^{\infty} e^{-e^{-1}t} \pi^*(z_t^*) \lambda \bar{F}(\pi^*(z_t^*)) dt \middle| \tau, x_\tau^* = x-1, z_0^* = z \right] \\ &\leq \sup_{\pi_t: t \geq \tau} E \left[\int_{t=\tau}^{\infty} e^{-e^{-1}t} \pi_t \lambda \bar{F}(\pi_t) dt \middle| \tau, x_\tau^* = x-1, z_0^* = z \right] \\ &= e^{-e^{-1}\tau} J^*(x-1, a+1, b_\tau^{\text{db}}) \end{aligned}$$

¹This follows easily from the fact that for any positive constant k , X/k is distributed as a Gamma random variable with parameters (a, bk) if X is distributed as a Gamma random variable with parameters (a, b) .

and similarly

$$\begin{aligned}
& E \left[\int_{t=\tau}^{\infty} e^{-e^{-1}t} \pi^*(z_t^*) \lambda \bar{F}(\pi^*(z_t^*)) dt \middle| \tau, x_\tau^* = x, z_0^* = z \right] \\
& \leq \sup_{\pi_t: t \geq \tau} E \left[\int_{t=\tau}^{\infty} e^{-e^{-1}t} \pi_t \lambda \bar{F}(\pi_t) dt \middle| \tau, x_\tau^* = x, z_0^* = z \right] \\
& = e^{-e^{-1}\tau} J^*(x, a+1, b_\tau^{\text{db}})
\end{aligned}$$

This yields the result. \square

Lemma 12. For $x > 1, a > 1, b > 0$, $J^*(x, a, b) \leq 2.05J^*(x-1, a, b)$.

Proof: We establish this result for the case where $\alpha = e^{-1}$. This is without loss since by Lemma 5 we know that for all $x > 1, a > 1, b > 0$, $J^{*,\alpha}(x, a, b) \leq 2.05J^{*,\alpha}(x-1, a, b) \Leftrightarrow J^{*,e^{-1}}(x, a, \alpha b/e^{-1}) \leq 2.05J^{*,e^{-1}}(x-1, a, \alpha b/e^{-1})$.

Let $\tau_1 = \inf\{t : n^*(t) = x-1\}$, and define

$$J^{*,\tau_1}(z) = E_{z, \pi^*} \left[\sum_{k=1}^{x-1} e^{-e^{-1}t_k} \pi_{t_k}^- \right].$$

Now,

$$(8) \quad J^*(z) = J^{*,\tau_1}(z) + E \left[e^{-e^{-1}\tau_1} J^*(1, a+x-1, b_{\tau_1}) \right]$$

We will show that $E \left[e^{-e^{-1}\tau_1} J^*(1, a+x-1, b_{\tau_1}) \right] \leq 1.05J^*(x-1, a, b)$. Since we know by definition that $J^*(x-1, a, b) \geq J^{*,\tau_1}(z)$, the result will then follow immediately from (8).

To show $E \left[e^{-e^{-1}\tau_1} J^*(1, a+x-1, b_{\tau_1}) \right] \leq 1.05J^*(x-1, a, b)$, we will first establish a lower bound on

$$\pi^*(2, a+x-2, b_{\tau_1}) / J^*(1, a+x-1, b_{\tau_1}).$$

Let $a+x-2 \equiv k, a+x-1 \equiv k'$. Certainly, $k' \leq 2k$ since $a > 1$. Now,

$$\pi^*(2, k, b) = 1 + \log k/b - \log J^*(2, k, b) \geq 1 + \log k/b - \log J_{k/b}^*(2)$$

and $J^*(1, k', b) \leq J^*(1, 2k, b) \leq J_{2k/b}^*(1)$ so that

$$\frac{\pi^*(2, k, b)}{J^*(1, k', b)} \geq \frac{1 + \log k/b - \log J_{k/b}^*(2)}{J_{2k/b}^*(1)}$$

But,

$$\inf_{y \in (0, \infty)} \frac{1 + \log y - \log J_y^*(2)}{J_{2y}^*(1)} = \inf_{y \in (0, \infty)} \frac{1 + \log y - \log W(ye^{W(y)})}{W(2y)} \geq 0.96$$

recalling the expression for $J_y^*(x)$ from Section 3.1.

