Asymptotically Optimal Dynamic Pricing in Observable Queues

Jeunghyun Kim, Ramandeep Randhawa

Asymptotically Optimal Dynamic Pricing in Observable Queues

February 23, 2015

Abstract

We study optimal dynamic pricing to maximize revenues in queueing systems with price and delay sensitive customers. A key feature of our model is that the system congestion is visible so that upon arrival, customers decide to join the system based on the congestion and the price at that time. We analyze this problem in the typical asymptotic regime of large customer market size and capacity. This asymptotic analysis involves solving a first order or fluid optimization problem that ignores stochastic variability and then refining it by minimizing the revenue loss that occurs due to stochasticity. Denoting the market size by n, one expects the revenue loss due to stochasticity to be on the scale of n raised to the power one-half. However, surprisingly, we find that the optimal dynamic pricing leads to an order improvement and the loss due to stochasticity is on the scale of n raised to the power one-third. The corresponding asymptotic control problem also turns out to be non-conventional. We solve this problem to obtain a near-optimal dynamic pricing policy, and further, we show that a simple policy of using only two prices can achieve most of the benefits of dynamic pricing.

Keywords: revenue management, dynamic pricing, lead-time quotation, asymptotic analysis, diffusion analysis

1 Introduction

Time is an important attribute of many products and services with customers valuing quick access at a premium. Given that the ability of a firm to provide quick access changes with the congestion in the system, pricing dynamically as a function of the congestion seems to be a good strategy. It has been observed (see, for example, Biller et al., 2005) that employing dynamic pricing has improved profits and other supply chain metrics in the automotive industry. Indeed, it is now standard for make-to-order manufacturers to change prices dynamically as a means to manage demand and supply. Dynamic pricing has also been implemented in road-tolls, where toll-paying single drivers gain access to high occupancy carpool lanes, and more recently by firms such as Uber that use surge pricing to manage congestion between customers and available taxis. The goal of this paper is to address the questions: *What is the value of changing prices dynamically and what is the optimal dynamic pricing strategy?*

We study this question in the context of a monopolistic firm processing price and delay sensitive jobs (or, equivalently customers) who arrive as a Poisson process. Customers are heterogeneous in their valuation and have linear disutility in their waiting times. At the time of arrival, the customers observe the posted or quoted waiting time and compare their value with the sum total of the price and waiting costs and then either join if the net value is positive or leave the system otherwise. The firm has limited and fixed capacity and we model the system as a single-server queue. The firm's decision is to set prices as a function of the congestion or queue-length in order to maximize its long-term average revenue rate. As we will discuss in the literature review section, static versions of this problem, in which the firm sets a single price using steady-state queueing behavior, are well understood. However, the literature on dynamic pricing is quite sparse and treats only some special cases. A priori, one expects some similarity between dynamic pricing and static pricing especially for large systems. This is so because for large systems both pricing methods, loosely speaking, attempt to minimize the system variability, or rather the revenue loss due to variability.

Our first result is that this intuition doesn't quite play out and in fact there is a fundamental difference between the two pricing methods in large systems; dynamic pricing can lead to an "order improvement" in performance relative to static pricing. Formally, we consider the typical large system asymptotics in which the potential customer arrival rate and processing capacity are both large, and increasing without bound. In this regime, static pricing, with linear customer disutility for waits, has been established to follow the conventional square-root behavior, i.e., the loss in revenue due to variability is of $\mathcal{O}(\sqrt{n})$, where n denotes the system size. Somewhat surprisingly, we find that dynamic pricing can mitigate the variability to the extent that the corresponding loss in revenue is of a lower order. We identify this order to be the system scale raised to the power one-third, i.e., $\mathcal{O}(n^{1/3})$. Further, we prove that a simple two-price policy, that sets a high price when the queues are large and a low price otherwise can reap most of the benefits of dynamic pricing. In particular, we prove that our proposed two-price policy achieves the $\mathcal{O}(n^{1/3})$-scale up to logarithmic terms.

Intuitively, the benefit from dynamic pricing that we find arises because a dynamic pricing policy uses larger price refinements (of a larger order) compared with that in static pricing and can maintain congestion at a lower order by increasing prices when queues are long, but at the same time when the queues are short, the prices can be decreased to increase volume, and thus, the revenue. In this sense, dynamic pricing provides an order of magnitude improvement over static pricing. A mathematical explanation for this order improvement is as follows: the second-order optimization problem trades-off the expected steady-state price refinement with the queue-length scaled by system capacity (n). The conventional logic suggests that both of these should be on the same order, which then would be the square-root order. However, by changing prices dynamically especially by introducing negative price corrections, one can maintain the expected steady-state price refinement at a much lower level, and hence balancing the trade-off with the expected queue-length leads to a lower order of revenue loss. Technically, the optimization problem also contains the second moment of the steady-state price refinement, which is typically of a lower order than the expected steady-state price refinement. However, by changing the price dynamically, the first moment of the steady-state price refinement is lowered to an extent that the second moment of the steady-state price refinement becomes important and balancing this with the scaled expected steady-state queue-length yields the one-third order.

From an analysis perspective, the conventional approaches of formulating the optimization problem using a limiting process do not work in our setting. Instead, we directly analyze the underlying Markov Chain and work with approximations to the steady-state probabilities. These allow us to formulate a control problem, which interestingly is not separable in the system scale, and in this sense differentiates our work from antecedent literature as well. We solve this control problem to propose a near-optimal pricing policy. This policy is intricate (as one would expect) and changes prices with every arrival and departure to the system. Given such an intricate near-optimal dynamic pricing policy, we feel our finding that a two-price policy can perform extremely well, practically appealing.

2 Literature Review

This paper studies dynamic pricing to maximize revenues in queueing systems with observable congestion under large market asymptotics. To place this paper's model and results in perspective, it is useful to view the related literature in terms of three dimensions: (a) *Pricing method*: do the prices change with congestion (dynamic pricing) or remain fixed (static pricing); (b) *Observability of congestion*: do the customers observe the congestion at time of joining the system or not; and (c) *Mode of analysis*: is the analysis exact or asymptotic.

The literature that studies static pricing is much more extensive compared with that on dynamic pricing. One of the first papers in this literature is Naor (1969). That paper studies the optimal static price to be set when queues are observable. Another influential paper that considers static pricing is Mendelson and Whang (1990), which studies differentiation between different customer types when maximizing social welfare, but when the queues are not observable. Unobservable queues lend a certain simplicity to the analysis because they allow using steady-state congestion formulas directly, rather than dealing with the underlying Markov chain. More recent papers that study static prices with unobservable queues are Cachon and Feldman (2011), that compares subscription with pay-per-use, and Haviv and Randhawa (2014), that shows that a fixed price can perform very well without any knowledge of the overall demand. All these papers take an exact analysis approach. Maglaras and Zeevi (2003) is the first paper that studied the pricing problem asymptotically and characterized the optimality of the "square-root" regime when the customer delay sensitivity is linear. That paper considers the case of unobservable queues. In that asymptotic framework, Kumar and Randhawa (2010) studies static pricing and capacity sizing and reveals that optimality scales depend on the curvature of the customer delay cost function near the origin. It is worth mentioning that in this literature, there are also papers that consider settings in which the firm announces lead-times. Plambeck and Ward (2008) considers such a model and uses an asymptotic analysis to characterize the optimal static price and dynamic sequencing policies when lead-times are quoted dynamically. Asymptotically, quoting lead-times is quite similar to simply making the queue-length visible. One key difference is that in manufacturing settings, expediting orders is allowed to ensure quoted lead-times are always met, which provides some additional flexibility. Another related paper Lee and Ward (2014) considers customer abandonments while studying the asymptotically optimal static pricing and capacity sizing decisions.

Turning to the literature on dynamic pricing. Early papers in this area assume that customers are sensitive to only prices and not delay, and the firm changes prices dynamically because there is a cost to the firm from having high congestion. Examples of such papers are Low (1974) and Paschalidis and Tsitsiklis (2000). These two papers consider finite buffer systems. Low (1974) proves that prices are non-decreasing in the number of customers in the system, whereas Paschalidis and Tsitsiklis (2000) considers a multi-class system and numerically shows that static pricing can perform quite well. More recent papers that study related problems are Yoon and Lewis (2004) and Maglaras (2006): Yoon and Lewis (2004) assumes deterministic customer valuation but allows non-stationarity in arrival and service rates and establishes interesting structural properties of the optimal policy, and propose a practical point-wise stationary approximation; and Maglaras (2006) proposes tractable capacity sizing, dynamic pricing, and sequencing for a multi-class system based on a fluid approximation. Ata and Shneorson (2006) incorporates customer delay sensitivity and observable queues into the dynamic pricing problem to maximize social welfare; the capacity is also controllable. In this literature, perhaps the paper closest to ours is Celik and Maglaras (2008). That paper uses an asymptotic approach to study dynamic pricing and scheduling to maximize revenues when the firm quotes lead-times. A key difference is that we focus on settings in which capacity is constraining. We feel that this approach provides flexibility in modeling situations in which capacity decisions are made over a long horizon and further helps study the cases in which there are mismatches between supply and demand. Interestingly, this model change leads to a completely different structure of the asymptotic problem. Another related paper in this domain is Ata and Olsen (2013) which studies asymptotically optimal dynamic pricing and lead-time quotation in a setting with two customer classes, and the customers have convex-concave delay costs; the customer type is assumed not known to the manager.

While dynamic pricing is useful in dealing with congestible systems, its usefulness can be enhanced when there is parameter uncertainty or changes in underlying parameters. For instance, Afeche and Ata (2013) uses dynamic pricing to learn customer delay sensitivity and Besbes and Maglaras (2010) studies dynamic pricing when the demand is time-varying and stochastic. The latter paper uses an observable queue framework with a customer model that is identical to ours. However, the authors use an asymptotic fluid approach to capture the changes in market size, whereas we use a diffusion based approach to capture the changes in congestion.

Our paper also relates to the literature on dynamic pricing in inventory systems. In particular, our discussion of simple dynamic pricing policies relates to Netessine (2006), which optimizes the number of price changes when dealing with a non-stationary arrival process. Other papers that use dynamic pricing to learn demand characteristics in such systems are Farias and Van Roy (2010), Eren and Maglaras (2010), Besbes and Zeevi (2012), and Harrison et al. (2012).

3 Model

We model the firm as a single-server queueing system that process jobs (or customers) which are price and delay sensitive. We assume customer processing times are independent and identically distributed according to an exponential distribution with unit mean and that the server processes work at a fixed rate of n. Customers are differentiated on their valuation, which we assume is independent and identically distributed across customers with a distribution whose cumulative distribution and density functions are denoted by F and f, respectively. We assume that f is continuously differentiable and its derivative is denoted by f'. We also assume that the customer valuation distribution has a non-decreasing hazard rate, i.e., $H(x) := \frac{f(x)}{1-F(x)}$ is non-decreasing.

Customers are homogeneous on their sensitivity to delay and we use h to denote the per unit time cost of waiting. Upon arrival at time t, a customer observes the posted price $p(t)$ and the current queue-length $Q(t)$. The customer joins the system if his randomly drawn valuation V exceeds the total expected cost of joining the system. Because the firm's processing rate is n, it follows that a customer arriving at time t joins the queue if $V > p(t) + hQ(t)/n$. Alternatively, instead of assuming the queue-length is visible, one can consider the case in which the firm quotes the current lead-time, $Q(t)/n$. We further discuss lead-time quotations in Section 6.

We assume that potential customers arrive according to a Poisson process with an arrival rate of $n\lambda$, which represents the market size. Thus, the effective arrival rate of customers who actually join the system at time t is given by

$$n\lambda\mathbb{P}\left(V > p(t) + h\frac{Q(t)}{n}\right) = n\lambda\bar{F}\left(p\left(t\right) + h\frac{Q\left(t\right)}{n}\right), \tag{1}$$

where the tail distribution function $\bar{F}\left(\cdot\right) := 1 - F\left(\cdot\right)$.

The firm's decision is to select the optimal dynamic pricing strategy, i.e., the function $p(\cdot)$, in order to maximize the long-term average revenue. For a pricing strategy p, using (1), we can write firm's rate of revenue accrual at time t as

$$p\left(t\right)n\lambda\bar{F}\left(p\left(t\right) + h\frac{Q\left(t\right)}{n}\right).$$

So, the firm's optimization problem is

$$\sup_{p}\limsup_{T\to\infty}\frac{1}{T}\int_{0}^{T}p\left(t\right)n\lambda\bar{F}\left(p\left(t\right) + h\frac{Q\left(t\right)}{n}\right)dt. \tag{2}$$

In this paper, we focus on stationary pricing policies so that $p(t)$ is in fact $p(Q(t))$, and henceforth we interpret the pricing function p as a function of the queue-length rather than the time explicitly. So, (2) can be cast in the following steady-state formulation:

$$\sup_{p\in\pi}R_n(p) := \mathbb{E}\left[p\left(Q\right)n\lambda\bar{F}\left(p\left(Q\right) + h\frac{Q}{n}\right)\right], \tag{3}$$

where the expectation is with respect to the steady-state queue-length distribution and π represents the set of stationary pricing policies that are non-anticipating. We denote the optimal objective function value by $R_n^\star$, and the optimizing price function by $p_n^\star$.

