

# Simple Policies for Dynamic Pricing with Imperfect Forecasts

Yiwei Chen      Vivek F. Farias \*

This Version: May 16, 2011

## Abstract

We consider the ‘classical’ single product dynamic pricing problem allowing the ‘scale’ of demand intensity to be modulated by an an exogenous ‘market size’ stochastic process. This is a natural model of dynamically changing market conditions. We show that for a broad family of Gaussian market size processes, simple dynamic pricing rules that are essentially agnostic to the specification of this market size process perform provably well. The pricing policies we develop are shown to compensate for forecast imperfections (or a lack of forecast information altogether) by frequent re-optimization and re-estimation of the ‘instantaneous’ market size.

## 1. Introduction

The following is one of the central (and perhaps, simplest) problems in the theory of revenue management: A vendor is endowed with some finite inventory that he must sell over some fixed sales horizon; no inventory replenishment is permitted. The vendor’s customers are price sensitive and arrive randomly over time. The vendor is thus faced with the task of dynamically adjusting prices over time so as to maximize expected revenue earned over the course of the selling season. With a view to providing managers with implementable prescriptions, this problem has been studied in many different guises. Central to the theoretical study of this dynamic pricing problem is a landmark paper by Gallego and van Ryzin [1994]. That paper studied a model wherein potential customers arrived at a rate whose magnitude as a function of time and posted prices was *known in advance*. Given this knowledge, the elegant and practical insight from that work was simply this: by posting a *fixed* price over the course of the selling season, the vendor was guaranteed to earn close to the maximum revenue achievable under a dynamic pricing policy, especially in ‘high volume’ settings<sup>1</sup>.

In reality, it is typically not the case that a vendor has access to a reliable prediction of how customer demand will evolve over the course of the season. In particular, the very nature of the product being sold may preclude the possibility of coming up with accurate predictions, or any prediction whatsoever: fashion items, or novelty luxury goods are good examples of such products. More to the point, given that valuable information is revealed over time, the simple ‘fixed price’ prescriptions alluded to above are unlikely to be sufficient in the face of uncertainty in the evolution of customer demand. Faced with such uncertainty, a natural alternative is to consider

---

\*YC is with the Sloan School of Management, MIT. VF is with the Operations Research Center and Sloan School of Management, MIT. Emails: {ywchen, vivekf@mit.edu}

<sup>1</sup>Regimes where the initial inventory and scale of demand grow large simultaneously.

building *stochastic* ‘forecast’ models for how consumer demand might evolve. This alternative has its own perils: building a useful model of this type is far from trivial and the predictive power of such models constructed in practice is frequently questionable. Moreover, the implementability (or even, computability) of an optimal pricing scheme incorporating such a stochastic model is unlikely to be as simple or clean as in the case where predictions made at the start of the selling season are perfect.

What is needed at this juncture is a simple to interpret and implement prescription for the above dynamic pricing problem. This prescription should rely only on data that a manager can easily access or calculate which in the real world is essentially just sales information over time. Of course, simplicity in itself is not sufficient; our prescription needs to provide compelling performance. Given the restrictions on information about the market size process it is not clear what an appropriate benchmark is. A ‘gold standard’ benchmark is, of course, the revenue under a ‘clairvoyant’ optimal policy computed assuming perfect knowledge of the realization of customer demand over time. This paper takes a first step towards constructing such a prescription. In brief, we consider a dynamic pricing model wherein arriving customers are price sensitive. These potential customers arrive over time at some potentially non-stationary rate. However, as opposed to being known in advance, this rate process is *stochastic, un-modeled and unobservable*. This is representative of a volatile demand environment and the reality that initial, pre-selling season demand predictions are often very crude (or often, not even available). We make several contributions relative to the dynamic pricing model above:

**A New Prescription:** Optimal dynamic pricing in our setting is challenging from a computational and implementational perspective. We propose a sub-optimal heuristic that accounts for the stochasticity in the market-size (demand) process while preserving much of the simplicity and implementability of the Gallego and van Ryzin [1994] policy. The policy we propose, the ‘Re-optimized Fixed Price Policy’, or RFP- $\Delta$  policy in short, is akin to repeatedly applying the fixed price policy at discrete epochs in time  $\Delta$  apart, with updated values for market size and inventory. The proxy for ‘market size’ used at a given epoch is computed from sales over the preceding epoch in time. In the event that the manager has access to side information or wishes to hedge against some specific realization of demand, this estimate is ‘tuned’ by a certain hedging parameter. The RFP- $\Delta$  policy is attractive from a practical perspective for two reasons:

- It is easy to interpret: Indeed, the price posted at each price revision may be interpreted as the optimal ‘fixed price’ in response to the inventory level at that point in time and the estimate of demand computed by the scheme at that point in time. This interpretability is valuable not just in and of itself, but also because a number of ‘legacy’ dynamic pricing systems already rely on fixed price logic.
- The variant of our policy that we predominantly study requires absolutely *no* information pertaining to the market size process. Our policy never requires that the market size process be directly observed. In particular, we present a family of schemes one member of which can be run in an entirely mechanical fashion given the ability to observe sales and nothing else.

**Performance Analysis:** In spite of its simplicity, we show the RFP- $\Delta$  policy is competitive with a *clairvoyant* pricing algorithm with access not just to a probabilistic model of demand evolution but *realizations* of how demand will evolve over the course of the selling season. We establish this by showing that the performance loss in using our prescription relative to an optimal clairvoyant pricing strategy is *uniformly* bounded for a broad class of Gaussian demand or ‘market-size’ processes in the high volume setting; in particular this bound holds for arbitrarily volatile market-size processes. In

addition, we present a parametric performance analysis that succinctly describes the performance of our prescription as a function of key market-size process parameters, and obtain correspondingly tighter bounds. Key to our analysis is a certain ‘inventory balancing’ property inherent to the RFP- $\Delta$  policy which mitigates the need for a model of the market size process. In addition, the frequency with which we review prices must clearly impact performance; it stands to reason that as  $\Delta$  grows large, performance will suffer. As such, we provide an analysis of the ‘price’ of such discrete price reviews that isolates the key factors that influence performance loss. We believe that these results are potentially of broader independent interest.

**Computational Evidence:** We present a computational study that delves into the implementation of our prescription and the performance we might expect in practice. In our experiments, we model market-size processes as OU processes (which are continuous time analogues to the moving average processes that serve as canonical stochastic forecast models). Our numerical results suggest that the RFP- $\Delta$  policy performs consistently well (i.e well within 90% of an optimal pricing policy) for a wide range of market-size volatilities and inventory levels or ‘load factors’. We show that these gains can be achieved with a relatively small number of price adjustments. Finally, these experiments show that the use of the RFP- $\Delta$  policy yields valuable gains over price updates that account for inventory shocks but do not update demand forecasts, using instead some initial forecast.

## 1.1. Literature Review

There are several streams of literature that are relevant to our work. The most relevant perhaps is the large literature on dynamic pricing: Gallego and van Ryzin [1994] formulated an elegant model in which a vendor starts with a finite number of identical indivisible units of inventory and is allowed to adjust prices over time. Customers arrive according to a Poisson process, with independent, identically distributed reservation prices and make purchases if and only if their reservation prices exceed the posted price. The primary insight in this work was that fixed price policies are essentially optimal if the vendor has an accurate forecast of customer demand over time. Bitran and Mondschein [1997] study optimal pricing policies in a periodic-review model. They derive structural properties of the optimal pricing policy and show that it is consistent with observations in practice. Zhao and Zheng [2000] specialize the model formulated in Gallego and van Ryzin [1994] and explicitly model the evolution of the customer reservation price distribution over time. They derive structural properties for this interesting setting. Comprehensive literature reviews on dynamic pricing can be found in McGill and van Ryzin [1999], Bitran and Caldentey [2003], Elmaghraby and Keskinocak [2003], and Talluri and van Ryzin [2004]. A related stream of literature considers learning issues that arise in the above dynamic pricing setting. The work here typically considers (relatively simple) market size processes parameterized by some un-observable parameter that must be learned over time. Optimal policies are developed in some contexts (see for example, Besbes and Zeevi [2009], Xu and Hopp [2004]), and sub-optimal heuristics in others (see for example, Aviv and Pazgal [2005b], Aviv and Pazgal [2005a], Araman and Caldentey [2009] and Farias and Van Roy [2010]).

In an important departure from the models above, Akan and Ata [2009] considered a model for network revenue management wherein the relevant market size processes are allowed to be arbitrary diffusions. In a tour-de-force analysis, that work showed that optimal policies in that model (in high volume settings) were of the well-studied ‘bid-price’ type, and moreover that these bid prices could be computed via the solution of certain PDE’s derived from the diffusions describing the market size process. The present work can be seen to complement that line of literature in the sense that it asks what one may do if the diffusions describing the underlying market-size process

are unspecified or only partially specified to the seller. In other words, what can be done when a perfectly specified forecast model is unavailable to the seller? The present work also complements a recent paper by Besbes and Maglaras [2009] which studies issues similar to the ones here albeit in the context of admission control to a queue via modulating prices. Both of the above papers study the relevant systems in a limiting regime that produces a *stochastic* fluid model. This appears to be the right regime for the issues at hand wherein the time scale at which customers arrive is substantially ‘faster’ than that at which one sees shocks in the aggregate arrival rate. Of course, the use of *deterministic* fluid models in RM contexts is relatively common; see for instance, Gallego and van Ryzin [1994], Gallego and van Ryzin [1997], Bitran and Caldentey [2003], Maglaras [2006], and Maglaras and Meissner [2006]. Finally, it is worth noting that outside of dynamic pricing, inter-temporal correlation in the customer arrival process is frequently modeled by assuming that customer arrival rates are driven by some autoregressive integrated moving average (ARIMA) process in the supply chain/ inventory management literature. See for instance the work by Lee et al. [2000], Raghunathan [2001], Gaur et al. [2005] and Graves [1999].

The remainder of this paper is organized as follows: In Section 2, we formulate our model and define the vendor’s optimization problem. Section 3 introduces ‘re-optimized fixed price’ (RFP) policies which are the subject of this paper. Section 4 is devoted to a theoretical performance analysis of an ‘idealized’ RFP policy in a closely related *fluid* system, which retains a stochastic market-size process, but assumes that demand is infinitely divisible and the customer arrival process is thus a ‘fluid’ whose rate is modulated by the underlying market-size process. In Section 5, we discuss the impact of using the RFP- $\Delta$  policy (which has discrete reviews) as opposed to the idealized RFP policy which is allowed to update price continuously and is the subject of the preceding section. The analysis remains in the fluid model. Finally, Section 6 shows that the performance analysis of the preceding section is also valid in the stochastic model introduced in Section 2 provided we operate in an appropriate ‘high volume’ regime. This section also retracts the ability to observe the market size process directly, relying instead on empirical estimates from past sales. The results in Section 6 thus close the loop on our analysis and provide a relatively complete performance analysis of a practically implementable variant of our policy. Section 7 presents a numerical investigation of the RFP- $\Delta$  policy. Section 8 concludes with thoughts on future directions.

## 2. Problem Formulation

We consider a vendor who begins a selling season of length  $T$  with  $x_0$  units of inventory of some given product. The vendor posts a per-unit price  $p_t \in \mathbb{R}_+ \cup \{\infty\}$  at time  $t$  and is allowed to dynamically adjust this price to compensate for demand shocks he may experience. Potential customers arrive according to a point process with rate  $\Lambda_t$ ; this instantaneous demand rate is itself determined by a ‘market size’ stochastic process  $\{\Lambda_t : t \geq 0\}$ . The market-size process is exogenous and independent of everything else. Arriving customers are endowed with a reservation price drawn independently from a fixed cumulative distribution  $F(\cdot)$ . For a customer arriving at time  $t$ , the customer chooses to purchase a single unit of the product if her reservation price exceeds the price posted at that time,  $p_t$ ; otherwise she is lost to the system. The vendor’s goal is to dynamically adjust prices in a manner that maximizes expected revenue. The remainder of this section is dedicated to formalizing this problem.

**Reservation Prices:** As stated above, we assume that each customer’s reservation price is an

independent random variable with cumulative distribution  $F(\cdot)$ . Letting  $\bar{F}(p) = 1 - F(p)$ , we have that the probability an arriving customer will choose to make a purchase when the posted price is  $p$  is simply  $\bar{F}(p)$ . We will make the following assumptions on the reservation price distribution  $F(\cdot)$ :

**Assumption 1.**

1.  $F(\cdot)$  has a continuous density  $f(\cdot)$  with support  $\mathbb{R}_+$ .<sup>2</sup>
2.  $F$  has a non-decreasing hazard rate on  $\mathbb{R}_+$ . That is,  $f(p)/\bar{F}(p)$  is non-decreasing in  $p$  on  $\mathbb{R}_+$ .
3.  $p\bar{F}(p)$  is concave and has a unique maximizer  $p^*$ .

The first assumption guarantees that  $F(\cdot)$  is invertible. Many commonly used distribution functions, such as the exponential, logistic and Weibull, satisfy the second assumption (see Farias and Van Roy [2010]). The third assumption is also a standard regularity assumption in the revenue management literature (see Talluri and van Ryzin [2004]). While each of these assumptions have an economic interpretation (see Farias and Van Roy [2010]), we do not dwell on such interpretations here since they are well studied in the extant literature. We note simply that the assumptions will permit us to use first order conditions to guarantee the optimality of various quantities in the sequel, and are thus made for convenience.

**2.1. The Market-Size Process:**

In order to capture shocks to aggregate demand, we assume that the instantaneous demand rate (or market size) is itself determined by an exogenous stochastic process,  $\{\Lambda_t : t \geq 0\}$ . In positing such a process we seek to model inter-temporal correlations in demand in addition to potential non-stationarity. Here we will restrict ourselves to a special class of market-size processes that while being broad, are sufficiently well-behaved to admit a number of useful pricing strategies. Our restrictions can be viewed as natural *structural* assumptions on the nature of the process driving aggregate demand as we will see shortly. In particular, we will make the following assumptions on  $\{\Lambda_t : t \geq 0\}$ .

**Assumption 2.**

1.  $\Lambda_t = (\bar{\Lambda}_t)^+$  where  $\bar{\Lambda}_t$  is a Gaussian process with continuous sample paths.
2.  $E[\bar{\Lambda}_t] \triangleq \lambda_t$  is positive.
3. The variance of the random variable  $\bar{\Lambda}_t$ ,  $\sigma_t^2$ , is increasing as a function of  $t$  and concave.

Several of the above restrictions are rather benign. Allowing continuity in the sample paths of the underlying process permits us with a minimal degree of ‘predictability’ for the market-size process<sup>3</sup>. The second assumption is partially justified by the fact that it is meaningless to anticipate negative demand. Finally requiring that  $\sigma_t^2$  be increasing in  $t$  simply captures the fact that as one looks further into the future the uncertainty in market size at those times increases

---

<sup>2</sup>We extend the domain of  $\bar{F}$  to define  $\bar{F}(\infty) = 0$  and  $\infty.\bar{F}(\infty) = 0$ ; these formal definitions agree with the limiting values of  $\bar{F}(p)$  and  $p\bar{F}(p)$  under our assumptions.

<sup>3</sup>Indeed, a large useful class of Gaussian processes have continuous modifications, including, for instance the generalized moving average processes we introduce.

and its concavity will be seen as a natural outcome of assuming that the impact of a shock on the instantaneous market size process diminishes over time in the class of processes we will discuss next. A more precise justification of the above assumptions is provided by exhibiting a class of stochastic processes satisfying these assumptions and are simultaneously seen to be good models of reality as we now describe.

A natural class of processes satisfying the above assumptions are what we dub *Generalized Moving Average Processes*. In particular, define:

$$\bar{\Lambda}_t = \lambda_t + \int_0^t \phi(t-s) dZ_s$$

where we assume that  $\lambda_t > 0$ ,  $\phi(\cdot) \in \mathcal{C}_1$  and  $dZ_s$  is an increment of Brownian motion. We think of  $\{\lambda_t\}$  as a deterministic forecast that the vendor may or may not possess and  $\int_0^t \phi(t-s) dZ_s$  as a ‘shock’ term that is difficult to model; indeed depending on the precise definition of  $\phi$  this term could behave in drastically different ways. In addition we assume that the function  $|\phi|$  is non-increasing; it is simple to verify that the resulting market-size process does indeed satisfy the stipulations of Assumption 2. Our labeling of such processes, and indeed the reason we believe they are interesting follow from the fact that the discrete time analogue of such a process is given by

$$\bar{\Lambda}_{i\Delta} = \lambda_{i\Delta} + \sum_{k=0}^{i-1} \phi((i-k)\Delta) \epsilon_k$$

where the  $\epsilon_k$  are i.i.d  $\mathcal{N}(0, \Delta)$  random variables. This is precisely a moving average process and forecast models employed in practice are likely to be of this type (see, for instance, Chapter 9 in Taluri and van Ryzin [2004]). Moreover, this analogy makes our assumptions on  $\phi$  fairly transparent: the requirement that  $|\phi|$  be non-increasing implies that demand shocks today have a diminishing influence on aggregate demand in the future. Finally, we note that a number of well-studied continuous time stochastic processes are of this type: Basic examples include the Wiener process with drift  $\{\mu_t : t \geq 0\}$  (which is recovered by setting  $\lambda_t = \lambda_0 + \int_0^t \mu_s ds$  and  $\phi(t) = \sigma$  for arbitrary  $\sigma > 0$ ) and the Ornstein-Uhlenbeck (OU) process with mean  $\lambda$  and  $\Lambda_0 = \lambda$  (which is recovered by setting  $\lambda_t = \lambda$  and  $\phi(t) = \sigma \exp(-\beta t)$  for arbitrary  $\sigma > 0$ , and  $\beta > 0$ ).

We end this section with an attractive property describing the *modulus of continuity* of the samples paths of generalized moving average processes; in addition to providing us with a concrete sense of the role of volatility in movements of the market size process, this property will serve as a valuable guide to understanding precisely how frequently one might want to review pricing decisions in a setting where prices can only be changed at discrete intervals. In particular, we will show that

$$|\Lambda_{t+\tau} - \Lambda_t| = O\left(\sigma \sqrt{2\tau \log(1/\tau)}\right)$$

where the constant in the big-Oh notation is  $\omega$  dependent. More precisely, we have:

**Lemma 1.** (*Sample Path Modulus of Continuity*) Assume that  $\bar{\Lambda}_t$  is a generalized moving average process with  $\phi \in \mathcal{C}_2$  and  $\lambda_t \in \mathcal{C}_1$ . Then, for  $\tau > 0$ , and any  $t \in [0, T)$ , we have:

$$\limsup_{\tau \rightarrow 0} \frac{|\Lambda_{t+\tau} - \Lambda_t|}{\sigma \sqrt{2\tau \log(1/\tau)}} \leq 1 \text{ a.s.}$$

## 2.2. Dynamics and the Revenue Optimization Problem for an Idealized Vendor

We begin with defining the sales/ inventory process: Let us denote a sample path of the market-size process up to time  $t$  by  $\Lambda^t \triangleq \{\Lambda_s : s \in [0, t]\}$  and similarly denote the price history up to time  $t$  by  $p^t \triangleq \{p_s : s \in [0, t]\}$ . Let  $\mathcal{F}_t = \sigma(\Lambda^t, p^t)$  and  $\mathcal{G}_t = \sigma(\{\Lambda_s : 0 \leq s \leq T\}, p^t)$ . We then define the sales process,  $N_t$ , as a point process with  $\mathcal{G}_t$ (also,  $\mathcal{F}_t$ )-intensity  $\Lambda_t \bar{F}(p_t)$ ; see Definition D7 in Bremaud [1980].  $N_t$  is simply interpreted as the cumulative sales up to time  $t$ ; we denote the corresponding inventory process by  $X_t = x_0 - N_t$ . We will allow for optimal pricing policies  $p_t$  that are progressively measurable with respect to the filtration  $\sigma(\Lambda^t, N^t)$ , so that  $\mathcal{F}_t = \sigma(\Lambda^t, N^t)$ <sup>4</sup>. In particular, such policies *cannot* see the future evolution of the market size process  $\{\Lambda_t\}$ . In fact, practical policies are likely to be even *further* restricted – in particular, we will eventually consider policies that are  $\mathcal{F}_t^{N_t}$ -progressive, i.e. policies that can simply observe the sales process  $N_t$  and have no knowledge of the measure over sample paths of the market size process (i.e. events in  $\mathcal{G}_0$ ) nor observe  $\Lambda_t$ .

**The ‘Idealized’ Vendor’s Revenue Optimization Problem.** We now discuss a revenue optimization problem faced by a vendor with perfect knowledge of the specification of the market-size process (i.e. knowledge of a probability distribution over sample paths of the  $\Lambda_t$  process), potentially unlimited computational power and the ability to monitor market size. All of these assumptions are objectionable and we will eventually seek an implementable prescription that requires *none* of these assumptions.

We require that the idealized vendor be restricted to causal pricing policies that respect the inventory constraint. More precisely, let  $\Pi$  denote the family of all  $\mathbb{R}_+ \cup \{\infty\}$ -valued price processes (‘policies’)  $\{\pi_t : t \geq 0\}$ , that are  $\mathcal{F}_t$ -progressive and in addition satisfy

$$\int_0^T dN_t \leq x_0 \text{ a.s.}$$

where  $p_t = \pi_t$ . The vendors objective is to find a pricing policy that maximizes expected revenue. In particular, define the expected revenue under a policy  $\pi$  according to <sup>5</sup>

$$(1) \quad J^\pi(x^0, \lambda^0, 0) = \mathbb{E} \left[ \int_0^T \pi_t dN_t \middle| X_0 = x_0, \Lambda_0 = \lambda_0 \right].$$

The vendor then seeks to find a policy  $\pi^*$  that achieves

$$\sup_{\pi \in \Pi} J^\pi(x^0, \lambda^0, 0) \triangleq J^*(x^0, \lambda^0, 0).$$

Let us try to understand the challenges associated with the problem above. First, we observe that solving the dynamic optimization problem implicit in finding a revenue maximizing policy above entails having to deal with a function valued state space (corresponding to histories of the market-size process), and this is likely to be difficult. At this juncture, one may observe that the size of the state-space can be reduced by restricting attention to special market-size processes and exploiting, for example, Markov structure; as an example if the market-size process is an OU process, one must simply track the current market size yielding a three dimensional state space.