so that

$$\frac{\pi^*(2, a+x-2, b_{\tau_1})}{J^*(1, a+x-1, b_{\tau_1})} \geq 0.96$$

It follows that

$$\begin{aligned}
J^*(x-1, a, b) &\geq J^{*,\tau_1}(z) \\
&\geq E[e^{-e^{-1}\tau_1} \pi^*(2, a+x-2, b_{\tau_1})] \\
&\geq 0.96 E[e^{-e^{-1}\tau_1} J^*(1, a+x-1, b_{\tau_1})]
\end{aligned}$$

Substituting in (8), we have the result. \square

A Remark on the Proof of Lemma 12.

The infimum in Lemma 12 is computed as follows. We first observe that

$$\frac{1 + \log y - \log W(ye^{W(y)})}{W(2y)} \geq \frac{1 + \log y - \log 2W(y)}{W(2y)}.$$

Some simple algebra establishes that

$$\frac{1 + \log y - \log 2W(y)}{W(2y)} = \frac{1 - \log 2 + W(y)}{W(2y)} \geq \frac{1 - \log 2 + W(y)}{2W(y)} \geq 1$$

for $y < 0.1$ using the fact that $W(\cdot)$ is concave increasing and $W(0.1) < 0.092$. In addition, using the fact that $W(x)/W(2x)$ is increasing in x and by evaluating $W(2 \times 10^8)/W(4 \times 10^8) > 0.961$, we can conclude that

$$\frac{1 - \log 2 + W(y)}{W(2y)} \geq \frac{1 - \log 2 + W(y)}{1.041W(y)} \geq 0.961$$

for $y > 2 \times 10^8$. It is then straightforward to numerically minimize $\frac{1 + \log y - \log W(ye^{W(y)})}{W(2y)}$ over the compact interval $[0.1, 2 \times 10^8]$ to any finite precision since it is Lipschitz over that interval.

D Auxiliary Results for Section 6

In what follows we derive an approximation bound for decay balancing prices when reservation prices satisfy the following assumption in addition to Assumption 1:

Assumption 3.

1. $\frac{\rho(p)}{F(p)}$ is a differentiable, convex function of p with support \mathbb{R}_+ .
2. There exists a unique static revenue maximizing price $p^* > 0$ with $\left. \frac{d}{dp} \frac{\rho(p)}{F(p)} \right|_{p=p^*} \geq 1/\overline{F}(p^*)p^{*2}$.

Corollary D.1. For all $z \in \mathcal{S}$, and reservation price distributions satisfying Assumptions 1 and 3

$$\frac{1}{\kappa(a)} \leq \frac{\pi_{\text{db}}(z)}{\pi^*(z)} \leq 1$$

Proof: Recall that the decay balance equation implies that $\frac{\overline{F}(p^*)p^*\rho(\pi^*(z))}{F(\pi^*(z))} = \frac{\overline{F}(p^*)p^*a}{J^*(z)b\alpha} \equiv r^*$. Let $\tilde{r} = \frac{\overline{F}(p^*)p^*a}{J(z)b\alpha}$. Lemma 3 implies that $r^* \geq \tilde{r} \geq 1$.

Define a function $g : [p^*, \pi^*(z)] \rightarrow [1, r^*]$ according to $g(p) = \frac{\bar{F}(p^*)p^*\rho(p)}{\bar{F}(p)}$. Observe that $g(p^*) = 1$, $g(\pi^*(z)) = r^*$ and further by Assumptions 1 and 3, $g(\cdot)$ is an increasing convex function of p on $[p^*, \pi^*(z)]$ with range $[1, r^*]$. It follows that the inverse function g^{-1} is a concave increasing function on $[1, r^*]$ with range $[p^*, \pi^*(z)]$.

Now we have that $\pi^{\text{db}}(z) = g^{-1}(\tilde{r}) = p^* + \frac{\pi^{\text{db}}(z) - p^*}{\tilde{r} - 1}(\tilde{r} - 1)$ and by the concavity of g^{-1} , we have $\pi^*(z) = g^{-1}(r^*) \leq g^{-1}(\tilde{r}) + \frac{g^{-1}(\tilde{r}) - g^{-1}(1)}{\tilde{r} - 1}(r^* - \tilde{r}) = p^* + \frac{\pi^{\text{db}}(z) - p^*}{\tilde{r} - 1}(r^* - 1)$.

