

MIT Open Access Articles

Simple Policies for Dynamic Pricing with Imperfect Forecasts

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

Citation: Chen, Yiwei, and Vivek F. Farias. "Simple Policies for Dynamic Pricing with Imperfect Forecasts." *Operations Research* 61, no. 3 (June 2013): 612–624.

As Published: <http://dx.doi.org/10.1287/opre.2013.1166>

Publisher: Institute for Operations Research and the Management Sciences (INFORMS)

Persistent URL: <http://hdl.handle.net/1721.1/87676>

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

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



Simple Policies for Dynamic Pricing with Imperfect Forecasts'

Yiwei Chen

Vivek F. Farias *

Abstract

We consider the ‘classical’ single product dynamic pricing problem allowing the ‘scale’ of demand intensity to be modulated by 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’ settings¹.

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 building *stochastic* ‘forecast’ models for how consumer demand might evolve. This alternative has

*YC is with the Cheung Kong Graduate School of Business and Sloan School of Management, MIT. VF is with the Operations Research Center and Sloan School of Management, MIT. Emails: {ywchen@ckgsb.edu.cn, vivekf@mit.edu}

¹Regimes where the initial inventory and scale of demand grow large simultaneously.

its own perils: building a model of this type with predictive power is far from trivial. 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- Δ policy in short, is akin to repeatedly applying the fixed price policy at discrete epochs in time spaced Δ 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- Δ 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- Δ 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- Δ 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 Δ 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- Δ 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- Δ 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 and one operates in a ‘high volume’ regime where inventory and demand grow large simultaneously.

In related pieces of work, 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 fairly recent 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(s) 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]). Boyaci and Ozer [2010] study related issues in the context of capacity planning via advance selling.

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. Our analysis will be in a similar regime – we ignore fluctuations due to ‘point process noise’. 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 Section 5, we discuss the impact of using a discrete review version of the RFP policy (the RFP- Δ policy) and exhibit that the performance achieved by the idealized RFP policy can, in essence, be obtained with this discrete review policy. We characterize the performance loss due to discrete price reviews. Section 6 presents a numerical investigation of the RFP- Δ policy. Section 7 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 at rate Λ_t . These customers are assumed infinitely divisible and the 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. The corresponding sales rate at time t is $\Lambda_t \bar{F}(p_t)$ where p_t is the price posted by the seller at time t and $\bar{F}(\cdot)$ is a modulating function that we will specify shortly. The vendors revenue rate at time t is then simply $\Lambda_t \bar{F}(p_t) p_t$. The vendor’s goal is to dynamically adjust prices in a manner that maximizes expected revenue while respecting an inventory constraint.

Reservation Prices: To motivate the choice of the modulating function $\bar{F}(\cdot)$, consider a model with indivisible customers where 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. 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 assumptions on $\bar{F}(\cdot)$ that, by the above analogy, may be viewed as assumptions on customer reservation price distributions:

Assumption 1.

1. $F(\cdot)$ has a continuous density $f(\cdot)$ with support \mathbb{R}_+ .²
2. F has a non-decreasing hazard rate on \mathbb{R}_+ . That is, $f(p)/\overline{F}(p)$ is non-decreasing in p on \mathbb{R}_+ .
3. $p\overline{F}(p)$ 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. In particular, define $\Lambda_t = (\overline{\Lambda}_t)^+$:

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

where we assume that $\lambda_t > 0$ and $\lambda_t \in \mathcal{C}_0$; $\phi(\cdot) \in \mathcal{C}_1$ and non-increasing in absolute value; and dZ_s is an increment of Brownian motion. We dub these processes *Generalized Moving Average Processes*³. 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 ϕ this term could behave in drastically different ways. We use σ_t^2 to denote the variance of $\overline{\Lambda}_t$, where $\sigma_t^2 = \int_0^t \phi^2(s) ds$, which follows from Ito’s isometry.

Our labeling of such processes, and indeed the reason we believe they are interesting follow from the fact that when evaluated at discrete times $i\Delta$, we have:

$$\overline{\Lambda}_{n\Delta} = \lambda_{n\Delta} + \sum_{k=0}^{n-1} \theta_{n-k} \epsilon_k$$

where the ϵ_k are standard normal random variables and $\theta_j = \sqrt{\int_{(j-1)\Delta}^{j\Delta} \phi^2(s) ds}$. 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 Talluri and van Ryzin [2004]). Moreover, this analogy makes our assumptions on ϕ 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

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

³we will abuse terminology and refer to $\overline{\Lambda}_t$ as a generalized moving average process as well, when this does not risk confusion

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 λ and $\Lambda_0 = \lambda$ (which is recovered by setting $\lambda_t = \lambda$ and $\phi(t) = \sigma \exp(-\beta t)$ for arbitrary $\sigma > 0$, and $\beta > 0$). Further, as it turns out, we could consider an even more general class of processes; see Appendix 2.

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 $p^t \triangleq \{p_s : s \in [0, t]\}$. We define the sales process encoding the number of sales up to time t , N_t^4 , according to

$$N_t = \int_0^t \Lambda_t \bar{F}(p_t) dt,$$

and denote by $X_t = x_0 - N_t$ the corresponding inventory process. We denote by \mathcal{F}_t the filtration generated by the history of the market-size and sales processes up to time t , $\sigma(\Lambda^t, N^t)$, and allow for optimal pricing policies p_t that are progressively measurable with respect to this history.

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 Λ_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 Π 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 $N_T \leq x_0$ 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 π according to ⁵

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

The vendor then seeks to find a policy π^* that achieves

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

The problem above raises some serious challenges:

1. In the absence of making restrictions on the market size process, the optimization problem is an intractable, infinite dimensional one.
2. Even if one were to assume a relatively ‘tractable’ model for the market size process – in that it induced a low dimensional optimization problem – it is unclear whether the specification of this process is clear at the outset. To be concrete, the *parameters* specifying a tractable market size process model will themselves need to be learned and this poses a further challenge to the sellers optimization problem. Indeed, even very simple forms of such parametric uncertainty are difficult to resolve; see Araman and Caldentey [2009], Aviv and Pazgal [2005a,b], Farias and Van Roy [2010].

⁴Notice that we are not specifying a point process but a fluid sales process. It is not difficult, but rather technical, to augment our model with N_t defined as an appropriate point process.