Notice that in this problem, the queue-length's steady-state distribution is intertwined with the choice of p, which makes solving (3) exactly difficult and not amenable to generating insights. Therefore, we perform an asymptotic large system analysis in which the firm's capacity n is large, and correspondingly, the market size $n\lambda$ is also large.

4 Asymptotic Analysis: Preliminaries and Static Pricing

Our asymptotic approach proceeds in the conventional manner. In Section 4.1, We first analyze the system under a fluid approximation by taking a rate-based approach. Then, in Section 4.2, we perform some preliminary analysis for refining this approximation by incorporating the inter-temporal fluctuations associated with customer arrivals and departures. Section 4.3 then characterizes the asymptotically optimal static price that also serves as a benchmark for our analysis of dynamic pricing in the next section.

4.1 Fluid analysis

In the fluid model of the system, customers are processed at a fixed rate of n as long as there is work in the system, and customers arrive deterministically at the rate of $n\lambda\bar{F}\left(p\left(Q(t)\right)+h\frac{Q(t)}{n}\right)$ at time t. Given the deterministic system behavior, it follows that the optimal pricing strategy for the fluid model is to maintain the queue-length at a zero level by ensuring that the customer arrival rate never exceeds the processing capacity. Thus, optimizing the fluid system entails solving

$$\begin{aligned} &\sup_p \; pn\lambda\bar{F}(p) \\ &s.t. \; n\lambda\bar{F}(p) \le n. \end{aligned} \tag{4}$$

We denote the unconstrained maximizer of the above program by $p^\star := \arg\max_p p\bar{F}(p)$. Because the customer valuation distribution has a non-decreasing hazard rate, $p^\star$ is unique, and further the objective function in (4) is concave. Thus, (4) is solved by the price

$$\bar{p} := \max \left\{ \bar{F}^{-1}\left(\frac{1}{\lambda}\right), p^{\star} \right\}.$$

We denote the optimal fluid objective function by $\bar{R}_n^{\star}$.

The following result formally establishes a bound on the performance of dynamic pricing, in particular, that the optimal revenue is bounded above by the optimal fluid objective. This bound is quite useful as it allows us to focus on the revenue loss due to stochasticity that we define as the gap between revenue obtained under a pricing policy and that obtained using the fluid model (that has no stochasticity). The following result also establishes that the simple pricing strategy of pricing at the fluid optimal price $\bar{p}$, that is, $p(q) = \bar{p}$ for all $q \geq 0$, leads to $o(n)$-revenue loss. That is, the revenue loss as a fraction of the system scale converges to zero.

Proposition 1.

(a) The optimal fluid objective value is an upper bound for the revenue obtained under any pricing policy, i.e., for any $n > 0$, we have

$$R_n^{\star} \leq \bar{R}_n^{*}.$$

(b) The static price $\bar{p}$ leads to $o(n)$-revenue loss due to stochasticity, i.e., we have

$$R_n(\bar{p}) = \bar{R}_n^{\star} - o(n).$$

We would like to highlight that though the naive policy of pricing at the fixed level $\bar{p}$ is optimal on the fluid scale, it ignores the queueing aspect of the problem completely. Hence, we only expect this policy to perform well if the system is not capacity constrained or is extremely large in scale. Because we anticipate capacity constraints to arise in practice, we focus on this setting by making the following assumption.

Assumption 1. *The system is capacity constrained so that the unconstrained optimizer of the fluid optimization program is not achievable, i.e., we have $\lambda\bar{F}(p^{\star}) > 1$.*

We briefly discuss what happens if this assumption does not hold in Section 6. We next analyze price refinements in the next section.

4.2 Refining the fluid approximation

In this section, we refine the fluid optimal price $\bar{p}$ by considering prices of the form

$$p_n(q) = \bar{p} + \theta_n(q) \text{ for } q \geq 0,$$

for some function θ_n. Because the static price $\bar{p}$ is optimal on the fluid-scale, we expect the refinement θ_n to take small values. Thus, we focus on price refinements that are asymptotically zero, i.e., $\theta_n(q) \to 0$ as $n \to \infty$.

We proceed with an informal argument that characterizes the revenue loss due to stochasticity, which will play a key role in our analysis. Our formal results will make this analysis precise. Denoting Q_n as the steady-state queue-length, we expect the queue-length to be small relative to n for large n. So, we approximate the expected steady-state revenue by applying the Taylor series expansion to the term $\bar{F}(p_n(Q_n) + hQ_n/n)$ around $\bar{p}$ as follows:

$$\begin{aligned}
R_n(p_n) =& \mathbb{E}\left[\Big(\bar{p} + \theta_n\left(Q_n\right)\Big) n\lambda \bar{F}\left(\bar{p} + \theta_n\left(Q_n\right) + h\frac{Q_n}{n}\right)\right] \\
\overset{(a)}{\approx}& n\lambda \mathbb{E}\left[\Big(\bar{p} + \theta_n\left(Q_n\right)\Big)\left(\bar{F}(\bar{p}) - f(\bar{p})\left(\theta_n\left(Q_n\right) + h\frac{Q_n}{n}\right) - f'(\bar{p})\frac{\left(\theta_n\left(Q_n\right) + h\frac{Q_n}{n}\right)^2}{2}\right)\right] \\
\overset{(b)}{\approx}& n\lambda\bar{p}\bar{F}(\bar{p}) - \left(-r'(\bar{p})n\lambda\mathbb{E}[\theta_n(Q_n)] + \left(f\left(\bar{p}\right) + \frac{\bar{p}f'\left(\bar{p}\right)}{2}\right)n\lambda\mathbb{E}[\theta_n(Q_n)^2] + h\bar{p}f(\bar{p})\lambda\mathbb{E}[Q_n]\right) \\
=& \bar{R}_n^\star - \Big(\alpha n\mathbb{E}[\theta_n(Q_n)] + \beta n\mathbb{E}[\theta_n(Q_n)^2] + \gamma\mathbb{E}[Q_n]\Big). \qquad (5)
\end{aligned}$$

where

$$r(p) := p\bar{F}(p), \ \alpha := -r'(\bar{p})\lambda, \ \beta := \big(f\left(\bar{p}\right) + \bar{p}f'\left(\bar{p}\right)/2\big)\,\lambda, \text{ and } \gamma := h\bar{p}f(\bar{p})\lambda.$$

In (5), we obtain (a) using the first two terms of the Taylor series expansion, and we obtain (b) by ignoring the lower order terms, specifically we ignore all terms of the form $\mathbb{E}[\theta_n(Q_n)^i Q_n^j]$ for $i + j \geq 2$ except the term $\mathbb{E}[\theta_n(Q_n)^2]$, which is the second moment of the pricing refinement. In the next section, we will prove that this second moment term (surprisingly) plays an important role in characterizing the optimal dynamic price.

The term in parenthesis in (5) represents the revenue loss relative to the fluid problem arising due to stochasticity, and so our revenue maximization problem for the refined problem can be restated as minimizing this revenue loss, i.e., we have

$$\inf_{\theta_n} \Big(\alpha n \mathbb{E}[\theta_n(Q_n)] + \beta n \mathbb{E}[\theta_n(Q_n)^2] + \gamma \mathbb{E}[Q_n] \Big). \tag{6}$$

Thus, this refinement depends on the first and second moments of the steady-state price refinement and the expected steady-state queue-length. We now proceed by solving (6). We begin by analyzing the static pricing problem to illustrate the mode of analysis and because it serves as a good benchmark policy to compare dynamic pricing with.

4.3 Static pricing benchmark

A static or fixed price is independent of the queue-length, and so we abuse notation and set the price $p_n(q) = \bar{p} + \theta_n$ for all $q \geq 0$. Denoting the optimal static price by $p^\star_{n,S}$, we have the following asymptotic characterization.

Proposition 2. *If Assumption 1 holds, then static pricing leads to $\mathcal{O}(\sqrt{n})$-revenue loss due to stochasticity, i.e., for some constant $\Pi_S > 0$, we have*

$$R_n(p^\star_{n,S}) = \bar{R}^\star_n - \Pi_S \sqrt{n} + o(\sqrt{n}).$$

Further, the asymptotically optimal static price is

$$p^\star_{n,S} = \bar{p} + \frac{1}{\sqrt{n}} \pi_S \quad \textit{for some } \pi_S \in \mathbb{R}.$$

We next provide an intuitive derivation of this result, which will be useful to illustrate the key value of dynamic pricing in this problem. We focus on positive price refinements. Noting that for large n, the term θ_n^2 is much smaller, or more precisely, of a lower order compared with θ_n, we can ignore it and this simplifies (6) to the following program:

$$\inf_{\theta_n} \alpha n\theta_n + \gamma\mathbb{E}\left[Q_n\right]. \tag{7}$$

Next, consider the expected steady-state queue-length. We can obtain an upper bound for this term by ignoring the customers' delay sensitivity, i.e., by letting the customer arrival rate depend only on the price. This upper bounding system is an $M/M/1$ queue with arrival rate $n\lambda\bar{F}(\bar{p}+\theta_n)$. So, we have

$$\mathbb{E}\left[Q_n\right] \leq \frac{n}{n - n\lambda\bar{F}\left(\bar{p}+\theta_n\right)} = \frac{1}{\lambda f(\bar{p})\theta_n} + o\left(\frac{1}{\theta_n}\right).$$

In fact, for price refinements of interest, we have $\mathbb{E}\left[Q_n\right] = \frac{\delta}{\theta_n} + o\left(\frac{1}{\theta_n}\right)$, where $\delta > 0$ is some constant. So, (6) reduces to

$$\inf_{\theta_n} \alpha n\theta_n + \frac{\gamma\delta}{\theta_n}.$$

Note that $\lambda\bar{F}(p^\star) > 1$ implies that $\alpha = -r'(\bar{p})\lambda > 0$. Thus, we obtain that $\theta_n^\star = n^{-\frac{1}{2}}\sqrt{\frac{\gamma\delta}{\alpha}}$ and the value of the objective function equals $2\sqrt{n}\sqrt{\alpha\gamma\delta}$.

5 Dynamic Pricing: Asymptotic Analysis

In this section, we analyze dynamic pricing policies. First, in Section 5.1, we provide some intuitive reasoning behind the order-improvement in revenue loss that occurs when using dynamic pricing. Then, in Section 5.2, we analyze a simple dynamic pricing policy that utilizes only two price levels and provide its asymptotic characterization. The two-price policy provides a lower-bound on the performance of optimal dynamic pricing. In Section 5.3, we compute an asymptotic upper bound on this performance. Finally, in Section 5.4, we formulate and solve a drift control problem to characterize asymptotically near-optimal dynamic prices. Section 5.5 contains a numerical study that illustrates the key insights of this section.

5.1 An intuitive argument

Clearly, the analysis of dynamic pricing is more complicated than that for static pricing. However, it turns out that beyond complexity of analysis, there is also a fundamental benefit to dynamic pricing that leads to an "order improvement" over static pricing. We proceed by first illustrating this benefit using an informal argument. Unlike static pricing, in which the fixed pricing refinement is positive, under dynamic pricing this refinement can be negative for small queue-lengths and positive for long queue-lengths. If these prices are chosen properly, one can potentially ensure that $n\mathbb{E}\theta_n \approx 0$ so that (6) reduces to

$$\inf_{\theta_n} \Big(\beta n \mathbb{E}[\theta_n(Q_n)^2] + \gamma \mathbb{E}[Q_n] \Big). \tag{8}$$

Similar to the static pricing case, one expects that the expected steady-state queue-length should be inversely proportional to the price refinement, in this case, the positive price refinement. That is, $\mathbb{E}Q_n \approx \frac{\phi}{\mathbb{E}[\theta_n(Q_n)^+]}$ for some $\phi > 0$. Then, (8) reduces to

$$\inf_{\theta_n} \Big(\beta n \mathbb{E}[\theta_n(Q_n)^2] + \gamma \frac{\phi}{\mathbb{E}[\theta_n(Q_n)^+]} \Big). \tag{9}$$

Expecting θ_n^+ and θ_n^- functions to be of similar order, we further expect that $\mathbb{E}[\theta_n(Q_n)^2] = \mathcal{O}\left(\mathbb{E}[\theta_n(Q_n)^+]^2\right)$ and hence, ignoring the constants, the trade-off in (9) becomes that between the terms

$$n\mathbb{E}[\theta_n(Q_n)^+]^2 \text{ and } \frac{1}{\mathbb{E}[\theta_n(Q_n)^+]}.$$

It follows that the optimal refinement should set

$$\mathbb{E}[\theta_n(Q_n)^+] = \mathcal{O}\left(n^{-\frac{1}{3}}\right).$$

This choice of price refinement also results in the objective (8) being $\mathcal{O}(n^{1/3})$, i.e., the total revenue under such a dynamic pricing scheme is $\bar{R}_n^\star - \mathcal{O}(n^{1/3})$ as compared with $\bar{R}_n^\star - \mathcal{O}(n^{1/2})$, which one obtains under static pricing. Thus, dynamic pricing is able to mitigate the revenue loss that occurs due to stochasticity by varying the price. In particular, it uses a larger value of price refinement compared with the static pricing and when queue-lengths are small, the refinement is negative to prevent idleness of the server and as the queue-length increases, the refinement becomes positive to lower the queue-length. This leads to lower congestion as well, with the expected steady-state queue-length being of $\mathcal{O}\left(n^{1/3}\right)$ under dynamic pricing compared with $\mathcal{O}\left(n^{1/2}\right)$ under static pricing.