<sup>4</sup>We will in the sequel consider regimes wherein  $\{\Lambda_t\}$  is replaced by the process  $\{\Lambda_{t,(n)}\}$  where  $\Lambda_{t,(n)} = n\Lambda_t$ . Denoting the corresponding sales process by  $N_{t,(n)}$  we will define  $\mathcal{F}_t^{(n)} = \sigma(\Lambda_{t,(n)}, N_{t,(n)})$  and similarly define  $\mathcal{G}_t^{(n)}$ .

<sup>5</sup>We will frequently omit the conditioning in the sequel when this is apparent from context.

If we were willing to ignore the obvious restrictions this imposes on modeling the market-size process, the vendor will still require knowledge of the specification of such a market-size process in order to solve such a problem; this requirement itself is unrealistic. Finally, in the interest of implementability, simple, easy to understand policies are highly desirable, and even if an optimal policy were computable, its complexity may preclude an easy implementation. Thus motivated, we will in the next section introduce a simple dynamic pricing strategy that we will: (a) Require little or no knowledge of the specification of the market-size process and (b) Remain a ‘good’ alternative to optimal pricing policies.

### 3. Re-optimized Fixed Price (RFP) Policies

To motivate the policies that we will introduce in this section, we begin with a simpler scenario that was the subject of the landmark paper by Gallego and van Ryzin [1994]: we consider the case where  $\{\Lambda_t : t \geq 0\}$  is in fact a deterministic process, so that  $\Lambda_t$  is identically equal to  $\lambda_t$ . Gallego and van Ryzin [1994] proposed the following *static* fixed price policy: At time 0, one selects a price,  $p_{\text{FP}}$  as follows:

$$p_{\text{FP}} = \begin{cases} p^*, & \text{if } \int_0^T \bar{F}(p^*) \lambda_s ds \leq x_0; \\ \bar{F}^{-1} \left( x_0 / \int_0^T \lambda_s ds \right), & \text{otherwise.} \end{cases}$$

This price is left unchanged over the course of the selling season. Gallego and van Ryzin [1994] showed that in a ‘high volume’ regime where the initial inventory level and market size are scaled simultaneously<sup>6</sup>, this fixed price policy is near-optimal. The intuition underlying this result is that in such a high volume regime, the dynamic pricing policy at hand is an essentially deterministic problem and the above solution is easily recognized as being the optimal solution of this limiting deterministic problem.

It is easy to see that if one considers an analogous regime for the case where the market-size process is *not* deterministic, the limiting optimization problem is no longer a deterministic problem, and there is no reason to believe that a fixed price policy such as the one above might work well. The following example illustrates what might happen:

**Example 1.** Consider the market-size process  $\Lambda_t = (\lambda + \sigma Z_t)^+$  where  $Z_t$  is Brownian motion. Let the initial inventory level,  $x_0$  and the length of the horizon  $T$  satisfy  $\lambda T = x_0$ . Further, assume that customers’ reservation prices are exponentially distributed with parameter 1, i.e.,  $\bar{F}(p) = \exp(-p)$ . We consider a sequence of ‘scaled’ problems; the  $n$ th problem has initial inventory  $x_{0,(n)} = nx_0$ , market-size process  $\{\Lambda_{t,(n)}\}$  (where  $\Lambda_{t,(n)} = n\Lambda_t$ ) and a fixed selling season of length  $T$ . Consider a static pricing policy  $\pi_{\text{FP}}$  that employs a constant selling price of  $p^*$  over the sales season. When  $\sigma = 0$ , we are left with a deterministic market-size process so that as  $n$  gets large, the results of Gallego and van Ryzin [1994] imply that:

$$\lim_{n \rightarrow \infty} \frac{J^{\pi_{\text{FP}}}(x_{(n)}^0, \lambda_{(n)}^0, 0)}{J^*(x_{(n)}^0, \lambda_{(n)}^0, 0)} = 1.$$

However, if  $\sigma > 0$ , we can show (see Appendix E.2) that

$$\limsup_{n \rightarrow \infty} \frac{J^{\pi_{\text{FP}}}(x_{(n)}^0, \lambda_{(n)}^0, 0)}{J^*(x_{(n)}^0, \lambda_{(n)}^0, 0)} = O((\log T)^{-1}).$$

---

<sup>6</sup>That is, we consider a sequence of problems indexed by  $n$  where in the  $n$ th problem, initial inventory is equal to  $nx_0$  and the market-size process is  $\{n\lambda_t : t \geq 0\}$ .

The point here is that the natural fixed price policy one constructs with an incorrect specification of the ‘shock’ term in the market-size process (in the case above, we simply ignored it) can be arbitrarily bad. Of course, with more information about the market size process one could construct better fixed price policies – for instance, if one knew the realization of the shocks at time 0, we are back to the Gallego and van Ryzin [1994] setting where fixed price policies are essentially optimal. If we knew enough about the market size process to compute an unbiased estimate of total demand realized over  $[0, T]$  (i.e. the quantity  $\int_0^T \Lambda_t dt$ ) at time 0, one may again potentially obtain an improvement over the setting above. Ideally, however, we wish to operate in a setting where the seller knows nothing about the market size process outside of perhaps  $\lambda$ ; even this quantity may be elusive.

The fixed price policy is attractive from a managerial perspective for its simplicity, implementability and ease of interpretation; unfortunately the negative result above points out that this policy is, not surprisingly, unsuitable for situations where the seller is unable to model the ‘shock’ in the market size process. As it turns out, we will be able to get away with the use of a policy that is only slightly more sophisticated than the above policy: the policy we consider will be akin to repeated application of the policy above with updated values for market size and inventory along with an intuitive ‘hedging’ adjustment to account for predictable variability in the market size process.

### 3.1. An Idealized Re-optimized Fixed Price Policy

This section introduces a simple dynamic pricing policy motivated by the simple fixed price policy above. The policy will require that one monitor inventory levels and the current market size process continuously <sup>7</sup>. In addition should deterministic forecast information be available, the policy will require as input an additional ‘hedging’ parameter  $\alpha$  (in subsequent sections we will discuss an oblivious choice of  $\alpha$ ). The dynamic pricing policy we propose is defined according to:

$$(2) \quad \pi_{\text{RFP}}(X^t, \Lambda^t, t) = \bar{F}^{-1} \left( \min \left\{ \bar{F}(p^*), \frac{X_t h(t, \alpha)}{\Lambda_t (T - t)} \right\} \right),$$

where  $\alpha \in [0, 1]$  is a parameter whose choice we will discuss shortly, and  $h(t, \alpha)$  is defined according to

$$h(t, \alpha) = \frac{\alpha(1 - \frac{t}{T}) + (1 - \alpha) \frac{\lambda_t (T-t)}{\int_0^T \lambda_s ds}}{\alpha(1 - \frac{t}{T}) + (1 - \alpha) \frac{\int_t^T \lambda_s ds}{\int_0^T \lambda_s ds}}.$$

for  $t < T$ ;  $h(T, \alpha) = 1$ . This policy is idealized for the simple reason that it requires that prices be updated continuously and moreover that the value of  $\Lambda_t$  be known to the manager at time  $t$ . We will address these shortcomings and provide an implementable prescription shortly; in the interim we use this idealized policy to build intuition for our eventual prescription.

The policy above is easy to interpret. To understand the general structure of the policy, let us begin with setting  $\alpha = 1$  so that  $h(t, \alpha) = 1$ . This results in a simple policy that is easy to interpret, has minimal information requirements, and that will eventually translate into a practical variant of the idealized policy here. Here the policy reduces to

$$\pi_{\text{RFP}}(X^t, \Lambda^t, t) = \bar{F}^{-1} \left( \min \left\{ \bar{F}(p^*), \frac{X_t}{\Lambda_t (T - t)} \right\} \right).$$

---

<sup>7</sup>Formally, this policy will be an  $\mathcal{F}_t$ -progressive process

Notice that when  $\Lambda_t = \lambda$  the near-optimal fixed price policy of Gallego and van Ryzin [1994] introduced earlier in this section is equivalently re-written as:

$$p_{\text{FP}} = \bar{F}^{-1} \left( \min \left\{ \bar{F}(p^*), \frac{x_0}{\lambda T} \right\} \right).$$

Thus, the RFP policy here may simply be re-interpreted as repeated application of the above fixed price policy where at every point in time  $t$  the policy uses an updated value for the inventory level, and assumes that the prevailing market-size  $\Lambda_t$  will remain constant for the remainder of the horizon. We will show that in this setting, one is guaranteed that under the RFP policy,

$$X_t \geq \frac{T-t}{T} x_0$$

so that the RFP policy can be seen as ‘rationing’ inventory uniformly over time in the face of complete uncertainty in the market-size process. We next adjust the above policy for scenarios wherein the manager wishes to ‘hedge’ his performance losses relative to some known *deterministic* forecast in addition to random realizations of the market size process.

### 3.1.1. Incorporating Deterministic Forecast Information: The multiplier $h(t, \alpha)$

The multiplier  $h(t, \alpha)$ , and in particular, the parameter  $\alpha$  allows us to ‘hedge’ between a deterministic forecast of market evolution and the possibility of demand shocks. In particular, a manager might have available forecast information on the evolution of the market size process. In the event that demand shocks are small, it stands to reason that there is inherent value to incorporating this information into any pricing decision made. On the other hand in the event that one anticipates large shocks in demand, one might want to place less of an emphasis on these forecasts.

To develop a qualitative understanding of this multiplier and the parameter  $\alpha$ , we consider here the two extreme cases where  $\alpha = 0$  and 1. In particular, setting  $\alpha = 0$  may be seen as an appropriate choice when our forecasts of market evolution are perfect so that  $\Lambda_t = \lambda_t$ ; i.e. there are *no* demand shocks. In this case, the policy above reduces to

$$\pi_{\text{RFP}}(X^t, \Lambda^t, t) = \bar{F}^{-1} \left( \min \left\{ \bar{F}(p^*), \frac{X_t}{\int_t^T \lambda_s ds} \right\} \right),$$

which is essentially the fixed price policy for deterministic market-size processes. Next, consider a scenario wherein the manager has essentially no forecast information available. Here one may consider setting  $\alpha = 1$ . This may also be seen as an appropriate choice when our forecasts are ‘swamped’ by volatility in the market-size process. In particular, for  $\alpha = 1$ , our policy reduces to

$$\pi_{\text{RFP}}(X^t, \Lambda^t, t) = \bar{F}^{-1} \left( \min \left\{ \bar{F}(p^*), \frac{X_t}{\Lambda_t(T-t)} \right\} \right).$$

In the absence of no predictions on how the market will evolve, the policy above prices as though the market size at time  $t$  will prevail over the remainder of the selling season. In reality, we will operate in an intermediate regime, where the market size process will consist of a predictable component (corresponding to initial forecasts e.g. a well established seasonality pattern) in addition to an un-modeled stochastic process component (that one would otherwise attempt to capture via a stochastic forecast model); intermediate choices of  $\alpha$  allow us to hedge appropriately between the two extreme scenarios described above. As we will show later, the interpolation between the two

extreme policies enforces an intuitive convex combination of the inventory levels one may expect under either of the two extreme policies at any given point in time. In particular, we will show that under the policy above, one must have:

$$X_t \geq \left[ \alpha \left( 1 - \frac{t}{T} \right) + (1 - \alpha) \frac{\int_t^T \lambda_s ds}{\int_0^T \lambda_s ds} \right] x_0$$

where one recognizes  $x_0(\int_t^T \lambda_s ds)/(\int_0^T \lambda_s ds)$  as the inventory on hand under the optimal policy in the event that the shock term were 0, and  $x_0(T - t)/T$  as a lower bound on the inventory on hand in the event that one employed the RFP policy without forecast information and thereby chose to ration inventory uniformly over time.

In the sequel, we will provide a ‘universally good’ choice of  $\alpha$  that one may select in the absence of any information about the market size process whatsoever. We will in addition show how one may optimally select  $\alpha$  given information about the drift and volatility of the market-size process.

### 3.2. An Implementable Prescription

We now consider an implementable variant of the policies described above. In particular, we consider a pricing policy that recognizes the following restrictions on the vendor: First, the vendor can only update prices at discrete, preferably infrequent, intervals. Second, the vendor is *never* able to observe the market size process.

As such, our scheme will require that the vendor set a single parameter  $\Delta$ , specifying how frequently prices need to be updated. In Section 5.1, we discuss the choice to  $\Delta$  and how it depends on volatility in the market size process. In addition, should the vendor have available forecast information on the market size process, he will be required to select the  $\alpha$  parameter described above; should this need arise, Section 4.5 discusses the selection of  $\alpha$ ; we anticipate that the most accessible choice in practice will be setting  $\alpha = 1$ . The pricing scheme will proceed as follows:

---

#### The RFP- $\Delta$ Policy

---

1. Over the interval  $[0, \Delta)$ , post the price  $p^*$  (the static revenue maximizing price).
2. At times  $i\Delta$  (where  $i \in \{1, \dots, \lfloor T/\Delta \rfloor\}$ ) for which  $X_{i\Delta} > 0$ , estimate current market size according to:

$$\hat{\Lambda}_{i\Delta} = \frac{N_{i\Delta} - N_{(i-1)\Delta}}{\bar{F}(p_{(i-1)\Delta})\Delta}$$

where  $p_{(i-1)\Delta}$  is understood to be the price posted over the interval  $[i - 1\Delta, i\Delta)$ .<sup>8</sup>

3. Over the period  $[i\Delta, (i + 1)\Delta)$ , post the price:

$$p_{i\Delta} = \bar{F}^{-1} \left( \min \left\{ \bar{F}(p^*), \frac{X_{i\Delta}}{\hat{\Lambda}_{i\Delta}(T - i\Delta)} \right\} \right).$$

---

<sup>8</sup>It is simple to establish inductively that so long as  $X_{i\Delta} > 0$ , we must have  $p_{(i-1)\Delta} < \infty$ .

In the event that deterministic forecast information is available and one wishes to incorporate this information into the pricing decision, post the price:

$$p_{i\Delta} = \bar{F}^{-1} \left( \min \left\{ \bar{F}(p^*), \frac{X_{i\Delta} h(i\Delta, \alpha)}{\hat{\Lambda}_{i\Delta}(T - i\Delta)} \right\} \right).$$

The first choice is consistent with setting  $\alpha = 1$ ; from here on we will refer to this as an ‘oblivious’ selection of  $\alpha$ .

4. As an exception to the above pricing rules, if at any time  $t$ , inventory hits zero (i.e  $X_t = 0$ ), we immediately set the price to  $\infty$ .<sup>9</sup>

The RFP policy we have proposed above is attractive from a practical perspective for several reasons:

- It can be implemented in an entirely mechanical fashion with absolutely no knowledge of the underlying market size process; this corresponds to selecting  $\alpha = 1$ . In the event that deterministic forecast information is available to the retailer, this information is easily incorporated into the pricing decision.
- It is easy to interpret as a discrete review policy, where price updates are made based on unexpected shocks in the sales process. The updated prices reflect a belief that the extant state of the world will prevail over the remainder of the selling season.
- It uses easy to understand proxies of the prevailing market size process. The only data required at each price update is the number of sales since the previous update and the price posted over that period.

In fact, the general policy described here is reminiscent of what sophisticated retailers employ in practice; see Talluri and van Ryzin [2004]. What we have been careful to specify is the precise forecasting that one may use. Moreover, as we will see in the sequel, we will show that with the appropriate forecasting rule and choice of review frequency  $\Delta$ , the above pricing prescription enjoys attractive theoretical performance guarantees as well as excellent numerical performance.

We end this section with a succinct formalization of the above pricing policy, which we refer to as the RFP- $\Delta$  policy. Define  $t(\Delta) = \max_i \{i\Delta : i\Delta \leq t\}$ . Then the RFP- $\Delta$  policy is given by:

$$\pi_{\text{RFP}}^\Delta(X^t, t) = \begin{cases} \bar{F}^{-1} \left( \min \left\{ \bar{F}(p^*), \frac{X_{t(\Delta)} h(t(\Delta), \alpha)}{\hat{\Lambda}_{t(\Delta)}(T - t(\Delta))} \right\} \right) & \text{if } X_t > 0 \\ \infty & \text{if } X_t = 0 \end{cases}$$

where the estimated market size  $\hat{\Lambda}_{i\Delta}$  assumed over times  $t \in [i\Delta, i + 1\Delta)$  where  $X_{i\Delta} > 0$  is defined as

$$\hat{\Lambda}_{i\Delta} = \begin{cases} \frac{X_{(i-1)\Delta} - X_{i\Delta}}{\bar{F}(\pi_{\text{RFP}}^\Delta(X^{(i-1)\Delta}, (i-1)\Delta))\Delta} & \text{if } i \in \{1, 2, 3, \dots\} \\ \frac{x_0}{\bar{F}(p^*)} & \text{if } i = 0. \end{cases}$$

<sup>9</sup>This is a standard formalism of the notion that one cannot sell more than the initial endowment of inventory.

In the context of generalized moving average processes, the above policy is oblivious to  $\phi$ , and in the event one set  $\alpha = 1$ , it is oblivious to  $\lambda_t$  as well. Moreover, it can only observe *sales* and make a limited number of pricing updates. The next three sections are devoted to a performance analysis of this policy. In contrast with the performance degradation illustrated for the fixed price policy in Example 1 (which was also oblivious to the structure of the market size process), we will demonstrate an *uniform* upper bound on performance loss relative to a *clairvoyant optimal policy* with advance knowledge of the realized sample path of the market-size process. We will in addition present stronger parametric guarantees for certain special market-size processes. The analysis will proceed in steps: we will first examine the performance of the *idealized* RFP policy presented in Section 3.1 in a fluid regime. We will next move to understanding the impact of discrete pricing updates via the RFP- $\Delta$  policy in the fluid regime. Finally, we will consider the stochastic model we began with and present performance guarantees for the RFP- $\Delta$  policy in that setting.

## 4. Performance Analysis I: The Idealized RFP Policy In A Fluid Model

This Section takes the first of three steps towards understanding the performance of the RFP- $\Delta$  policy that was the outcome of our discussion in the previous Section. In particular, we will first conduct a performance analysis for the *idealized* RFP policy, namely (2), which we recall can update prices continuously in a closely related *fluid* system. In particular, we consider a system wherein the inventory process follows the ‘fluid’ dynamic:

$$d\tilde{X}_t = -\Lambda_t \bar{F}(p_t) dt, \tilde{X}_0 = x_0.$$

This is in contrast to the model proposed in Section 2, wherein the inventory at time  $t$ ,  $X_t = x_0 - N_t$ , where  $N_t$  is a stochastic point process with intensity  $\Lambda_t \bar{F}(p_t)$ .

Such a dynamic arises by viewing individual customers as infinitesimally small (a notion we will make precise in a subsequent section). The notion of a policy remains unchanged and we define the fluid value of a policy  $\pi$  according to:

$$\tilde{J}^\pi(x^0, \lambda^0, 0) = \mathbb{E} \left[ \int_0^T \pi(\tilde{X}^t, \Lambda^t, t) \Lambda_t \bar{F}(\pi(\tilde{X}^t, \Lambda^t, t)) dt \middle| \tilde{X}_0 = x_0, \Lambda_0 = \lambda_0 \right]$$

where the expectation is now only over the sample paths of the  $\Lambda_t$  process. Similarly, we define the optimal fluid value function

$$\tilde{J}^*(x^0, \lambda^0, 0) = \sup_{\pi \in \Pi} \tilde{J}^\pi(x^0, \lambda^0, 0).$$

Our goal for the remainder of this section will be to produce a lower bound on the quantity  $\tilde{J}^{\pi_{\text{RFP}}}(x^0, \lambda^0, 0) / \tilde{J}^*(x^0, \lambda^0, 0)$ . We begin with an overview and discussion of the results in this section.

### 4.1. Performance Guarantees for the Idealized RFP Policy

This section will establish several performance guarantees for the idealized RFP policy in the fluid regime. In particular, we will establish the following results:

1. For the RFP policy with  $\alpha = 1$  (i.e. the oblivious choice) and market-size processes satisfying Assumption 2, and the additional requirement that  $\lambda_t = \lambda$  for all  $t$ , we will show that

$$\frac{\tilde{J}^{\pi_{\text{RFP}}}(x^0, \lambda^0, 0)}{\tilde{J}^*(x^0, \lambda^0, 0)} \geq 0.342.$$

2. If in addition to the assumptions above, we have  $\sigma_T/\lambda \leq \sqrt{2\pi}\lambda B$  for some constant  $B > 0$  then we will show that

$$\frac{\tilde{J}^{\pi_{\text{RFP}}}(x^0, \lambda^0, 0)}{\tilde{J}^*(x^0, \lambda^0, 0)} \geq \max \left\{ 0.342, \frac{1}{1+B} - \frac{B}{1+B} \left( \exp(-1/4\pi B^2) + 0.853 \right) \right\}.$$

3. Finally, for the RFP policy with  $\alpha = 0.594$  and *all* market-size processes satisfying Assumption 2, we will show that

$$\frac{\tilde{J}^{\pi_{\text{RFP}}}(x^0, \lambda^0, 0)}{\tilde{J}^*(x^0, \lambda^0, 0)} \geq 0.203.$$

The guarantees of the first two results apply to the setting where the vendor sets  $\alpha = 1$ . In particular, this is a setting wherein the manager knows absolutely nothing about the specification of the market size process and as such, we believe this is a highly relevant setting. We see that with frequent re-optimization, a well studied fixed price policy suffices to combat uncertainty in market-size. In particular, while a fixed price policy with *no* re-optimization can be arbitrarily bad (example 1), the RFP policy will have uniformly bounded relative performance losses. These bounds hold for *arbitrarily* volatile market-size processes. Simultaneously, we see via the second guarantee that in a regime where the volatility is low, the RFP policy proposed is near optimal.

The third guarantee applies to a setting wherein the manager is aware of  $\{\lambda_t\}$ , i.e. has access to a deterministic forecast. In that setting he may wish to incorporate this information into pricing decisions, and the performance result provides a choice of the parameter  $\alpha$  that allows him to do so while preserving a uniform performance guarantee. Additional knowledge of the market-size process will allow for a problem specific selection of the  $\alpha$  parameter. This will be evident in the proof of Theorem 3. We will also examine this issue from a computational perspective in our numerical experiments and show that, in fact, the theoretically ideal choice of  $\alpha$  performs quite well, while the oblivious choice of  $\alpha$  performs adequately as well.

As it will turn out, all of these guarantees hold relative to a ‘clairvoyant’ optimal policy that has full knowledge of the realized sample path at time 0. This allows us to interpret our bounds as stating that the value of an accurate stochastic forecasting model is mitigated by frequent re-optimization. We now turn to establishing these guarantees.