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

3. 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 is closely related to that studied in a 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 Λ_t is identically equal to λ_t . Here one may show (see Gallego and van Ryzin [1994]) that an optimal pricing policy selects a *fixed* price p_{FP} at time 0 and keeps this price fixed over the length of the selling season. This fixed price is given by

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

Gallego and van Ryzin [1994] showed that this policy remained near optimal in models wherein customers were *indivisible* and arrived according to independent increment point processes, as one scaled the initial inventory and arrival rate process to grow large.

It is easy to see that if one considers the case where the market-size process is *not* deterministic, 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., $\bar{F}(p) = \exp(-p)$. Consider a static pricing policy π_{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 the results of Gallego and van Ryzin [1994] imply that:

$$\frac{J^{\pi_{\text{FP}}}(x^0, \lambda^0, 0)}{J^*(x^0, \lambda^0, 0)} = 1.$$

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

$$\frac{J^{\pi_{\text{FP}}}(x^0, \lambda^0, 0)}{J^*(x^0, \lambda^0, 0)} = O((\log T)^{-1}).$$

The fixed price policy is attractive from a managerial perspective for its simplicity, implementability and ease of interpretation. In the above example this policy is aware of λ but does not know ϕ . It is also disadvantaged relative to the optimal policy in that it is allowed no price changes. In what follows we will propose a policy that is given *no* more information at the outset (it will not even know λ) but *is allowed* to update prices in a causal fashion and thus respond to shocks. We will see that allowing this update recovers a substantial portion of the gap between the fixed price and optimal policies above. This policy will be only slightly more sophisticated than the above fixed price policy. It will be akin to repeated application of the fixed price policy 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 continuous-time dynamic pricing policy motivated by the fixed price policy above. In addition should deterministic forecast information be available, the policy will require as input an additional ‘hedging’ parameter α (in subsequent sections we will discuss an oblivious choice of α). 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 Λ_t be known to the manager at time t . These shortcomings are not fundamental; we will address them comprehensively in Section 5 with a simple revision to the above policy.

To develop a qualitative understanding of the RFP policy, 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 both cases, we see that the form taken by the policy closely resembles repeated application of the GVR fixed price rule with an appropriate forecast for expected demand over the remainder of the selling season. This is an immensely attractive feature as it agrees with existing real-world pricing practice.

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 α 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 α that one may select in the absence of any information about the market size process whatsoever and also argue that selecting $\alpha = 1$ (which clearly results in a policy that is highly simple and intuitive in form) will frequently be desirable and suffice.

4. Performance Analysis for the Idealized RFP Policy

This Section aims to understand the performance of the RFP policy presented in the previous Section. Our goal will be to produce a lower bound on the quantity $J^{\pi_{\text{RFP}}}(x^0, \lambda^0, 0)/J^*(x^0, \lambda^0, 0)$. We begin with an overview and discussion of our results in this regard:

4.1. Performance Guarantees for the Idealized RFP Policy

This section will establish several performance guarantees for the idealized RFP policy. In particular, we will establish the following result:

Theorem 1. *Assume Λ_t is a generalized moving average process. Then,*

1. *For the RFP policy with $\alpha = 1$, we have*

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

where $B \triangleq \sigma_T/\sqrt{2\pi}\lambda^2$, and we assume $\lambda_t = \lambda$ for all t .

2. *For the RFP policy with $\alpha = 0.594$, and arbitrary forecasts $\{\lambda_t\}$, we have:*

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

The proof can be found in the appendix. The guarantees of the first result 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. The two salient features of this guarantee are:

1. The performance loss relative to an optimal policy remains bounded for *arbitrarily* volatile market-size processes.
2. In a regime where the volatility is low, the *same* RFP policy is near optimal.

We see that with continuous 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 which is allowed to adjust prices will have uniformly bounded relative performance losses.

The second 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 α 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 α parameter. This will be evident in the proof of Theorem 1. We will also examine this issue from a computational perspective in our numerical experiments and show that, in fact, the theoretically ideal choice of α performs quite well, while the oblivious choice of $\alpha = 1$ 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 (in this section, continuous) re-optimization. Finally, we note that the analysis here could be extended to other market-size process families; Appendix A.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). In the remainder of this section, we provide a summary of the key steps required in proving the above theorem.

4.2. An Upper Bound on the Optimal Value function

We follow a familiar path to deriving an upper bound on the optimal value function. In particular, we consider a problem wherein the entire sample path of the market-size process, that is $\{\Lambda_t : t \geq 0\}$, is available at time 0. The optimal value function for this ‘clairvoyant’ problem is easy to derive and provides an intuitive upper bound. The revenue optimization problem here is a deterministic one. In particular, let us denote by $J_{\{\Lambda_t\}}^*(x_0, 0)$, the optimal value of the revenue maximization problem:

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

Proposition 2 of Gallego and van Ryzin [1994] establishes that $J_{\{\Lambda_t\}}^*(x_0, 0) = x_0 g(\int_0^T \Lambda_t dt / x_0)$, where the *unit revenue function* $g : \mathbb{R}_+ \rightarrow \mathbb{R}_+$ is defined according to

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

The properties of this unit revenue function are useful tools in our analysis and are established in the appendix. Now since any policy $\pi \in \Pi$ describes a feasible solution to the above problem, we immediately have that

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

In addition, we define the *certainty equivalent* value function J_{CE}^* according to

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

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

Lemma 1.

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

4.3. A Lower Bound on $J^{\pi_{\text{RFP}}}$

The more difficult challenge in establishing our performance guarantees is deriving a useful *lower* bound on the value of the RFP policy. In order to establish our lower bound we will first demonstrate that the RFP policy possesses an intuitive ‘balancing’ property that yields useful *uniform* lower bounds on the inventory process under *any* market-size sample path.

Lemma 2. (*Inventory Balancing*) *The inventory at time t under the RFP policy, X_t , satisfies*

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

Proof. By the definition of π_{RFP} ,

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

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

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

Thus,

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

Integrating on both sides and using the initial value $X_0 = x_0$, yields

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

■

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

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

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. Moreover, when $\alpha = 0$, we see that

$$X_t \geq \frac{\int_t^T \lambda_s ds}{\int_0^T \lambda_s ds} x_0$$

which is easily recognized as the inventory on hand under the optimal policy in the event that the shock term were 0. As discussed earlier an appropriate choice of α allows the policy designer to chose to ration inventory as a convex combination of these two levels irrespective of the sample path.