We formalize this argument in the next section. The key departure of the analysis from the intuitive discussion is that we find that we cannot set $\mathbb{E}\left[\theta_n\left(Q_n\right)\right]$ to be arbitrarily small. In fact, the best we can do is to have $n\mathbb{E}\left[\theta_n\left(Q_n\right)\right] = \mathcal{O}(n^{1/3})$ so that although this term plays a role in the refined optimization problem, it does not impact its scale or order. The fact that both the first and second moments of the price refinement play important roles in the optimization problem complicates matters because this means that we need to consider both the first-order and second order terms in the Taylor series expansion that we computed in (5).

5.2 The optimal two-price policy

We begin our formal analysis of dynamic pricing by studying a class of policies that uses only two price levels and in this sense we refer to it as two-price (TP) policies. Such a policy sets one price for small queue-lengths and another price for large queue-lengths and can be characterized as follows:

$$p_{n,TP}\left(q\right) = \begin{cases} \bar{p} - \theta_n^- & \text{if } q \leq \tau_n, \\ \bar{p} + \theta_n^+ & \text{otherwise,} \end{cases} \tag{10}$$

for some non-negative constants θ_n^-, θ_n^+ and τ_n.

Defining

$$\phi = \frac{1}{2}\left(H\left(\bar{p}\right) + \frac{H'(\bar{p})}{H(\bar{p})}\right) \tag{11}$$

(recall that H denotes the hazard rate function of the valuation distribution) and denoting the optimal revenue under two-price policies by $R^\star_{n,TP}$, the following proposition provides an asymptotic characterization of the optimal two-price policy and its performance.

Proposition 3. *If Assumption 1 holds, then:*

(a) The optimal two-price policy leads to $\mathcal{O}\left(n^{1/3}\left(\log n\right)^{1/3}\right)$-revenue loss due to stochasticity, i.e., we have

$$R^\star_{n,TP} = \bar{R}^\star_n - n^{1/3}\left(\log n\right)^{1/3}\Pi_{TP} + o\left(n^{1/3}\left(\log n\right)^{1/3}\right),$$

where $\Pi_{TP} := \phi^{1/3}\left(\frac{3h}{\lambda f(\bar{p})}\right)^{2/3}$.

(b) The asymptotically optimal two-price policy is:

$$p^{\star}_{n,TP}(q) = \begin{cases} \bar{p} - \frac{\log n}{(n\log n)^{1/3}}\pi & \textit{if } q \le \frac{(n\log n)^{1/3}}{3\lambda f(\bar{p})\pi}, \\ \bar{p} + \frac{3}{(n\log n)^{1/3}}\pi & \textit{otherwise,} \end{cases}$$

where $\pi := \frac{1}{3}\left(\frac{3h}{\lambda f(\bar{p})\phi}\right)^{1/3}$.

This result formalizes the intuitive reasoning of the previous section to prove that the optimal two-price policy leads to a revenue-loss due to stochasticity of $\mathcal{O}(n^{1/3}(\log n)^{1/3})$ compared with $\mathcal{O}(n^{1/2})$ that we obtain with static pricing. Notice that compared with the intuitive reasoning, the performance of the two-price policy has an additional $(\log n)^{1/3}$ term in the revenue loss. This term arises from the inability to keep the expected price refinement arbitrarily small because of the discontinuity of the pricing function. A brief explanation for the presence of the $(\log n)^{1/3}$ term is as follows: in the proof of the result, we show that the revenue can be written as $R_n(p_{n,TP}) = \bar{R}^{\star}_n - n^{1/3}\Pi^n + o(n^{1/3})$, where defining $\hat{\pi}^i_n := \theta^i_n n^{1/3}$ for $i \in \{-,+\}$ and $\hat{\tau}_n := \tau_n n^{-1/3}$, we have

$$\Pi^n := \min_{\hat{\pi}^-_n,\hat{\pi}^+_n,\hat{\tau}_n \ge 0} \frac{n^{1/3}}{\left(\lambda\bar{F}\left(\bar{p} - \hat{\pi}^-_n/n^{1/3}\right)\right)^{n^{1/3}\hat{\tau}_n}} \frac{\alpha\hat{\pi}^-_n\hat{\pi}^+_n}{\left(\hat{\pi}^-_n + \hat{\pi}^+_n\right)} + \phi\hat{\pi}^-_n\hat{\pi}^+_n + h\left(\frac{1}{\lambda f(\bar{p})}\left(\frac{1}{\hat{\pi}^+_n} - \frac{1}{\hat{\pi}^-_n}\right) + \hat{\tau}_n\right). \tag{12}$$

A careful examination of (12) shows that Π^n must be $\mathcal{O}((\log n)^{1/3})$ and further that the optimizers $(\hat{\pi}^-_n, \hat{\pi}^+_n, \hat{\tau}_n)$ must converge in the following fashion: $\hat{\pi}^-_n(\log n)^{-2/3} \to \pi$, $\hat{\pi}^+_n(\log n)^{1/3} \to \pi/3$, and $\hat{\tau}_n(\log n)^{-1/3} \to 1/(3\lambda f(\bar{p})\pi)$, as in Proposition 3(b). The details are in the proof of the result.

We would like to comment that it is somewhat surprising that the asymptotically optimal two-price policy can be easily characterized, and that it is of such a simple form. In contrast, the asymptotically-optimal static pricing cannot be characterized explicitly. This difference arises from the asymptotic analysis of the objective of (6). For instance, the expected queue-length term in the objective equals:

$$\mathbb{E}[Q_n] = \frac{\sum_{i=0}^{\infty} i \prod_{j=0}^{i-1} \lambda\bar{F}\left(\bar{p} + \theta_n(j) + h\frac{j}{n}\right)}{\sum_{i=0}^{\infty} \prod_{j=0}^{i-1} \lambda\bar{F}\left(\bar{p} + \theta_n(j) + h\frac{j}{n}\right)}. \tag{13}$$

In the two-price policy, the price refinement is of a larger order relative to the traditional square-root order, implying that the price refinement dominates the queue-length effect and asymptotically the arrival rates take two values: low and high. This makes (13) easy to analyze. However, in static pricing, the price refinement and queue-length effect are on the same order, and hence asymptotically, we need to deal with the state dependent argument in the arrival rate, and precludes a simple explicit characterization.

The two-price policy we have identified provides us with a lower bound on the performance of dynamic pricing, and we expect allowing for additional prices will lead to even better performance. This leads us to a natural question about how well can dynamic pricing be expected to perform if we allow general pricing policies.

5.3 An upper bound on the performance of dynamic pricing

We next establish an asymptotic upper bound on optimal revenue achievable using dynamic pricing. For asymptotic analysis, one typically introduces a scale factor that scales the decision variables as the system scale increases. The analysis of the asymptotically optimal two-price policy yielded that the positive and negative price refinements were scaled differently. Extending this to a fully dynamic price would entail an infinite number of scales, one for each price possible, and would be intractable. To make our asymptotic analysis tractable, we restrict attention to a smaller, yet fairly general class of dynamic pricing policies that has two scales: one for positive price refinements, and another for negative price refinements, and we establish the bound for this class of policies.

Dynamic pricing policy class. We focus on a class of sequences of dynamic pricing policies denoted by $\mathcal{P}$ such that a sequence $\{p_n\} \in \mathcal{P}$ is of the form $p_n(q) = \bar{p} + \theta_n(q)$, where

$$\theta_n(q) = \begin{cases} s_n^- \theta\left(\frac{q}{\tau_n}\right) & \text{if } \theta\left(\frac{q}{\tau_n}\right) \leq 0, \\ s_n^+ \theta\left(\frac{q}{\tau_n}\right) & \text{otherwise,} \end{cases} \tag{14}$$

with the following properties:

(a) s_n^-, s_n^+, τ_n are positive constants with $\lim_{n\to\infty} s_n^+ = \lim_{n\to\infty} s_n^- = 0$.

(b) θ is a non-decreasing, bounded function (possibly discontinuous).

Equation (14) is motivated by the typical approach in asymptotic analysis to separate the scale from the decision variable. In this case, s_n^- and s_n^+ are the scale parameters that multiply the decision variable θ when it is negative and positive, respectively. We would like to emphasize that this is unlike typical asymptotic analysis, in which there tends to be only one such multiplying scale. Further, because we expect the queue-length to be asymptotically large, we use τ_n to scale down the queue-length in the argument of θ.

Before presenting the result, we would like to point out that the requirements that the price refinement function θ be non-decreasing and bounded are technical conditions that we require for our analysis to work. Notice that a non-decreasing θ implies that the prices we consider are non-decreasing in the queue-length. Intuitively, one does expect a firm to charge higher prices as the congestion increases, so this requirement does not seem too restrictive. Nevertheless, in the appendix, we formulate the exact (non-asymptotic) dynamic program that the firm faces and prove that the optimal price $p_n^\star$ has the property that $p_n^\star(q) + h\frac{q}{n}$, a customer's total cost of joining the system, is non-decreasing in q (see Lemma 1 in Appendix A.1). Because we expect $p_n^\star(q) = \bar{p} + \theta_n^\star(q)$ with $|\theta_n^\star(q)| \gg q/n$,[1] the optimal prices should indeed be asymptotically non-decreasing. So, restricting attention to this class of policies seems reasonable. We would also like to point out that we do not require any continuity or differentiability conditions on θ.

Result. The following result proves that under any dynamic pricing policy from the class $\mathcal{P}$, the smallest possible revenue-loss due to stochasticity is $\mathcal{O}(n^{1/3})$, and thus establishes a formal limit on the achievable performance of dynamic pricing.

Proposition 4. *If Assumption 1 holds, then for any sequence of dynamic pricing policy* $\{p_n\}_{n\geq 1} \in \mathcal{P}$*, there exists a constant* $K > 0$ *such that*

$$R_n(p_n) \leq \bar{R}_n^\star - Kn^{1/3}.$$

This result shows that the intuitive reasoning of Section 5.1 is tight with respect to the order of optimality and dynamic pricing (within a large class of policies) cannot reduce the revenue loss

[1] The relation $|\theta_n^\star(q)| \gg q/n$, which is our convention for $q/n = o\left(|\theta_n^\star(q)|\right)$, is expected because under the optimal pricing policy, we expect $|\theta_n^\star(q)| \gg \mathcal{O}(\frac{1}{\sqrt{n}})$, and further under such pricing policy, we must have $\mathcal{O}(\frac{1}{\sqrt{n}}) \gg Q_n/n$.

to a value smaller than $Kn^{1/3}$.

5.4 Computing asymptotically optimal dynamic prices

In this section, we formulate a drift control problem (DCP) in the diffusion limit to propose the asymptotically optimal pricing policy. Our analysis thus far suggests that under the optimal policy, the system queue-length should be on the $\mathcal{O}(n^{1/3})$ scale. So, in order to reach a proper limiting problem, we need to spatially scale the system, i.e., scale the state-space by $n^{1/3}$. Further, to be consistent with the newly scaled state-space, we modify the price refinement function to take the scaled queue-length as an argument, i.e., we replace $\theta_n(Q_n(t))$ by $\theta_n(Q_n(t)n^{-1/3})$.