## 4.2. Preliminaries: The Unit Revenue Function

Define the *unit revenue function*  $g : \mathbb{R}_+ \rightarrow \mathbb{R}_+$  according to

$$g(y) = \begin{cases} p^* \bar{F}(p^*) y & \text{if } y \leq 1/\bar{F}(p^*); \\ \bar{F}^{-1}(1/y) & \text{otherwise.} \end{cases}$$

Notice that  $g(y)$  is the optimal value of the optimization problem <sup>10</sup>

$$\begin{aligned} & \underset{p_t \in \mathbb{R}_+, t \geq 0}{\text{maximize}} && \int_0^1 p_t y \bar{F}(p_t) dt \\ & \text{subject to} && \int_0^1 y \bar{F}(p_t) dt \leq 1. \end{aligned}$$

From the results of Gallego and van Ryzin [1994], it follows that in the limit as  $n$  grows large, the unit revenue function  $g(y)$  corresponds to the (normalized) optimal expected revenue in a system

<sup>10</sup>see Proposition 2 of Gallego and van Ryzin [1994].

with initial inventory  $n$  and instantaneous market size  $\Lambda_t = ny/T$ . The following lemma, proved in the appendix, establishes a number of properties for  $g$  that we will have opportunity to use in the sequel.

**Lemma 2.**

1.  $g(\cdot)$  is a non-negative, continuous, non-decreasing, and concave function on  $\mathbb{R}_+$ , with  $g(0) = 0$ .
2.  $yg(1/y)$  is non-decreasing and concave on  $\mathbb{R}_{++}$ .
3.  $g(y)/y$  is non-increasing on  $\mathbb{R}_+$ .
4. If  $u, v > 0$ , then  $\frac{g(u)}{g(v)} \geq \min(\frac{u}{v}, 1)$ ,  $\frac{1}{u} \int_0^u g(v)dv \leq g(u/2)$ .

**4.3. An Upper Bound on the Fluid Optimal Value function**

Consider the fluid system, but with the entire sample path of the market-size process, that is  $\{\Lambda_t : t \geq 0\}$ , available at time 0. The revenue optimization problem here is a deterministic one: In particular, let us denote by  $\tilde{J}_{\{\Lambda_s\}}^*(x_0, 0)$ , the optimal value of the revenue maximization problem:

$$\begin{aligned} & \underset{p_t \in \mathbb{R}_+, t \geq 0}{\text{maximize}} && \int_0^T p_t \Lambda_t \bar{F}(p_t) dt \\ & \text{subject to} && \int_0^T \Lambda_t \bar{F}(p_t) dt \leq x_0. \end{aligned}$$

Now since any policy  $\pi \in \Pi$  describes a feasible solution to the above problem, we immediately have that

$$\tilde{J}^*(x^0, \lambda^0, 0) \leq \mathbf{E} \left[ \tilde{J}_{\{\Lambda_s\}}^*(x_0, 0) \right].$$

In addition, we define the *certainty equivalent* fluid value function  $\tilde{J}_{\text{CE}}^*$  according to

$$\tilde{J}_{\text{CE}}^*(x_0, \lambda_0, 0) \triangleq \begin{aligned} & \underset{p_t \in \mathbb{R}_+, t \geq 0}{\text{maximize}} && \int_0^T p_t \mathbf{E}[\Lambda_t] \bar{F}(p_t) dt \\ & \text{subject to} && \int_0^T \mathbf{E}[\Lambda_t] \bar{F}(p_t) dt \leq x_0. \end{aligned}$$

We then have the following upper bounds on the fluid optimal value function  $\tilde{J}^*$ . The proof relies essentially on establishing the appropriate convexity and applications of Jensen's inequality; it may be found in the appendix:

**Lemma 3.**

$$\begin{aligned} \tilde{J}^*(x^0, \lambda^0, 0) & \leq \mathbf{E} \left[ \tilde{J}_{\{\Lambda_t\}}^*(x_0, 0) \right] \\ & \leq \tilde{J}_{\text{CE}}^*(x_0, \lambda_0, 0) \\ & \leq x_0 g \left( \frac{\int_0^T (\lambda_t + \sigma_t / \sqrt{2\pi}) dt}{x_0} \right) \\ & \leq x_0 g \left( \frac{\int_0^T \lambda_t dt}{x_0} \right) + x_0 g \left( \frac{\int_0^T \sigma_t dt}{x_0 \sqrt{2\pi}} \right). \end{aligned}$$

#### 4.4. A Lower Bound on $\tilde{J}^{\pi_{\text{RFP}}}$

Here we establish a lower bound on the value of the RFP policy in the fluid regime. With the upper bound we have established on the optimal fluid value function, this will lead to a performance bound in the fluid regime. In order to establish our bound we will first demonstrate that the RFP policy has a useful ‘balancing’ property that yields useful *uniform* lower bounds on the inventory process under *any* market-size sample path.

**Lemma 4.** (*Inventory Balancing*) *The inventory at time  $t$  under the RFP policy in the fluid model,  $\tilde{X}_t$ , satisfies*

$$(3) \quad \tilde{X}_t \geq \left[ \alpha \left( 1 - \frac{t}{T} \right) + (1 - \alpha) \frac{\int_t^T \lambda_s ds}{\int_0^T \lambda_s ds} \right] x_0.$$

**Proof.** By the definition of  $\pi_{\text{RFP}}$ ,

$$\bar{F}(\pi_{\text{RFP}}(\tilde{X}^t, \Lambda^t, t)) = \min \left\{ \bar{F}(p^*), \frac{\tilde{X}_t h(t, \alpha)}{\Lambda_t(T-t)} \right\}.$$

Consider an arbitrary sample path of the market-size process,  $\{\Lambda_t\}$ , we have

$$\begin{aligned} d\tilde{X}_t &= -\Lambda_t \bar{F}(\pi_{\text{RFP}}(\tilde{X}^t, \Lambda^t, t)) dt \\ &\geq -\frac{\tilde{X}_t h(t, \alpha)}{T-t} dt \\ &= \tilde{X}_t d \log \left( \alpha \left( 1 - \frac{t}{T} \right) + (1 - \alpha) \frac{\int_t^T \lambda_s ds}{\int_0^T \lambda_s ds} \right). \end{aligned}$$

Thus,

$$d \log(\tilde{X}_t) \geq d \log \left( \alpha \left( 1 - \frac{t}{T} \right) + (1 - \alpha) \frac{\int_t^T \lambda_s ds}{\int_0^T \lambda_s ds} \right).$$

Integrating on both sides and using the initial value  $\tilde{X}_0 = x_0$ , yields

$$\tilde{X}_t \geq \left[ \alpha \left( 1 - \frac{t}{T} \right) + (1 - \alpha) \frac{\int_t^T \lambda_s ds}{\int_0^T \lambda_s ds} \right] x_0. \quad \blacksquare$$

The above result reflects a natural ‘balancing’ property of the RFP policy. To see this consider the oblivious choice of  $\alpha = 1$ . Here, the above bound reduces to

$$\tilde{X}_t \geq \left( 1 - \frac{t}{T} \right) x_0,$$

so that when  $\alpha = 1$ , the RFP policy enforces a ‘balanced’ allocation of inventories across all time intervals irrespective of the actual realization of the market-size sample path. We next use this balancing property in a crucial way to establish a useful intermediate lower bound on the fluid value of the RFP policy. The proof of this bound, and all subsequent results may be found in the appendix.

**Lemma 5.**

$$\tilde{J}^{\pi_{\text{RFP}}}(x^0, \lambda^0, 0) \geq \alpha \frac{x_0}{T} \mathbb{E} \left[ \int_0^T g \left( \frac{\Lambda_t T}{x_0} \right) dt \right] + (1 - \alpha) \frac{x_0}{\int_0^T \lambda_s ds} \mathbb{E} \left[ \int_0^T \lambda_t g \left( \frac{\Lambda_t \int_0^T \lambda_s ds}{x_0 \lambda_t} \right) dt \right].$$

We next derive lower bounds on each of the summands in the lower bound above in the interest of deriving quantities comparable to the upper bound of Lemma 3. The proofs rely essentially on properties of the fluid optimization problem (as reflected in properties of the unit revenue function), and distributional properties of the marginals of the market size process. In particular, it is here that we exploit the Gaussianity of the marginals of our process via the results of Lemma 14. The other assumptions we made on the market-size process are also material for these bounds. In particular, we have:

**Lemma 6.**

$$\frac{1}{T} \mathbb{E} \left[ \int_0^T g \left( \frac{\Lambda_t T}{x_0} \right) dt \right] \geq 0.342g \left( \frac{\int_0^T \sigma_t dt}{x_0 \sqrt{2\pi}} \right).$$

and

**Lemma 7.**

$$\frac{1}{\int_0^T \lambda_s ds} \mathbb{E} \left[ \int_0^T \lambda_t g \left( \frac{\Lambda_t \int_0^T \lambda_s ds}{x_0 \lambda_t} \right) dt \right] \geq \frac{1}{2} g \left( \frac{\int_0^T \lambda_t dt}{x_0} \right).$$

The lower bounds we have established will allow us to establish performance bounds for the RFP policy in the fluid model that we provide next.

#### 4.5. Performance Guarantees in the Fluid Model

We begin with a guarantee relevant to the setting where the seller has absolutely *no* information about the market size process. He consequently sets  $\alpha = 1$ . Remarkably, we have:

**Theorem 1.** *Consider the RFP policy for an arbitrary  $\alpha$ . We then have for all market-size processes satisfying Assumption 2, and the additional requirement that  $\lambda_t = \lambda$  for all  $t$ ,*

$$\frac{\tilde{J}^{\pi_{\text{RFP}}}(x^0, \lambda^0, 0)}{\tilde{J}^*(x^0, \lambda^0, 0)} \geq \frac{\tilde{J}^{\pi_{\text{RFP}}}(x^0, \lambda^0, 0)}{\tilde{J}_{\text{CE}}^*(x^0, \lambda^0, 0)} \geq 0.342.$$

The guarantee above holds for arbitrarily volatile market size processes; stronger guarantees are available in an environment wherein volatility is low. In particular, we have:

**Theorem 2.** *Consider a market-size process with no drift and bounded volatility. That is assume that  $\lambda_t = \lambda$  and moreover,  $\sigma_T/\lambda \leq \sqrt{2\pi}\lambda B$  for some constant  $B > 0$ . Then,*

$$\frac{\tilde{J}^{\pi_{\text{RFP}}}(x^0, \lambda^0, 0)}{\tilde{J}^*(x^0, \lambda^0, 0)} \geq \frac{\tilde{J}^{\pi_{\text{RFP}}}(x^0, \lambda^0, 0)}{\tilde{J}_{\text{CE}}^*(x^0, \lambda^0, 0)} \geq \max \left\{ 0.342, \frac{1}{1+B} - \frac{B}{1+B} \left( \exp(-1/4\pi B^2) + 0.853 \right) \right\}.$$

In order to get a stronger sense for the parameter  $B$  in the parametric guarantee above, consider the setting where the market size process is a generalized moving average process. In this case if  $\phi(t) \leq \sigma$ , we automatically have  $B \leq \sigma T^{1/2}/\lambda$ ; a natural measure of the ‘relative’ uncertainty in the process. Similarly, if  $\phi(t) \leq \sigma e^{-\beta t}$  (so that we face an Ornstein-Uhlenbeck market size process), we

have  $B \leq \sigma/\lambda\beta^{1/2}$ , which again shows a natural dependence on what might consider the inherent volatility in the process.

In an ideal world, we would like the result of Theorems 1 and 2 which apply to the RFP policy with the oblivious choice of  $\alpha = 1$  (i.e. the policy which uses absolutely no information about the market size process) to hold for arbitrary  $\{\lambda_t\}$ ; this would unfortunately be asking for too much – in particular, in a setting where volatility were identically zero this would be akin to asking for guarantees relative to some *arbitrary, fixed sample path* of the market size process with no a-priori knowledge of this sample path. Nonetheless, we present a guarantee relevant for the RFP policy and *any* market size process satisfying Assumption 2 next. We assume for this next result that the seller *knows*  $\{\lambda_t\}$  (recall we think of this as a deterministic forecast) but no other information about the market size process and selects  $\alpha = 0.594$ ; as is evident from the proof of the result that follows, this selection of  $\alpha$  can be fine-tuned with additional information about the market size process.

**Theorem 3.** *Consider the RFP policy with  $\alpha = 0.594$ . We then have for all market-size processes satisfying Assumption 2,*

$$\frac{\tilde{J}^{\pi_{\text{RFP}}}(x^0, \lambda^0, 0)}{\tilde{J}^*(x^0, \lambda^0, 0)} \geq \frac{\tilde{J}^{\pi_{\text{RFP}}}(x^0, \lambda^0, 0)}{\tilde{J}_{\text{CE}}^*(x^0, \lambda^0, 0)} \geq 0.203.$$

The guarantees above were for the *idealized* RFP policy in a fluid model. The remainder of our analysis (i.e. the two sections that follow) will be devoted to showing that under appropriate conditions the guarantees here extend to the implementable RFP- $\Delta$  policy applied to the (stochastic) model we began with. We end this Section with noting that the analysis here could be extended to other market-size process families; Appendix B.1 establishes versions of Theorem 1 for reflected generalized moving average processes and also for Cox-Ingersoll-Ross (CIR) processes (which are an important example of an affine process).

## 5. Performance Analysis II: The Impact of Discretization

The previous section demonstrated performance guarantees for the idealized RFP policy in a fluid model. We continue to consider performance in the fluid model in this section, but now move to understanding the impact of discrete pricing updates. In particular, this section will consider the use of the RFP- $\Delta$  policy in the natural fluid model for our system. The goal of this section is two-fold. First, we must establish the natural fact that as pricing updates get more frequent, the RFP- $\Delta$  policy yields performance on par with that of the idealized RFP policy. Second, we provide operational insights for the choice of  $\Delta$  by demonstrating the precise rate at which the performance of the RFP- $\Delta$  policy approaches that of the RFP policy; in particular, we provide an estimate of the ‘price’ of such discrete reviews.

We begin by introducing some notation: We denote by  $\tilde{X}_t^\Delta$  the inventory process under the RFP- $\Delta$  policy in the fluid regime. Recall that at the  $i$ th price update, the RFP- $\Delta$  policy uses an estimate of the market size process based on sales in the preceding interval of time,  $\hat{\Lambda}_{i\Delta}$ . We denote the analogue to this quantity in the fluid model by  $\hat{\Lambda}_{i\Delta}^{\text{FL}}$ . We define:

$$\hat{\Lambda}_{i\Delta}^{\text{FL}} = \frac{1}{\Delta} \int_{i-1\Delta}^{i\Delta} \Lambda_s ds.$$

We begin with showing that the inventory process under the RFP- $\Delta$  policy approaches that under the idealized RFP policy. In particular:

**Lemma 8.** For all  $t \in [0, T)$ , we have

$$\lim_{\Delta \rightarrow 0} \tilde{X}_{t(\Delta)}^\Delta = \tilde{X}_t.$$

Now define the revenue garnered under the RFP on a specific sample path of the market size process according to the  $\mathcal{G}_0$ -measurable random variable

$$\tilde{J}^{\pi_{\text{RFP}}}(x^0, \lambda^0, 0 | \mathcal{G}_0) = \int_0^T \pi_{\text{RFP}}(\tilde{X}^t, \Lambda^t, t) \bar{F}(\pi_{\text{RFP}}(\tilde{X}^t, \Lambda^t, t)) \Lambda_t dt$$

and similarly define  $\tilde{J}^{\pi_{\text{RFP}}^\Delta}(x^0, \lambda^0, 0 | \mathcal{G}_0)$ . We then have that in the fluid model on almost every sample path of the market size process, the revenue garnered under the RFP- $\Delta$  policy approach those garnered under the RFP policy. In particular, we have:

**Lemma 9.**

$$\lim_{\Delta \rightarrow 0} \tilde{J}^{\pi_{\text{RFP}}^\Delta}(x^0, \lambda^0, 0 | \mathcal{G}_0) = \tilde{J}^{\pi_{\text{RFP}}}(x^0, \lambda^0, 0 | \mathcal{G}_0).$$

Via an appeal to the dominated convergence theorem, this result yields as an easy corollary (whose proof is omitted):

**Corollary 1.**

$$\lim_{\Delta \rightarrow 0} \tilde{J}^{\pi_{\text{RFP}}^\Delta}(x^0, \lambda^0, 0) = \tilde{J}^{\pi_{\text{RFP}}}(x^0, \lambda^0, 0).$$

## 5.1. The Price of Discrete Reviews

This section explores the *precise* loss in revenue due to the fact that prices can only be updated at discrete points in time. It stands to reason that this loss will depend first on the frequency of updates (in a limiting sense this is already clear by our results thus far), and second on the volatility of the underlying market size process. In particular, one expects that if the market size process is not too volatile then we can get away with relatively fewer price updates and conversely for the scenario where volatility is high. This sort of question has broad relevance to retail in general, and this section will yield a relatively simple rule of thumb to judge re-optimization frequency.

Our study here will restrict attention to generalized moving average processes with  $\lambda_t = \lambda$  and  $\phi(0) = \sigma > 0$ ; an extension to the more general case is straightforward but tedious and not terribly insightful. To begin, we note that the modulus of continuity established for sample paths of our market size process (Lemma 1) and the proof of Lemma 8 yield as an immediate corollary the following estimate whose proof is omitted:

**Corollary 2.** Assume that  $\bar{\Lambda}_t$  is a generalized moving average process with  $\phi \in \mathcal{C}_2$  and  $\lambda_t \in \mathcal{C}_1$ . Then, for any  $t \in [0, T)$ , we have almost surely:

$$\limsup_{\Delta \rightarrow 0} \frac{|\hat{\Lambda}_{t(\Delta)}^{\text{FL}} - \Lambda_t|}{\sigma \sqrt{\Delta \log(1/\Delta)}} \leq 2$$

and further,

$$\limsup_{\Delta \rightarrow 0} \frac{|\tilde{X}_{t(\Delta)}^\Delta - \tilde{X}_t|}{\sigma \sqrt{\Delta \log(1/\Delta)}} \leq 4$$

We use these results to establish a precise rate at which the revenue under the RFP- $\Delta$  policy approach those under the RFP policy. We do so for the ‘oblivious’ policy with  $\alpha = 1$ . Define  $B \triangleq \max_{t \in [0, T]} \Lambda_t$ , and  $\eta(\Delta) = \sqrt{\Delta \log(1/\Delta)}$ . We will then show that

$$\left| \tilde{J}_{\text{RFP}}^{\Delta}(x^0, \lambda^0, 0 | \mathcal{G}_0) - \tilde{J}_{\text{RFP}}(x^0, \lambda^0, 0 | \mathcal{G}_0) \right| = O(\sigma \eta(\Delta) \log(1/\eta(\Delta)))$$

where the only random term in the constant within the big-Oh notation is  $B$ . More precisely, we have the following theorem:

**Theorem 4.** (*The Price of Discretization*) *Under the assumptions on the market size process above, we have almost surely:*

$$\limsup_{\Delta \rightarrow 0} \frac{\left| \tilde{J}_{\text{RFP}}^{\Delta}(x^0, \lambda^0, 0 | \mathcal{G}_0) - \tilde{J}_{\text{RFP}}(x^0, \lambda^0, 0 | \mathcal{G}_0) \right|}{\sigma \eta(\Delta) \log(1/\eta(\Delta))} \leq \frac{4p^* \bar{F}(p^*) B^2 T^2}{x_0^2}$$

An application of the reverse Fatou Lemma and Jensen’s inequality immediately yields:

$$\limsup_{\Delta \rightarrow 0} \frac{\left| \tilde{J}_{\text{RFP}}(x^0, \lambda^0, 0) - \tilde{J}_{\text{RFP}}^{\Delta}(x^0, \lambda^0, 0) \right|}{\eta(\Delta) \log(1/\eta(\Delta))} \leq 4p^* \bar{F}(p^*) \sigma \frac{\mathbb{E} B^2 T^2}{x_0^2}$$

Given that the RFP- $\Delta$  policy with  $\alpha = 1$  or similar variants thereof already finds use among sophisticated retailers in pricing products with high demand volatility, this theorem provides a great deal of valuable intuition. Precisely, it shows how the frequency with which prices are updated should be adjusted relative to various intuitive problem parameters; the relationships are intuitive but *non-linear*.

1. Volatility: Ignoring logarithmic factors, if volatility  $\sigma$  is halved, one may shrink the frequency with which prices are updated by a factor of  $\sqrt{2}$  and maintain the same additive revenue loss. As one might anticipate, low volatility consequently calls for lower price update frequencies.
2. Load Factor: The quantity  $\sqrt{\mathbb{E} B^2 T} / x_0$  is a measure of load, i.e. demand relative to inventory, or at least an upper bound thereof. It stands to reason that when inventory is scarce relative to demand one might wish to update prices frequently; indeed that is precisely what our result suggests – the revenue loss due to discretization scales like the square of this measure of load. In particular, the more inventory one has relative to demand, the less important it is to update price. Again ignoring logarithmic factors, doubling the amount of inventory will permit cutting the price review frequency by a factor of 4.

## 6. Performance Analysis III: The High Volume Regime

The previous sections analyzed the performance of the RFP (and RFP- $\Delta$ ) policies in a stochastic fluid model. Whereas the fluid model retained a stochastic market-size process, the customer arrival process was no longer a stochastic point process; rather demand was assumed to be infinitely divisible and the customer arrival process was thus a ‘fluid’ whose rate was modulated by the underlying market-size process. This section establishes that the fluid model of the previous section can be viewed as a limit of the stochastic model we proposed in Section 2, in an appropriate ‘high volume’ regime characterized by large initial inventory levels and proportionately scaled market-size processes. In particular, allowing ourselves some imprecision for the purpose of exposition, this

will imply the following high-level result: Provided one adjusts prices sufficiently frequently, the RFP- $\Delta$  policy (our implementable prescription) inherits, in an appropriately high volume regime, the performance guarantees established for the *idealized* RFP policy in the fluid model. In particular, this section will close the loop in our analysis and provide performance guarantees for the implementable RFP- $\Delta$  policy in our original stochastic model.

Before describing an overview of the results in detail, we first describe the high volume regime precisely.

**The ‘High Volume’ Regime:** We assume we are given a dynamic pricing problem of the type described in Section 2, with initial inventory level  $x_0$ , and market-size process  $\{\Lambda_t : t \geq 0\}$ . We next consider an  $n$ th ‘scaled’ version of this problem with initial inventory level  $nx_0 \triangleq x_{0,(n)}$  and market-size process  $\{\Lambda_{t,(n)}\}$ , where  $\Lambda_{t,(n)} = n\Lambda_t$ . All other problem parameters (namely, the time horizon,  $T$ , and the reservation price distribution of an individual arriving customer,  $F(\cdot)$ ) remain unchanged in this  $n$ th problem. By ‘high volume’ we understand the regime where  $n$  grows large.