We next use this balancing property in a crucial way to establish a useful lower bound on the value of the RFP policy. The proofs also rely essentially on 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 structural properties of our market size process (see Lemma 8). In particular, we have:

Lemma 3.

$$J^{\pi_{\text{RFP}}}(x^0, \lambda^0, 0) \geq 0.342\alpha x_0 g \left(\frac{\int_0^T \sigma_t dt}{x_0 \sqrt{2\pi}} \right) + \frac{1-\alpha}{2} x_0 g \left(\frac{\int_0^T \lambda_t dt}{x_0} \right).$$

The Performance Guarantee: The upper and lower bounds derived thus far allow us to provide a uniform bound on the performance loss of the RFP policy relative to an optimal clairvoyant pricing strategy. In particular, observe that by Lemma 1 we have that $J^*(x^0, \lambda^0, 0) \leq J_{\text{CE}}^*(x^0, \lambda^0, 0)$. Consequently, we have:

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

The second inequality above is via Lemmas 1 and 3. Setting $\alpha = 0.594$ maximizes the lower bound established above, yielding the second claim in the Theorem. The first claim follows along similar lines but is somewhat more tedious; the proof can be found in the appendix.

5. A Discrete Review RFP Policy

We now consider a variant of the idealized RFP policy that is potentially better suited to implementation. 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 Δ , specifying how frequently prices need to be updated. In addition, should the vendor have available forecast information on the market size process, he will be required to select the α parameter as well; we anticipate that the most accessible choice in practice will be setting $\alpha = 1$. The pricing scheme will proceed as follows:

The RFP- Δ 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 $[i-1\Delta, i\Delta)$.⁶

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

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

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 ∞ .⁷

The above policy can be formalized succinctly as follows: Define $t(\Delta) = \max_i \{i\Delta : i\Delta \leq t\}$. Then the RFP- Δ 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

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

The RFP- Δ policy 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.

⁶It is simple to establish inductively that so long as $X_{i\Delta} > 0$, we must have $p_{(i-1)\Delta} < \infty$.

⁷This is a standard formalism of the notion that one cannot sell more than the initial endowment of inventory.

- 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. As it turns out, the RFP- Δ policy enjoys similar performance guarantees as the idealized policy presented and analyzed in previous sections. In particular, let $K \triangleq \max_{t \in [0, T]} \Lambda_t$, and $\eta(\Delta) = \sqrt{\Delta \log(1/\Delta)}$. We will establish the following approximation result that, together with the results of the previous section, theoretically establishes the virtues of the RFP- Δ policy.

Theorem 2. (*The Price of Discretization*) *For generalized moving average processes and an RFP- Δ policy with $\alpha = 1$, we have:*

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

where $\sigma \triangleq \phi(0)$, and we assume $\lambda_t = \lambda$ for all t .

This theorem provides valuable intuition. Precisely, it shows how the frequency with which prices are updated should be adjusted relative to key problem parameters; the relationships are intuitive but *non*-linear:

1. Volatility: Ignoring logarithmic factors, if volatility σ 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 relative to the policy that updates prices continuously. As one might anticipate, low volatility consequently calls for lower price update frequencies.
2. Load Factor: The quantity $\sqrt{\mathbb{E}K^2}T/x_0$ 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.

Our study in this section will restrict attention to generalized moving average processes with $\lambda_t = \lambda$; an extension to the more general case is straightforward but tedious. The remainder of this section outlines the proof of this approximation result.

5.1. An Outline of the Analysis

Our approximation result will naturally require some regularity in the market size process. In particular, the regularity we exploit is characterized by the following estimate of the *modulus of*

continuity of the samples paths of generalized moving average processes:

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

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

Lemma 4. (*Sample Path Modulus of Continuity*) Assume that $\bar{\Lambda}_t$ is a generalized moving average process with $\phi \in \mathcal{C}_2$ and $\lambda_t = \lambda$. Then, for $\Delta > 0$, and any $t \in [0, T)$, we have:

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

Now define

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

Observe that from the definition of the estimated market size, we have $\hat{\Lambda}_{i\Delta} = \tilde{\Lambda}_{i\Delta}$ if $X_{i\Delta}^\Delta > 0$ where we denote by X_t^Δ and X_t the inventory processes under the RFP- Δ and RFP policies respectively. We begin by estimating the deviation of the estimated market size process and the inventory process under the RFP policy from the true market size process and the inventory process under the idealized RFP policy; these results rely crucially on the sample path regularity result above. In particular, we show:

Lemma 5. Assume that $\bar{\Lambda}_t$ is a generalized moving average process with $\phi \in \mathcal{C}_2$ and $\lambda_t = \lambda$. Then, for any $t \in [0, T)$, we have almost surely:

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

and further,

$$\limsup_{\Delta \rightarrow 0} \frac{|X_{t(\Delta)}^\Delta - X_t|}{\sigma T\sqrt{\Delta\log(1/\Delta)}} \leq 4$$

We use these results to establish a precise rate at which the revenue under the RFP- Δ policy approaches that under the RFP policy, as stated in Theorem 2. The proofs for all of these results can be found in the appendix.

6. Numerical Validation

This section is dedicated to a computational exploration of the RFP- Δ 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, it is intuitive to expect that 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- Δ policy across the scenarios described above in parameter regimes that include realistic scenarios.

2. Tuning Δ : 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 λ_t and the belief that ϕ is identically 0) involves a re-optimization scheme which in our language would translate to the RFP- Δ scheme with α 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 ϕ is zero and λ_t is correctly known), we wish to explore whether the RFP- Δ scheme with $\alpha = 1$ is viable in this setting. We also explore our robust choice of α designed for a setting where λ_t is known but one wishes to be oblivious to the nature of ϕ .

6.1. Problem Specifications

We examine RFP- Δ 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 λ_t to be a constant; the RFP- Δ policy will neither be aware of this constant nor know the specification of ϕ . We will also explore the case where λ_t is not constant in Section 6.4. 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 upper bound:

$$\mathbb{E} \left[J_{\{\Lambda_t\}}^*(x_0, 0) \right] = x_0 \mathbb{E} \left[g \left(\frac{\int_0^T \Lambda_t dt}{x_0} \right) \right].$$

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 Λ_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 , λ_t , σ and T will vary across our experiments, while β will be normalized to 1. Moving forward, we will often be interested in the following summary statistics about a particular ensemble of problem instances:

- Coefficient of Variation: We define

$$\text{CV} \triangleq \frac{\sqrt{\text{var}[\int_0^T \bar{\Lambda}_t dt]}}{\mathbb{E}[\int_0^T \bar{\Lambda}_t dt]}.$$

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.