Approximating the system dynamics. In order to write out the limiting DCP, we need to approximate the system dynamics with an appropriate diffusion process. To do so, we first write out the exact system dynamics using the following notation: we define N_a and N_s as two independent unit rate Poisson processes, I_n as the cumulative server idle time process, and we use Q_n to denote the queue-length process. Then, noting that the effective arrival rate at time t is

$$\Lambda_n(t) := n\lambda\bar{F}\left(\bar{p} + \theta_n\left(\frac{Q_n\,(t)}{n^{1/3}}\right) + h\frac{Q_n(t)}{n}\right)$$

and the amount of time the server is busy until time t is

$$T_n(t) := t - I_n(t),$$

we can write

$$Q_n\,(t) = N_a\Big(\int_0^t \Lambda_n(s)ds\Big) - N_s\Big(nT_n(t)\Big).$$

We now apply the strong approximation to the arrival and the job completion processes. In particular, we assume X_a and X_s are two independent standard Brownian motions so that $N_i(t) =$

$t + X_i(t) + o(\sqrt{t})$ for $i \in \{a, s\}$. Then, we have

$$N_a\Big(\int_0^t \Lambda_n(s)ds\Big) = \int_0^t \Lambda_n(s)ds + \sqrt{n}X_a\left(\int_0^t \frac{\Lambda_n(s)}{n}ds\right) + o\left(\sqrt{nt}\right)$$
$$N_s\left(nT_n(t)\right) = nt + \sqrt{n}X_s\left(T_n\left(t\right)\right) - nI_n\left(t\right) + o\left(\sqrt{nt}\right),$$

implying that

$$Q_n\left(t\right) = \int_0^t (\Lambda_n(s) - n)ds + \sqrt{n}X_a\left(\int_0^t \frac{\Lambda_n(s)}{n}ds\right) - \sqrt{n}X_s\left(T_n\left(t\right)\right) + nI_n\left(t\right) + o\left(\sqrt{nt}\right).$$

Define

$$\Delta_n\left(q\right) := \lambda\left(\bar{F}\left(\bar{p} + \theta_n\left(q\right) + h\frac{q}{n^{2/3}}\right) - \bar{F}\left(\bar{p}\right)\right), \text{ for all } q \geq 0,$$

so that $n\Delta_n\left(\frac{Q_n(s)}{n^{1/3}}\right) = \Lambda_n(s) - n$ is the drift of the system at time s. Then, we can write

$$Q_n\left(t\right) = n\int_0^t \Delta_n\left(\frac{Q_n(s)}{n^{1/3}}\right)ds + \sqrt{n}X_a\left(\int_0^t \frac{\Lambda_n(s)}{n}ds\right) - \sqrt{n}X_s\left(T_n\left(t\right)\right) + nI_n\left(t\right) + o\left(\sqrt{nt}\right).$$

We next consider the terms X_a and X_s. Because we expect $\theta_n \to 0$ and $Q_n/n \to 0$ as $n \to \infty$, we have

$$\int_0^t \frac{\Lambda_n(s)}{n}ds = t + o(t)$$

and hence

$$X_a\left(\int_0^t \frac{\Lambda_n(s)}{n}ds\right) = X_a\left(t\right) + o(\sqrt{t}).$$

Also, because the server should be busier when operating under this policy than when operating under the optimal static pricing policy, we should have $T_n(t) = t\left(1 - o\left(\frac{1}{\sqrt{n}}\right)\right)$, and so

$$X_s\left(T_n\left(t\right)\right) = X_s\left(t\right) + o\left(n^{-1/4}\sqrt{t}\right).$$

Putting all these components together, we obtain the following approximation of the system dynamics:

$$Q_n\left(t\right) = n\int_0^t \Delta_n\left(\frac{Q_n\left(s\right)}{n^{1/3}}\right)ds + \sqrt{2n}X\left(t\right) + nI_n\left(t\right) + o\left(\sqrt{nt}\right), \tag{15}$$

where X is another standard Brownian motion.

By further scaling the time dimension, equation (15) motivates the use of a diffusion process $Z(u)$ to approximate the queue-length process $\frac{1}{n^{1/3}}Q_n\left(\frac{u}{n^{1/3}}\right)$ for large n, where Z is the (weak) solution of the following stochastic differential equation:

$$Z\left(u\right) = \int_0^u \Delta\left(Z\left(s\right)\right)ds + \sqrt{2}B\left(u\right) + I\left(u\right), \tag{16}$$

where I is a non-decreasing process such that $\int_0^u I\left(s\right)dZ\left(u\right) = 0$, and B is another independent standard Brownian motion. Note that $Z(u)$ approximates $\frac{1}{n^{1/3}}Q_n\left(\frac{u}{n^{1/3}}\right)$ rather than $\frac{Q_n(u)}{n^{1/3}}$. This additional scaling in the time dimension is an artifact of the optimal dynamic pricing policy under which the expected steady-state queue-length is of $\mathcal{O}\left(n^{1/3}\right)$ rather than $\mathcal{O}\left(\sqrt{n}\right)$. We refer readers to Appendix A.2 for additional details on how Z is derived from (15). In (16), Δ is the drift of the diffusion process. Also, the non-decreasing process I relates well to I_n in (16) because at any time u, I can increase if and only if $Z(u) = 0$ and similarly I_n can increase if and only if $Q_n(u) = 0$.

Solving the limiting DCP. With the state-dependent drift Δ in (16) as the decision variable, we now write out the DCP to propose an asymptotically optimal pricing policy. To do so, it will be convenient to rewrite the revenue-loss due to stochasticity in terms of the drift rather than the price refinement in (5). Straightforward application of the Taylor series expansion yields

$$\bar{R}_n^\star - R_n\left(p_n\right) \approx n\psi\mathbb{E}\left[\Delta_n\left(\frac{Q_n}{n^{1/3}}\right)\right] + n\frac{\phi}{\left(\lambda f(\bar{p})\right)^2}\mathbb{E}\left[\Delta_n\left(\frac{Q_n}{n^{1/3}}\right)^2\right] + h\mathbb{E}\left[Q_n\right], \tag{17}$$

where $\psi := r'(\bar{p})/f(\bar{p}) < 0$ and ϕ is defined in (11).

Recall that the process Z approximates the time and spatially scaled queue-length and Δ approximates the drift of the queue-length process in the pre-limit. Hence, we correspondingly scale the objective in (17) to obtain the following DCP:

$$\inf_{\Delta} n^{1/3}\psi\mathbb{E}\left[\Delta\left(Z\right)\right] + \frac{\phi}{\left(\lambda f(\bar{p})\right)^2}\mathbb{E}\left[\Delta\left(Z\right)^2\right] + h\mathbb{E}\left[Z\right], \tag{18}$$

where the expectation is with respect to the steady-state distribution of Z. In solving (18), we restrict attention to increasing drift functions, i.e., $\Delta(q)$ is increasing in q. This is not restrictive, because Lemma 1 in Appendix A.1 proves that the analogous drift term in the pre-limit queueing system is also increasing.

We denote the optimizer and optimal objective value of (18) by $\hat{\Delta}_n^\star$ and $\hat{R}_n^\star$, respectively. Notice that unlike typical asymptotic control problems that involve diffusion processes, our optimization problem (18) has an "order-inconsistency" because it includes the term $n^{1/3}$ even after we have scaled the problem to the correct order. In particular, the first term in the objective, $n^{1/3}\Delta(Z)$, dominates the other two terms for each Z. However, when we solve this optimization problem, we find that the expected steady-state drift under the optimal solution is of a lower magnitude (in order sense) than its point-wise value, i.e., $\mathbb{E}[\hat{\Delta}_n^\star(Z(t))] \ll \hat{\Delta}_n^\star(Z(t))$, so that the optimal objective value of (18) remains bounded. The following result characterizes the solution to the DCP (18).

Proposition 5. *For each $n \geq 0$, there exists a constant $\kappa^\star$ and a continuously differentiable function $g^\star$ such that the optimal control that solves* (18) *is*

$$\hat{\Delta}_n^\star(q) = -\left(\psi n^{1/3} + g^\star(q)\right)\frac{(\lambda f(\bar{p}))^2}{2\phi}, \text{ for all } q \geq 0,$$

where the pair $(g^\star, \kappa^\star)$ solves

$$g^{\star\prime}(q) + hq - \left(\psi n^{1/3} + g^\star(q)\right)^2 \frac{(\lambda f(\bar{p}))^2}{4\phi} = \kappa^\star \tag{19}$$

with $g^\star(0) = 0$ and $g^{\star\prime}(q) > 0$ and $g^{\star\prime}(q) \leq C\sqrt{q}$ for all $q > 0$ for some positive constant C. Further, the optimal objective value of (18) *$\hat{R}_n^\star = \kappa^\star$.*

Notice that for any κ, (19) is an ordinary differential equation that can be solved explicitly. Combining this solution with the conditions $g^{\star\prime}(q) > 0$ and $g^{\star\prime}(q) \leq C\sqrt{q}$ for all $q > 0$, allows us to nail down the value of κ, and thus easily obtain the optimal drift $\hat{\Delta}_n^\star$.

Proposed policy. We now "un-scale" the solution to the asymptotic DCP $\Delta_n^\star(q)$ to propose a pricing policy for the actual system. We first set the drift in the pre-limit system as

$$\Delta_n(q) = n^{-1/3}\hat{\Delta}_n^\star\left(\frac{q}{n^{1/3}}\right)$$

to obtain the price refinement as

$$\theta_n(q) = -\frac{1}{\lambda f(\bar{p})}\Delta_n(q) = -n^{-1/3}\frac{1}{\lambda f(\bar{p})}\hat{\Delta}_n^\star\left(\frac{q}{n^{1/3}}\right) \text{ for } q \geq 0.$$

So, the proposed pricing policy is to post the following price to a customer who arrives when the queue-length is q:

$$\hat{p}_n^\star(q) = \bar{p} - n^{-1/3}\frac{1}{\lambda f(\bar{p})}\hat{\Delta}_n^\star\left(\frac{q}{n^{1/3}}\right), \tag{20}$$

In the next section, we numerically compare the performance of this policy with that of the asymptotically optimal static price and two-price policies, and also with that of the optimal dynamic price obtained from exact analysis.

5.5 Numerical study

We use numerical experiments to illustrate two points. First, we will show that the solution obtained using the approximate DCP in Section 5.4 well approximates the exact optimal solution obtained by solving the underlying Markov Decision Process (MDP). Second, we will compare the performance of the asymptotically optimal static and two-price schemes with the approximating DCP solution. We verify that static pricing exhibits a greater (and on a higher order) loss in revenue compared with other dynamic pricing schemes, and further that the two-price scheme has near-optimal performance.

To illustrate our numerical observations, we pick one set of parameters, in particular, we fix the customer valuation distribution as a unit mean exponential, the customer delay sensitivity as $h = 1$ and the potential demand $\lambda = 2e$ so that the "load" on the system $\lambda\bar{F}(p^\star) = 2$. We would like to point out that the results presented in the following paragraphs well represent the results obtained from other settings in our numerical experiments that include $\lambda\bar{F}(p^\star) = 1.1, 2, 5$, and

Weibull distribution with shape parameter being 2 and scale parameter being 1 and Uniform $[0,1]$ distribution.

Accuracy of approximating DCP. Figure 1a illustrates the accuracy of the DCP in approximating the actual revenue by comparing the scaled revenue-loss of the DCP objective (18) with the actual scaled revenue-loss $\left(\frac{\bar{R}_n^\star - R_n(\hat{p}_n^\star)}{n^{1/3}}\right)$ for the price function $\hat{p}_n^\star$ given in (20) that solves the DCP. We see that the objective function of our DCP indeed well approximates the exact objective. Figure 1b compares the scaled revenue-loss obtained from implementing the DCP solution $\hat{p}_n^\star$ with the optimal solution obtain from solving the MDP. From the figure, we observe that indeed the price obtained by solving the DCP has excellent performance relative to the optimal. Note that because the complexity of solving the exact MDP increases very quickly with system size, we are only able to solve it exactly for system sizes up to $n = 10^5$.