The eventual goal of this section will be to show that the performance guarantees we demonstrated for the RFP- $\Delta$  policy in the fluid model also hold when the RFP- $\Delta$  policy as described in Section 3 is employed in the high volume regime. Notice that this is *distinct* from analyzing the use of the fluid policy (i.e. a policy that anticipates that the inventory process behaves according to the fluid model) in the stochastic system. The analysis will proceed in two broad steps:

- We show that the inventory process in the stochastic model under the RFP- $\Delta$  policy is close to the inventory process in the fluid model in the high-volume regime; along the way we must establish that the market-size forecasts used are also close to their fluid analogues.
- We use the above result to show that revenue under the RFP- $\Delta$  policy in the fluid model are close to revenues under the RFP- $\Delta$  policy in the stochastic model (Corollary 1). We use this result in turn to show our main result for this section, Lemma 12 which shows that the performance guarantees we established in the fluid model remain for the *idealized* RFP policy remain valid in the high volume regime under the use of the RFP- $\Delta$  policy.

## 6.1. The Inventory Process and Convergence to the Fluid Value

We begin with some notation relevant to this section. We denote by  $X_{t,(n)}^\Delta$  the inventory at time  $t$  under the RFP- $\Delta$  policy in the  $n$ th (stochastic) problem. Similarly, we denote by  $\hat{\Lambda}_{t(\Delta),(n)}$  the inventory forecast in place over the interval  $[t(\Delta), t(\Delta) + \Delta)$  in the  $n$ th problem. We then have:

**Lemma 10.** *For any given  $\Delta > 0$ , we have that the inventory process and the forecast process in the  $n$ th system converge to their fluid analogues. In particular:*

1. For any  $t$ ,

$$\frac{X_{t,(n)}^\Delta}{n} \rightarrow \tilde{X}_t^\Delta \quad \text{a.s.}$$

2. For any  $t$  such that  $\tilde{X}_t^\Delta > 0$ ,

$$\frac{\hat{\Lambda}_{t(\Delta),(n)}}{n} \rightarrow \hat{\Lambda}_{t(\Delta)}^{\text{FL}} \quad \text{a.s.}$$

This result suffice for us to show that on every realization of the market size process, we have in the high-volume regime that the revenue yielded in the stochastic model converges to the revenue yielded in the fluid model under RFP –  $\Delta$  policy for any fixed  $\Delta$ . In particular, we have:

**Lemma 11.**

$$\liminf_n \frac{1}{n} J^{\pi_{\text{RFP}}^\Delta} \left( x_{(n)}^0, \lambda_{(n)}^0, 0 | \mathcal{G}_0 \right) \geq \tilde{J}^{\pi_{\text{RFP}}^\Delta} \left( x^0, \lambda^0, 0 | \mathcal{G}_0 \right) \text{ a.s.}$$

As a corollary to this result we have:

**Corollary 3.** *For  $\Delta$  sufficiently small, we have:*

$$\liminf_n \frac{J^{\pi_{\text{RFP}}^\Delta} \left( x_{(n)}^0, \lambda_{(n)}^0, 0 \right)}{\tilde{J}^{\pi_{\text{RFP}}^\Delta} \left( x_{(n)}^0, \lambda_{(n)}^0, 0 \right)} \geq 1.$$

## 6.2. Closing the Loop: Performance Guarantees in the High Volume Regime

We finally present a result that yields performance guarantees for our implementable prescription, the RFP- $\Delta$  policy. Our result will imply that the guarantees we established for the *idealized* RFP policy in the fluid model, apply just as well to the RFP- $\Delta$  policy when employed in the actual stochastic setting under appropriate conditions. In particular, we consider a setting where:

1. The re-optimization frequency is high (i.e.  $\Delta$  is small).
2. The volume of sales is high (i.e.  $n$  is large).

Our work thus far, and in particular, Corollaries 1 and 3 allows us to obtain the following natural result:

**Lemma 12.** *We have*

$$\liminf_{\Delta \rightarrow 0} \liminf_{n \rightarrow \infty} \frac{J^{\pi_{\text{RFP}}^\Delta} \left( x_{(n)}^0, \lambda_{(n)}^0, 0 \right)}{J^* \left( x_{(n)}^0, \lambda_{(n)}^0, 0 \right)} \geq \frac{\tilde{J}^{\pi_{\text{RFP}}^\Delta} \left( x^0, \lambda^0, 0 \right)}{\tilde{J}_{\text{CE}}^* \left( x^0, \lambda^0, 0 \right)}.$$

This Lemma serves to bridge the results we established for the *idealized* RFP policy in the *fluid* model to results for the RFP- $\Delta$  policy in the model with point process arrivals. More specifically, the guarantees established in Theorems 1, 2, and 3 now apply to the RFP- $\Delta$  policy in the high volume regime. As an example, one has for instance, using Theorem 2 and Lemma 12 that for the RFP- $\Delta$  policy using  $\alpha = 1$

$$(4) \liminf_{\Delta \rightarrow 0} \liminf_{n \rightarrow \infty} \frac{J^{\pi_{\text{RFP}}^\Delta} \left( x_{(n)}^0, \lambda_{(n)}^0, 0 \right)}{J^* \left( x_{(n)}^0, \lambda_{(n)}^0, 0 \right)} \geq \max \left\{ 0.342, \frac{1}{1+B} - \frac{B}{1+B} \left( \exp(-1/4\pi B^2) + 0.853 \right) \right\}.$$

when  $\lambda_t = \lambda$  and  $\sigma_T/\lambda \leq \sqrt{2\pi}\lambda B$ . One obtains similar corollaries to Theorem 1 and 3.

In addition, the proof of Lemma 12 reveals that under the conditions of Theorem 4, we have in fact that almost surely

$$\liminf_{n \rightarrow \infty} \frac{J^{\pi_{\text{RFP}}^\Delta} \left( x_{(n)}^0, \lambda_{(n)}^0, 0 \right)}{J^* \left( x_{(n)}^0, \lambda_{(n)}^0, 0 \right)} \geq \frac{\tilde{J}^{\pi_{\text{RFP}}^\Delta} \left( x^0, \lambda^0, 0 \right)}{\tilde{J}_{\text{CE}}^* \left( x^0, \lambda^0, 0 \right)} - O(\sigma\eta(\Delta) \log(1/\eta(\Delta)))$$

In particular, applying Theorem 2 one then has:

$$(5) \quad \liminf_{n \rightarrow \infty} \frac{J^{\pi_{\text{RFP}}^{\Delta}}(x_{(n)}^0, \lambda_{(n)}^0, 0)}{J^*(x_{(n)}^0, \lambda_{(n)}^0, 0)} \geq \max \left\{ 0.342, \frac{1}{1+B} - \frac{B}{1+B} \left( \exp(-1/4\pi B^2) + 0.853 \right) \right\} \\ - O(\sigma\eta(\Delta) \log(1/\eta(\Delta))),$$

an apparent refinement to the guarantee (4). Of course, with the use of the appropriate fluid model guarantee a number of other performance guarantees are possible; as opposed to listing these exhaustively, we stop here and reflect on the guarantee (5). The guarantee exposes a number of factors that influence the design of the RFP policy and influence its performance:

1. Volatility: Volatility influences performance loss in two spots – first, low volatility implies, not surprisingly, that loss relative to the certainty equivalent policy is small (as clarified by the first term in the bound of (5) where  $B$  shrinks linearly with  $\sigma_T$ ) and moreover that the loss due to discrete price reviews is also small (the second term in the bound). Conversely, high volatility has adverse effects. The component of the loss relative to the certainty equivalent policy grows, but there is a uniform upper bound to the relative loss (namely,  $\sim 34\%$  in this setting). There is an additional additive loss component due to discrete price reviews and in order to combat this loss one must increase review frequency.
2. Review Frequency: The fact that the performance loss of the idealized RFP policy is uniformly bounded in the face of arbitrary volatility can, at its root, be traced to the simple inventory balancing property demonstrated for that policy. This property, in turn, relied crucially on the continuous price updates inherent to the idealized policy. We see that when one moves away from this idealized policy one incurs an additive loss that depends on volatility and the frequency of re-optimization. This additive loss can be controlled by adjusting re-optimization frequency appropriately in response to the volatility one anticipates.

## 7. Numerical Validation

This section is dedicated to a computational exploration of the RFP- $\Delta$  policy that we have studied carefully to this point. We are particularly interested in the policy with the oblivious choice of  $\alpha = 1$ ; this is the policy one may implement with absolutely no information about the market size process and thereby of potentially greatest interest to practitioners. We will be interested in exploring the following questions:

1. What are ‘difficult’ (or conversely, ‘easy’) regimes for our problem? For example, we find it natural to conjecture that as the relative volatility grows large we anticipate a degradation in performance. As another example, as one increases inventory available while leaving the market size process fixed, the problem at hand becomes easier. We will carefully explore the performance of the RFP- $\Delta$  policy across the scenarios described above in regimes that straddle reality.
2. Tuning  $\Delta$ : In our analysis of the price of discretization, we established various qualitative dependencies of the loss in revenue due to our inability to update prices continuously as a function of various problem parameters of interest. We are interested in seeing these insights reflected in our computational experiments.

3. **Naive Re-Optimization:** A natural scheme involving re-optimization in the *presence* of a forecast believed to be ‘good’ (i.e. knowledge of  $\lambda_t$  and the belief that  $\phi$  is identically 0) involves a re-optimization scheme which in our language would translate to the RFP- $\Delta$  scheme with  $\alpha$  set to zero. Since there may indeed be settings where this is reasonable (indeed, consider the setting where it is in fact the case that  $\phi$  is zero and  $\lambda_t$  is correctly known), we wish to explore whether the RFP- $\Delta$  scheme with  $\alpha = 1$  is viable in this setting. We also explore our robust choice of  $\alpha$  designed for a setting where  $\lambda_t$  is known but one wishes to be oblivious to the nature of  $\phi$ .

## 7.1. Problem Specifications

We examine RFP- $\Delta$  policy performance numerically for a class of market-size processes that are O-U type processes and evolve according to

$$\bar{\Lambda}_t = \lambda_t + \sigma \int_0^t e^{-\beta(t-s)} dZ_s.$$

In the bulk of our experiments we will consider  $\lambda_t$  to be a constant; the RFP- $\Delta$  policy will neither be aware of this constant nor know the form or specification of  $\phi$ . We will also explore the case where  $\lambda_t$  is not constant in Section 7.4. We consider exponentially distributed customer reservation prices so that  $F(p) = 1 - e^{-p}$ . We will compare ourselves to a super-optimal policy that knows  $\{\Lambda_t : t \geq 0\}$  at time 0. This yields the (easy to compute) upper bound:

$$\mathbb{E} \left[ J_{\{\Lambda_t\}}^*(x_0, 0) \right] = \mathbb{E} \left[ \log \left( \sum_{n=0}^{x_0} \left( \frac{\int_0^T \Lambda_t dt}{e} \right)^n \frac{1}{n!} \right) \right].$$

While computationally convenient, this bound can be quite loose in high volatility regimes, and as such we will also consider another super-optimal bound corresponding to a policy that merely knows the specification of  $\Lambda_t$  and can monitor the market size process; we compute the corresponding optimal policy by numerically solving the associated HJB equation.

The specific choice of  $x_0$ ,  $\lambda_t$ ,  $\beta$  and  $\sigma$  will vary across our experiments, while  $T$  will be normalized to 5. Moreover, we will keep the quantity  $n\Delta$  fixed and equal to 100 in all of our experiments; this ensures an ‘apples to apples’ comparison that adjusts for forecast in-accuracies as we vary  $\Delta$ . Moving forward, we will often be interested in the following summary statistics about a particular ensemble of problem instances:

- **Coefficient of Variation:** We define

$$\text{CV} \triangleq \limsup_{t \rightarrow \infty} \frac{\sqrt{\text{var} [\bar{\Lambda}_t]}}{\mathbb{E} [\bar{\Lambda}_t]}.$$

In the special case where  $\lambda_t = \lambda$  the limit above exists and is  $\sigma/\lambda\sqrt{2\beta}$ . This is a natural measure of the relative volatility in the underlying market size process. An example of a ‘high’ CV in practice is typically on the order of 1 to 2; we will go as far as 5 in our experiments.

- **Load Factor:** We define load factor as the quantity  $x_0/\lambda T$ . This is a measure of the abundance of inventory relative to demand. Large values typically signal easier problems, and in reality it is fair to anticipate load factors close to 1; we will consider load factors as low as 0.3 in our experiments.

We next set out to investigate each of the issues outlined at the outset of this section.

## 7.2. Performance Across Varying Volatility and Load Regimes

Here we seek to understand how the RFP- $\Delta$  policy performs across varying problem regimes. In particular, we take  $\lambda_t = \lambda = e$  and  $\beta = 1$  and vary  $x_0$  and  $\sigma$  so as to create various combinations of relative volatility (the ‘CV’ measure) and load factor. Recall, that we expect problems with high CV and low load factor (i.e scarce inventory) to be the most challenging. In the experiments below, we take  $\Delta = 0.1$ .

**Table 1:** A Lower Bound on Relative Optimality (i.e.  $J^{\pi_{\text{RFP}}^\Delta}/J^*$ ) across various CV/load factor combinations. Common parameters across problem instances:  $\lambda = e, \beta = 1, n = 1000, T = 5, \Delta = 0.1$ .

Initial Inventory $x_0^{(n)}/n$	Load Factor $x_0/\lambda T$	Relative Optimality			
		CV=0.5	1	2.5	5
4	0.294	0.967	0.943	0.880	0.824
8	0.589	0.995	0.989	0.938	0.879
12	0.883	1.000	0.999	0.972	0.914
16	1.177	1.000	1.000	0.989	0.937
20	1.472	1.000	1.000	0.996	0.958

The results here are very encouraging with performance generally within 95% for parameter regimes one might encounter in practice. For extremely high CV and low load factor, the performance is as bad as 82%; while this performance loss occurs in a regime quite far away from what one encounters in practice, it is worth examining this issue further. As it turns out, it is not really the case that the RFP- $\Delta$  policy degrades in this setting but rather that the upper bound we use on the optimal value function is weak. In particular, if one computed a certain tighter but substantially more difficult to compute super-optimal policy in this setting (by considering the optimal policy that was allowed to observe  $\Lambda_t$  causally and knew the distribution over its sample paths) and compared performance against *this* policy, one again obtains performance figures essentially within 95%; see Appendix E.3.

## 7.3. Selecting A Re-Optimization Frequency

Section 5.1 developed theory around the ‘price’ of discretization, i.e. the revenue loss inherent in the fact that we permitted a limited number of price updates. To summarize that theory, we characterized precisely how volatility impacted this revenue loss, all else being the same, and further how greater amounts of inventory permitted more infrequent discretization, all else being the same. Here we try to understand these tradeoffs in a more concrete setting. In particular, we examine performance loss as a function of discretization frequency across various combinations of load factor (which implicitly translates to varying levels of initial inventory) and CV (which translates to varying  $\sigma$ ). The results are summarized in the following tables:

**Table 2:** A Lower Bound on Relative Optimality (i.e.  $J^{\pi_{\text{RFP}}^{\Delta}}/J^*$ ) as a function of review frequency across varying values of load factor and market size volatility. Common parameters across problem instances:  $\lambda = e, \beta = 1, T = 5, n\Delta = 100$ .

	$x_0/\lambda T = 0.368$				0.736				1.104			
CV	$\Delta=0.1$	0.5	1	2.5	0.1	0.5	1	2.5	0.1	0.5	1	2.5
0.1	0.999	0.999	0.997	0.992	1.000	1.000	1.000	1.000	1.000	1.000	1.000	1.000
0.5	0.984	0.980	0.973	0.948	1.000	1.000	1.000	1.000	1.000	1.000	1.000	1.000
1	0.958	0.948	0.933	0.877	0.998	0.997	0.997	0.995	1.000	1.000	1.000	0.999
2	0.915	0.886	0.847	0.748	0.975	0.972	0.967	0.950	0.997	0.996	0.995	0.993

What Table 2 shows is entirely in line with our theoretical development. What this adds to the theoretical development is the surprising fact that in absolute terms one needs relatively little re-optimization to capture most of the gains of the RFP- $\Delta$  policy in practical regimes. In particular, under the majority of circumstances, *ten or even five* price updates suffices to get within 90% (and frequently, 95%) of our loose upper bound; in light of Appendix E.3, this is likely to be essentially optimal given how loose the upper bound is.

#### 7.4. The Hedging Parameter $\alpha$ and the Gain Over Re-optimization Without Forecasting

The literature abounds with example of re-optimization schemes *without* forecast updates; in the language of the RFP policy this corresponds to setting  $\alpha = 0$ . In particular, this is a setting where the manager has a forecast  $\{\lambda_t\}$  which he believes to be perfect. Moreover, he believes and that  $\phi$  is identically zero. This section seeks to answer the following questions:

1. How might the RFP- $\Delta$  policy with forecast updates but oblivious to any knowledge of  $\{\lambda_t\}$  perform in this setting, i.e. how does the RFP- $\Delta$  policy with  $\alpha = 1$  perform here?
2. Moreover, if the manager did indeed have knowledge of  $\{\lambda_t\}$  can he hedge between this perfect forecast and a scenario where his forecast is corrupted by noise that is difficult to model, i.e. how does our robust choice of  $\alpha (= 0.594)$  for scenarios where  $\{\lambda_t\}$  is available fare?

We assume here that  $\lambda_t$  evolves according to<sup>11</sup>

$$\frac{\dot{\lambda}_t}{\lambda - \lambda_t} = p + q \frac{\lambda_t}{\lambda}, \lambda_0 = 0.$$

We measure relative uncertainty as  $CV = \sigma / \max_{0 \leq t \leq T} \lambda_t \sqrt{2\beta}$ . We fix problem scale (at  $n = 50$ ) and consider three sets of experiments, each corresponding to different levels of CV (0.5, 2.5 and 5). In each set of experiments we consider the performance of the RFP policy for various settings of  $\alpha$ ; we are most interested in the setting where  $\alpha = 1$  and  $\alpha = 0.594$  which are respectively settings appropriate to no knowledge of any specification of the market size process whatsoever, and knowledge of  $\{\lambda_t\}$ . In order to tease apart the effect of discrete reviews and errors in estimating the current market size, we consider the idealized RFP policy here. The results can be found in Tables 3. We draw the following principle conclusions:

<sup>11</sup>This is the so-called Bass model (Bass [1969]) which is widely used to characterize how a new product or service grows after it is introduced to the market.  $p$  and  $q$  are termed the coefficient of innovation and imitation respectively,  $\lambda$  represents potential market size

1. Except for scenarios where CV is very low, re-optimization without forecast updates can be improved upon dramatically.
2. The oblivious choice of  $\alpha = 1$ , wherein the manager requires absolutely no knowledge of the market size process performs surprisingly well; notice that this is a scenario outside of the purview of our analysis since  $\lambda_t$  is no longer constant.
3. The robust choice of  $\alpha$  appears to provide the hedging suggested by the theory providing intermediate performance; it is closer to the naive scheme when forecasts are exact and closer to the oblivious scheme when they are not. That said, even at its worst, the oblivious scheme with  $\alpha = 1$  incurs a marginal loss relative to the best choice of  $\alpha$ .

**Table 3:** Lower Bounds on Relative Optimality (i.e.  $J^{\pi_{RFP}}/J^*$ ) across varying CV for different settings of  $\alpha$  in the idealized RFP policy. Common parameters across problem instances:  $\lambda = e, \beta = 1, n = 50, T = 20, p = 0.03, q = 0.5$ .

CV	Load Factor	Naive Re-opt.	Oblivious/ Robust Choice	
	$x_0/\int_0^T \lambda_t dt$	$\alpha = 0$	$\alpha = 1$	$\alpha = 0.594$
.5	0.250	0.987	0.971	0.983
	0.417	0.991	1.000	1.000
	0.583	0.996	1.000	1.000
	0.750	0.997	1.000	1.000
2.5	0.250	0.822	0.914	0.893
	0.417	0.831	0.939	0.915
	0.583	0.850	0.963	0.939
	0.750	0.871	0.981	0.961
5	0.250	0.704	0.820	0.783
	0.417	0.697	0.825	0.785
	0.583	0.696	0.833	0.791
	0.750	0.699	0.842	0.799

## 7.5. Summary of Experimental Conclusions

Our computational experiments provide valuable insights on the questions we set out to answer. In particular:

1. The performance of the RFP- $\Delta$  policy with  $\alpha$  set to 1 (so that no market size information whatsoever is required) is robust across a broad swathe of parameter regimes that control volatility and the scarcity of inventory. Performance well within 95% can be expected for parameter regimes of practical interest. Degradation beyond this point can be attributed largely to the fact that the clairvoyant upper bound we compare ourselves against gets exceedingly loose
2. The revenue loss due to discrete reviews behaves largely as predicted by the theory; in particular, high volatility and low inventory levels call for higher review frequencies. Surprisingly, in absolute terms, this frequency was quite low; ten, and frequently just five price updates tended to suffice.

3. The oblivious setting of  $\alpha = 1$  appears to provide excellent performance even in settings where  $\lambda_t$  is itself time varying and unknown to the manager. A robust selection of the  $\alpha$  parameter provides added value in settings where  $\lambda_t$  is time varying and known and volatility is low.

## 8. Concluding Remarks

The present paper revisited a classical model of dynamic pricing in order to address an important reality: forecasts are rarely accurate, and retailers frequently witness demand shocks that are material to their revenues. The natural cure for such issues is typically the incorporation of *stochastic* forecast models into the relevant dynamic pricing problem; such a cure is best avoidable if possible: such forecast models are difficult to calibrate and their predictive power in practice is questionable. Fortunately for us, it appears that at least in the context of single product dynamic pricing one may well be able to deal with the issue of imperfect forecasts and potentially large demand shocks with a combination of re-optimization and ‘running-average’ type forecasts. In particular, we presented a simple dynamic pricing policy (the RFP- $\Delta$  policy) that can in fact be shown to be competitive with a clairvoyant policy with a-priori access to information about demand evolution over the course of the selling season. The policy we presented is easy to implement: in the guise we focused on primarily (namely, the setting where  $\alpha = 1$ ), the policy required *no* initial information about the market size process whatsoever, nor the ability to monitor it over time. The policy was simply allowed a finite number of price adjustments that were made on the basis of observed sales.