6.2. Performance Across Varying Volatility and Load Regimes

Here we seek to understand how the RFP- Δ policy performs across varying problem regimes. In particular, we take $\lambda_t = \lambda = e$ and $\beta = 1$ and vary x_0 and σ 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: $\lambda = e, \beta = 1, T = 5, \Delta = 0.1$.

Initial Inventory	Load Factor	Relative Optimality			
x_0	$x_0/\lambda T$	(CV, σ)=(0.5, 3.63)	(1, 7.25)	(2.5, 18.13)	(5, 36.25)
4	0.294	0.947	0.902	0.830	0.768
8	0.589	0.991	0.958	0.886	0.828
12	0.883	1.000	0.987	0.922	0.861
16	1.177	1.000	0.997	0.949	0.887
20	1.472	1.000	0.999	0.968	0.908

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 77%; 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- Δ 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 Λ_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 C.3.

6.3. Selecting A Re-Optimization Frequency

Section 5 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 σ). 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: $\lambda = e, \beta = 1, T = 5$.

(CV, σ)	$x_0/\lambda T=0.368$				0.736				1.104			
	$\Delta=0.1$	0.5	1	2.5	0.1	0.5	1	2.5	0.1	0.5	1	2.5
(0.1, 0.73)	0.997	0.996	0.994	0.985	1.000	1.000	1.000	1.000	1.000	1.000	1.000	1.000
(0.5, 3.63)	0.962	0.953	0.937	0.884	0.998	0.998	0.998	0.997	1.000	1.000	1.000	1.000
(1, 7.25)	0.919	0.894	0.856	0.761	0.975	0.975	0.972	0.957	0.995	0.996	0.996	0.995
(2, 14.5)	0.847	0.762	0.669	0.576	0.905	0.884	0.850	0.775	0.943	0.940	0.928	0.892

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- Δ 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.

6.4. The Hedging Parameter α 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 (i.e. he believes that ϕ is identically zero). This section seeks to answer the following questions:

1. How might the RFP- Δ policy with forecast updates but oblivious to any knowledge of $\{\lambda_t\}$ perform in this setting, i.e. how does the RFP- Δ 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 λ_t evolves according to⁸

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

We consider four sets of experiments, each corresponding to different levels of CV (0.1, 0.5, 2.5 and 5). In each set of experiments we consider the performance of the RFP policy for various settings of α ; 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:

1. Except for scenarios where CV is very low, re-optimization without forecast updates can be improved upon dramatically.

⁸This 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, λ represents potential market size

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 λ_t is no longer constant.
3. The robust choice of α 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 α .

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

(CV, σ)	Load Factor	Naive Re-opt.	Oblivious/ Robust Choice	
	$x_0/\int_0^T \lambda_t dt$	$\alpha = 0$	$\alpha = 1$	$\alpha = 0.594$
(0.1, 0.90)	0.250	0.974	0.975	0.980
	0.417	0.985	0.999	0.998
	0.583	0.993	1.000	1.000
	0.750	0.995	1.000	1.000
(0.5, 4.49)	0.250	0.853	0.905	0.894
	0.417	0.873	0.938	0.924
	0.583	0.904	0.971	0.959
	0.750	0.927	0.990	0.981
(2.5, 22.45)	0.250	0.756	0.820	0.801
	0.417	0.748	0.821	0.799
	0.583	0.748	0.828	0.805
	0.750	0.753	0.840	0.815
(5, 44.90)	0.250	0.755	0.815	0.796
	0.417	0.742	0.809	0.789
	0.583	0.736	0.809	0.786
	0.750	0.733	0.810	0.787

6.5. 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- Δ policy with α 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 λ_t is itself time varying and unknown to the manager. A robust selection of the α parameter provides added value in settings where λ_t is time varying and known and volatility is low.

7. 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- Δ 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. *Oper. Res.*, 57(5):1169–1188, 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.
- T. Boyaci and O. Ozer. Information acquisition for capacity planning via pricing and advance selling: When to stop and act? *Oper. Res.*, 58(5):1328–1349, 2010.
- 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.

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

A. Proofs for Section 4

We begin with establishing properties of the unit revenue function, $g(\cdot)$.

Lemma 6.

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 non-decreasing 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 7.

■

Lemma 1.

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

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

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

By the concavity of $g(\cdot)$ established in Lemma 6 and Jensen's inequality, we immediately have:

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

which is the second inequality. The fact that $J_{\{\Lambda_t\}}^*(x_0, 0) = x_0 g \left(\frac{\int_0^T \Lambda_t dt}{x_0} \right)$ 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 μ and variance σ^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 6, $g(\cdot)$ is non-decreasing, it then follows that

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

The sub-additivity of $g(\cdot)$ from the fourth part of Lemma 6 then yields the final inequality. ■

Lemma 3.

$$J^{\pi_{\text{RFP}}}(x^0, \lambda^0, 0) \geq 0.342\alpha x_0 g \left(\frac{\int_0^T \sigma_t dt}{x_0 \sqrt{2\pi}} \right) + \frac{1-\alpha}{2} x_0 g \left(\frac{\int_0^T \lambda_t dt}{x_0} \right).$$

Proof. We have:

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

$$(6) \quad \geq \alpha \frac{x_0}{T} \mathbb{E} \left[\int_0^T g \left(\frac{\Lambda_t T}{x_0} \right) dt \right] + (1 - \alpha) \frac{x_0}{\int_0^T \lambda_s ds} \mathbb{E} \left[\int_0^T \lambda_t g \left(\frac{\Lambda_t \int_0^T \lambda_s ds}{x_0 \lambda_t} \right) dt \right].$$

where the second equality holds by the definition of $g(\cdot)$, the first inequality follows from the lower bound on X_t established in Lemma 2 on the inventory balancing property and the property that $zg(1/z)$ is a non-decreasing function. The final inequality holds because $zg(1/z)$ is a concave function.

Next, we prove the lower bounds of two terms in 6 respectively. For the first term, 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 6, and the positivity of λ_t . The second inequality holds due to Property 4 in Lemma 6. The final inequality was derived as a property of the class of market-size processes we consider in Property 3 in Lemma 8.

For the second term in 6, 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 non-decreasing. ■

Theorem 1. Assume Λ_t is a generalized moving average process. Then,

1. For the RFP policy with $\alpha = 1$, we have

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

where $B \triangleq \sigma_T / \sqrt{2\pi} \lambda^2$, and we assume $\lambda_t = \lambda$ for all t .

2. For the RFP policy with $\alpha = 0.594$, and arbitrary forecasts $\{\lambda_t\}$, we have:

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

Proof. We provide a proof of the first part of the theorem; the second part is proved in Section 4.