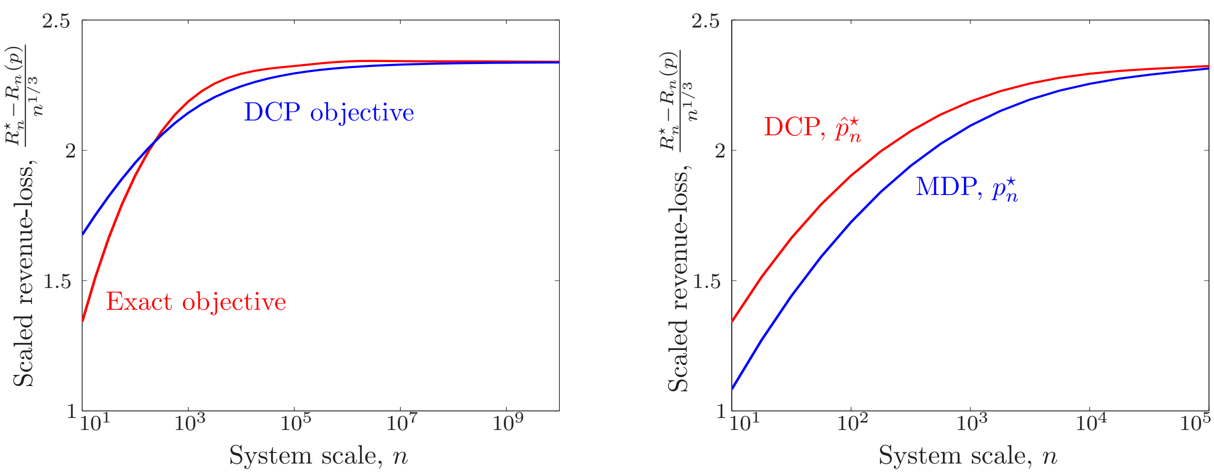

(a) Accuracy of the DCP objective in approximating the actual objective

(b) Performance of the DCP solution, $\hat{p}_n^\star$ compared with optimal MDP solution

Figure 1: Accuracy of drift control problem.

Performance of two-price policy. We next compare the performance of the asymptotically optimal two-price policy characterized in Section 5.2 with the asymptotically optimal static pricing and the solution to the DCP. Figure 2 plots the scaled revenue-loss for each of these policies for different system sizes. The figure clearly illustrates the $\mathcal{O}(\sqrt{n})$ revenue-loss of static pricing, which when scaled by $n^{1/3}$ grows without bound as n increases. We also observe that the two-price policy performs very well and has a very small gap relative to the solution of the DCP.

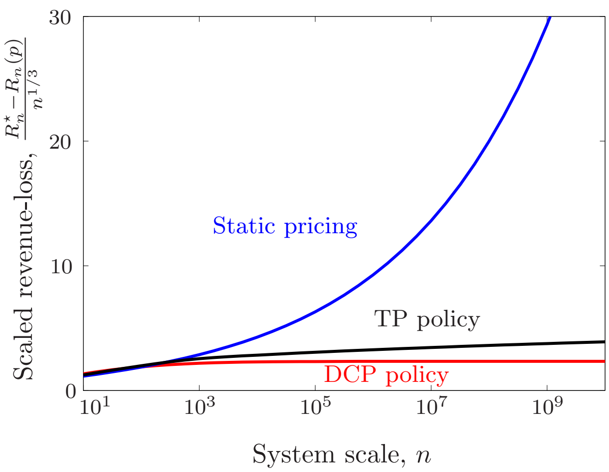

Figure 2: Performance of two-price (TP) policy relative to static pricing and the DCP-based policy.

6 Conclusion

In this paper, we study optimal dynamic pricing to maximize revenue in a queueing system when the system congestion is observable by arriving customers. We take an asymptotic approach and find that there is a fundamental benefit to dynamic pricing that results in an order improvement in the revenue-loss due to stochasticity relative to static pricing. In particular, under the optimal dynamic pricing scheme, the queue-length is maintained at $\mathcal{O}(n^{1/3})$-scale relative to the traditional $\mathcal{O}(\sqrt{n})$-scale that one expects. We formulate an approximating DCP, solving which yields near-optimal performance. Further, we propose a simple two-price policy that sets a low price when system congestion is low and a high price when the congestion is high. We prove that this policy has $\mathcal{O}((\log n)^{1/3} n^{1/3})$ revenue loss, i.e., within a logarithm term of the optimal scale. Our numerical experiments show that this policy performs very close to optimal.

Our result that the optimal operating scale is $\mathcal{O}(n^{1/3})$ relates to the recent paper Kumar and Randhawa (2010), which studies static pricing in unobservable queues. In that paper, the non-square optimality scale is obtained due to the curvature of the customer delay costs in the vicinity of the fluid operating point. Interestingly, in our paper, the customer delay sensitivity is linear, and it is the dynamic pricing that leads to the non-conventional operating scale. We would also like to mention that such non-conventional asymptotic behavior has also been observed recently in Lee and Ward (2014) in a setting with customer abandonments.

In the paper, we focus on the capacity-constrained setting which is characterized by the condition $\lambda\bar{F}(p^\star) > 1$, so that the unconstrained optimal price is not optimal for the fluid optimization problem. If this condition does not hold, the value of dynamic pricing is asymptotically quite limited. In particular, if $\lambda\bar{F}(p^\star) = 1$, then static pricing by itself generates a revenue-loss of $\mathcal{O}(n^{1/3})$. To see this, notice that we would have $\alpha = 0$ in the Taylor series expansion (5). Hence, the revenue-loss of static pricing would be on the same order as that of the corresponding optimal dynamic pricing policy. If $\lambda\bar{F}(p^\star) < 1$, then simply pricing at $p^\star$ leads to an under-loaded situation, i.e., we obtain an $M/M/1$ queueing system with utilization $\rho = \lambda\bar{F}(p^\star) < 1$. Hence, as n grows without bound, the queueing fluctuations will be extremely small, in particular of $\mathcal{O}(1)$-scale, and will effectively play no role in the optimal solution.

Finally, we would like to emphasize that our focus on observable queues is quite analogous to settings of lead-time quotations. In particular, asymptotically, our setting is identical to one in which the firm truthfully announces the current expected delay at each time instant. There are several papers in the literature that consider static lead-time quotations (that do not change with congestion) with an option of expediting orders to ensure that the quotes are always met (see, for example, Celik and Maglaras, 2008, and Plambeck and Ward, 2008). Our mode of analysis can be extended to cover this case. Incorporating expediting adds a "pushing boundary" to the drift control problem, but does not affect the nature of the solutions, and the basic insights still carry over. That is, our analysis based on Taylor series expansion remains essentially the same except an additional constraint is needed on the quoted lead-time. In fact, the additional constraint makes solving the DCP easier.

References

Afeche, P. and Ata, B. (2013), 'Bayesian dynamic pricing in queueing systems with unknown delay cost characteristics', *Manufacturing and Service Operations Management* **15**(2), 292–304.

Ata, B. and Olsen, T. (2013), 'Congestion-based leadtime quotation and pricing for revenue maximization with heterogeneous customers', *Queueing Systems* **73**(1), 35–78.

Ata, B. and Shneorson, S. (2006), 'Dynamic control of an M/M/1 service system with adjustable arrival and service rates', *Management Science* **52**(11), 1778–1791.

Besbes, O. and Maglaras, C. (2010), 'Revenue optimization for a make-to-order queue in an uncertain market environment', *Operations Research* **57**(6), 1438–1450.

Besbes, O. and Zeevi, A. (2012), 'Blind network revenue management', *Operations Research* **60**(6), 1537–1550.

Biller, S., Chan, L. M. A., Simchi-Levi, D. and Swann, J. (2005), 'Dynamic pricing and the direct-to-customer model in the automotive industry', *Electronic Commerce Research* **5**(2), 309–334.

Cachon, G. and Feldman, P. (2011), 'Pricing services subject to congestion: Charge per-use fees or sell subscriptions?', *Manufacturing and Service Operations Management* **13**(2), 244–260.

Celik, S. and Maglaras, C. (2008), 'Dynamic pricing and lead-time quotation for a multiclass make-to-order queue', *Management Science* **54**(6), 1132–1146.

Eren, S. and Maglaras, C. (2010), 'Monopoly pricing with limited demand information', *Journal of Revenue and Pricing Management* **9**, 23–48.

Farias, V. and Van Roy, B. (2010), 'Dynamic pricing with a prior on market response', *Operations Research* **58**(1), 16–29.

Harrison, J. M., Keskin, N. B. and Zeevi, A. (2012), 'Bayesian dynamic pricing policies: Learning and earning under a binary prior distribution', *Management Science* **58**(3), 570–586.

Haviv, M. and Randhawa, R. (2014), 'Pricing in queues without demand information', *Manufacturing and Service Operations Management* **16**(3), 401–411.

Kim, J. and Ward, A. (2013), 'Dynamic scheduling of a GI/GI/1+GI queue with multiple customer classes', *Queueing Systems* **75**(2), 339–284.

Kumar, S. and Randhawa, R. S. (2010), 'Exploiting market size in service systems', *Manufacturing & Service Operations Management* **12**(3), 511–526.

Lee, C. and Ward, A. (2014), 'Optimal pricing and capacity sizing when customers abandon'. Working Paper.

Low, D. (1974), 'Optimal dynamic pricing policies for an M/M/s queue', *Operations Research* **22**(3), 545–561.

Maglaras, C. (2006), 'Revenue management for a multiclass single-server queue via a fluid model analysis', *Operations Research* **54**(5), 914–932.

Maglaras, C. and Zeevi, A. (2003), 'Pricing and capacity sizing for systems with shared resources: Approximate solutions and scaling relations', *Management Science* **49**(8), 1018–1038.

Mendelson, H. and Whang, S. (1990), 'Optimal incentive-compatible priority pricing for the M/M/1 queue', *Operations Research* **38**(5), 870–883.

Naor, P. (1969), 'The regulation of queue size by levying tolls', *Econometrica* **37**(1), 15–24.

Netessine, S. (2006), 'Dynamic pricing of inventory/capacity with infrequent price changes', *European Journal of Operations Research* **174**, 553–580.

Paschalidis, I. and Tsitsiklis, J. (2000), 'Congestion-dependent pricing of network services', *IEEE/ACM Transactions on Networking* **8**(2), 171–184.

Plambeck, E. and Ward, A. (2008), 'Optimal control of a high volume assemble-to-order system with leadtime quotation and expediting', *Queueing Systems* **60**(1), 1–69.

Puterman, M. (2009), *Markov decision processes: discrete stochastic dynamic programming*, Vol. 414, John Wiley & Sons.

Yoon, S. and Lewis, M. (2004), 'Optimal pricing and admission control in a queueing system with periodically varying parameters', *Queueing Systems* **47**, 177–199.

Appendices

The appendix is organized as follows. In Appendix A, we provide some supporting materials that are referred to in the main text. In Appendix A.1, we formally prove that the solution to the exact MDP satisfies $p_n^\star(q) + h\frac{q}{n}$ is non-decreasing in q. In Appendix A.2, we present the spatial and time scaling used to derive (16) from (15). Appendix B provides the proofs of all the results in the body of the paper. Appendix C contains proofs for technical lemmas that are used in Appendix B.

A Supporting Materials

A.1 The MDP solution

We analyze the exact (non-asymptotic) optimization problem (3). Using the uniformization technique (Puterman, 2009, p. 562), we can derive the following set of equations from our model:

$$\gamma_n = \sup_{p\in\mathcal{F}} \left\{p\lambda\bar{F}\left(p\right) - \lambda\bar{F}\left(p\right)y_n\left(0\right)\right\} \tag{21}$$

$$\gamma_n = \sup_{p\in\mathcal{F}} \left\{p\lambda\bar{F}\left(p + h\frac{i}{n}\right) - \lambda\bar{F}\left(p + h\frac{i}{n}\right)y_n\left(i\right)\right\} + y_n\left(i-1\right) \text{ for } i \geq 1, \tag{22}$$

where $\mathcal{F}$ is a set of feasible prices, γ_n is a "guess" for $\frac{R_n^\star}{n}$, and y_n is an associated value function. The following lemma formalizes the "non-decreasing" property that is mentioned in Section 5.3.

Lemma 1. *If $\mathcal{F} = [0, p_{max}]$ for some $p_{max} \in (\bar{p}, \infty)$ and $f(q) > 0$ for all $q \geq 0$, there exists a unique pair of γ_n and y_n that jointly solves* (21) *and* (22)*. Furthermore, $R_n^\star = n\gamma_n$ and $\lambda\bar{F}\left(p_n^\star(q) + h\frac{q}{n}\right)$ is non-increasing in q, i.e., $p_n^\star(q) + h\frac{q}{n}$ is non-decreasing in q.*

Note that (21) and (22) are of the same form as (12) and (13) of Ata and Shneorson (2006) with a caveat that (22) has an infinite number of equations whereas (13) of Ata and Shneorson (2006) has a finite number of equations. Indeed, the proof of Lemma 1 is quite similar to that of its counterpart in Ata and Shneorson (2006) (a combination between Proposition 1 and Corollary 1 therein), and so is omitted here.