Consequently,

$$\begin{aligned} \frac{\pi^{\text{db}}(z)}{\pi^*(z)} &\geq \frac{p^* + \frac{\pi^{\text{db}}(z) - p^*}{\tilde{r} - 1}(\tilde{r} - 1)}{p^* + \frac{\pi^{\text{db}}(z) - p^*}{\tilde{r} - 1}(r^* - 1)} \\ &\geq \frac{p^* + \frac{\pi^{\text{db}}(z) - p^*}{\tilde{r} - 1}(\tilde{r} - 1)}{p^* + \frac{\pi^{\text{db}}(z) - p^*}{\tilde{r} - 1}(\kappa(a)\tilde{r} - 1)} \\ &\geq \frac{p^* + (\tilde{r} - 1)/(\bar{F}(p^*)p^* \frac{d}{dp} \frac{\rho(p)}{\bar{F}(p)} \Big|_{p=p^*})}{p^* + (\kappa(a)\tilde{r} - 1)/(\bar{F}(p^*)p^* \frac{d}{dp} \frac{\rho(p)}{\bar{F}(p)} \Big|_{p=p^*})} \\ &\geq \frac{1}{\kappa(a)} \end{aligned}$$

where the second inequality follows from Theorem 1. The third inequality follows from the fact that by Assumption 3, $\frac{\tilde{r} - 1}{\pi^{\text{db}}(z) - p^*} \geq g'(p)|_{p=p^*} = \bar{F}(p^*)p^* \frac{d}{dp} \frac{\rho(p)}{\bar{F}(p)} \Big|_{p=p^*}$. The final inequality follows from part 2 of Assumption 3: $\bar{F}(p^*)p^* \frac{d}{dp} \frac{\rho(p)}{\bar{F}(p)} \Big|_{p=p^*} \geq 1/p^*$. That $\frac{\pi^{\text{db}}(z)}{\pi^*(z)} \leq 1$ is immediate from the fact that $J^*(z) \leq \tilde{J}(z)$. \square

Armed with this result, we can derive a performance bound analogous to Theorem 2, but for general reservation price distributions:

Theorem D.1. *For all $z \in \mathcal{S}$, and reservation price distributions satisfying Assumptions 1 and 3,*

$$\frac{1}{\kappa(a)} \leq \frac{J^{\pi^{\text{db}}}(z)}{J^*(z)} \leq 1.$$

E Existence and Uniqueness of solutions to the HJB equation

Our analysis thus far has been predicated on using the HJB equation to characterize the optimal value function J^* . This section makes this argument rigorous for the case of a Gamma prior (which is the focus of our analysis). In particular, we establish the following theorems for this special case:

Theorem E.1. *The value function J^* is the unique solution in \mathcal{J} to $HJ = 0$.*

Theorem E.2. *A policy $\pi \in \Pi$ is optimal if and only if $H^\pi J^* = 0$.*

Our proofs to both Theorems E.1 and E.2 will rely on showing the existence of a bounded solution to the HJB Equation $(HJ)(z) = 0$ for $z \in \mathcal{S}_{\tilde{x}, \tilde{a}, \tilde{b}}$. We restrict attention to exponential reservation prices (which are the primary focus of our analysis). All of the arguments that follow are easily extended to the case of general reservation prices satisfying Assumption 1, but doing so is notationally quite cumbersome.

E.1 Existence of Solutions to the HJB Equation

We will demonstrate the existence of a solution to the HJB Equation wherein price is restricted to some bounded interval. We will later show that the solution obtained is in fact a solution to the original HJB Equation. Throughout, this section, we will let r denote the mean of the reservation price.

Define $B = r + \frac{r}{\tilde{b}} \left(1 + \frac{e^{-1}(\tilde{a} + \tilde{x})}{\tilde{a}\alpha} + \frac{e^{-1}(\tilde{a} + \tilde{x})}{\alpha} \right)$. Let Π_B be the set of admissible price functions bounded by B , and define the Dynamic programming operator

$$(H^B J)(z) = \sup_{\pi \in \Pi_B} (H^\pi J)(z)$$

We will first illustrate the existence of a bounded solution to the HJB Equation:

$$(9) \quad (H^B J)(z) = 0$$

for $z \in \mathcal{S}_{\tilde{x}, \tilde{a}, \tilde{b}}$.