By Lemma 1 we have that $J^*(x^0, \lambda^0, 0) \leq J_{\text{CE}}^*(x^0, \lambda^0, 0)$. Consequently, if $\lambda_t = \lambda$ for all t , then for an arbitrary α ,

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

Now, we have:

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

The first equality holds by definition of $g(\cdot)$ and Lemma 1. The second inequality follows by applying the inventory balancing Lemma (Lemma 2) to obtain a lower bound on X_t along with the property that $zg(1/z)$ is a non-decreasing function, which is established in Lemma 6. The third

inequality follows by Property 4 in Lemma 6. The final inequality follows by Property 3 in Lemma 8.

In addition, we have:

$$\begin{aligned}
\frac{J^{\pi_{RFP}}(x^0, \lambda^0, 0)}{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 Φ 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 8, 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

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

■

A.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 3. *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:*

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

Proof. Now, we have:

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

$$\begin{aligned}
&\geq \frac{1}{T} \int_0^T \left[1 - \Phi \left(\frac{\bar{\sigma}_{T,1}\sqrt{2}}{\sigma_t\sqrt{\pi}} \right) + \int_0^{\bar{\sigma}_{T,1}\sqrt{2/\pi}} \frac{y}{\bar{\sigma}_{T,1}\sqrt{2/\pi}} \frac{\exp(-y^2/2\sigma_t^2)}{\sqrt{2\pi\sigma_t^2}} dy \right] dt \\
&\geq 0.243.
\end{aligned}$$

The first equality holds by definition of $g(\cdot)$ and Lemma 1. In addition, we use the fact that for a Normal random variable X with mean μ and variance σ^2 , we know that $\mathbb{E}[|X|] \leq \mu + \sqrt{2}\sigma/\sqrt{\pi}$ so that $\mathbb{E}[\Lambda_t] = \mathbb{E}[\bar{\Lambda}_t] \leq \lambda + \sqrt{2}\sigma_t/\sqrt{\pi}$. The second inequality follows by applying the inventory balancing Lemma (Lemma 2) to obtain a lower bound on X_t along with the property that $zg(1/z)$ is a non-decreasing function, which is established in Lemma 6. The fourth inequality follows by Property 4 in Lemma 6. Finally, by Lemma 1 we have that $J^*(x^0, \lambda^0, 0) \leq J_{\text{CE}}^*(x^0, \lambda^0, 0)$ so that

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

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, θ controls the speed of market-size adjustment, λ and σ 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 Λ_0 is distributed according to this stationary distribution and define $\lambda = \Lambda_0$.

Theorem 4. *Consider the RFP policy with $\alpha = 0$. Then if Λ_t is driven by the CIR process above, we have:*

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

Proof. Now, we have:

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

$$\geq 0.632.$$

The second inequality follows by applying the inventory balancing Lemma (Lemma 2) to obtain a lower bound on X_t along with the property that $zg(1/z)$ is a non-decreasing function, which is established in Lemma 6. The third inequality follows by Property 4 in Lemma 6. $\Gamma(\cdot, \cdot)$ 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 1 we have that $J^*(x^0, \lambda^0, 0) \leq J_{\text{CE}}^*(x^0, \lambda^0, 0)$ so that

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

and the guarantee follows. ■

B. Proofs for Section 5

Lemma 4. (*Sample Path Modulus of Continuity*) Assume that $\bar{\Lambda}_t$ is a generalized moving average process with $\phi \in \mathcal{C}_2$ and $\lambda_t = \lambda$. Then, for $\Delta > 0$, and any $t \in [0, T)$, we have:

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

Proof. For any $0 \leq t \leq T - \tau$ and $0 \leq \tau \leq \Delta$, we have

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

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

Now, we have

$$\limsup_{\Delta \rightarrow 0} \sup_{0 \leq t \leq T-\tau, 0 \leq \tau \leq \Delta} \frac{|\Lambda_{t+\tau} - \Lambda_t|}{\sqrt{2\Delta \log(1/\Delta)}} \leq \limsup_{\Delta \rightarrow 0} \sup_{0 \leq t \leq T-\tau, 0 \leq \tau \leq \Delta} \frac{\sigma |Z_{t+\tau} - Z_t| + (3L_{\phi_1} B + L_{\phi_2} B t) \tau}{\sqrt{2\Delta \log(1/\Delta)}}$$

$$\begin{aligned}
&= \limsup_{\Delta \rightarrow 0} \sup_{0 \leq t \leq T-\tau, 0 \leq \tau \leq \Delta} \frac{\sigma |Z_{t+\tau} - Z_t|}{\sqrt{2\Delta \log(1/\Delta)}} \\
&= \sigma,
\end{aligned}$$

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

Lemma 5. *Assume that $\bar{\Lambda}_t$ is a generalized moving average process with $\phi \in \mathcal{C}_2$ and $\lambda_t = \lambda$. Then, for any $t \in [0, T)$, we have almost surely:*

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

and further,

$$\limsup_{\Delta \rightarrow 0} \frac{|X_{t(\Delta)}^\Delta - X_t|}{\sigma T \sqrt{\Delta \log(1/\Delta)}} \leq 4$$

Proof. First, we prove the convergence rate of the estimated market size. We have

$$\begin{aligned}
|\tilde{\Lambda}_{t(\Delta)} - \Lambda_t| &= \left| \frac{1}{\Delta} \int_{t(\Delta)-\Delta}^{t(\Delta)} \Lambda_s ds - \Lambda_t \right| \\
&\leq \sup_{0 \leq s \leq 2\Delta} |\Lambda_{t-s} - \Lambda_t|.
\end{aligned}$$

Therefore,

$$\begin{aligned}
\limsup_{\Delta \rightarrow 0} \frac{|\tilde{\Lambda}_{t(\Delta)} - \Lambda_t|}{\sigma \sqrt{\Delta \log(1/\Delta)}} &\leq \limsup_{\Delta \rightarrow 0} \frac{\sup_{0 \leq s \leq 2\Delta} |\Lambda_{t-s} - \Lambda_t|}{\sigma \sqrt{\Delta \log(1/\Delta)}} \\
&\leq \limsup_{\Delta \rightarrow 0} \frac{\sup_{0 \leq t \leq T-s, 0 \leq s \leq 2\Delta} |\Lambda_{t-s} - \Lambda_t|}{\sigma \sqrt{\Delta \log(1/\Delta)}} \\
&= \limsup_{\Delta \rightarrow 0} \frac{\sup_{0 \leq t \leq T-s, 0 \leq s \leq 2\Delta} |\Lambda_{t-s} - \Lambda_t|}{\sigma \sqrt{4\Delta \log(1/2\Delta)}} \frac{\sqrt{4\Delta \log(1/2\Delta)}}{\sqrt{\Delta \log(1/\Delta)}} \\
&\leq 2.
\end{aligned}$$

The last inequality follows from Lemma 4.