A.2 Additional steps in derivation of DCP of Section 5.4

We provide the time scaling steps that we use to derive (16) from (15). Dividing both sides of (15) by $n^{1/3}$, we obtain:

$$\frac{1}{n^{1/3}}Q_n\left(t\right) = n^{2/3}\int_0^t \Delta_n\left(\frac{Q_n\left(s\right)}{n^{1/3}}\right)ds + \sqrt{2}n^{1/6}X\left(t\right) + n^{2/3}I_n\left(t\right) + o\left(\sqrt{n^{1/3}t}\right). \tag{23}$$

Observe that $\sqrt{2}n^{1/6}X(t) \stackrel{d}{=} \sqrt{2}X(n^{1/3}t)$. We next scale the time index by $n^{1/3}$, i.e., we define a new time variable $u = n^{1/3}t$ and use this in (23) to obtain

$$\frac{1}{n^{1/3}}Q_n\left(\frac{u}{n^{1/3}}\right) = n^{2/3}\int_0^{\frac{u}{n^{1/3}}} \Delta_n\left(\frac{Q_n(s)}{n^{1/3}}\right)ds + \sqrt{2}n^{1/6}X\left(\frac{u}{n^{1/3}}\right) + n^{2/3}I_n\left(\frac{u}{n^{1/3}}\right) + o\left(\sqrt{u}\right)$$
$$\stackrel{(a)}{=} \int_0^u \Delta\left(\frac{1}{n^{1/3}}Q_n(s)\left(\frac{s}{n^{1/3}}\right)\right)ds + \sqrt{2}\hat{X}(u) + n^{2/3}I_n\left(\frac{u}{n^{1/3}}\right) + o\left(\sqrt{u}\right), \quad (24)$$

where $\hat{X}$ is another independent standard Brownian motion, (a) is obtained by the change of variables in the integration and by setting $\Delta = n^{1/3}\Delta_n$ motivated from $\Delta_n = \mathcal{O}\left(n^{-1/3}\right)$. As I in (16) increases only when $Z(u) = 0$, it is well matched with I_n in (24) and this motivates the approximation of $\frac{1}{n^{1/3}}Q_n\left(\frac{u}{n^{1/3}}\right)$ in (24) by $Z(u)$ in (16).

B Proofs of Main Results

Notation. In the following proofs, we will use the following convention: for any two sequences of real numbers $\{a_n\}_{n\geq 1}$ and $\{b_n\}_{n\geq 1}$, we use the notation $a_n \sim b_n$, $a_n \gtrsim b_n$, $a_n \gg b_n$, and $a_n = \Theta(b_n)$ to represent

$$\lim_{n\to\infty}\frac{a_n}{b_n} = 1,\ \liminf_{n\to\infty}\frac{a_n}{b_n} \geq 1,\ \lim_{n\to\infty}\frac{b_n}{a_n} = 0,\ \text{and } 0 < \liminf_{n\to\infty}\frac{|a_n|}{|b_n|} \leq \limsup_{n\to\infty}\frac{|a_n|}{|b_n|} < \infty, \text{ respectively.}$$

Also, we use $x \downarrow a$ for $a \in \mathbb{R}$ to denote "as x approaches to a from the right". Finally, it is useful to set $\mathcal{P}_{TP}$ as a class of sequences of two-price policies defined in Section 5.2.

Proof of Proposition 1:

Since part (b) of the proposition is subsumed by the results in Proposition 2, we only prove part (a) here. For $p \geq 0$, let $r(p) = p\lambda\bar{F}(p)$. Let $\bar{F}^{-1}$ be the inverse of $\bar{F}$ so that $p = \bar{F}^{-1}(q)$ for each $q \in [0,1]$. Let $\tilde{r}(q) = \lambda q\bar{F}^{-1}(q)$. Because the distribution has an increasing hazard rate function, it is straightforward to check $\tilde{r}''(q) < 0$, so that $\tilde{r}$ is concave in $q \in [0,1]$ and it is maximized at

$q = \bar{F}(p^\star)$. Observe that for any pricing policy p_n, we have

$$\begin{aligned} R_n(p_n) &= n\lambda \mathbb{E}\left[p_n(Q_n)\bar{F}\left(p_n(Q_n) + h\frac{Q_n}{n}\right)\right] \\ &\leq n\lambda \mathbb{E}\left[\left(p_n(Q_n) + h\frac{Q_n}{n}\right)\bar{F}\left(p_n(Q_n) + h\frac{Q_n}{n}\right)\right] \\ &\overset{(a)}{\leq} n\tilde{r}\left(\mathbb{E}\left[\bar{F}\left(p_n(Q_n) + h\frac{Q_n}{n}\right)\right]\right) \end{aligned} \tag{25}$$

where (a) follows by Jensen's inequality. We first consider the case when $\bar{F}^{-1}\left(\frac{1}{\lambda}\right) \leq p^\star$ so that $\bar{R}_n^\star = n\lambda p^\star \bar{F}(p^\star)$. In this case we have

$$R_n(p_n) \leq n\tilde{r}\left(\mathbb{E}\left[\bar{F}\left(p_n(Q_n) + h\frac{Q_n}{n}\right)\right]\right) \leq \bar{R}_n^\star$$

as desired because $\tilde{r}(q) \leq \lambda p^\star \bar{F}(p^\star)$ for any $q \in [0,1]$. On the other hand, suppose $p^\star < \bar{F}^{-1}\left(\frac{1}{\lambda}\right)$, so that $\bar{R}_n^\star = \bar{p}n = n\tilde{r}\left(\frac{1}{\lambda}\right)$ (recall that $\lambda > 1$ because $\lambda\bar{F}(\bar{p}) = 1$). Because $\tilde{r}(q)$ is increasing in $q \in \left[0, \bar{F}(p^\star)\right]$ and $\frac{1}{\lambda} < \bar{F}(p^\star)$, we have $\tilde{r}(q) \leq \tilde{r}\left(\frac{1}{\lambda}\right)$ for $q \in \left[0, \frac{1}{\lambda}\right]$. Observe that $n\lambda\mathbb{E}\left[\bar{F}\left(p_n(Q_n) + h\frac{Q_n}{n}\right)\right]$ is the expected steady-state arrival rate and therefore it cannot exceed the service capacity n, i.e,

$$\mathbb{E}\left[\bar{F}\left(p_n(Q_n) + h\frac{Q_n}{n}\right)\right] \leq \frac{1}{\lambda}. \tag{26}$$

Putting (25) and (26) together, we have $R_n(p_n) \leq n\tilde{r}\left(\frac{1}{\lambda}\right) = \bar{R}_n^\star$. $\square$

Proof of Proposition 2:

First, we show that for any static price $p_{n,S}(q) = \bar{p} + \theta_{n,S}$ for $q \geq 0$, where $\theta_{n,S}$ is some constant such that $\theta_{n,S} \ll 1$,

$$\liminf_{n\to\infty} \frac{\bar{R}_n^\star - R_n(p_{n,S})}{\sqrt{n}} > 0. \tag{27}$$

We next show that there exists a sequence $\{p_{n,S}\}_{n\geq 1}$ under which

$$\limsup_{n\to\infty} \frac{\bar{R}_n^\star - R_n(p_{n,S})}{\sqrt{n}} < \infty. \tag{28}$$

Then, the first part of the proposition is established by combining (27) with (28). The second part is established using a portion of the argument used to establish (27) because for any $\theta_{n,S} \gg \frac{1}{\sqrt{n}}$ we prove

$$\liminf_{n\to\infty} \frac{\bar{R}_n^\star - R_n\left(p_{n,S}\right)}{\sqrt{n}} = \infty. \tag{29}$$

The following lemma is useful in establishing (27) - (29) (and many results in the sequel).

Lemma 2. *For any sequence of prices $\{p_n\}$ in $\mathcal{P}$ or $\mathcal{P}_{TP}$, we have*

$$\begin{aligned}
&\bar{R}_n^\star - R_n\left(p_n\right) \gtrsim \phi n \mathbb{E}\left[\theta_n\left(Q_n\right)^2\right] + h\mathbb{E}\left[Q_n\right],\\
&\bar{R}_n^\star - R_n\left(p_n\right) \sim \alpha n \mathbb{E}\left[\theta_n\left(Q_n\right)\right] + \beta n \mathbb{E}\left[\theta_n\left(Q_n\right)^2\right] + \gamma\mathbb{E}\left[Q_n\right].
\end{aligned}$$

All lemmas are proved in Appendix C. To show (27), first, consider the case $\frac{1}{\sqrt{n}} \ll \theta_{n,S}$ and $\theta_{n,S} > 0$. By Lemma 2, we have

$$\bar{R}_n^\star - R_n\left(p_{n,S}\right) \sim \alpha n\theta_{n,S} + \beta n\theta_{n,S}^2 + \gamma\mathbb{E}\left[Q_n\right] \overset{(a)}{\sim} \alpha n\theta_{n,S} + \gamma\mathbb{E}\left[Q_n\right] \geq \alpha n\theta_{n,S}, \tag{30}$$

where (a) follows because $\theta_{n,S} \ll 1$. Since $\frac{1}{\sqrt{n}} \ll \theta_{n,S}$ and $\theta_{n,S} > 0$, (30) implies (27) by (29).

Next, consider the case $\frac{1}{\sqrt{n}} \ll \theta_{n,S}$ and $\theta_{n,S} < 0$. Let $\hat{t}_n := \frac{n\left|\theta_{n,S}\right|}{h}$ so that $\theta_{n,S} + h\frac{\hat{t}_n}{n} = 0$. Consider an $M/M/1/\hat{t}_n$ system in which the static price is $\bar{p}$ and customers are only price-sensitive. Let $Q_{n,lb}$ be the steady-state queue-length of this system. It is obvious that $Q_{n,lb} \leq Q_n$ a.s. and $\mathbb{E}\left[Q_{n,lb}\right] \leq \mathbb{E}\left[Q_n\right]$. Also by the expected steady-state queue-length formula for an $M/M/1$ system with a finite buffer, we can easily obtain that $\sqrt{n} \ll \mathbb{E}\left[Q_{n,lb}\right]$, so that (29) is established by Lemma 2.

For the cases of $\theta_{n,S} = \mathcal{O}\left(\frac{1}{\sqrt{n}}\right)$, we can prove that $\sqrt{n} = \mathcal{O}\left(\mathbb{E}\left[Q_n\right]\right)$, so that (27) holds, by considering an $M/M/1/\sqrt{n}$ system in which the static price is $\bar{p}+\max\{0, \theta_{n,S}\}+h\frac{1}{\sqrt{n}}$ and customers are only price sensitive.

Finally, to show that there exists a sequence of static pricing policies under which (28) is achieved, set $\theta_{n,S} = 0$. Then, by Lemma 2, we have $\bar{R}_n^\star - R_n\left(p_{n,S}\right) \sim \gamma\mathbb{E}\left[Q_n\right]$. So, it suffices to show that $\mathbb{E}\left[Q_n\right] = \mathcal{O}\left(\sqrt{n}\right)$. This can be established by considering an $M/M/1$ queue with a two-price policy of $\bar{p} + \theta_{n,TP}$, where $\theta_{n,TP}\left(q\right)$ is 0 if $q \leq \sqrt{n}$ and $h\frac{1}{\sqrt{n}}$ otherwise and customers are only price-sensitive. $\square$

Proof of Proposition 3:

Define $\mathcal{P}^{\star}_{TP}$ to be a class of sequences of two-price policies defined in (10) such that

$$\theta_n^- > 0,\ \theta_n^+ > 0,\ \frac{1}{\sqrt{n}} \ll \min\left\{\theta_n^-, \theta_n^+\right\},\ \text{and},\ \tau_n \ll \sqrt{n}. \tag{31}$$

We will use the following three lemmas to prove the proposition.