There are several extensions to the present work possible along the lines of extending the scope of the market size processes the analysis applies to, and incorporating learning of the reservation price distribution. By far, the most interesting direction to pursue perhaps is an understanding of what can be done in the multi-dimensional setting (i.e. the setting in Gallego and van Ryzin [1997]). Doing so would require that we first understand how one might accomplish the requisite ‘inventory balancing’ in that setting. Finally, it is worth noting that many retailers employ a pricing strategy closely related to the RFP policy in practice (especially towards the end of the selling season); what is typically missing is a careful understanding of what sort of forecast to use. A real-world study with the present policy would thus not require a big departure from current practice and would be of great value.

**Acknowledgments:** Ben Van Roy gave us the idea that using re-optimization in a dynamic pricing context might obviate the need to know the specification of what we have called the market-size process here.

## References

- M. Akan and B. Ata. Bid-price controls for network revenue management: Martingale characterization of optimal bid prices. *Math. Oper. Res.*, 34 (4):912–936, 2009.
- V. Araman and R. Caldentey. Dynamic pricing for non-perishable products with demand learning. *Operations Research*, 2009.
- Y. Aviv and A. Pazgal. Pricing of short life-cycle products through active learning. *Working Paper*, 2005a.
- Y. Aviv and A. Pazgal. A partially observed Markov decision process for dynamic pricing. *Management Sci.*, 51 (8):1400–1416, 2005b.

- F. Bass. A new product growth for model consumer durables. *Management Sci.*, 15 (5):215–227, 1969.
- O. Besbes and C. Maglaras. Revenue optimization for a make-to-order queue in an uncertain market environment. *Oper. Res.*, 57 (6):1438–1450, 2009.
- O. Besbes and A. Zeevi. Dynamic pricing without knowing the demand function: Risk bounds and near-optimal algorithms. *Oper. Res.*, 57 (6):1407–1420, 2009.
- G. Bitran and R. Caldentey. An overview of pricing models for revenue management. *Manufacturing Service Oper. Management*, 5 (3):203–229, 2003.
- G. Bitran and S. Mondschein. Periodic pricing of seasonal products in retailing. *Management Sci.*, 43 (1):64–79, 1997.
- P. Bremaud. *Point Processes and Queues, Martingale Dynamics*. Springer-Verlag, NY, 1980.
- W. Elmaghraby and P. Keskinocak. Dynamic pricing: Research overview, current practices and future directions. *Management Sci.*, 49 (10):1287–1309, 2003.
- V. Farias and B. Van Roy. Dynamic pricing with a prior on market response. *Oper. Res.*, 58 (1):16–29, 2010.
- G. Gallego and G. van Ryzin. Optimal dynamic pricing of inventories with stochastic demand over finite horizons. *Management Sci.*, 40 (8):999–1020, 1994.
- G. Gallego and G. van Ryzin. A multiproduct dynamic pricing problem and its applications to network yield management. *Oper. Res.*, 45 (1):24–41, 1997.
- V. Gaur, A. Giloni, and S. Seshadri. Information sharing in a supply chain under ARMA demand. *Management Sci.*, 51 (6):961–969, 2005.
- S. Graves. A single-item inventory model for a nonstationary demand process. *Manufacturing Service Oper. Management*, 1 (1):50–61, 1999.
- H. Lee, K. So, and C. Tang. The value of information sharing in a two-level supply chain. *Management Sci.*, 46 (5):626–643, 2000.
- C. Maglaras. Revenue management for a multiclass single-server queue via a fluid model analysis. *Oper. Res.*, 54 (5):914–932, 2006.
- C. Maglaras and J. Meissner. Dynamic pricing strategies for multiproduct revenue management problems. *Manufacturing Service Oper. Management*, 8 (2):136–148, 2006.
- J. McGill and G. van Ryzin. Revenue management: Research overview and prospects. *Transportation Sci.*, 33 (2):233–256, 1999.
- S. Raghunathan. Information sharing in a supply chain: A note on its value when demand is nonstationary. *Management Sci.*, 47 (4):605–610, 2001.
- K. Talluri and G. van Ryzin. *The Theory and Practice of Revenue Management*. Springer Science+Business Media, 2004.

X. Xu and W. Hopp. Dynamic pricing and inventory control: the value of demand learning. *Working paper*, 2004.

W. Zhao and Y. Zheng. Optimal dynamic pricing for perishable assets with nonhomogeneous demand. *Management Sci.*, 46 (3):375–388, 2000.

## A. Proofs for Section 2

**Lemma 1.** (*Sample Path Modulus of Continuity*) Assume that  $\bar{\Lambda}_t$  is a generalized moving average process with  $\phi \in \mathcal{C}_2$  and  $\lambda_t \in \mathcal{C}_1$ . Then, for  $\tau > 0$ , and any  $t \in [0, T)$ , we have:

$$\limsup_{\tau \rightarrow 0} \frac{|\Lambda_{t+\tau} - \Lambda_t|}{\sigma \sqrt{2\tau \log(1/\tau)}} \leq 1 \text{ a.s.}$$

**Proof.** We have

$$\begin{aligned} |\Lambda_{t+\tau} - \Lambda_t| &= \left| \left( \lambda_{t+\tau} + \int_0^{t+\tau} \phi(t+\tau-s) dZ_s \right)^+ - \left( \lambda_t + \int_0^t \phi(t-s) dZ_s \right)^+ \right| \\ &\leq |\lambda_{t+\tau} - \lambda_t| + \left| \int_t^{t+\tau} \phi(t+\tau-s) dZ_s \right| + \left| \int_0^t (\phi(t-s) - \phi(t+\tau-s)) dZ_s \right| \\ &= |\lambda_{t+\tau} - \lambda_t| + \left| \sigma Z_{t+\tau} - \phi(\tau) Z_t + \int_t^{t+\tau} \phi'(t+\tau-s) Z_s ds \right| \\ &\quad + \left| (\phi(0) - \phi(\tau)) Z_t + \int_0^t (\phi'(t-s) - \phi'(t+\tau-s)) Z_s ds \right| \\ &\leq |\lambda_{t+\tau} - \lambda_t| + \sigma |Z_{t+\tau} - Z_t| + (\phi(0) - \phi(\tau)) |Z_t| + \left| \int_t^{t+\tau} \phi'(t+\tau-s) Z_s ds \right| \\ &\quad + (\phi(0) - \phi(\tau)) |Z_t| + \left| \int_0^t (\phi'(t-s) - \phi'(t+\tau-s)) Z_s ds \right| \\ &\leq L_\lambda \tau + \sigma \sup_{t \leq s \leq t+1-\tau} |Z_{s+\tau} - Z_s| + L_{\phi_1} \tau B + L_{\phi_1} \tau B + L_{\phi_1} \tau B + L_{\phi_2} \tau B t \\ &= \sigma \sup_{t \leq s \leq t+1-\tau} |Z_{s+\tau} - Z_s| + (L_\lambda + 3L_{\phi_1} B + L_{\phi_2} B t) \tau \end{aligned}$$

where  $B \triangleq \sup_{0 \leq t \leq T} Z_t$ . The first inequality follows property that  $|(A+B)^+ - (C+D)^+| \leq |A-C| + |B-D|$ , the second equality follows from the integration by parts formulas for stochastic integrals, the third inequality follows from the assumed differentiability properties of  $\lambda_t$  and  $\phi(t)$  (the constants correspond to bounds on the appropriate differentials) and the definition of  $B$ .

Now, we have

$$\begin{aligned} \limsup_{\tau \rightarrow 0} \frac{|\Lambda_{t+\tau} - \Lambda_t|}{\sqrt{2\tau \log(1/\tau)}} &\leq \limsup_{\tau \rightarrow 0} \frac{\sigma \sup_{t \leq s \leq t+1-\tau} |Z_{s+\tau} - Z_s| + (L_\lambda + 3L_{\phi_1} B + L_{\phi_2} B t) \tau}{\sqrt{2\tau \log(1/\tau)}} \\ &= \limsup_{\tau \rightarrow 0} \frac{\sigma \sup_{t \leq s \leq t+1-\tau} |Z_{s+\tau} - Z_s|}{\sqrt{2\tau \log(1/\tau)}} \\ &= \sigma, \end{aligned}$$

where the first inequality follows from the first part of our argument, and the second inequality is Levy's theorem on the modulus of continuity of sample paths of Brownian motion.  $\blacksquare$

## B. Proofs for Section 4

### Lemma 2.

1.  $g(\cdot)$  is a non-negative, continuous, non-decreasing, and concave function on  $\mathbb{R}_+$ , with  $g(0) = 0$ .
2.  $yg(1/y)$  is non-decreasing and concave on  $\mathbb{R}_{++}$ .
3.  $g(y)/y$  is non-increasing on  $\mathbb{R}_+$ .
4. If  $u, v > 0$ , then  $\frac{g(u)}{g(v)} \geq \min(\frac{u}{v}, 1)$ ,  $\frac{1}{u} \int_0^u g(v)dv \leq g(u/2)$ .

### Proof.

1. That  $g(\cdot)$  is non-negative, continuous and non-decreasing with  $g(0) = 0$  follows by definition. We show  $g(\cdot)$  is a concave function. In the remainder of the proof, we use the fact that  $(p\bar{F}(p))'|_{p=p^*} = \bar{F}(p^*) - p^*f(p^*) = 0$ . We know that on  $y \leq 1/\bar{F}(p^*)$ ,  $g'(y) = p^*\bar{F}(p^*)$ . Now on  $y \geq 1/\bar{F}(p^*)$ ,  $g(y)$  is increasing in  $y$  and we have  $g'(y) = \bar{F}^2(g(y))/f(g(y))$ , which in turn must be non-increasing following the second part of Assumption 1 that  $\bar{F}(p)/f(p)$  is non-increasing. Finally,

$$\bar{F}^2(g(1/\bar{F}(p^*))) / f(g(1/\bar{F}(p^*))) = \bar{F}^2(p^*) / f(p^*) = p^*\bar{F}(p^*).$$

so that  $g(\cdot)$  is continuously differentiable on  $\mathbb{R}_+$  with a non-increasing derivative. Thus,  $g(\cdot)$  is concave on  $\mathbb{R}_+$ .

2. Note that

$$yg(1/y) = \begin{cases} p^*\bar{F}(p^*) & \text{if } y \geq \bar{F}(p^*); \\ y\bar{F}^{-1}(y) & \text{otherwise.} \end{cases}$$

It follows that  $g(y)' = 0$  on  $y \geq \bar{F}(p^*)$ . On the domain  $(0, \bar{F}(p^*)]$ , define the function  $p(y) = \bar{F}^{-1}(y)$ ;  $p(y)$  is decreasing in  $y$ . On  $(0, \bar{F}(p^*)]$ , we have  $(yg(1/y))' = p(y) - \bar{F}(p(y))/f(p(y))$ , which is non-increasing in  $y$  following the second part of Assumption 1 that  $\bar{F}(p)/f(p)$  is non-increasing, and the fact that  $p(y)$  is decreasing in  $y$ . Moreover, on  $(0, \bar{F}(p^*)]$ ,

$$(yg(1/y))' \geq (yg(1/y))'|_{y=\bar{F}(p^*)} = p^* - \bar{F}(p^*)/f(p^*) = 0.$$

It follows that  $yg(1/y)$  is non-decreasing and concave on  $\mathbb{R}_{++}$ .

3. That  $g(y)/y$  is non-increasing on  $\mathbb{R}_+$  follows directly from property (2) above.
4. Since  $g(\cdot)$  is a non-decreasing and concave function on  $\mathbb{R}_+$ , this property holds due to Lemma 13.

■

**Lemma 3.**

$$\begin{aligned}
\tilde{J}^*(x^0, \lambda^0, 0) &\leq \mathbf{E} \left[ \tilde{J}_{\{\Lambda_t\}}^*(x_0, 0) \right] \\
&\leq \tilde{J}_{\text{CE}}^*(x_0, \lambda_0, 0) \\
(6) \quad &\leq x_0 g \left( \frac{\int_0^T (\lambda_t + \sigma_t / \sqrt{2\pi}) dt}{x_0} \right) \\
(7) \quad &\leq x_0 g \left( \frac{\int_0^T \lambda_t dt}{x_0} \right) + x_0 g \left( \frac{\int_0^T \sigma_t dt}{x_0 \sqrt{2\pi}} \right).
\end{aligned}$$

**Proof.** The first inequality is evident by definition. Now, by definition of the unit revenue function  $g(\cdot)$  and Section 5.2 of Gallego and van Ryzin [1994], we have that

$$\tilde{J}_{\text{CE}}^*(x_0, \lambda_0, 0) = x_0 g \left( \frac{\int_0^T \mathbf{E}[\Lambda_t] dt}{x_0} \right)$$

By the concavity of  $g(\cdot)$  established in Lemma 2 and Jensen's inequality, we immediately have:

$$\mathbf{E} \left[ \tilde{J}_{\{\Lambda_t\}}^*(x_0, 0) \right] = \mathbf{E} \left[ x_0 g \left( \frac{\int_0^T \Lambda_t dt}{x_0} \right) \right] \leq x_0 g \left( \frac{\int_0^T \mathbf{E}[\Lambda_t] dt}{x_0} \right) = \tilde{J}_{\text{CE}}^*(x_0, \lambda_0, 0)$$

which is the second inequality. The fact that  $\tilde{J}_{\{\Lambda_t\}}^*(x_0, 0) = x_0 g \left( \frac{\int_0^T \Lambda_t dt}{x_0} \right)$  follows from the definition of  $g(\cdot)$  and Section 5.2 in Gallego and van Ryzin [1994].

Now for a Normal random variable  $X$  with mean  $\mu$  and variance  $\sigma^2$ , we know that  $\mathbf{E}[X^+] \leq \mu + \sigma/\sqrt{2\pi}$ . Thus,  $\mathbf{E}[\Lambda_t] = \mathbf{E}[\bar{\Lambda}_t^+] \leq \lambda_t + \sigma_t/\sqrt{2\pi}$ . Since, by Lemma 2,  $g(\cdot)$  is non-decreasing, it then follows that

$$\tilde{J}_{\text{CE}}^*(x_0, \lambda_0, 0) = x_0 g \left( \frac{\int_0^T \mathbf{E}[\Lambda_t] dt}{x_0} \right) \leq x_0 g \left( \frac{\int_0^T (\lambda_t + \sigma_t/\sqrt{2\pi}) dt}{x_0} \right).$$

The sub-additivity of  $g(\cdot)$  from the fourth part of Lemma 2 then yields the final inequality. ■

**Lemma 5.**

$$\tilde{J}^{\pi_{\text{RFP}}}(x^0, \lambda^0, 0) \geq \alpha \frac{x_0}{T} \mathbf{E} \left[ \int_0^T g \left( \frac{\Lambda_t T}{x_0} \right) dt \right] + (1 - \alpha) \frac{x_0}{\int_0^T \lambda_s ds} \mathbf{E} \left[ \int_0^T \lambda_t g \left( \frac{\Lambda_t \int_0^T \lambda_s ds}{x_0 \lambda_t} \right) dt \right].$$

**Proof.** We have:

$$\begin{aligned}
\tilde{J}^{\pi_{\text{RFP}}}(x^0, \lambda^0, 0) &= \mathbf{E} \left[ \int_0^T \pi_{\text{RFP}}(\tilde{X}^t, \Lambda^t, t) \bar{F}(\pi_{\text{RFP}}(\tilde{X}^t, \Lambda^t, t)) \Lambda_t dt \right] \\
&= \mathbf{E} \left[ \int_0^T \frac{\tilde{X}_t h(t, \alpha)}{\Lambda_t (T-t)} g \left( \frac{\Lambda_t (T-t)}{\tilde{X}_t h(t, \alpha)} \right) \Lambda_t dt \right] \\
&\geq x_0 \mathbf{E} \left[ \int_0^T \left( \frac{\alpha}{T} + (1 - \alpha) \frac{\lambda_t}{\int_0^T \lambda_s ds} \right) g \left( \Lambda_t / x_0 \left( \frac{\alpha}{T} + (1 - \alpha) \frac{\lambda_t}{\int_0^T \lambda_s ds} \right) \right) dt \right] \\
&\geq \alpha \frac{x_0}{T} \mathbf{E} \left[ \int_0^T g \left( \frac{\Lambda_t T}{x_0} \right) dt \right] + (1 - \alpha) \frac{x_0}{\int_0^T \lambda_s ds} \mathbf{E} \left[ \int_0^T \lambda_t g \left( \frac{\Lambda_t \int_0^T \lambda_s ds}{x_0 \lambda_t} \right) dt \right].
\end{aligned}$$

where the second equality holds by the definition of  $g(\cdot)$ , the first inequality follows from the lower bound on  $\tilde{X}_t$  established in Lemma 4 on the inventory balancing property and the property that  $zg(1/z)$  is an increasing function. The final inequality holds because  $zg(1/z)$  is a concave function. ■

**Lemma 6.**

$$\frac{1}{T} \mathbb{E} \left[ \int_0^T g \left( \frac{\Lambda_t T}{x_0} \right) dt \right] \geq 0.342g \left( \frac{\int_0^T \sigma_t dt}{x_0 \sqrt{2\pi}} \right).$$

**Proof.** We have:

$$\begin{aligned} \frac{\mathbb{E} \left[ \int_0^T g \left( \frac{\Lambda_t T}{x_0} \right) dt \right]}{Tg \left( \frac{\int_0^T \sigma_t dt}{x_0 \sqrt{2\pi}} \right)} &= \frac{\int_0^T \int_{-\infty}^{\infty} g \left( \frac{T(\lambda_t + y)^+}{x_0} \right) \frac{\exp(-y^2/2\sigma_t^2)}{\sqrt{2\pi\sigma_t^2}} dy dt}{Tg \left( \frac{\int_0^T \sigma_t dt}{x_0 \sqrt{2\pi}} \right)} \\ &\geq \frac{\int_0^T \int_{-\infty}^{\infty} g \left( \frac{T y^+}{x_0} \right) \frac{\exp(-y^2/2\sigma_t^2)}{\sqrt{2\pi\sigma_t^2}} dy dt}{Tg \left( \frac{\int_0^T \sigma_t dt}{x_0 \sqrt{2\pi}} \right)} \\ &\geq \frac{1}{T} \int_0^T \int_{-\infty}^{\infty} \min \left\{ 1, \frac{y^+}{\int_0^T \sigma_t dt / T \sqrt{2\pi}} \right\} \frac{\exp(-y^2/2\sigma_t^2)}{\sqrt{2\pi\sigma_t^2}} dy dt \\ &= \frac{1}{T} \int_0^T \left[ 1 - \Phi \left( \frac{\bar{\sigma}_{T,1}}{\sigma_t \sqrt{2\pi}} \right) + \int_0^{\bar{\sigma}_{T,1}/\sqrt{2\pi}} \frac{y}{\bar{\sigma}_{T,1} \sigma_t} \exp(-y^2/2\sigma_t^2) dy \right] dt \\ &\geq 0.342. \end{aligned}$$

The first inequality holds due to Property 1 in Lemma 2, and the positivity of  $\lambda_t$ . The second inequality holds due to Property 4 in Lemma 2. The final inequality was derived as a property of the class of market-size processes we consider in Property 3 in Lemma 14. ■

**Lemma 7.**

$$\frac{1}{\int_0^T \lambda_s ds} \mathbb{E} \left[ \int_0^T \lambda_t g \left( \frac{\Lambda_t \int_0^T \lambda_s ds}{x_0 \lambda_t} \right) dt \right] \geq \frac{1}{2} g \left( \frac{\int_0^T \lambda_t dt}{x_0} \right).$$

**Proof.** We have:

$$\begin{aligned}
& \frac{1}{\int_0^T \lambda_s ds} \mathbb{E} \left[ \int_0^T \lambda_t g \left( \frac{\Lambda_t \int_0^T \lambda_s ds}{x_0 \lambda_t} \right) dt \right] \\
&= \frac{1}{\int_0^T \lambda_s ds} \int_0^T \lambda_t \int_{-\infty}^{\infty} g \left( \frac{(\lambda_t + y)^+ \int_0^T \lambda_s ds}{x_0 \lambda_t} \right) \frac{\exp(-y^2/2\sigma_t^2)}{\sqrt{2\pi\sigma_t^2}} dy dt \\
&\geq \frac{1}{\int_0^T \lambda_s ds} \int_0^T \lambda_t \int_0^{\infty} g \left( \frac{(\lambda_t + y) \int_0^T \lambda_s ds}{x_0 \lambda_t} \right) \frac{\exp(-y^2/2\sigma_t^2)}{\sqrt{2\pi\sigma_t^2}} dy dt \\
&\geq \frac{1}{\int_0^T \lambda_s ds} \int_0^T \lambda_t \int_0^{\infty} g \left( \frac{\lambda_t \int_0^T \lambda_s ds}{x_0 \lambda_t} \right) \frac{\exp(-y^2/2\sigma_t^2)}{\sqrt{2\pi\sigma_t^2}} dy dt \\
&= \frac{1}{\int_0^T \lambda_s ds} \int_0^T \lambda_t g \left( \frac{\int_0^T \lambda_s ds}{x_0} \right) \int_0^{\infty} \frac{\exp(-y^2/2\sigma_t^2)}{\sqrt{2\pi\sigma_t^2}} dt \\
&= \frac{1}{2} g \left( \frac{\int_0^T \lambda_t dt}{x_0} \right).
\end{aligned}$$

The first and second inequalities follow respectively from the fact that  $g(\cdot)$  is non-negative and increasing.  $\blacksquare$

**Theorem 1.** Consider the RFP policy for an arbitrary  $\alpha$ . We then have for all market-size processes satisfying Assumption 2, and the additional requirement that  $\lambda_t = \lambda$  for all  $t$ ,

$$\frac{\tilde{J}^{\pi_{\text{RFP}}}(x^0, \lambda^0, 0)}{\tilde{J}^*(x^0, \lambda^0, 0)} \geq \frac{\tilde{J}^{\pi_{\text{RFP}}}(x^0, \lambda^0, 0)}{\tilde{J}_{\text{CE}}^*(x^0, \lambda^0, 0)} \geq 0.342.$$

**Proof.** By Lemma 3 we have that  $\tilde{J}^*(x^0, \lambda^0, 0) \leq \tilde{J}_{\text{CE}}^*(x^0, \lambda^0, 0)$ . Consequently,

$$\frac{\tilde{J}^{\pi_{\text{RFP}}}(x^0, \lambda^0, 0)}{\tilde{J}^*(x^0, \lambda^0, 0)} \geq \frac{\tilde{J}^{\pi_{\text{RFP}}}(x^0, \lambda^0, 0)}{\tilde{J}_{\text{CE}}^*(x^0, \lambda^0, 0)}.$$