For some arbitrary $N > \tilde{b}$ we first obtain a solution on the compact set $\mathcal{S}_{\tilde{x}, \tilde{a}, \tilde{b}}^N \equiv \{(x, a, b) \in \mathcal{S} : x + a = \tilde{x} + \tilde{a}; \tilde{b} \leq b \leq N\}$ with the boundary conditions $J(x, a, N) = 0$ and $J(0, a, b) = 0$:

Lemma E.1. (9) has a unique bounded solution on $\mathcal{S}_{\tilde{x}, \tilde{a}, \tilde{b}}^N$ satisfying $J(x, a, N) = 0$ and $J(0, a, b) = 0$.

The proof is analogous to that of Theorem VII.T3 in Bremaud (1981); upon setting $J(0, a, b) = 0$, (9) can be interpreted as an initial value problem of the form $\dot{J} = f(J, b)$ with $J(N) = 0$, in the space $\mathbb{R}^{\tilde{x}-1}$ equipped with the max-norm.

The following two Lemma's construct a solution to (9) on $\mathcal{S}_{\tilde{x}, \tilde{a}, \tilde{b}}$ using solutions constructed on $\mathcal{S}_{\tilde{x}, \tilde{a}, \tilde{b}}^N$.

Lemma E.2. Let J^N be the unique solution to (9) on $\mathcal{S}_{\tilde{x}, \tilde{a}, \tilde{b}}^N$ with $J(x, a, N) = 0$ and $J(0, a, b) = 0$. Moreover, let $J^{N'}$ be the unique solution to (9) on $\mathcal{S}_{\tilde{x}, \tilde{a}, \tilde{b}}^{N'}$ for some $N' > N$ with $J(x, a, N') = 0$ and $J(0, a, b) = 0$. Then, for $(x, a, b) \in \mathcal{S}_{\tilde{x}, \tilde{a}, \tilde{b}}^N$,

$$|J^N(x, a, b) - J^{N'}(x, a, b)| \leq r \frac{\tilde{a} + \tilde{x}}{\tilde{b}} \exp(-\alpha(N - b))$$

Moreover, $J^N(x, a, b) \leq \frac{r e^{-1}(\tilde{a} + \tilde{x})}{\alpha \tilde{b}}$

Proof: Define $\tau_N = \inf\{t : n_t = x\} \wedge \inf\{t : b_t = N\}$. Similarly, define $\tau_{N'}$. Let $\pi^{*,N}(\cdot)$, defined on $\mathcal{S}_{\tilde{x}, \tilde{a}, \tilde{b}}^N$, be the greedy price with respect to J^N . Finally, define the 'revenue' function $r_t^{*,N} = \frac{a_t e^{-\pi_t^{*,N}/r} / \pi_t^{*,N}}{b_t}$. We then have, via an application of Lemma 6,

$$\begin{aligned} J^N(x, a, b) &= E_{z, \pi^{*,N}} \left[\int_0^{\tau_N} e^{-\alpha t} r_t^{*,N} dt \right] + E_{z, \pi^{*,N}} \left[e^{-\alpha \tau_N} J^N(x_{\tau_N}, a_{\tau_N}, b_{\tau_N}) \right] \\ &= E_{z, \pi^{*,N}} \left[\int_0^{\tau_N} e^{-\alpha t} r_t^{*,N} dt \right] \end{aligned}$$

Note that this immediately yields:

$$J^N(x, a, b) \leq J^*(x, a, b) \leq J_{a/b}^*(x) \leq \frac{re^{-1}(\tilde{a} + \tilde{x})}{\alpha \tilde{b}}.$$