Lemma 3. *For any* $\{p_{n,TP}\}_{n\geq 1} \in \mathcal{P}_{TP} \backslash \mathcal{P}^{\star}_{TP}$, *we have*

$$\liminf_{n\to\infty} \frac{\bar{R}^{\star}_n - R_n\left(p_{n,TP}\right)}{\left(n\log n\right)^{1/3}} = \infty.$$

Lemma 4. *For any* $\{p_{n,TP}\}_{n\geq 1} \in \mathcal{P}^{\star}_{TP}$, *we have*

$$\bar{R}^{\star}_n - R_n\left(p_{n,TP}\right) \sim L_n\left(\theta_n^-, \theta_n^+, \tau_n\right),$$

where

$$\begin{aligned} L_n\left(\theta^-, \theta^+, \tau\right) := {} & \alpha n \frac{-\theta^- D_1\left(\theta^-, \tau\right) + \theta^+ D_2\left(\theta^+\right)}{D_1\left(\theta^-, \tau\right) + D_2\left(\theta^+\right)} + \beta n \frac{\left(\theta^-\right)^2 D_1\left(\theta^-, \tau\right) + \left(\theta^+\right)^2 D_2\left(\theta^+\right)}{D_1\left(\theta^-, \tau\right) + D_2\left(\theta^+\right)} \\ & + h\left(\frac{\left(\lambda\bar{F}\left(\bar{p}-\theta^-\right)\right)^{-\tau}}{\left(D_1\left(\theta^-, \tau\right) + D_2\left(\theta^+\right)\right)\left(\lambda\bar{F}\left(\bar{p}-\theta^-\right)-1\right)^2} + \tau - \frac{1}{\lambda\bar{F}\left(\bar{p}-\theta^-\right)-1} + D_2\left(\theta^+\right)\right) \end{aligned} \tag{32}$$

for

$$D_1\left(\theta, \tau\right) := \frac{1-\left(\lambda\bar{F}\left(\bar{p}-\theta\right)\right)^{-\tau-1}}{\lambda\bar{F}\left(\bar{p}-\theta\right)-1}\ \text{ and }\ D_2\left(\theta\right) := \frac{1}{1-\lambda\bar{F}\left(\bar{p}+\theta\right)}.$$

Lemma 5. *For any* $\{p_{n,TP}\}_{n\geq 1} \in \mathcal{P}^{\star}_{TP}$, *we have*

$$\liminf_{n\to\infty} \frac{L\left(\theta_n^-, \theta_n^+, \tau_n\right)}{\left(n\log n\right)^{1/3}} \geq \Pi_{TP}.$$

Furthermore, if we let

$$\theta_n^{\star,-} := \frac{\log n}{(n \log n)^{1/3}}\pi, \theta_n^{\star,+} := \frac{3}{(n \log n)^{1/3}}\pi, \tau_n^{\star} := \frac{(n \log n)^{1/3}}{3\lambda f(\bar{p})\pi},$$

then

$$\limsup_{n\to\infty} \frac{L\left(\theta_n^{\star,-}, \theta_n^{\star,+}, \tau_n^{\star}\right)}{(n \log n)^{1/3}} = \Pi_{TP}.$$

Given Lemma 3-5, we now prove the proposition. Lemma 3 proves that the sequence of asymptotically optimal two-price policies lies in $\mathcal{P}_{TP}^{\star}$. Lemma 4 provides a tractable characterization of the revenue gap under a sequence in $\mathcal{P}_{TP}^{\star}$. This characterization is used in Lemma 5 to show that the sequence defined in part (b) of the proposition is asymptotically optimal and its performance is as given in part (a) of the proposition. □

Proof of Proposition 4:

We define $x_\theta := \inf\{q \geq 0 : \theta(q) \geq 0\}$. So, if $x_\theta \in (0, \infty)$, then we can think of $x_\theta \tau_n$ as the "switch-point," with $\theta_n(q) \leq 0$ for $q \leq x_\theta \tau_n$ and $\theta_n(q) \geq 0$ otherwise. The following lemma will be useful in the proof of the proposition.

Lemma 6. *For any sequence* $\{p_n\}_{n\geq 1} \in \mathcal{P}$ *such that* $\lim_{q\downarrow 0} \theta(q) < 0$*, we have* $\min\{\sqrt{n}, x_\theta \tau_n\} = \mathcal{O}(\mathbb{E}[Q_n])$*.*

Notice that if $x_\theta = \infty$, then the proposition is proved by applying Lemma 2 and Lemma 6. So, in what follows, we only consider $x_\theta \in [0, \infty)$ and formally consider two cases: Case 1, in which $x_\theta > 0$ and Case 2, in which $x_\theta = 0$. Furthermore, when $x_\theta > 0$, we will fix $x_\theta = 1$. This is only to simplify notation and is without loss of generality.

Case 1: $\lim_{q\downarrow 0} \theta(q) < 0$**, i.e.,** $x_\theta > 0$**.** Suppose that $\theta(q) \leq 0$ for $q \geq 0$. Then, by considering an $M/M/1/\sqrt{n}$ system with the arrival rate of $n\lambda \bar{F}\left(\bar{p} + h\frac{1}{\sqrt{n}}\right)$ and the service rate of n, we can show that $\sqrt{n} = \mathcal{O}(\mathbb{E}[Q_n])$. Therefore, using Lemma 2, the result follows. To proceed, consider the following three cases: $s_n^+ \tau_n \ll 1$, $s_n^+ \tau_n = \Theta(1)$, and $s_n^+ \tau_n \gg 1$.

Case 1(a): $s_n^+\tau_n = \Theta(1)$**.** We already argued that there must exist $\epsilon > 0$ such that $\theta(1+\epsilon) > 0$. Observe that

$$\mathbb{E}\left[\theta_n\left(Q_n\right)^2\right] \geq \left(s_n^+\theta\left(1+\epsilon\right)\right)^2 \mathbb{P}\left(Q_n \geq \left(1+\epsilon\right)\tau_n\right)$$
$$\mathbb{E}\left[Q_n\right] \geq \left(1+\epsilon\right)\tau_n \mathbb{P}\left(Q_n \geq \left(1+\epsilon\right)\tau_n\right).$$

We will next prove

$$\liminf_{n\to\infty} \mathbb{P}\left(Q_n \geq \left(1+\epsilon\right)\tau_n\right) > 0. \tag{33}$$

The proof of the result for this case will then follow by applying Lemma 2 because

$$n\left(s_n^+\right)^2 + \tau_n = \mathcal{O}\left(\phi n \mathbb{E}\left[\theta_n\left(Q_n\right)^2\right] + h\mathbb{E}\left[Q_n\right]\right) \text{ and } n^{1/3} = \mathcal{O}\left(n\left(s_n^+\right)^2 + \tau_n\right),$$

when $s_n^+\tau_n = \Theta\left(1\right)$. To show (33), fix $N > 1+\epsilon$ and let $Q_{n,lb}$ be the steady-state queue-length of an $M/M/1/N\tau_n$ system with the arrival rate of $n\lambda\bar{F}\left(\bar{p} + \theta_n\left(N\tau_n\right) + h\frac{N\tau_n}{n}\right)$ and the service rate of n. Then, it is straightforward to check

$$\mathbb{P}\left(Q_n \geq \left(1+\epsilon\right)\tau_n\right) \geq \mathbb{P}\left(Q_{n,lb} \geq \left(1+\epsilon\right)\tau_n\right) \text{ and } \liminf_{n\to\infty} \mathbb{P}\left(Q_{n,lb} \geq \left(1+\epsilon\right)\tau_n\right) > 0,$$

which establishes (33), and consequently the result for this case.

Case 1(b): $s_n^+\tau_n \gg 1$**.** Suppose $\lim_{q\downarrow 1}\theta(q) > 0$. Observe that

$$\mathbb{E}\left[\theta_n\left(Q_n\right)^2\right] \geq \left(s_n^+\theta\left(1+\frac{1}{\tau_n}\right)\right)^2 \mathbb{P}\left(Q_n \geq \tau_n + 1\right),$$

Suppose we establish that

$$\frac{1}{s_n^+\tau_n} = \mathcal{O}\left(\mathbb{P}\left(Q_n \geq \tau_n + 1\right)\right). \tag{34}$$

Then, because $\theta\left(1+\frac{1}{\tau_n}\right) \geq \lim_{x\downarrow 1}\theta\left(x\right) > 0$ and $\frac{1}{\tau_n} = o\left(s_n^+\right)$, using (34), we have

$$\frac{1}{\left(\tau_n\right)^2} = \mathcal{O}\left(\mathbb{E}\left[\theta_n\left(Q_n\right)^2\right]\right). \tag{35}$$

By Lemma 2, Lemma 6, and (35), the result is then proved for this case. So, we focus on proving (34). To do so, we fix $N > 1$ and let $Q_{n,lb}$ be the steady-state queue-length of an $M/M/1/N\tau_n$ system with the service rate of n and the arrival rate of $n \exp\left(-\lambda f\left(\bar{p}\right) ch\frac{\tau_n}{n}\right)$, if $q \leq \tau_n$,, and $n \exp\left(-\lambda f\left(\bar{p}\right) c\left(\theta_n\left(N\tau_n\right) + h\frac{N\tau_n}{n}\right)\right)$, otherwise. for some positive constant $c > 0$ that satisfies

$$\exp\left(-\lambda f\left(\bar{p}\right) ch\frac{\tau_n}{n}\right) \leq \lambda\bar{F}\left(\bar{p}\right)\bar{F}\left(\bar{p} + h\frac{\tau_n}{n}\right), \text{ and}$$
$$\exp\left(-\lambda f\left(\bar{p}\right) c\left(\theta_n\left(N\tau_n\right) + h\frac{N\tau_n}{n}\right)\right) \leq \lambda\bar{F}\left(\bar{p} + \theta_n\left(N\tau_n\right) + h\frac{N\tau_n}{n}\right)$$

for sufficiently large n. Then, it is straightforward to check that

$$\mathbb{P}\left(Q_n \geq \tau_n + 1\right) \geq \mathbb{P}\left(Q_{n,lb} \geq \tau_n + 1\right) \text{ and } \frac{1}{s_n^+\tau_n} = \mathcal{O}\left(\mathbb{P}\left(Q_{n,lb} \geq \tau_n + 1\right)\right),$$

establishing (34).

We now consider the case when $\lim_{q\downarrow 1}\theta\left(q\right) = 0$. If there exists $\zeta > 0$ such that $\theta\left(1+\zeta\right) = 0$ but $\lim_{q\downarrow 1+\zeta}\theta\left(q\right) > 0$, then we can repeat the same arguments in the previous case to complete the proof of the proposition. Therefore, we only need to consider the case $\theta\left(q\right) > 0$ for $q > q_0$ where $q_0 := \sup\left\{q \geq 1 : \theta\left(q\right) = 0\right\}$ while $\lim_{q\downarrow q_0}\theta\left(q\right) = 0$. Observe that we only need to consider the case when $q_0 < \infty$ because otherwise, we already know that $\sqrt{n} = \mathcal{O}\left(\mathbb{E}\left[Q_n\right]\right)$. Let $\nu \geq 0$ be some constant such that $q_0 = 1 + \nu$. To proceed, define $\nabla^{(1)}$ and $\nabla^{(k)}$ for $k = 2, 3, \ldots$, as follows:

$$\nabla^{(1)}\left(\theta\right)\left(q\right) = \lim_{x\downarrow q}\frac{\theta\left(x\right) - \theta\left(q\right)}{x - q},\ \nabla^{(k)}\left(\theta\right)\left(q\right) = \lim_{x\downarrow q}\frac{\nabla^{k-1}\left(\theta\right)\left(x\right) - \nabla^{k-1}\left(\theta\right)\left(q\right)}{x - q}.$$

Then, we must have $k_\theta < \infty$ where $k_\theta := \inf\left\{k \geq 1 : \nabla^{(k)}\left(\theta\right)\left(1+\nu\right) > 0\right\}$.

Let $\hat{\tau}_{n,1} := \tau_n\left(1 + \nu + \frac{1}{\left(s_n^+\tau_n\right)^{1/k_\theta}}\right)$, By the definition of k_θ and $\nabla^{(k_\theta)}$, we have

$$\theta_n\left(\hat{\tau}_{n,1}\right) = \Theta\left(\frac{1}{\tau_n}\right) = \Theta\left(\frac{1}{\hat{\tau}_{n,1}}\right) \tag{36}$$

where the last inequality follows because $s_n^+\tau_n \gg 1$. Let $d_n = \left(\frac{\tau_n^{k_\theta}}{s_n^+}\right)^{\frac{1}{k_\theta+1}}$. Then, because $s_n^+\tau_n \gg 1$, we have

$$d_n \ll \tau_n \text{ and } \frac{1}{s_n} \ll d_n. \tag{37}$$

Define $\hat{\tau}_{n,2} := \tau_n \left(1 + \nu + \frac{d_n}{\tau_n}\right)$. Then, by (37) and by the definition of k_θ and $\nabla^{(k_\theta)}$, we have $\hat{\tau}_{n,1} < \hat{\tau}_{n,2}$ for large n and

$$\hat{\tau}_{n,2} - \hat{\tau}_{n,1} = \Theta\left(d_n\right) \text{ and } \theta_n\left(\hat{\tau}_{n,2}\right) = \Theta\left(\frac{1}{d_n}\right). \tag{38}$$