Next, we prove the convergence rate of the inventory process under the RFP- Δ policy. Now for $i > 0$, we have

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

We have that for $\epsilon > 0$, there exist numbers $C(\epsilon), D(\epsilon) < \infty$, such that $1/(T-s) \leq C(\epsilon)$ and

$|\frac{d}{ds} \frac{1}{T-s}| \leq D(\epsilon)$ for all $s < T - \epsilon$. Now, for any s, Δ such that $2\Delta \leq s < T - \epsilon$, we have:

$$\begin{aligned}
(7) \quad & \left| \min \left\{ \bar{F}(p^*), \frac{X_s}{\Lambda_s(T-s)} \right\} \Lambda_s - \min \left\{ \bar{F}(p^*), \frac{X_{s(\Delta)}}{\tilde{\Lambda}_{s(\Delta)}(T-s(\Delta))} \right\} \Lambda_s \right| \\
& \leq \left| \min \left\{ \bar{F}(p^*) \Lambda_s, \frac{X_s}{T-s} \right\} - \min \left\{ \bar{F}(p^*) \tilde{\Lambda}_{s(\Delta)}, \frac{X_{s(\Delta)}}{T-s(\Delta)} \right\} \right| \\
& \quad + \min \left\{ \bar{F}(p^*), \frac{X_{s(\Delta)}}{\tilde{\Lambda}_{s(\Delta)}(T-s(\Delta))} \right\} |\Lambda_s - \tilde{\Lambda}_{s(\Delta)}| \\
& \leq \bar{F}(p^*) |\Lambda_s - \tilde{\Lambda}_{s(\Delta)}| + \left| \frac{X_s}{T-s} - \frac{X_{s(\Delta)}}{T-s(\Delta)} \right| + \bar{F}(p^*) |\Lambda_s - \tilde{\Lambda}_{s(\Delta)}| \\
& \leq 2 \sup_{0 \leq \tau \leq 2\Delta} |\Lambda_s - \Lambda_{s-\tau}| + \left| \frac{X_s}{T-s} - \frac{X_{s(\Delta)}}{T-s} \right| + \left| \frac{X_{s(\Delta)}}{T-s} - \frac{X_{s(\Delta)}}{T-s(\Delta)} \right| \\
& \leq 2 \sup_{0 \leq s \leq T-\tau, 0 \leq \tau \leq 2\Delta} |\Lambda_s - \Lambda_{s-\tau}| + (C(\epsilon)K + x_0 D(\epsilon))\Delta,
\end{aligned}$$

where $K \triangleq \sup_{t \in [0, T]} \Lambda_t$. 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 $(i+1)\Delta < T - \epsilon$,

$$\begin{aligned}
|X_{(i+1)\Delta}^\Delta - X_{(i+1)\Delta}| &= \left| \left(X_{i\Delta}^\Delta - \int_{i\Delta}^{(i+1)\Delta} \min \left\{ \bar{F}(p^*), \frac{X_{i\Delta}^\Delta}{\tilde{\Lambda}_{i\Delta}(T-i\Delta)} \right\} \Lambda_s ds \right)^+ \right. \\
& \quad \left. - \left(X_{i\Delta} - \int_{i\Delta}^{(i+1)\Delta} \min \left\{ \bar{F}(p^*), \frac{X_s}{\Lambda_s(T-s)} \right\} \Lambda_s ds \right)^+ \right| \\
& \leq \left| \left(X_{i\Delta}^\Delta - \int_{i\Delta}^{(i+1)\Delta} \min \left\{ \bar{F}(p^*), \frac{X_{i\Delta}^\Delta}{\tilde{\Lambda}_{i\Delta}(T-i\Delta)} \right\} \Lambda_s ds \right)^+ \right. \\
& \quad \left. - \left(X_{i\Delta} - \int_{i\Delta}^{(i+1)\Delta} \min \left\{ \bar{F}(p^*), \frac{X_{i\Delta}}{\tilde{\Lambda}_{i\Delta}(T-i\Delta)} \right\} \Lambda_s ds \right)^+ \right| \\
& \quad + \left| \int_{i\Delta}^{(i+1)\Delta} \min \left\{ \bar{F}(p^*), \frac{X_s}{\Lambda_s(T-s)} \right\} \Lambda_s ds \right. \\
& \quad \left. - \int_{i\Delta}^{(i+1)\Delta} \min \left\{ \bar{F}(p^*), \frac{X_{i\Delta}}{\tilde{\Lambda}_{i\Delta}(T-i\Delta)} \right\} \Lambda_s ds \right| \\
& \leq |X_{i\Delta}^\Delta - X_{i\Delta}| \\
& \quad + \int_{i\Delta}^{(i+1)\Delta} \left| \min \left\{ \bar{F}(p^*), \frac{X_s}{\Lambda_s(T-s)} \right\} - \min \left\{ \bar{F}(p^*), \frac{X_{s(\Delta)}}{\tilde{\Lambda}_{s(\Delta)}(T-s(\Delta))} \right\} \right| \Lambda_s ds \\
& \leq |X_{i\Delta}^\Delta - X_{i\Delta}| + 2 \sup_{0 \leq s \leq T-\tau, 0 \leq \tau \leq 2\Delta} |\Lambda_s - \Lambda_{s-\tau}| \Delta + (C(\epsilon)K + x_0 D(\epsilon))\Delta^2,
\end{aligned}$$

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 (7). Moreover, since trivially $|X_\Delta^\Delta - X_\Delta| \leq \int_0^\Delta \Lambda_s ds \leq$

$K\Delta$, we must have for any positive integer i with $i\Delta < T - \epsilon$,

$$\begin{aligned} |X_{i\Delta}^\Delta - X_{i\Delta}| &\leq 2T \sup_{0 \leq s \leq T-\tau, 0 \leq \tau \leq 2\Delta} |\Lambda_s - \Lambda_{s-\tau}| + (C(\epsilon)K + x_0 D(\epsilon))\Delta^2(i-1) + K\Delta \\ &\leq 2T \sup_{0 \leq s \leq T-\tau, 0 \leq \tau \leq 2\Delta} |\Lambda_s - \Lambda_{s-\tau}| + (C(\epsilon)KT + x_0 D(\epsilon)T + K)\Delta. \end{aligned}$$