Now, for an arbitrary $\pi \in \Pi^B$, and the corresponding revenue function r , we have (again, via Lemma 6)

$$\begin{aligned} J^{N'}(x, a, b) &\geq E_{z, \pi} \left[\int_0^{\tau_{N'}} e^{-\alpha t} r_t dt \right] + E_{z, \pi} \left[e^{-\alpha \tau_{N'}} J^{N'}(x_{\tau_{N'}}, a_{\tau_{N'}}, b_{\tau_{N'}}) \right] \\ &= E_{z, \pi} \left[\int_0^{\tau_{N'}} e^{-\alpha t} r_t dt \right] \end{aligned}$$

In particular, using the price function $\pi = \pi^{*,N}$ for $b \leq N$ and 0 otherwise, yields,

$$(10) \quad J^{N'}(x, a, b) \geq E_{z, \pi^{*,N}} \left[\int_0^{\tau_N} e^{-\alpha t} r_t^{*,N} dt \right] = J^N(x, a, b)$$

The same argument, applied to J^N , with the price function $\pi^{*,N'}$, yields

$$E_{z, \pi^{*,N'}} \left[\int_0^{\tau_N} e^{-\alpha t} r_t^{*,N'} dt \right] \leq J^N(x, a, b)$$

Finally, noting that on $\{\tau_{N'} > \tau_N\}$, $\tau_N \geq N - b$, we have

$$E_{z, \pi^{*,N'}} \left[\int_{\tau_N}^{\tau_{N'}} e^{-\alpha t} r_t^{*,N'} dt \right] \leq r \frac{\tilde{a} + \tilde{x}}{\tilde{b}} \exp(-\alpha(N - b))$$

Adding the two preceding inequalities, yields

$$J^{N'}(x, a, b) - r \frac{\tilde{a} + \tilde{x}}{\tilde{b}} \exp(-\alpha(N - b)) \leq J^N(x, a, b).$$

Since $J^{N'}(x, a, b) \geq J^N(x, a, b)$ by (10), the result follows. □

This yields as a corollary the following result:

Lemma E.3. $\lim_{N \rightarrow \infty} J^N$ exists on $\mathcal{S}_{\tilde{x}, \tilde{a}, \tilde{b}}$, is bounded, and solves system (9)

Proof: From Lemma E.2, we have $\lim_{N \rightarrow \infty} J^N(x, a, b)$ exists and is bounded for all $(x, a, b) \in \mathcal{S}$. We posit that this limit is a solution to system (9). First note that by the continuity of

$$f(x, a, J, b) \equiv \inf_{p \in [0, B]} \left[e^{\gamma p} \alpha J(x, a) - \frac{a}{b} p + \frac{a}{b} (J(x - 1, a + 1) - J(x, a)) \right]$$

in J , we have:

$$\lim_{N \rightarrow \infty} f(x, a, J^N, b) = f(x, a, \lim_{N \rightarrow \infty} J^N, b)$$

for each x, a, b . It remains for us to show that

$$\lim_{\delta \rightarrow 0} \lim_{N \rightarrow \infty} \frac{J^N(x, a, b + \delta) - J^N(x, a, b)}{\delta}$$

exists and equals $\lim_{N \rightarrow \infty} dJ^N(x, a, b)/db$. Note however by the Mean Value Theorem that

$$J^N(x, a, b + \delta) - J^N(x, a, b)/\delta = dJ^N(x, a, b)/db + R_N$$

where

$$\begin{aligned} |R_N| &\leq \sup_{b' \in [b, b + \delta]} dJ^N(x, a, y)/dy|_{y=b'} - \inf_{b' \in [b, b + \delta]} dJ^N(x, a, y)/dy|_{y=b'} \\ &= \sup_{b' \in [b, b + \delta]} f(x, a, J^N(x, a, b'), b') - \inf_{b' \in [b, b + \delta]} f(x, a, J^N(x, a, b'), b') \end{aligned}$$

But $J^N(x, a, b)$ converges uniformly to its limit on $[b, b + \delta]$ by Lemma E.2, and f is uniformly continuous on $[b, b + \delta]$ being a continuous function restricted to a compact set, so that

$$\limsup_N |R_N| \leq \sup_{b' \in [b, b + \delta]} f(x, a, J^*(x, a, b'), b') - \inf_{b' \in [b, b + \delta]} f(x, a, J^*(x, a, b'), b')$$

Finally, by the continuity of J^* in b ,

$$\lim_{\delta \rightarrow 0} \limsup_N |R_N| = 0$$

Similarly,

$$\lim_{\delta \rightarrow 0} \liminf_N |R_N| = 0$$

This completes the proof. \square

The previous Lemma constructs a bounded solution to (9). We now show that this solution is in fact a solution to the original HJB Equation $(HJ)(z) = 0$ for $z \in \mathcal{S}_{\tilde{x}, \tilde{a}, \tilde{b}}$.