Now, we have:

$$\begin{aligned}
\frac{\tilde{J}^{\pi_{\text{RFP}}}(x^0, \lambda^0, 0)}{\tilde{J}_{\text{CE}}^*(x^0, \lambda^0, 0)} &\geq \frac{\mathbb{E} \left[ \int_0^T \frac{\tilde{X}_t}{\Lambda_t(T-t)} g \left( \frac{\Lambda_t(T-t)}{\tilde{X}_t} \right) \Lambda_t dt \right]}{x_0 g \left( \frac{\lambda T + \int_0^T \sigma_t dt / \sqrt{2\pi}}{x_0} \right)} \\
&\geq \frac{\mathbb{E} \left[ \int_0^T g \left( \frac{\Lambda_t T}{x_0} \right) dt \right]}{T g \left( \frac{\lambda T + \int_0^T \sigma_t dt / \sqrt{2\pi}}{x_0} \right)} \\
&= \frac{\int_0^T \int_{-\infty}^{\infty} g \left( \frac{T(\lambda+y)^+}{x_0} \right) \frac{\exp(-y^2/2\sigma_t^2)}{\sqrt{2\pi\sigma_t^2}} dy dt}{T g \left( \frac{\lambda T + \int_0^T \sigma_t dt / \sqrt{2\pi}}{x_0} \right)} \\
&\geq \frac{1}{T} \int_0^T \int_{-\infty}^{\infty} \min \left\{ 1, \frac{(\lambda+y)^+}{\lambda + \int_0^T \sigma_t dt / T \sqrt{2\pi}} \right\} \frac{\exp(-y^2/2\sigma_t^2)}{\sqrt{2\pi\sigma_t^2}} dy dt \\
&= \frac{1}{T} \int_0^T \left[ 1 - \Phi \left( \frac{\bar{\sigma}_{T,1}}{\sigma_t \sqrt{2\pi}} \right) + \int_{-\lambda}^{\bar{\sigma}_{T,1}/\sqrt{2\pi}} \frac{\lambda+y}{\lambda + \bar{\sigma}_{T,1}/\sqrt{2\pi}} \frac{\exp(-y^2/2\sigma_t^2)}{\sqrt{2\pi\sigma_t^2}} dy \right] dt \\
&\geq \frac{1}{T} \int_0^T \left[ 1 - \Phi \left( \frac{\bar{\sigma}_{T,1}}{\sigma_t \sqrt{2\pi}} \right) + \int_0^{\bar{\sigma}_{T,1}/\sqrt{2\pi}} \frac{y}{\bar{\sigma}_{T,1}/\sqrt{2\pi}} \frac{\exp(-y^2/2\sigma_t^2)}{\sqrt{2\pi\sigma_t^2}} dy \right] dt \\
&\geq 0.342.
\end{aligned}$$

The first equality holds by definition of  $g(\cdot)$  and Lemma 3. The second inequality follows by applying the inventory balancing Lemma (Lemma 4) to obtain a lower bound on  $\tilde{X}_t$  along with the property that  $zg(1/z)$  is an increasing function, which is established in Lemma 2. The third inequality follows by Property 4 in Lemma 2. The final inequality follows by Property 3 in Lemma 14.  $\blacksquare$

**Theorem 2.** *Consider a market-size process with no drift and bounded volatility. That is assume that  $\lambda_t = \lambda$  and moreover,  $\sigma_T/\lambda \leq \sqrt{2\pi}\lambda B$  for some constant  $B > 0$ . Then,*

$$\frac{\tilde{J}^{\pi_{\text{RFP}}}(x^0, \lambda^0, 0)}{\tilde{J}^*(x^0, \lambda^0, 0)} \geq \frac{\tilde{J}^{\pi_{\text{RFP}}}(x^0, \lambda^0, 0)}{\tilde{J}_{\text{CE}}^*(x^0, \lambda^0, 0)} \geq \max \left\{ 0.342, \frac{1}{1+B} - \frac{B}{1+B} \left( \exp(-1/4\pi B^2) + 0.853 \right) \right\}.$$

**Proof.** We have:

$$\begin{aligned}
\frac{\tilde{J}^{\pi_{RFP}}(x^0, \lambda^0, 0)}{\tilde{J}^*(x^0, \lambda^0, 0)} &\geq \frac{1}{T} \int_0^T \left[ 1 - \Phi \left( \frac{\bar{\sigma}_{T,1}}{\sigma_t \sqrt{2\pi}} \right) + \int_{-\lambda}^{\bar{\sigma}_{T,1}/\sqrt{2\pi}} \frac{\lambda + y}{\lambda + \bar{\sigma}_{T,1}/\sqrt{2\pi}} \frac{\exp(-y^2/2\sigma_t^2)}{\sqrt{2\pi\sigma_t^2}} dy \right] dt \\
&= \frac{1}{T} \int_0^T \left[ 1 - \Phi \left( \frac{\bar{\sigma}_{T,1}}{\sigma_t \sqrt{2\pi}} \right) + \frac{\lambda}{\lambda + \bar{\sigma}_{T,1}/\sqrt{2\pi}} \left( \Phi \left( \frac{\bar{\sigma}_{T,1}}{\sigma_t \sqrt{2\pi}} \right) - \Phi \left( -\frac{\lambda}{\sigma_t} \right) \right) \right. \\
&\quad \left. + \frac{\sigma_t}{\sqrt{2\pi}\lambda + \bar{\sigma}_{T,1}} \left( \exp(-\lambda^2/2\sigma_t^2) - \exp(-\bar{\sigma}_{T,1}^2/4\pi\sigma_t^2) \right) \right] dt \\
&\geq \frac{1}{T} \int_0^T \left[ 1 - \Phi \left( \frac{\bar{\sigma}_{T,1}}{\sigma_t \sqrt{2\pi}} \right) + \frac{1}{1+B} \left( \Phi \left( \frac{\bar{\sigma}_{T,1}}{\sigma_t \sqrt{2\pi}} \right) - \Phi \left( -\frac{1}{\sqrt{2\pi}B} \right) \right) \right. \\
&\quad \left. - \frac{\sigma_t}{\sqrt{2\pi}\lambda + \bar{\sigma}_{T,1}} \exp(-\bar{\sigma}_{T,1}^2/4\pi\sigma_t^2) \right] dt \\
&= \frac{1}{T} \int_0^T \left[ 1 - \Phi \left( \frac{\bar{\sigma}_{T,1}}{\sigma_t \sqrt{2\pi}} \right) + \frac{1}{1+B} \left( \Phi \left( \frac{\bar{\sigma}_{T,1}}{\sigma_t \sqrt{2\pi}} \right) - \Phi \left( -\frac{1}{\sqrt{2\pi}B} \right) \right) \right. \\
&\quad \left. - \frac{\bar{\sigma}_{T,1}}{\sqrt{2\pi}\lambda + \bar{\sigma}_{T,1}} \frac{\sigma_t}{\bar{\sigma}_{T,1}} \exp(-\bar{\sigma}_{T,1}^2/4\pi\sigma_t^2) \right] dt \\
&\geq \frac{1}{T} \int_0^T \left[ 1 - \Phi \left( \frac{\bar{\sigma}_{T,1}}{\sigma_t \sqrt{2\pi}} \right) + \frac{1}{1+B} \left( \Phi \left( \frac{\bar{\sigma}_{T,1}}{\sigma_t \sqrt{2\pi}} \right) - \Phi \left( -\frac{1}{\sqrt{2\pi}B} \right) \right) \right. \\
&\quad \left. - \frac{B}{1+B} \frac{\sigma_t}{\bar{\sigma}_{T,1}} \exp(-\bar{\sigma}_{T,1}^2/4\pi\sigma_t^2) \right] dt \\
&= \frac{1}{T} \int_0^T \left[ \frac{1}{1+B} \Phi \left( \frac{1}{\sqrt{2\pi}B} \right) + \frac{B}{1+B} \left( 1 - \Phi \left( \frac{\bar{\sigma}_{T,1}}{\sigma_t \sqrt{2\pi}} \right) - \frac{\sigma_t}{\bar{\sigma}_{T,1}} \exp(-\bar{\sigma}_{T,1}^2/4\pi\sigma_t^2) \right) \right] dt \\
&\geq \frac{1}{1+B} \Phi \left( \frac{1}{\sqrt{2\pi}B} \right) \\
&\quad + \frac{B}{1+B} \frac{1}{T} \int_0^T \left[ 1 - \Phi \left( \sqrt{\frac{\max\left\{\frac{T}{2t}, 1\right\}}{2\pi}} \right) - \frac{1}{1-t/3T} \exp(-(1-t/3T)^2/4\pi) \right] dt \\
&= \frac{1}{1+B} \Phi \left( \frac{1}{\sqrt{2\pi}B} \right) \\
&\quad + \frac{B}{1+B} \int_0^1 \left[ 1 - \Phi \left( \sqrt{\frac{\max\left\{\frac{1}{2v}, 1\right\}}{2\pi}} \right) - \frac{1}{1-v/3} \exp(-(1-v/3)^2/4\pi) \right] dv \\
&= \frac{1}{1+B} \Phi \left( \frac{1}{\sqrt{2\pi}B} \right) - 0.853 \frac{B}{1+B} \\
&\geq \frac{1}{1+B} - \frac{B}{1+B} \left( \exp(-1/4\pi B^2) + 0.853 \right).
\end{aligned}$$

Here  $\Phi$  is the C.D.F of a standard normal random variable. The first inequality follows from the third inequality in the proof of Theorem 1, the second and third inequalities hold because  $\sigma_t \vee \bar{\sigma}_{T,1} \leq \sigma_T \leq \sqrt{2\pi}\lambda B$ , the fourth inequality follows Property 2 in Lemma 14, and the last inequality is derived from the fact that  $1 - \Phi(x) \leq \exp(-x^2/2)/x\sqrt{2\pi}$  for  $x > 0$ .

Combined with the lower bound derived in Theorem 1, we have

$$\frac{\tilde{J}^{\pi_{\text{RFP}}}(x^0, \lambda^0, 0)}{\tilde{J}^*(x^0, \lambda^0, 0)} \geq \max \left\{ 0.342, \frac{1}{1+B} - \frac{B}{1+B} \left( \exp(-1/4\pi B^2) + 0.853 \right) \right\}.$$

■

**Theorem 3.** *Consider the RFP policy with  $\alpha = 0.594$ . We then have for all market-size processes satisfying Assumption 2,*

$$\frac{\tilde{J}^{\pi_{\text{RFP}}}(x^0, \lambda^0, 0)}{\tilde{J}^*(x^0, \lambda^0, 0)} \geq \frac{\tilde{J}^{\pi_{\text{RFP}}}(x^0, \lambda^0, 0)}{\tilde{J}_{\text{CE}}^*(x^0, \lambda^0, 0)} \geq 0.203.$$

**Proof.** By Lemma 3 we have that  $\tilde{J}^*(x^0, \lambda^0, 0) \leq \tilde{J}_{\text{CE}}^*(x^0, \lambda^0, 0)$ . Consequently,

$$\begin{aligned} \frac{\tilde{J}^{\pi_{\text{RFP}}}(x^0, \lambda^0, 0)}{\tilde{J}^*(x^0, \lambda^0, 0)} &\geq \frac{\tilde{J}^{\pi_{\text{RFP}}}(x^0, \lambda^0, 0)}{\tilde{J}_{\text{CE}}^*(x^0, \lambda^0, 0)} \\ &\geq \frac{\alpha \frac{x_0}{T} \mathbb{E} \left[ \int_0^T g \left( \frac{\Lambda_t T}{x_0} \right) dt \right] + (1 - \alpha) \frac{x_0}{\int_0^T \lambda_s ds} \mathbb{E} \left[ \int_0^T \lambda_t g \left( \frac{\Lambda_t \int_0^T \lambda_s ds}{x_0 \lambda_t} \right) dt \right]}{x_0 g \left( \frac{\int_0^T \sigma_t dt}{x_0 \sqrt{2\pi}} \right) + x_0 g \left( \frac{\int_0^T \lambda_t dt}{x_0} \right)} \\ &\geq \frac{0.342 \alpha x_0 g \left( \frac{\int_0^T \sigma_t dt}{x_0 \sqrt{2\pi}} \right) + \frac{1-\alpha}{2} x_0 g \left( \frac{\int_0^T \lambda_t dt}{x_0} \right)}{x_0 g \left( \frac{\int_0^T \sigma_t dt}{x_0 \sqrt{2\pi}} \right) + x_0 g \left( \frac{\int_0^T \lambda_t dt}{x_0} \right)} \\ &\geq \min \left\{ 0.342 \alpha, \frac{1-\alpha}{2} \right\}. \end{aligned}$$

The second inequality above is via Lemmas 3 and 5. The third inequality follows from the lower bounds in Lemmas 6 and 7. Setting  $\alpha = 0.594$  maximizes the lower bound established above, yielding the first claim in the theorem. ■

## B.1. Performance Guarantees for Alternate Market Size Processes

While we focused on providing performance guarantees for market size processes satisfying Assumption 2, our analysis is easily extended to a number of distinct classes of market size processes. The analysis schema is essentially identical to what we have seen thus far, except for the final steps of the analysis where one must specialize to properties of the marginals of the market size process in question. To illustrate this, we present analogues to Theorem 1 for two market size processes outside of those specified by Assumption 2. The first class of processes we consider are ‘reflected’ generalized moving average processes, where as opposed to considering  $\Lambda_t = (\bar{\Lambda}_t)^+$  we consider  $\Lambda_t = |\bar{\Lambda}_t|$  where  $\bar{\Lambda}_t$  is constructed as before. Here we have:

**Theorem 5.** *Consider the RFP policy with  $\alpha = 0$ . Let  $\bar{\Lambda}_t$  satisfy the requirements of Assumption 2. Moreover, assume that  $\lambda_t = \lambda$  for all  $t$ . Then, if  $\Lambda_t = |\bar{\Lambda}_t|$ , we must have:*

$$\frac{\tilde{J}^{\pi_{\text{RFP}}}(x^0, \lambda^0, 0)}{\tilde{J}^*(x^0, \lambda^0, 0)} \geq 0.243.$$

**Proof.** Now, we have:

$$\begin{aligned}
\frac{\tilde{J}^{\pi_{\text{RFP}}}(x^0, \lambda^0, 0)}{\tilde{J}_{\text{CE}}^*(x^0, \lambda^0, 0)} &\geq \frac{\mathbb{E} \left[ \int_0^T \frac{\tilde{X}_t}{\Lambda_t(T-t)} g \left( \frac{\Lambda_t(T-t)}{\tilde{X}_t} \right) \Lambda_t dt \right]}{x_0 g \left( \frac{\lambda T + \int_0^T \sqrt{2}\sigma_t dt / \sqrt{\pi}}{x_0} \right)} \\
&\geq \frac{\mathbb{E} \left[ \int_0^T g \left( \frac{\Lambda_t T}{x_0} \right) dt \right]}{T g \left( \frac{\lambda T + \int_0^T \sqrt{2}\sigma_t dt / \sqrt{\pi}}{x_0} \right)} \\
&= \frac{\int_0^T \int_{-\infty}^{\infty} g \left( \frac{T|\lambda+y|}{x_0} \right) \frac{\exp(-y^2/2\sigma_t^2)}{\sqrt{2\pi\sigma_t^2}} dy dt}{T g \left( \frac{\lambda T + \int_0^T \sqrt{2}\sigma_t dt / \sqrt{\pi}}{x_0} \right)} \\
&\geq \frac{\int_0^T \int_{-\infty}^{\infty} g \left( \frac{T(\lambda+y)^+}{x_0} \right) \frac{\exp(-y^2/2\sigma_t^2)}{\sqrt{2\pi\sigma_t^2}} dy dt}{T g \left( \frac{\lambda T + \int_0^T \sqrt{2}\sigma_t dt / \sqrt{\pi}}{x_0} \right)} \\
&\geq \frac{1}{T} \int_0^T \int_{-\infty}^{\infty} \min \left\{ 1, \frac{(\lambda+y)^+}{\lambda + \int_0^T \sqrt{2}\sigma_t dt / T\sqrt{\pi}} \right\} \frac{\exp(-y^2/2\sigma_t^2)}{\sqrt{2\pi\sigma_t^2}} dy dt \\
&= \frac{1}{T} \int_0^T \left[ 1 - \Phi \left( \frac{\bar{\sigma}_{T,1}\sqrt{2}}{\sigma_t\sqrt{\pi}} \right) + \int_{-\lambda}^{\bar{\sigma}_{T,1}\sqrt{2}/\pi} \frac{\lambda+y}{\lambda + \bar{\sigma}_{T,1}\sqrt{2}/\pi} \frac{\exp(-y^2/2\sigma_t^2)}{\sqrt{2\pi\sigma_t^2}} dy \right] dt \\
&\geq \frac{1}{T} \int_0^T \left[ 1 - \Phi \left( \frac{\bar{\sigma}_{T,1}\sqrt{2}}{\sigma_t\sqrt{\pi}} \right) + \int_0^{\bar{\sigma}_{T,1}\sqrt{2}/\pi} \frac{y}{\bar{\sigma}_{T,1}\sqrt{2}/\pi} \frac{\exp(-y^2/2\sigma_t^2)}{\sqrt{2\pi\sigma_t^2}} dy \right] dt \\
&\geq 0.243.
\end{aligned}$$

The first equality holds by definition of  $g(\cdot)$  and Lemma 3. In addition, we use the fact that for a Normal random variable  $X$  with mean  $\mu$  and variance  $\sigma^2$ , we know that  $\mathbb{E}[|X|] \leq \mu + \sqrt{2}\sigma/\sqrt{\pi}$  so that  $\mathbb{E}[\Lambda_t] = \mathbb{E}[|\tilde{\Lambda}_t|] \leq \lambda + \sqrt{2}\sigma_t/\sqrt{\pi}$ . The second inequality follows by applying the inventory balancing Lemma (Lemma 4) to obtain a lower bound on  $\tilde{X}_t$  along with the property that  $zg(1/z)$  is an increasing function, which is established in Lemma 2. The fourth inequality follows by Property 4 in Lemma 2. Finally, by Lemma 3 we have that  $\tilde{J}^*(x^0, \lambda^0, 0) \leq \tilde{J}_{\text{CE}}^*(x^0, \lambda^0, 0)$  so that

$$\frac{\tilde{J}^{\pi_{\text{RFP}}}(x^0, \lambda^0, 0)}{\tilde{J}^*(x^0, \lambda^0, 0)} \geq \frac{\tilde{J}^{\pi_{\text{RFP}}}(x^0, \lambda^0, 0)}{\tilde{J}_{\text{CE}}^*(x^0, \lambda^0, 0)},$$

and the guarantee follows. ■

As a second example of an alternate market size process, we consider a market-size process specified by the Cox-Ingersoll-Ross (CIR) process

$$d\Lambda_t = \theta(\lambda - \Lambda_t)dt + \sigma\sqrt{\Lambda_t}dZ_t,$$

where  $\theta, \lambda, \sigma > 0$ . As is customary for the use of this process in applications we consider the regime where  $2\theta\lambda > \sigma^2$  wherein the process above becomes an example of a strictly positive and ergodic

affine process. In this model,  $\theta$  controls the speed of market-size adjustment,  $\lambda$  and  $\sigma$  corresponds to mean and volatility of the process respectively. The stationary distribution for this process is Gamma distributed with shape parameter  $2\theta\lambda/\sigma^2$  and scale parameter  $\sigma^2/2\theta$ . We assume  $\Lambda_0$  is distributed according to this stationary distribution and define  $\lambda = \Lambda_0$ .