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

$$\begin{aligned} |X_{t(\Delta)}^\Delta - X_t| &\leq |X_{t(\Delta)}^\Delta - X_{t(\Delta)}| + |X_{t(\Delta)} - X_t| \\ &\leq 2T \sup_{0 \leq s \leq T-\tau, 0 \leq \tau \leq 2\Delta} |\Lambda_s - \Lambda_{s-\tau}| + (C(\epsilon)KT + x_0 D(\epsilon)T + K)\Delta + K\Delta \end{aligned}$$

Therefore,

$$\begin{aligned} \limsup_{\Delta \rightarrow 0} \frac{|X_{t(\Delta)}^\Delta - X_t|}{\sigma T \sqrt{\Delta \log(1/\Delta)}} &\leq \limsup_{\Delta \rightarrow 0} \frac{2T \sup_{0 \leq s \leq T-\tau, 0 \leq \tau \leq 2\Delta} |\Lambda_s - \Lambda_{s-\tau}| + (C(\epsilon)BT + x_0 D(\epsilon)T + B)\Delta + B\Delta}{\sigma T \sqrt{\Delta \log(1/\Delta)}} \\ &= \limsup_{\Delta \rightarrow 0} \frac{2 \sup_{0 \leq s \leq T-\tau, 0 \leq \tau \leq 2\Delta} |\Lambda_s - \Lambda_{s-\tau}|}{\sigma \sqrt{\Delta \log(1/\Delta)}} \\ &= \limsup_{\Delta \rightarrow 0} \frac{2 \sup_{0 \leq s \leq T-\tau, 0 \leq \tau \leq 2\Delta} |\Lambda_s - \Lambda_{s-\tau}|}{\sigma \sqrt{4\Delta \log(1/2\Delta)}} \frac{\sqrt{4\Delta \log(1/2\Delta)}}{\sqrt{\Delta \log(1/\Delta)}} \\ &\leq 4, \end{aligned}$$

for all $t < T - \epsilon$. The last inequality follows from Lemma 4. Since our choice of $\epsilon > 0$ was arbitrary, the result follows. \blacksquare

Theorem 2. (*The Price of Discretization*) For generalized moving average processes and an RFP- Δ policy with $\alpha = 1$, we have:

$$\limsup_{\Delta \rightarrow 0} \frac{|J^{\pi_{\text{RFP}}}(x^0, \lambda^0, 0) - J^{\pi_{\text{RFP}}^\Delta}(x^0, \lambda^0, 0)|}{\eta(\Delta) \log(1/\eta(\Delta))} \leq 4p^* \bar{F}(p^*) \sigma \frac{\mathbb{E}K^2 T^3}{x_0^2}$$

where $\sigma \triangleq \phi(0)$, and we assume $\lambda_t = \lambda$ for all t .

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

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

Using this fact with Lemma 5 allows us to conclude after some algebraic manipulation that for any $t < T$ that

$$\limsup_{\Delta \rightarrow 0} \frac{1}{\sigma \eta(\Delta)} \left| \frac{\tilde{\Lambda}_{t(\Delta)}(T-t(\Delta))}{X_{t(\Delta)}^\Delta} - \frac{\Lambda_t(T-t)}{X_t} \right| \leq \frac{2T}{x_0} + \frac{4KT^3}{x_0^2(T-t)}$$

Let $\kappa(\Delta) \triangleq 8T^2 \sigma \eta(\Delta) / x_0$. Observe that on $t < T - \kappa(\Delta)$, we must have by the Balancing Lemma that $X_t \geq 8T \sigma \eta(\Delta)$, so that for Δ sufficiently small, Lemma 5 guarantees that $X_{t(\Delta)}^\Delta > 0$ as well.

Consequently, we have that for Δ sufficiently small:

$$\begin{aligned}
(8) \quad & |J^{\pi_{\text{RFP}}^\Delta}(x^0, \lambda^0, 0) - J^{\pi_{\text{RFP}}}(x^0, \lambda^0, 0)| \\
&= \left| \mathbb{E} \left[\int_0^T \pi_{\text{RFP}}^\Delta(X^{\Delta,t}, t) \bar{F} \left(\pi_{\text{RFP}}^\Delta(X^{\Delta,t}, t) \right) \Lambda_t dt \right] - \mathbb{E} \left[\int_0^T \pi_{\text{RFP}}(X^t, t) \bar{F} \left(\pi_{\text{RFP}}(X^t, t) \right) \Lambda_t dt \right] \right| \\
&\leq \mathbb{E} \left[\left| \int_0^T \pi_{\text{RFP}}^\Delta(X^{\Delta,t}, t) \bar{F} \left(\pi_{\text{RFP}}^\Delta(X^{\Delta,t}, t) \right) \Lambda_t dt - \int_0^T \pi_{\text{RFP}}(X^t, t) \bar{F} \left(\pi_{\text{RFP}}(X^t, t) \right) \Lambda_t dt \right| \right] \\
&\leq \mathbb{E} \left[\left| \int_0^{T-\kappa(\Delta)} \pi_{\text{RFP}}^\Delta(X_{t(\Delta)}^\Delta, t) \bar{F} \left(\pi_{\text{RFP}}^\Delta(X_{t(\Delta)}^\Delta, t) \right) \Lambda_t dt - \int_0^{T-\kappa(\Delta)} \pi_{\text{RFP}}(X^t, t) \bar{F} \left(\pi_{\text{RFP}}(X^t, t) \right) \Lambda_t dt \right| \right] \\
&\quad + \mathbb{E} K \kappa(\Delta) p^* \bar{F}(p^*).
\end{aligned}$$