Lemma E.4. *Let \tilde{J} be a bounded solution to (9). Then, \tilde{J} is a solution to $(HJ)(z) = 0$ for $z \in \mathcal{S}_{\tilde{x}, \tilde{a}, \tilde{b}}$.*

Proof: We show the claim by demonstrating that the greedy price (in Π^B) with respect to \tilde{J} is in fact attained in $[0, B)$. We begin by proving a bound on such a greedy price. Let $\pi^b \in \Pi^B$ be the greedy price with respect to \tilde{J} , and $\tau = \inf\{t : N_t = x_0\}$. Letting $\tilde{r}_t = \frac{a_t e^{-\pi_t^b/r} \pi_t^b}{b_t}$, we have, via Lemma 6,

$$\begin{aligned} \tilde{J}(z) &= E_{z, \pi^b} \left[\int_0^\tau e^{-\alpha t} \tilde{r}_t dt \right] + E_{z, \pi^b} \left[e^{-\alpha \tau} \tilde{J}(z_\tau) \right] \\ &= E_{z, \pi^b} \left[\int_0^\tau e^{-\alpha t} \tilde{r}_t dt \right] \\ &\leq J^*(z) \\ &\leq \frac{r e^{-1} (\tilde{a} + \tilde{x})}{\alpha \tilde{b}}. \end{aligned}$$

Now let \tilde{J}^δ be the solution to (9) when the discount factor is $\alpha(1+\delta/b)$. Let $\pi^{b,\delta}$ be the corresponding greedy price and $\tilde{r}_t^\delta = \frac{a_t e^{-\pi_t^{b,\delta}/r} \pi_t^{b,\delta}}{b_t}$. We then have from Lemma 6 and using the fact that $\tilde{J}(x, a, b + \delta) = \tilde{J}^\delta(x, a, b)$,

$$\begin{aligned}\tilde{J}(x, a, b + \delta) &= E_{z, \pi^{b,\delta}} \left[\int_0^{\tau^\delta} e^{-\alpha(1+\delta/b)t} \tilde{r}_t^\delta dt \right] \\ &\geq E_{z, \pi^b} \left[\int_0^\tau e^{-\alpha(1+\delta/b)t} \tilde{r}_t dt \right]\end{aligned}$$

It follows that

$$\begin{aligned}\tilde{J}(z) - \tilde{J}(x, a, b + \delta) &\leq E_{z, \pi^b} \left[\int_0^\tau (e^{-\alpha t} - e^{-\alpha(1+\delta/b)t}) \tilde{r}_t dt \right] \\ &\leq \int_0^\infty (e^{-\alpha t} - e^{-\alpha(1+\delta/b)t}) \frac{r e^{-1}(a+x)}{b} dt\end{aligned}$$

so that

$$\frac{d}{db} \tilde{J}(z) \geq -\frac{r\alpha e^{-1}(a+x)}{b b\alpha^2}$$

Putting the two bounds together yields

$$(11) \quad \tilde{J}(x-1, a+1, b) - \tilde{J}(z) + \frac{b}{a} \frac{d}{db} \tilde{J}(z) \geq -\frac{r e^{-1}(\tilde{a} + \tilde{x})}{\alpha \tilde{b}} - \frac{r e^{-1}(\tilde{a} + \tilde{x})}{\tilde{a} \tilde{b} \alpha}$$

Now observe that the greedy price $\pi^b \in \Pi$ with respect to \tilde{J} is given by

$$p = \left(r - \tilde{J}(x-1, a+1, b) + \tilde{J}(z) - \frac{b}{a} \frac{d}{db} \tilde{J}(z) \right)^+$$

which by (11) is in $[0, B)$, so that we have that \tilde{J} is, in fact, a solution to $(HJ)(z) = 0$ for $z \in \mathcal{S}_{\tilde{x}, \tilde{a}, \tilde{b}}$. \square