**Theorem 6.** *Consider the RFP policy with  $\alpha = 0$ . Then if  $\Lambda_t$  is driven by the CIR process above, we have:*

$$\frac{\tilde{J}^{\pi_{\text{RFP}}}(x^0, \lambda^0, 0)}{\tilde{J}^*(x^0, \lambda^0, 0)} \geq 0.632.$$

**Proof.** Now, we have:

$$\begin{aligned} \frac{\tilde{J}^{\pi_{\text{RFP}}}(x^0, \lambda^0, 0)}{\tilde{J}_{\text{CE}}^*(x^0, \lambda^0, 0)} &\geq \frac{\mathbb{E} \left[ \int_0^T \frac{\tilde{X}_t}{\Lambda_t(T-t)} g \left( \frac{\Lambda_t(T-t)}{\tilde{X}_t} \right) \Lambda_t dt \right]}{x_0 g \left( \frac{\lambda T}{x_0} \right)} \\ &\geq \frac{\mathbb{E} \left[ \int_0^T g \left( \frac{\Lambda_t T}{x_0} \right) dt \right]}{T g \left( \frac{\lambda T}{x_0} \right)} \\ &\geq \frac{1}{T} \int_0^T \mathbb{E} [\min \{ \frac{\Lambda_t}{\lambda}, 1 \}] dt \\ &= \mathbb{E} \left[ \min \left\{ \frac{\Lambda_0}{\lambda}, 1 \right\} \right] \\ &= 1 - \frac{\Gamma(a+1, a)}{\Gamma(a+1)} + \frac{\Gamma(a, a)}{\Gamma(a)} \\ &\geq 0.632. \end{aligned}$$

The second inequality follows by applying the inventory balancing Lemma (Lemma 4) to obtain a lower bound on  $\tilde{X}_t$  along with the property that  $zg(1/z)$  is an increasing function, which is established in Lemma 2. The third inequality follows by Property 4 in Lemma 2.  $\Gamma(\cdot, \cdot)$  is an incomplete Gamma function and is given by  $\Gamma(x, y) = \int_y^\infty s^{x-1} e^{-s} ds$ , and  $a \triangleq 2\theta\lambda/\sigma^2 \geq 1$ . By Lemma 3 we have that  $\tilde{J}^*(x^0, \lambda^0, 0) \leq \tilde{J}_{\text{CE}}^*(x^0, \lambda^0, 0)$  so that

$$\frac{\tilde{J}^{\pi_{\text{RFP}}}(x^0, \lambda^0, 0)}{\tilde{J}^*(x^0, \lambda^0, 0)} \geq \frac{\tilde{J}^{\pi_{\text{RFP}}}(x^0, \lambda^0, 0)}{\tilde{J}_{\text{CE}}^*(x^0, \lambda^0, 0)}.$$

and the guarantee follows. ■

## C. Proofs for Section 5

**Lemma 8.** *For all  $t \in [0, T)$ , we have*

$$\lim_{\Delta \rightarrow 0} \tilde{X}_{t(\Delta)}^\Delta = \tilde{X}_t.$$

**Proof.** Now for  $i > 0$ , we have in the fluid model, that

$$\tilde{X}_{(i+1)\Delta}^\Delta = \left( \tilde{X}_{i\Delta}^\Delta - \bar{F}(\pi_{\text{RFP}}^\Delta(\tilde{X}^{\Delta, i\Delta}, i\Delta)) \int_{i\Delta}^{(i+1)\Delta} \Lambda_s ds \right)^+.$$

From the definition of  $h(s, \alpha)$ , we have that for  $\epsilon > 0$ , there exist numbers  $C(\epsilon), D(\epsilon) < \infty$ , such that  $h(s, \alpha)/(T - s) \leq C(\epsilon)$  and  $|\frac{d}{ds}(h(s, \alpha)/(T - s))| \leq D(\epsilon)$  for all  $s < T - \epsilon$ . Now, for any  $\Delta \leq s < T - \epsilon$ , we have:

$$\begin{aligned}
& \left| \min \left\{ \bar{F}(p^*), \frac{\tilde{X}_s h(s, \alpha)}{\Lambda_s(T - s)} \right\} \Lambda_s - \min \left\{ \bar{F}(p^*), \frac{\tilde{X}_{s(\Delta)} h(s(\Delta), \alpha)}{\hat{\Lambda}_{s(\Delta)}^{\text{FL}}(T - s(\Delta))} \right\} \Lambda_s \right| \\
& \leq \left| \min \left\{ \bar{F}(p^*) \Lambda_s, \frac{\tilde{X}_s h(s, \alpha)}{T - s} \right\} - \min \left\{ \bar{F}(p^*) \hat{\Lambda}_{s(\Delta)}^{\text{FL}}, \frac{\tilde{X}_{s(\Delta)} h(s(\Delta), \alpha)}{T - s(\Delta)} \right\} \right| \\
(8) \quad & + \min \left\{ \bar{F}(p^*), \frac{\tilde{X}_{s(\Delta)} h(s(\Delta), \alpha)}{\hat{\Lambda}_{s(\Delta)}^{\text{FL}}(T - s(\Delta))} \right\} |\Lambda_s - \hat{\Lambda}_{s(\Delta)}^{\text{FL}}| \\
& \leq \bar{F}(p^*) |\Lambda_s - \hat{\Lambda}_{s(\Delta)}^{\text{FL}}| + \left| \frac{\tilde{X}_s h(s, \alpha)}{T - s} - \frac{\tilde{X}_{s(\Delta)} h(s(\Delta), \alpha)}{T - s(\Delta)} \right| + \bar{F}(p^*) |\Lambda_s - \hat{\Lambda}_{s(\Delta)}^{\text{FL}}| \\
& \leq 2(\Lambda_{s(\Delta)}^{\max} - \Lambda_{s(\Delta)}^{\min}) + \left| \frac{\tilde{X}_s h(s, \alpha)}{T - s} - \frac{\tilde{X}_{s(\Delta)} h(s, \alpha)}{T - s} \right| + \left| \frac{\tilde{X}_{s(\Delta)} h(s, \alpha)}{T - s} - \frac{\tilde{X}_{s(\Delta)} h(s(\Delta), \alpha)}{T - s(\Delta)} \right| \\
& \leq 2(\Lambda_{s(\Delta)}^{\max} - \Lambda_{s(\Delta)}^{\min}) + (C(\epsilon)B + x_0 D(\epsilon))\Delta,
\end{aligned}$$

where  $B \triangleq \sup_{t \in [0, T]} \Lambda_t$ . Moreover, we have defined  $\Lambda_{s(\Delta)}^{\max} \triangleq \max_{\tau \in [s(\Delta) - \Delta, s(\Delta) + \Delta]} \Lambda_\tau$  and similarly for  $\Lambda_{s(\Delta)}^{\min}$ . The second inequality follows from the fact that  $|\min\{A, B\} - \min\{C, D\}| \leq |A - C| + |B - D|$ . Now, we have, for  $i \geq 1$  with  $(i + 1)\Delta < T - \epsilon$ ,

$$\begin{aligned}
|\tilde{X}_{(i+1)\Delta}^\Delta - \tilde{X}_{(i+1)\Delta}| &= \left| \left( \tilde{X}_{i\Delta}^\Delta - \int_{i\Delta}^{(i+1)\Delta} \min \left\{ \bar{F}(p^*), \frac{\tilde{X}_{i\Delta}^\Delta h(i\Delta, \alpha)}{\hat{\Lambda}_{i\Delta}^{\text{FL}}(T - i\Delta)} \right\} \Lambda_s ds \right)^+ \right. \\
& \quad \left. - \left( \tilde{X}_{i\Delta} - \int_{i\Delta}^{(i+1)\Delta} \min \left\{ \bar{F}(p^*), \frac{\tilde{X}_s h(s, \alpha)}{\Lambda_s(T - s)} \right\} \Lambda_s ds \right)^+ \right| \\
& \leq \left| \left( \tilde{X}_{i\Delta}^\Delta - \int_{i\Delta}^{(i+1)\Delta} \min \left\{ \bar{F}(p^*), \frac{\tilde{X}_{i\Delta}^\Delta h(i\Delta, \alpha)}{\hat{\Lambda}_{i\Delta}^{\text{FL}}(T - i\Delta)} \right\} \Lambda_s ds \right)^+ \right. \\
& \quad \left. - \left( \tilde{X}_{i\Delta} - \int_{i\Delta}^{(i+1)\Delta} \min \left\{ \bar{F}(p^*), \frac{\tilde{X}_{i\Delta} h(i\Delta, \alpha)}{\hat{\Lambda}_{i\Delta}^{\text{FL}}(T - i\Delta)} \right\} \Lambda_s ds \right)^+ \right| \\
& \quad + \left| \int_{i\Delta}^{(i+1)\Delta} \min \left\{ \bar{F}(p^*), \frac{\tilde{X}_s h(s, \alpha)}{\Lambda_s(T - s)} \right\} \Lambda_s ds \right. \\
& \quad \left. - \int_{i\Delta}^{(i+1)\Delta} \min \left\{ \bar{F}(p^*), \frac{\tilde{X}_{i\Delta} h(i\Delta, \alpha)}{\hat{\Lambda}_{i\Delta}^{\text{FL}}(T - i\Delta)} \right\} \Lambda_s ds \right| \\
& \leq |\tilde{X}_{i\Delta}^\Delta - \tilde{X}_{i\Delta}| \\
& \quad + \int_{i\Delta}^{(i+1)\Delta} \left| \min \left\{ \bar{F}(p^*), \frac{\tilde{X}_s h(s, \alpha)}{\Lambda_s(T - s)} \right\} - \min \left\{ \bar{F}(p^*), \frac{\tilde{X}_{s(\Delta)} h(s(\Delta), \alpha)}{\hat{\Lambda}_{s(\Delta)}^{\text{FL}}(T - s(\Delta))} \right\} \right| \Lambda_s ds \\
& \leq |\tilde{X}_{i\Delta}^\Delta - \tilde{X}_{i\Delta}| + 2(\Lambda_{s(\Delta)}^{\max} - \Lambda_{s(\Delta)}^{\min})\Delta + (C(\epsilon)B + x_0 D(\epsilon))\Delta^2,
\end{aligned}$$

where the first inequality follows from the property that  $|A^+ - (B + C)^+| \leq |A^+ - B^+| + |C|$ , the second inequality follows from the property that  $|(X - \min\{a, bX\})^+ - (Y - \min\{a, bY\})^+| \leq |X - Y|$

for  $b \geq 0$ , and the last inequality follows from (8). Moreover, since trivially  $|\tilde{X}_\Delta^\Delta - \tilde{X}_\Delta| \leq \int_0^\Delta \Lambda_s ds \leq B\Delta$ , we must have for any positive integer  $i$  with  $i\Delta < T - \epsilon$ ,

$$\begin{aligned} |\tilde{X}_{i\Delta}^\Delta - \tilde{X}_{i\Delta}| &\leq 2\Delta \sum_{m=1}^{i-1} (\Lambda_{m\Delta}^{\max} - \Lambda_{m\Delta}^{\min}) + (C(\epsilon)B + x_0D(\epsilon))\Delta^2(i-1) + B\Delta \\ &\leq 2\Delta \sum_{m=1}^{\lfloor T/\Delta \rfloor} (\Lambda_{m\Delta}^{\max} - \Lambda_{m\Delta}^{\min}) + (C(\epsilon)BT + x_0D(\epsilon)T + B)\Delta. \end{aligned}$$

Hence, for an  $t < T - \epsilon$ ,

$$\begin{aligned} |\tilde{X}_{t(\Delta)}^\Delta - \tilde{X}_t| &\leq |\tilde{X}_{t(\Delta)}^\Delta - \tilde{X}_{t(\Delta)}| + |\tilde{X}_{t(\Delta)} - \tilde{X}_t| \\ &\leq 2\Delta \sum_{m=1}^{\lfloor T/\Delta \rfloor} (\Lambda_{m\Delta}^{\max} - \Lambda_{m\Delta}^{\min}) + (C(\epsilon)BT + x_0D(\epsilon)T + B)\Delta + B\Delta \end{aligned}$$

But, since  $\Lambda_t$  is a continuous function of  $t$  and therefore Riemann integrable, the inequality above immediately yields.

$$\lim_{\Delta \rightarrow 0} \tilde{X}_{t(\Delta)}^\Delta(\Delta) = \tilde{X}_t.$$

for all  $t < T - \epsilon$ . Since our choice of  $\epsilon > 0$  was arbitrary, the result follows.  $\blacksquare$

**Lemma 9.**

$$\lim_{\Delta \rightarrow 0} \tilde{J}^{\pi_{\text{RFP}}^\Delta}(x^0, \lambda^0, 0 | \mathcal{G}_0) = \tilde{J}^{\pi_{\text{RFP}}}(x^0, \lambda^0, 0 | \mathcal{G}_0).$$

**Proof.** Consider some  $0 < t < T - \epsilon$ . From Lemmas 4 and 8 it follows that for  $\Delta$  sufficiently small (smaller than  $\bar{\Delta}(\epsilon)$ , say),  $\tilde{X}_t^\Delta > 0$ . It consequently follows from the definition of the RFP- $\Delta$  policy that:

$$\begin{aligned} \pi_{\text{RFP}}^\Delta(\tilde{X}^{\Delta,t}, t) &= \bar{F}^{-1} \left( \min \left\{ \bar{F}(p^*), \frac{\tilde{X}_{t(\Delta)}^\Delta h(t(\Delta), \alpha)}{\hat{\Lambda}_{t(\Delta)}^{\text{FL}}(T - t(\Delta))} \right\} \right) \\ &\rightarrow \bar{F}^{-1} \left( \min \left\{ \bar{F}(p^*), \frac{\tilde{X}_t h(t, \alpha)}{\Lambda_t(T - t)} \right\} \right) \\ &= \pi_{\text{RFP}}(\tilde{X}^t, t) \end{aligned}$$

where the limit follows from Lemma 8 and the fact that by construction,  $\hat{\Lambda}_{t(\Delta)}^{\text{FL}} \rightarrow \Lambda_t$  for  $t > 0$ .

We then have:

$$\begin{aligned} &\lim_{\Delta \rightarrow 0} \left| \int_0^{T-\epsilon} \pi_{\text{RFP}}^\Delta(\tilde{X}^{\Delta,t}, t) \bar{F} \left( \pi_{\text{RFP}}^\Delta(\tilde{X}^{\Delta,t}, t) \right) \Lambda_t dt - \int_0^{T-\epsilon} \pi_{\text{RFP}}(\tilde{X}^t, t) \bar{F} \left( \pi_{\text{RFP}}(\tilde{X}^t, t) \right) \Lambda_t dt \right| \\ &\leq \lim_{\Delta \rightarrow 0} \int_0^{T-\epsilon} \left| \pi_{\text{RFP}}^\Delta(\tilde{X}^{\Delta,t}, t) \bar{F} \left( \pi_{\text{RFP}}^\Delta(\tilde{X}^{\Delta,t}, t) \right) - \pi_{\text{RFP}}(\tilde{X}^t, t) \bar{F} \left( \pi_{\text{RFP}}(\tilde{X}^t, t) \right) \right| \Lambda_t dt \\ &= \int_0^{T-\epsilon} \lim_{\Delta \rightarrow 0} \left| \pi_{\text{RFP}}^\Delta(\tilde{X}^{\Delta,t}, t) \bar{F} \left( \pi_{\text{RFP}}^\Delta(\tilde{X}^{\Delta,t}, t) \right) - \pi_{\text{RFP}}(\tilde{X}^t, t) \bar{F} \left( \pi_{\text{RFP}}(\tilde{X}^t, t) \right) \right| \Lambda_t dt \\ &= 0 \end{aligned}$$

where the first equality follows from the bounded convergence theorem and the second from the limit established in the first part of the argument. Now, the above argument suffices to show that

$$\liminf_{\Delta \rightarrow 0} \tilde{J}^{\pi_{\text{RFP}}^\Delta}(x^0, \lambda^0, 0|\mathcal{G}_0) \geq \tilde{J}^{\pi_{\text{RFP}}}(x^0, \lambda^0, 0|\mathcal{G}_0) - 2Bp^*\bar{F}(p^*)\epsilon$$

and

$$\limsup_{\Delta \rightarrow 0} \tilde{J}^{\pi_{\text{RFP}}^\Delta}(x^0, \lambda^0, 0|\mathcal{G}_0) \leq \tilde{J}^{\pi_{\text{RFP}}}(x^0, \lambda^0, 0|\mathcal{G}_0) + 2Bp^*\bar{F}(p^*)\epsilon$$

where  $B \triangleq \max_{[0, T]} \Lambda_t$ . Since our choice of  $\epsilon > 0$  was arbitrary, the result follows.  $\blacksquare$

**Theorem 4.** (*The Price of Discretization*) *Under the assumptions on the market size process above, we have almost surely:*

$$\limsup_{\Delta \rightarrow 0} \frac{\left| \tilde{J}^{\pi_{\text{RFP}}^\Delta}(x^0, \lambda^0, 0|\mathcal{G}_0) - \tilde{J}^{\pi_{\text{RFP}}}(x^0, \lambda^0, 0|\mathcal{G}_0) \right|}{\sigma\eta(\Delta) \log(1/\eta(\Delta))} \leq \frac{4p^*\bar{F}(p^*)B^2T^2}{x_0^2}$$

**Proof.** Recall that by the inventory balancing property we have that:

$$\frac{\tilde{X}_t}{T-t} \geq \frac{x_0}{T}.$$

Using this fact with Corollary 2 allows us to conclude after some algebraic manipulation that for any  $t < T$  that

$$\limsup_{\Delta \rightarrow 0} \frac{1}{\sigma\eta(\Delta)} \left| \frac{\hat{\Lambda}_{t(\Delta)}^{\text{FL}}(T-t(\Delta))}{\tilde{X}_{t(\Delta)}^\Delta} - \frac{\Lambda_t(T-t)}{\tilde{X}_t} \right| \leq \frac{2T}{x_0} + \frac{4BT^2}{x_0^2(T-t)}$$

Let  $\kappa(\Delta) \triangleq 8T\sigma\eta(\Delta)/x_0$ . Observe that on  $t < T - \kappa(\Delta)$ , we must have by the balancing Lemma that  $\tilde{X}_t \geq 8\sigma\eta(\Delta)$ , so that for  $\Delta$  sufficiently small, Corollary 2 guarantees that  $\tilde{X}_t^\Delta > 0$  as well. Consequently, we have that for  $\Delta$  sufficiently small:

$$\begin{aligned} & \left| \tilde{J}^{\pi_{\text{RFP}}^\Delta}(x^0, \lambda^0, 0) - \tilde{J}^{\pi_{\text{RFP}}}(x^0, \lambda^0, 0) \right| \\ &= \left| \int_0^T \pi_{\text{RFP}}^\Delta(\tilde{X}^{\Delta, t}, t) \bar{F} \left( \pi_{\text{RFP}}^\Delta(\tilde{X}^{\Delta, t}, t) \right) \Lambda_t dt - \int_0^T \pi_{\text{RFP}}(\tilde{X}^t, t) \bar{F} \left( \pi_{\text{RFP}}(\tilde{X}^t, t) \right) \Lambda_t dt \right| \\ &= \left| \int_0^{T-\kappa(\Delta)} \pi_{\text{RFP}}^\Delta(\tilde{X}_{t(\Delta)}^\Delta, t) \bar{F} \left( \pi_{\text{RFP}}^\Delta(\tilde{X}_{t(\Delta)}^\Delta, t) \right) \Lambda_t dt - \int_0^{T-\kappa(\Delta)} \pi_{\text{RFP}}(\tilde{X}^t, t) \bar{F} \left( \pi_{\text{RFP}}(\tilde{X}^t, t) \right) \Lambda_t dt \right| \\ &\quad + B\kappa(\Delta)p^*\bar{F}(p^*) \\ &\leq Bp^*\bar{F}(p^*) \int_0^{T-\kappa(\Delta)} \left| \frac{\hat{\Lambda}_{t(\Delta)}^{\text{FL}}(T-t(\Delta))}{\tilde{X}_{t(\Delta)}^\Delta} - \frac{\Lambda_t(T-t)}{\tilde{X}_t} \right| dt + B\kappa(\Delta)p^*\bar{F}(p^*) \\ &\leq Bp^*\bar{F}(p^*) \left( \sigma\eta(\Delta) \int_0^{T-\kappa(\Delta)} \left( \frac{2T}{x_0} + \frac{4BT^2}{x_0^2(T-t)} \right) dt + \kappa(\Delta) \right) \\ &\leq Bp^*\bar{F}(p^*) \left( \sigma\eta(\Delta) \left( \frac{2T^2}{x_0} + \frac{4BT^2(\log T + \log(1/\kappa(\Delta)))}{x_0^2} \right) + \kappa(\Delta) \right) \end{aligned}$$

where the second equality follows from our choice of  $\kappa(\Delta)$  and the first inequality follows from the fact that the function  $g(y)/y$  has its first derivative bounded by  $p^*\bar{F}(p^*)$  on  $y \geq 1/\bar{F}(p^*)$ . The result follows upon dividing through by  $\sigma\eta(\Delta) \log(1/\eta(\Delta))$  and taking the limit supremum.  $\blacksquare$

## D. Proofs for Section 6

**Lemma 10.** *For any given  $\Delta > 0$ , we have that the inventory process and the forecast process in the  $n$ th system converge to their fluid analogues. In particular:*

1. For any  $t$ ,

$$\frac{X_{t,(n)}^\Delta}{n} \rightarrow \tilde{X}_t^\Delta \quad \text{a.s.}$$

2. For any  $t$  such that  $\tilde{X}_t^\Delta > 0$ ,

$$\frac{\hat{\Lambda}_{t(\Delta),(n)}}{n} \rightarrow \hat{\Lambda}_{t(\Delta)}^{\text{FL}} \quad \text{a.s.}$$

**Proof.** We denote  $I \triangleq \lfloor \frac{T}{\Delta} \rfloor$ . We begin with establishing both claims for times  $i\Delta$  for  $i = 0, 1, \dots, I$ . We proceed by induction. The claims are trivially true for  $i = 0$  by definition of  $\hat{\Lambda}_0^{\text{FL}}$ . Assume the claims true for some  $i < I$  and consider time  $(i+1)\Delta$ . Now, if  $\tilde{X}_{i\Delta}^\Delta = 0$ , or else if  $\hat{\Lambda}_{(i+1)\Delta} = 0$ , we are done, so let us assume that  $\tilde{X}_{i\Delta}^\Delta > 0$  and further that  $\hat{\Lambda}_{(i+1)\Delta} > 0$ . It must then be that

$$\hat{\lambda}_{i+1} \triangleq \min \left\{ \bar{F}(p^*), \frac{\tilde{X}_{i\Delta}^\Delta h(i\Delta, \alpha)}{\hat{\Lambda}_{i\Delta}^{\text{FL}}(T - i\Delta)} \right\} \int_{i\Delta}^{(i+1)\Delta} \Lambda_s ds > 0.$$

Moreover, we define for convenience,

$$\hat{\lambda}_{i+1,(n)} \triangleq \min \left\{ \bar{F}(p^*), \frac{X_{i\Delta,(n)}^\Delta h(i\Delta, \alpha)}{\hat{\Lambda}_{i\Delta,(n)}(T - i\Delta)} \right\} \int_{i\Delta}^{(i+1)\Delta} \Lambda_s ds.$$

Now, we have that

$$\frac{X_{(i+1)\Delta,(n)}^\Delta}{n} = \left( \frac{X_{i\Delta,(n)}^\Delta}{n} - \frac{1}{n} \sum_{m=1}^n Y_{m,(n)}^i \right)^+$$

where  $Y_{m,(n)}^i$  are i.i.d. Poisson random variables with parameter  $\hat{\lambda}_{i+1,(n)}$ . Now, define the events

$$C_n = \{\omega : |\hat{\lambda}_{i+1} - \hat{\lambda}_{i+1,(n)}| \geq \hat{\lambda}_{i+1}/2\},$$

and

$$B_n = \left\{ \omega : \left| \frac{1}{n} \sum_{m=1}^n Y_{m,(n)}^i - \hat{\lambda}_{i+1,(n)} \right| \geq \frac{1}{n^{1/4}} \right\}.$$

Now, define  $\hat{\mathcal{G}}_{i\Delta} = \sigma \left( \cup_{n \in \mathbb{N}} \mathcal{G}_{i\Delta}^{(n)} \right)$  and notice that

$$\begin{aligned} \mathbb{P} \left( B_n \text{ i.o.} \mid \hat{\mathcal{G}}_{i\Delta} \right) &\leq 1 - \mathbb{I}_{\{\sum_n \mathbb{P}(B_n \mid \hat{\mathcal{G}}_{i\Delta}) < \infty\}} \\ &\leq 1 - \mathbb{I}_{\{\sum_n 2 \exp(-\sqrt{n} \hat{\lambda}_{i+1,(n)}/2e) < \infty\}} \\ &\leq \mathbb{I}_{\{C_n \text{ i.o.}\}} \end{aligned}$$

where the first inequality follows from the first Borel-Cantelli lemma and the second from a Chernoff bound for the sum of independent poisson random variables. Taking expectations, and using the fact that by the induction hypothesis,  $\mathbb{P}(C_n \text{ i.o.}) = 0$ , we have that

$$\mathbb{P}(B_n \text{ i.o.}) = 0.$$

In other words, we have shown that almost surely,

$$\lim_n \left| \frac{1}{n} \sum_{m=1}^n Y_{m,(n)}^i - \hat{\lambda}_{i+1,(n)} \right| = 0,$$

so that with the triangle inequality and the induction hypothesis, we must have that

$$(9) \quad \lim_n \left| \frac{1}{n} \sum_{m=1}^n Y_{m,(n)}^i - \hat{\lambda}_{i+1} \right| = 0$$

almost surely. The induction hypothesis then must imply that

$$\frac{1}{n} X_{(i+1)\Delta,(n)}^\Delta \rightarrow \tilde{X}_{(i+1)\Delta}^\Delta.$$