Now, by our choice of $\kappa(\Delta)$, we have that for Δ sufficiently small that

$$\begin{aligned}
\mathcal{E}(\Delta) &\triangleq \left| \int_0^{T-\kappa(\Delta)} \pi_{\text{RFP}}^\Delta(X_{t(\Delta)}^\Delta, t) \bar{F} \left(\pi_{\text{RFP}}^\Delta(X_{t(\Delta)}^\Delta, t) \right) \Lambda_t dt - \int_0^{T-\kappa(\Delta)} \pi_{\text{RFP}}(X^t, t) \bar{F} \left(\pi_{\text{RFP}}(X^t, t) \right) \Lambda_t dt \right| \\
&\leq K p^* \bar{F}(p^*) \int_0^{T-\kappa(\Delta)} \left| \frac{\tilde{\Lambda}_{t(\Delta)}(T-t(\Delta))}{X_{t(\Delta)}^\Delta} - \frac{\Lambda_t(T-t)}{X_t} \right| dt \\
&\leq K p^* \bar{F}(p^*) \left(\sigma \eta(\Delta) \int_0^{T-\kappa(\Delta)} \left(\frac{2T}{x_0} + \frac{4KT^3}{x_0^2(T-t)} \right) dt \right) \\
&\leq K p^* \bar{F}(p^*) \sigma \eta(\Delta) \left(\frac{2T^2}{x_0} + \frac{4KT^3(\log T + \log(1/\kappa(\Delta)))}{x_0^2} \right)
\end{aligned}$$

where the first inequality follows from the fact that the function $g(y)/y$ has its first derivative bounded in absolute value by $p^* \bar{F}(p^*)$, and the second inequality was established at the start of the proof. It follows that

$$\limsup_{\Delta \rightarrow 0} \frac{\mathcal{E}(\Delta)}{\sigma \eta(\Delta) \log(1/\eta(\Delta))} \leq \frac{4K^2 T^3 p^* \bar{F}(p^*)}{x_0^2}$$

Using this inequality, (8) then yields

$$\begin{aligned}
\limsup_{\Delta \rightarrow 0} \frac{|J^{\pi_{\text{RFP}}}(x^0, \lambda^0, 0) - J^{\pi_{\text{RFP}}^\Delta}(x^0, \lambda^0, 0)|}{\eta(\Delta) \log(1/\eta(\Delta))} &\leq \limsup_{\Delta \rightarrow 0} \mathbb{E} \left[\frac{\mathcal{E}(\Delta) + \kappa(\Delta) p^* \bar{F}(p^*) K}{\eta(\Delta) \log(1/\eta(\Delta))} \right] \\
&\leq \frac{4\mathbb{E} K^2 T^3 p^* \bar{F}(p^*) \sigma}{x_0^2}
\end{aligned}$$

where the second inequality follows from Fatou's lemma. ■

C. Miscellaneous Results and Computations

C.1. Properties of the Market-Size Process

We present in this Section, a few technical results for a class of market size processes satisfying the assumption below. It is simple to check that this class subsumes the class of generalized moving average processes we have studied in this paper.

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 non-decreasing as a function of t and concave.

Indeed it is evident that our moving average processes satisfy the first two requirements; to see that the last requirement is satisfied, we observe that in the case of moving average processes, Ito isometry yields $\text{Var}(\bar{\Lambda}_t) = \int_0^t \phi^2(s) ds$ which is evidently non-decreasing and concave.

Lemma 7. 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

$$(9) \quad \frac{1}{x} \int_0^x f(t) dt \leq f\left(\frac{x}{2}\right).$$

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 (9) follows by Jensen's inequality. ■

Now, we use Lemma 7 to characterize properties of the volatility of the market-size process, σ_t^2 . Defining

$$\begin{aligned} \bar{\sigma}_{T,1} &\triangleq \int_0^T \sigma_t dt / T \text{ and} \\ \bar{\sigma}_{T,2} &\triangleq \int_0^T \sigma_t^2 dt / T. \end{aligned}$$

we have:

Lemma 8.

1. σ_t is non-decreasing 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. σ_t^2 is non-decreasing in t directly implies that σ_t is non-decreasing in t . Now,

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

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

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

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

where inequality (10) follows by Lemma 7 and the concavity of σ_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 7 yields $\bar{\sigma}_{T,2} \leq \sigma_{T/2}^2$, so that the first part of Lemma 7 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:

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

where the first inequality follows from the previous property (i.e. Lemma 8, 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. ■

C.2. Analysis for Example 1 in Section 3

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

$$\frac{J^{\pi_{\text{FP}}}(x^0, \lambda^0, 0)}{J^*(x^0, \lambda^0, 0)} \leq O((\log T)^{-1}),$$

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 Section 4, as we will use performance under this policy as a lower bound to performance under an optimal policy. Now, we have

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

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 ; in the case of our example, recall that $p^* = 1$. We now focus on the second inequality: Theorem 1 showed that

$$J^{\pi_{\text{RFP}}}(x^0, \lambda^0, 0) \geq 0.342J_{\text{CE}}^*(x^0, \lambda^0, 0),$$

while by the definition of the unit revenue function, $g(\cdot)$, in Section 4, we know that

$$J_{\text{CE}}^*(x^0, \lambda^0, 0) = x_0g \left(\frac{\int_0^T \mathbb{E}[\Lambda_t]dt}{x_0} \right).$$

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 6, it follows that

$$J^{\pi_{\text{RFP}}}(x^0, \lambda^0, 0) \geq 0.342x_0g \left(1 + \frac{2T^{3/2}\sigma}{3\sqrt{2\pi}x_0} - \frac{\lambda^2\sqrt{T}}{\sigma\sqrt{2\pi}x_0} \right).$$

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

In our computational experiments, we compared performance of the RFP- Δ 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 summarized (for an OU process) in Tables 4 and 5.

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

Initial Inventory	Load Factor	Relative Optimality		
x_0	$x_0/\lambda T$	$J^{\pi_{RFP}}/J^*$	$J^{\pi_{RFP}^\Delta}/J^*$	$J^{\pi_{RFP}^\Delta}/J^{UB}$
4	0.294	0.951	0.923	0.830
8	0.589	0.962	0.941	0.886
12	0.883	0.979	0.965	0.922
16	1.177	0.990	0.977	0.949
20	1.472	0.998	0.990	0.968

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

Initial Inventory	Load Factor	Relative Optimality		
x_0	$x_0/\lambda T$	$J^{\pi_{RFP}}/J^*$	$J^{\pi_{RFP}^\Delta}/J^*$	$J^{\pi_{RFP}^\Delta}/J^{UB}$
4	0.294	0.922	0.891	0.768
8	0.589	0.938	0.915	0.828
12	0.883	0.947	0.929	0.861
16	1.177	0.951	0.936	0.887
20	1.472	0.966	0.951	0.908

In the experiments above $J^{\pi_{RFP}^\Delta}/J^{UB}$ is the quantity reported for the bulk of our experiments – performance relevant to a clairvoyant upper bound. The quantity $J^{\pi_{RFP}^\Delta}/J^*$ 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 $J^{\pi_{RFP}}/J^*$ 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.