E.2 Proofs for Theorems E.1 and E.2

Lemma E.5. For $J \in \mathcal{J}$, and $\pi \in \Pi$, let

$$\mathcal{A}_{\pi, z} J(z) = \lim_{t>0, t \rightarrow 0} \frac{e^{-\alpha t} E_{z, \pi} [J(z(t))] - J(z)}{t}.$$

We have:

$$\mathcal{A}_{\pi, z} J(z) = e^{-\pi(z)/r} \frac{a}{b} \left(J(z') - J(z) + \frac{b}{a} \frac{d}{db} J(z) \right) - \alpha J(z)$$

Proof: As in Theorem T1 in Section VII.2 of Bremaud (1981), one may show for $J \in \mathcal{J}$, and an

arbitrary $z_0 \in \mathcal{S}_{\tilde{x}, \tilde{a}, \tilde{b}}$,

$$\begin{aligned} J(z_t) = & J(z_0) + \int_0^t \left[\frac{b_s}{a_s} \frac{d}{db_s} J(z_s) + J(x_s - 1, a_s + 1, b_s) - J(z_s) \right] \frac{a_s}{b_s} e^{-p_s/r} ds \\ & + \int_0^t [J(x_{s-} - 1, a_{s-} + 1, b_{s-}) - J(z_{s-})] (dN_s - \frac{a_s}{b_s} e^{-p_s/r} ds) \end{aligned}$$

It is not hard to show that that $N_s - \frac{a_s}{b_s} e^{-p_s/r}$ is a zero-mean $\sigma(z^s, p^s)$ martingale, so that we may conclude

$$\begin{aligned} e^{-\alpha t} E[J(z_t)] - J(z_0) = \\ e^{-\alpha t} E \left[\int_0^t \left[\frac{b_s}{a_s} \frac{d}{db_s} J(z_s) + J(x_s - 1, a_s + 1, b_s) - J(z_s) \right] \frac{a_s}{b_s} e^{-p_s/r} ds \right] + (e^{-\alpha t} - 1) J(z_0) \end{aligned}$$

Dividing by t and taking a limit as $t \rightarrow 0$ yields, via bounded convergence, the result. \square

Lemma E.6. (*Verification Lemma*) *If there exists a solution, $\tilde{J} \in \mathcal{J}$ to*

$$(HJ)(z) = 0$$

for all $z \in \mathcal{S}_{\tilde{x}, \tilde{a}, \tilde{b}}$, we have:

1. $\tilde{J}(\cdot) = J^*(\cdot)$
2. Let $\pi^*(\cdot)$ be the greedy policy with respect to \tilde{J} . Then $\pi^*(\cdot)$ is an optimal policy.

Proof:

Let $\pi \in \Pi$ be arbitrary. By Lemma 6,

$$(12) \quad \begin{aligned} J^\pi(z_0) - \tilde{J}(z_0) = & E \left[\int_0^\tau e^{-\alpha s} H^\pi \tilde{J}(z_s) ds \right] \\ & \leq 0 \end{aligned}$$

with equality for $\pi^*(\cdot)$, since $H^{\pi^*} \tilde{J}(z) = (H\tilde{J})(z) = 0$ for all $z \in \mathcal{S}_{\tilde{x}, \tilde{a}, \tilde{b}}$. \square

Now we have shown the existence of a bounded solution, \tilde{J} to $(HJ)(z) = 0$ on $\mathcal{S}_{\tilde{x}, \tilde{a}, \tilde{b}}$ in the previous section, so that the first conclusion of the Verification Lemma gives

Theorem D.1. *The value function J^* is the unique solution in \mathcal{J} to $HJ = 0$.*

The second conclusion and (12) in the Verification Lemma give

Theorem D.2. *A policy $\pi \in \Pi$ is optimal if and only if $H^\pi J^* = 0$.*

References

- Bremaud, P. 1981. *Point Processes and Queues: Martingale Dynamics*. 1st ed. Springer-Verlag.
- Rogers, L.C.G, David Williams. 2000. *Diffusions, Markov Processes, and Martingales: Volume 1, Foundations*. Cambridge University Press.