Now, assuming that  $\tilde{X}_{(i+1)\Delta}^\Delta > 0$ , we must have by the argument above, and in particular, (9) that

$$\frac{\hat{\Lambda}_{(i+1)\Delta,(n)}}{n} \rightarrow \frac{\min \left\{ \bar{F}(p^*), \frac{\tilde{X}_{i\Delta}^\Delta h(i\Delta, \alpha)}{\hat{\Lambda}_{i\Delta}(T-i\Delta)} \right\} \int_{i\Delta}^{(i+1)\Delta} \Lambda_s ds}{\min \left\{ \bar{F}(p^*), \frac{\tilde{X}_{i\Delta}^\Delta h(i\Delta, \alpha)}{\hat{\Lambda}_{i\Delta}(T-i\Delta)} \right\}} = \hat{\Lambda}_{(i+1)\Delta}^{\text{FL}}.$$

This establishes the second claim and the first claim for times  $i\Delta$ ,  $i = 0, \dots, I$ . To establish the first claim for general  $t$ , we appeal to the claim just established for the time  $t(\Delta)$ , and repeat the argument in the induction step above.  $\blacksquare$

**Lemma 11.**

$$\liminf_n \frac{1}{n} J^{\pi_{\text{RFP}}^\Delta} (x_{(n)}^0, \lambda_{(n)}^0, 0 | \mathcal{G}_0) \geq \tilde{J}^{\pi_{\text{RFP}}^\Delta} (x^0, \lambda^0, 0 | \mathcal{G}_0) \text{ a.s.}$$

**Proof.** Define  $\hat{t}^\Delta = \inf\{t : \tilde{X}_t^\Delta > 0\}$ . Then,

$$\begin{aligned} & \lim_n \frac{1}{n} \mathbb{E} \left[ \int_0^{\hat{t}^\Delta} \pi_{\text{RFP}}^\Delta (X_{(n)}^{\Delta,t}, t) dN_t^{(n)} \right] \\ &= \lim_n \frac{1}{n} \mathbb{E} \left[ \int_0^{\hat{t}^\Delta} \pi_{\text{RFP}}^\Delta (X_{(n)}^{\Delta,t}, t) \bar{F} \left( \pi_{\text{RFP}}^\Delta (X_{(n)}^{\Delta,t}, t) \right) \Lambda_t^{(n)} dt \right] \\ &= \lim_n \frac{1}{n} \int_0^{\hat{t}^\Delta} \mathbb{E} \left[ \pi_{\text{RFP}}^\Delta (X_{(n)}^{\Delta,t}, t) \bar{F} \left( \pi_{\text{RFP}}^\Delta (X_{(n)}^{\Delta,t}, t) \right) \right] \Lambda_t^{(n)} dt \\ &= \frac{1}{n} \int_0^{\hat{t}^\Delta} \lim_n \mathbb{E} \left[ \pi_{\text{RFP}}^\Delta (X_{(n)}^{\Delta,t}, t) \bar{F} \left( \pi_{\text{RFP}}^\Delta (X_{(n)}^{\Delta,t}, t) \right) \right] \Lambda_t^{(n)} dt \\ &= \int_0^{\hat{t}^\Delta} \pi_{\text{RFP}}^\Delta (\tilde{X}^{\Delta,t}, t) \bar{F} \left( \pi_{\text{RFP}}^\Delta (\tilde{X}^{\Delta,t}, t) \right) \Lambda_t dt \\ &= \tilde{J}^{\pi_{\text{RFP}}^\Delta} (x^0, \lambda^0, 0 | \mathcal{G}_0) \end{aligned}$$

The first equality here follows from the definition of  $N_t^{(n)}$  and Theorem II.T8 in Bremaud [1980], the second equality follows from Fubini's theorem, the third equality is via the bounded convergence theorem, and the penultimate equality follows from Lemma 10 and the bounded convergence theorem. ■

**Corollary 3.** *For  $\Delta$  sufficiently small, we have:*

$$\liminf_n \frac{J^{\pi_{\text{RFP}}^\Delta}(x_{(n)}^0, \lambda_{(n)}^0, 0)}{\tilde{J}^{\pi_{\text{RFP}}^\Delta}(x_{(n)}^0, \lambda_{(n)}^0, 0)} \geq 1.$$

**Proof.** Lemmas 5, 7 and Theorem 1 together establish that for any  $\alpha$  we must have under our assumptions on  $\{\Lambda_t\}$  that  $\tilde{J}^{\pi_{\text{RFP}}}(x_0, \lambda_0, 0) > 0$ . Consequently, by Corollary 1, for  $\Delta$  sufficiently small, we must have  $\tilde{J}^{\pi_{\text{RFP}}^\Delta}(x_0, \lambda_0, 0) > 0$ . But Fatou's lemma applied to the result of Lemma 11 immediately yields

$$\liminf_n \frac{1}{n} J^{\pi_{\text{RFP}}^\Delta}(x_{(n)}^0, \lambda_{(n)}^0, 0) \geq \tilde{J}^{\pi_{\text{RFP}}^\Delta}(x^0, \lambda^0, 0),$$

so that dividing by  $\tilde{J}^{\pi_{\text{RFP}}^\Delta}(x^0, \lambda^0, 0)$  on both sides and using the fact that  $n\tilde{J}^{\pi_{\text{RFP}}^\Delta}(x^0, \lambda^0, 0) = \tilde{J}^{\pi_{\text{RFP}}^\Delta}(x_{(n)}^0, \lambda_{(n)}^0, 0)$  yields the result. ■

**Lemma 12.** *We have*

$$\liminf_{\Delta \rightarrow 0} \liminf_{n \rightarrow \infty} \frac{J^{\pi_{\text{RFP}}^\Delta}(x_{(n)}^0, \lambda_{(n)}^0, 0)}{J^*(x_{(n)}^0, \lambda_{(n)}^0, 0)} \geq \frac{\tilde{J}^{\pi_{\text{RFP}}}(x^0, \lambda^0, 0)}{\tilde{J}_{\text{CE}}^*(x^0, 0)}.$$

**Proof.** We first establish a few preliminary inequalities. For an arbitrary deterministic process  $\{\bar{\lambda}_t : t \geq 0\}$ , let us denote by  $J_{\{\bar{\lambda}_{t,(n)}\}}^*(x_{(n)}^0, 0)$  the value of an optimal policy in a system where  $\Lambda_t = \bar{\lambda}_t$  for all  $t$ . Now since an optimal policy derived for an arbitrary market-size process, and restricted to be  $\mathcal{F}_t$ -progressive, is certainly feasible for this problem, we have  $J_{\{\Lambda_{t,(n)}\}}^*(x_{(n)}^0, 0) \geq$

$J^*(x_{(n)}^0, 0|\mathcal{G}_0)$ . We then have:

$$\begin{aligned}
(10) \quad \liminf_{\Delta} \liminf_n \frac{J^{\pi_{\text{RFP}}^{\Delta}}(x_{(n)}^0, \lambda_{(n)}^0, 0)}{J^*(x_{(n)}^0, \lambda_{(n)}^0, 0)} &\geq \liminf_{\Delta} \left( \liminf_n \frac{J^{\pi_{\text{RFP}}^{\Delta}}(x_{(n)}^0, \lambda_{(n)}^0, 0)}{\tilde{J}^{\pi_{\text{RFP}}^{\Delta}}(x_{(n)}^0, \lambda_{(n)}^0, 0)} \liminf_n \frac{\tilde{J}^{\pi_{\text{RFP}}^{\Delta}}(x_{(n)}^0, \lambda_{(n)}^0, 0)}{J^*(x_{(n)}^0, \lambda_{(n)}^0, 0)} \right) \\
&\geq \liminf_{\Delta} \liminf_n \frac{\tilde{J}^{\pi_{\text{RFP}}^{\Delta}}(x_{(n)}^0, \lambda_{(n)}^0, 0)}{J^*(x_{(n)}^0, \lambda_{(n)}^0, 0)} \\
&\geq \liminf_{\Delta} \liminf_n \frac{\tilde{J}^{\pi_{\text{RFP}}^{\Delta}}(x_{(n)}^0, \lambda_{(n)}^0, 0)}{\mathbb{E} \left[ J_{\{\Lambda_{t,(n)}^*\}}^*(x_{(n)}^0, 0) \right]} \\
&\geq \liminf_{\Delta} \liminf_n \frac{\tilde{J}^{\pi_{\text{RFP}}^{\Delta}}(x_{(n)}^0, \lambda_{(n)}^0, 0)}{\mathbb{E} \left[ \tilde{J}_{\{\Lambda_{t,(n)}^*\}}^*(x_{(n)}^0, 0) \right]} \\
&= \liminf_{\Delta} \frac{\tilde{J}^{\pi_{\text{RFP}}^{\Delta}}(x^0, \lambda^0, 0)}{\tilde{J}_{\text{CE}}^*(x^0, 0)} \\
&= \frac{\tilde{J}^{\pi_{\text{RFP}}}(x^0, \lambda^0, 0)}{\tilde{J}_{\text{CE}}^*(x^0, 0)}.
\end{aligned}$$

The first inequality follows from the fact that for non-negative sequences  $\{a_n\}, \{b_n\}$ ,  $\liminf_n a_n b_n \geq \liminf_n a_n \liminf_n b_n$ . The second inequality follows from Corollary 3. The third inequality follows from our observation that  $J_{\{\Lambda_{t,(n)}^*\}}^*(x_{(n)}^0, 0) \geq J^*(x_{(n)}^0, 0|\mathcal{G}_0)$ . The fourth inequality follows from the fact that  $J_{\{\Lambda_{t,(n)}^*\}}^*(x_{(n)}^0, 0) \leq \tilde{J}_{\{\Lambda_{t,(n)}^*\}}^*(x_{(n)}^0, 0)$  which is Theorem 2 in Gallego and van Ryzin [1994]. The final equality is Corollary 1.  $\blacksquare$

## E. Miscellaneous Results and Computations

### E.1. Properties of the Market-Size Process

**Lemma 13.** *Let  $f : \mathbb{R}_+ \rightarrow \mathbb{R}_+$  be a non-decreasing, concave function with  $f(0) = 0$ . Then for all  $0 < y \leq x$ ,*

$$1 \leq \frac{f(x)}{f(y)} \leq \frac{x}{y},$$

and

$$(11) \quad \frac{1}{x} \int_0^x f(t) dt \leq f\left(\frac{x}{2}\right).$$

**Proof.** By definition,  $f(x)/f(y) \geq 1$ . Moreover, the concavity of  $f$  yields  $f(x) = f(0 + \frac{x}{y}y) \leq \frac{x}{y}f(y)$ . Thus,  $f(x)/x \leq f(y)/y$ . Inequality (11) follows by Jensen's inequality.  $\blacksquare$

Now, we use Lemma 13 to characterize properties of the volatility of the market-size process,  $\sigma_t^2$ . Defining

$$\bar{\sigma}_{T,1} \triangleq \int_0^T \sigma_t dt / T \text{ and}$$

$$\bar{\sigma}_{T,2} \triangleq \int_0^T \sigma_t^2 dt / T.$$

we have:

**Lemma 14.**

1.  $\sigma_t$  is increasing and concave in  $t$ .
2.  $1 - t/3T \leq \bar{\sigma}_{T,1}/\sigma_t \leq \sqrt{\bar{\sigma}_{T,2}/\sigma_t^2} \leq \sqrt{\max\{T/2t, 1\}}$ .
3.  $\frac{1}{T} \int_0^T \left[ 1 - \Phi\left(\frac{\bar{\sigma}_{T,1}}{\sigma_t \sqrt{2\pi}}\right) + \int_0^{\bar{\sigma}_{T,1}/\sqrt{2\pi}} \frac{y}{\bar{\sigma}_{T,1}\sigma_t} \exp(-y^2/2\sigma_t^2) dy \right] dt \geq 0.342$ .

**Proof.**

1.  $\sigma_t^2$  is increasing in  $t$  directly implies that  $\sigma_t$  is increasing in  $t$ . Now,

$$(\sigma_t^2)'' = 2(\sigma_t')^2 + 2\sigma_t\sigma_t''$$

so that since  $\sigma_t^2$  is concave,  $\sigma_t'' \leq 0$  and the concavity of  $\sigma_t$  follows.

2. To establish the first inequality, we see that:

$$\begin{aligned} \frac{\bar{\sigma}_{T,1}}{\sigma_t} &= \frac{\int_0^T \sigma_s ds}{T\sigma_t} = \frac{1}{T} \int_0^T \sqrt{\frac{\sigma_s^2}{\sigma_t^2}} ds \\ (12) \quad &\geq \frac{1}{T} \int_0^T \sqrt{\min\left\{\frac{s}{t}, 1\right\}} ds \\ &= 1 - \frac{t}{3T}, \end{aligned}$$

where inequality (12) follows by Lemma 13 and the concavity of  $\sigma_t^2$ .

That  $\bar{\sigma}_{T,1}/\sigma_t \leq \sqrt{\bar{\sigma}_{T,2}/\sigma_t^2}$  is a direct consequence of Jensen's inequality.

The second part of Lemma 13 yields  $\bar{\sigma}_{T,2} \leq \sigma_{T/2}^2$ , so that the first part of Lemma 13 then yields:

$$\frac{\bar{\sigma}_{T,2}}{\sigma_t^2} \leq \frac{\sigma_{T/2}^2}{\sigma_t^2} \leq \max\left\{\frac{T}{2t}, 1\right\}.$$

3. We have:

$$\begin{aligned} &\frac{1}{T} \int_0^T \left[ 1 - \Phi\left(\frac{\bar{\sigma}_{T,1}}{\sigma_t \sqrt{2\pi}}\right) + \int_0^{\bar{\sigma}_{T,1}/\sqrt{2\pi}} \frac{y}{\bar{\sigma}_{T,1}\sigma_t} \exp(-\frac{y^2}{2\sigma_t^2}) dy \right] dt \\ &= \frac{1}{T} \int_0^T \left[ 1 - \Phi\left(\frac{\bar{\sigma}_{T,1}}{\sigma_t \sqrt{2\pi}}\right) + \frac{\sigma_t}{\bar{\sigma}_{T,1}} \left(1 - \exp(-(\bar{\sigma}_{T,1})^2/4\pi\sigma_t^2)\right) \right] dt \\ &\geq \frac{1}{T} \int_0^T \left[ 1 - \Phi\left(\sqrt{\frac{\max\left\{\frac{T}{2t}, 1\right\}}{2\pi}}\right) + \frac{1}{\sqrt{\max\left\{\frac{T}{2t}, 1\right\}}} \left(1 - \exp(- (1 - t/3T)^2/4\pi)\right) \right] dt \\ &= \int_0^1 \left[ 1 - \Phi\left(\sqrt{\frac{\max\left\{\frac{1}{2v}, 1\right\}}{2\pi}}\right) + \frac{1}{\sqrt{\max\left\{\frac{1}{2v}, 1\right\}}} \left(1 - \exp(- (1 - v/3)^2/4\pi)\right) \right] dv \\ &= 0.342, \end{aligned}$$

where the first inequality follows from the previous property (i.e. Lemma 14, Property 2); the penultimate equality follows by employing the change of variables  $v = t/T$ , and the final equality follows from numerical evaluation of the definite integral in the penultimate line. ■

## E.2. Analysis for Example 1 in Section 3

Recall, that our goal is to show that if  $\sigma > 0$ , then

$$\limsup_{n \rightarrow \infty} \frac{J^{\pi_{\text{RFP}}}(x_{(n)}^0, \lambda_{(n)}^0, 0)}{J^*(x_{(n)}^0, \lambda_{(n)}^0, 0)} \leq O((\log T)^{-1}),$$

for the dynamic pricing problem described in Example 1. To show this, we will find it convenient to use properties of the RFP policy established in Sections 4 and 6, as we will use performance under this policy as a lower bound to performance under an optimal policy. Now, we have

$$\begin{aligned} \limsup_{n \rightarrow \infty} \frac{J^{\pi_{\text{RFP}}}(x_{(n)}^0, \lambda_{(n)}^0, 0)}{J^*(x_{(n)}^0, \lambda_{(n)}^0, 0)} &\leq \lim_{n \rightarrow \infty} \frac{p^* x_{0,(n)}}{J^{\pi_{\text{RFP}}}(x_{(n)}^0, \lambda_{(n)}^0, 0)} \\ &\leq \lim_{n \rightarrow \infty} \frac{x_{0,(n)}}{\tilde{J}^{\pi_{\text{RFP}}}(x_{(n)}^0, \lambda_{(n)}^0, 0)} \\ &\leq \left[ 0.342g \left( 1 + \frac{2T^{3/2}\sigma}{3\sqrt{2\pi}x_0} - \frac{\lambda^2\sqrt{T}}{\sigma\sqrt{2\pi}x_0} \right) \right]^{-1} \\ &= \left[ 0.342 \log \left( 1 + \frac{2T^{3/2}\sigma}{3\sqrt{2\pi}x_0} - \frac{\lambda^2\sqrt{T}}{\sigma\sqrt{2\pi}x_0} \right) \right]^{-1} \\ &= O((\log T)^{-1}). \end{aligned}$$

The first inequality follows by the definition of  $J^*$  and also the fact that performance under the fixed price policy is trivially upper bounded by  $p^*x_{0,(n)}$ ; in the case of our example, recall that  $p^* = 1$ . The second inequality then follows from Corollary 3 that showed that as  $n$  grows, performance in the stochastic model is at least as good as the performance in the fluid model. We now focus on the third inequality: Theorem 1 showed that

$$\tilde{J}^{\pi_{\text{RFP}}}(x_{(n)}^0, \lambda_{(n)}^0, 0) \geq 0.342 \tilde{J}_{\text{CE}}^*(x_{(n)}^0, \lambda_{(n)}^0, 0),$$

while by the definition of the unit revenue function,  $g(\cdot)$ , in Section 4.2, we know that

$$\tilde{J}_{\text{CE}}^*(x_{(n)}^0, \lambda_{(n)}^0, 0) = x_{0,(n)} g \left( \frac{\int_0^T \mathbb{E}[\Lambda_t] dt}{x_0} \right).$$

Since here,  $\int_0^T \mathbb{E}[\Lambda_t] dt \geq \lambda T + \frac{2T^{3/2}\sigma}{3\sqrt{2\pi}} - \frac{\lambda^2\sqrt{T}}{\sigma\sqrt{2\pi}}$  and  $g$  is non-decreasing from Lemma 2, it follows that

$$\tilde{J}^{\pi_{\text{RFP}}}(x_{(n)}^0, \lambda_{(n)}^0, 0) \geq 0.342 x_{0,(n)} g \left( 1 + \frac{2T^{3/2}\sigma}{3\sqrt{2\pi}x_0} - \frac{\lambda^2\sqrt{T}}{\sigma\sqrt{2\pi}x_0} \right).$$

### E.3. Computational Experiments Relative to a Tighter Super-Optimal Policy

In our computational experiments, we compared performance of the RFP- $\Delta$  policy against a clairvoyant upper bound that was permitted to observe the entire realization of a sample path of the market size process at time 0. While this bound was cheap to compute, we observed that in certain cases performance relative to this upper bound was worse than 10%. We conjectured that this did not reflect our pricing policies performance per se but rather simply the fact that our upper bound was loose in settings with high volatility. As such, we compute a tighter upper bound here, namely the expected revenue under an *optimal* policy with knowledge of the specification of the market size process (i.e. a probability distribution over its sample paths) and the ability to monitor the process and update prices in continuous time. This is obviously still an upper bound on the optimal value function, but nonetheless tighter than the clairvoyant bound. The results are summarize (for an OU process) in Tables 4 and 5.

**Table 4:** Performance Relative to a Tighter Upper Bound. Common parameters across problem instances:  $\lambda = e, \beta = 1, n = 1000, T = 5, CV = 2.5, \Delta = 0.1$ .

Initial Inventory	Load Factor	Relative Optimality		
$x_0^{(n)}/n$	$x_0/\lambda T$	$J^{\pi_{\text{RFP}}}/J^*$	$J^{\pi_{\text{RFP}}^\Delta}/J^*$	$J^{\pi_{\text{RFP}}^\Delta}/J^{UB}$
4	0.294	0.973	0.962	0.880
8	0.589	0.986	0.978	0.938
12	0.883	0.993	0.989	0.972
16	1.177	0.997	0.996	0.989
20	1.472	1.000	0.999	0.996

**Table 5:** Performance Relative to a Tighter Upper Bound. Common parameters across problem instances:  $\lambda = e, \beta = 1, n = 1000, T = 5, CV = 5, \Delta = 0.1$ .

Initial Inventory	Load Factor	Relative Optimality		
$x_0^{(n)}/n$	$x_0/\lambda T$	$J^{\pi_{\text{RFP}}}/J^*$	$J^{\pi_{\text{RFP}}^\Delta}/J^*$	$J^{\pi_{\text{RFP}}^\Delta}/J^{UB}$
4	0.294	0.945	0.911	0.824
8	0.589	0.953	0.934	0.879
12	0.883	0.962	0.952	0.914
16	1.177	0.973	0.964	0.937
20	1.472	0.988	0.980	0.958

In the experiments above  $J^{\pi_{\text{RFP}}^\Delta}/J^{UB}$  is the quantity reported for the bulk of our experiments – performance relevant to a clairvoyant upper bound. The quantity  $J^{\pi_{\text{RFP}}^\Delta}/J^*$  reports performance relative to the tighter upper bounds. Since even this tighter upper bound is potentially loose (since it re-optimizes continuously, and is allowed to observe the monitor the market size process), the quantity  $J^{\pi_{\text{RFP}}}/J^*$  report performance of the *idealized* RFP policy (that is also allowed to re-optimize continuously and monitor the market size process directly) against the tighter upper bound. We see that the results bear substantial support to the fact that a large fraction of the performance losses reported in our computational study are potentially due to the fact that we compare ourselves against an upper bound that can be fairly loose. This is not surprising given the amount of information used by the policy implicit in the clairvoyant upper bound.