Define $\hat{\tau}_{n,3} := \tau_n \left(1 + \nu + \frac{2d_n}{\tau_n}\right)$. Then, we also have $\hat{\tau}_{n,2} < \hat{\tau}_{n,3}$ for large n and

$$\hat{\tau}_{n,3} - \hat{\tau}_{n,2} = \Theta\left(d_n\right) \text{ and } \theta_n\left(\hat{\tau}_{n,3}\right) = \Theta\left(\frac{1}{d_n}\right) \tag{39}$$

using the same reasoning as in (38). Let $Q_{n,lb}$ be the steady-state queue-length of an $M/M/1/\hat{\tau}_{n,3}$ system for which the service rate is n and the arrival rate equals $n \exp\left(-\lambda f\left(\bar{p}\right) c \left(\theta_n\left(\hat{\tau}_{n,i}\right) + h\frac{\hat{\tau}_{n,i}}{n}\right)\right)$, when there are $q \in \left(\hat{\tau}_{n,i-1}, \hat{\tau}_{n,i}\right]$ customers are in the system, for $i = 1, 2, 3$ and $\hat{\tau}_{n,0} = -1$ where $c > 0$ is some constant that satisfies

$$\exp\left(-\lambda f\left(\bar{p}\right) c \left(\theta_n\left(\hat{\tau}_{n,i}\right) + h\frac{\hat{\tau}_{n,i}}{n}\right)\right) \leq \lambda \bar{F}\left(\bar{p} + \theta_n\left(\hat{\tau}_{n,i}\right) + h\frac{\hat{\tau}_{n,i}}{n}\right)$$

for $i = 1, 2, 3$ and for large n. Then, it is straightforward to check that $\mathbb{P}\left(Q_n \geq \hat{\tau}_{n,1}\right) \geq \mathbb{P}\left(Q_{n,lb} \geq \hat{\tau}_{n,1}\right)$ and by (36), (38), and (39), we can derive that

$$\mathbb{P}\left(\hat{Q}_{n,2} \geq \hat{\tau}_{n,1}\right) = \Theta\left(\frac{d_n}{\tau_n}\right). \tag{40}$$

Therefore, by (40), we have

$$\mathbb{E}\left[\theta_n\left(Q_n\right)^2\right] \geq \theta_n\left(\hat{\tau}_{n,1}\right)^2 \mathbb{P}\left(Q_{n,lb} \geq \hat{\tau}_{n,1}\right) = \Theta\left(\frac{1}{\tau_n d_n}\right).$$

Since $\frac{1}{\tau_n} = o\left(\frac{1}{d_n}\right)$, we have $\frac{1}{\tau_n^2} = \mathcal{O}\left(\mathbb{E}\left[\theta_n\left(Q_n\right)^2\right]\right)$. So, applying Lemma 2, the result is established for this case.

Case 1(c): $s_n^+ \tau_n \ll 1$**.** In this case, if $s_n^+ = \mathcal{O}\left(\frac{1}{\sqrt{n}}\right)$, we can prove that $\sqrt{n} = \mathcal{O}\left(\mathbb{E}\left[Q_n\right]\right)$, which establishes the result by applying Lemma 2. Let us focus on the case $\frac{1}{\sqrt{n}} \ll s_n^+$. Observe that

$$\mathbb{E}\left[\theta_n \left(Q_n\right)^2\right] \geq \theta_n \left(\frac{1}{s_n^+}\right)^2 \mathbb{P}\left(Q_n \geq \frac{1}{s_n^+}\right) \text{ and } \mathbb{E}\left[Q_n\right] \geq \frac{1}{s_n^+}\mathbb{P}\left(Q_n \geq \frac{1}{s_n^+}\right).$$

Further, there exists $\epsilon > 0$ such that $\theta\left(1+\epsilon\right) > 0$ and for sufficiently large n, $\frac{1}{s_n^+ \tau_n} \geq 1+\varepsilon$, implying $\theta_n \left(\frac{1}{s_n^+}\right)^2 \geq \left(s_n^+\right)^2 \theta\left(1+\varepsilon\right)$. We now establish $\liminf_{n\to\infty} \mathbb{P}\left(Q_n \geq \frac{1}{s_n^+}\right) > 0$ because then the proof is complete by Lemma 2. For this, fix $N > 2$ and let $Q_{n,lb}$ be the steady-state queue-length of an $M/M/1/\frac{N}{s_n^+}$ queue with the service rate of n and the arrival rate of $n \exp\left(-\lambda f\left(\bar{p}\right) c s_n^+\right)$ for some constant $c > 0$ that satisfies

$$\exp\left(-\lambda f\left(\bar{p}\right) c s_n^+\right) < \lambda \bar{F}\left(\bar{p} + \theta_n \left(\frac{1}{s_n^+}\right) + h\frac{1}{n s_n^+}\right).$$

Such a c must exist because $\theta_n\left(\frac{1}{s_n^+}\right) + h\frac{1}{n s_n^+} = \Theta\left(s_n^+\right)$ and θ is bounded. It is straightforward to check that

$$\mathbb{P}\left(Q_n \geq \frac{1}{s_n^+}\right) \geq \mathbb{P}\left(Q_{n,lb} \geq \frac{1}{s_n^+}\right) \text{ and } \liminf_{n\to\infty} \mathbb{P}\left(Q_{n,lb} \geq \frac{1}{s_n^+}\right) > 0.$$

Hence, the result holds for this case.

Case 2: $\lim_{q\downarrow 0} \theta(q) \geq 0$**, i.e.,** $x_\theta = 0$**.** Recall that in the case when $\lim_{q\downarrow 0} \theta(q) < 0$, our analysis was focused on $\theta(q)$ for $q \geq 1$. In the current case, the same analysis can be applied with 0 replacing the role of 1 in the previous case. $\square$

Proof of Proposition 5:

Existence of the solution of the HJB equation. We first prove the existence of a pair $(g^\star, \kappa^\star)$ that solves (19). Similar to Kim and Ward (2013), we consider the following family of first order ODEs parameterized by $\kappa \geq 0$:

$$g_\kappa'\left(q\right) + hq - \left(\psi n^{1/3} + g_\kappa\left(q\right)\right)^2 \frac{\left(\lambda f\left(\bar{p}\right)\right)^2}{4\phi} = \kappa, \text{ with } g_\kappa(0) = 0. \tag{41}$$

With the same argument as in the proof of Lemma 4.1 of Kim and Ward (2013), we know that for each κ, there exist a unique g_κ that solves (41). Within this family, we will show that there exists a unique $\kappa^\star$ such that $(g_{\kappa^\star}, \kappa^\star)$ solves (19). To that end, for each κ, let $w_\kappa := -n^{1/3} + g_\kappa$ and define

$$q_{\infty,\kappa} := \inf \left\{ q \geq 0 : \lim_{x \uparrow q} w_\kappa(x) = \infty \right\}. \tag{42}$$

We also define sets $\mathcal{L}$ and $\mathcal{U}$ that bisect non-negative real numbers:

$$\begin{aligned} &\mathcal{L} := \{\kappa \geq 0 : \mathcal{S}_\kappa \neq \emptyset\} \text{ and } \mathcal{U} := \{\kappa \geq 0 : \mathcal{S}_\kappa = \emptyset\}, \text{ where} \\ &\mathcal{S}_\kappa := \left\{ q \in [0, q_{\infty,\kappa}] : w'_\kappa(q) \leq 0 \right\}. \end{aligned}$$

Lemma 7. *For any $\kappa_1 < \kappa_2$, we have $w_{\kappa_1}(q) < w_{\kappa_2}(q)$ for $q \in [0, \min\{q_{\infty,\kappa_1}, q_{\infty,\kappa_2}\}]$.*

Lemma 8. *Both $\mathcal{L}$ and $\mathcal{U}$ are non-empty.*

Lemma 9. *If $\kappa \in \mathcal{L}$, then $w_\kappa(q)$ is strictly quasi-concave in q and $\lim_{q\to\infty} w_\kappa(q) = -\infty$.*

Lemma 10. *$w_\kappa(q)$ is jointly continuous in $\kappa \geq 0$ and $q \geq 0$.*

Lemma 11. *Let $\kappa^\star := \sup\ \mathcal{L}$. Then, $\kappa^\star < \infty$ and $\kappa^\star \in \mathcal{U}$, i.e., $w'_{\kappa^\star}(q) > 0$ for $q \geq 0$.*

Lemma 12. *For any $\kappa \in \mathcal{L}$, $w_\kappa(q) \leq \sqrt{\kappa^\star + n^{2/3} + q}$.*

To complete the proof, suppose now $w_{\kappa^\star}$ does not satisfy the growth condition so that

$$\liminf_{q\to\infty} \frac{w_{\kappa^\star}(q)}{\sqrt{q}} = \infty.$$

Then, there must exist $\hat{q}$ such that

$$w_{\kappa^\star}(\hat{q}) \geq \sqrt{\kappa^\star + n^{2/3} + \hat{q}} + \epsilon,$$

which contradicts Lemma 10 because for any $\kappa < \kappa^\star$, we have

$$w_\kappa(\hat{q}) \leq \sqrt{\kappa^\star + n^{2/3} + \hat{q}}$$

by Lemma 12. Thus, the existence result follows by setting $g^\star = w_{\kappa^\star} + n^{1/3}$.

Optimality of $\hat{\Delta}_n^\star$. It remains to show that for any $\Delta(z)$ that is decreasing in z, we have

$$\liminf_{t\to\infty} \frac{1}{t}\mathbb{E}\left[\int_0^t \left\{n^{1/3}\psi\Delta(Z(s)) + \frac{\phi}{(\lambda f(\bar{p}))^2}\Delta(Z(s))^2 + hZ(s)\right\} ds\right] \geq \kappa^\star. \tag{43}$$

and the inequality is replaced by the equality if we replace Δ by $\hat{\Delta}_n^\star$.

If $\inf_{z\geq 0}\Delta(z) \geq 0$, then

$$\mathbb{E}\left[n^{1/3}\psi\Delta(Z(s)) + \frac{\phi}{(\lambda f(\bar{p}))^2}\Delta(Z(s))^2 + hZ(s)\right] \geq n^{1/3}\psi\Delta(0) + \mathbb{E}[hZ(s)]$$

because $\psi < 0$. However, in this case $\liminf_{t\to\infty}\frac{1}{t}\mathbb{E}\left[\int_0^t Z(s)\,ds\right] = \infty$ because Z is lower bounded by a reflected Brownian motion with zero drift. So, (43) follows.

Suppose now that $\inf_{z\geq 0}\Delta(z) := \delta < 0$. In this case, there must exist $\hat{z} < \infty$ such that $\Delta(z) \leq 0$ for all $z \geq \hat{z}$. To proceed, let $v^\star(q) := \int_0^q g^\star(x)\,dx$. Then, by the properties of $g^\star$ stated in the proposition, we have $v^\star(0) = 0$ and $v^\star(q) \leq Cq^{3/2}$. Furthermore, it is straightforward to check that $v^\star$ satisfies

$$v^{\star\prime\prime}(q) + hq + \inf_{\Delta}\left\{\left(n^{1/3}\psi + v^{\star\prime}(q)\right)\Delta(q) + \frac{\phi}{(\lambda f(\bar{p}))^2}\Delta(q)^2\right\} = \kappa^\star. \tag{44}$$

Recall that Z is the (weak) solution of

$$Z(t) = \int_0^t \Delta(Z(s))\,ds + \sqrt{2}B(t) + I(t),$$

where B is a standard Brownian motion and I is an increasing process such that $\int_0^t Z(u)\,dI(u) = 0$ for all $t \geq 0$ a.s. By Ito's lemma, we have

$$\begin{aligned}&v^\star(Z(t)) - v^\star(Z(0))\\ &= \int_0^t \left(v^{\star\prime}(Z(u))\Delta(Z(u)) + v^{\star\prime\prime}(Z(u))\right)du + \sqrt{2}\int_0^t v^{\star\prime}(Z(u))\,dB(u) + \int_0^t v^{\star\prime}(Z(u))\,dI(u).\end{aligned} \tag{45}$$

By adding

$$\int_0^t \left\{ n^{1/3}\psi\Delta\left(Z\left(s\right)\right) + \frac{\phi}{\left(\lambda f\left(\bar{p}\right)\right)^2}\Delta\left(Z\left(s\right)\right)^2 + hZ\left(s\right) \right\} ds$$

on both hand sides of (45), we obtain

$$\begin{aligned}
&\mathbb{E}\left[\frac{v^\star\left(Z\left(t\right)\right) - v^\star\left(Z\left(0\right)\right)}{t}\right] + \frac{1}{t}\mathbb{E}\left[\int_0^t \left\{ n^{1/3}\psi\Delta\left(Z\left(s\right)\right) + \frac{\phi}{\left(\lambda f\left(\bar{p}\right)\right)^2}\Delta\left(Z\left(s\right)\right)^2 + hZ\left(s\right) \right\} ds\right] \\
&= \frac{1}{t}\mathbb{E}\left[\int_0^t \left\{ v^{\star\prime\prime}\left(Z\left(s\right)\right) + hZ\left(s\right) + \left(v^{\star\prime}\left(Z\left(s\right)\right) + n^{1/3}\psi\right)\Delta\left(Z\left(s\right)\right) + \frac{\phi}{\left(\lambda f\left(\bar{p}\right)\right)^2}\Delta\left(Z\left(s\right)\right)^2 \right\} ds\right],
\end{aligned}$$

where the last equality follows because $\int_0^t v^{\star\prime}\left(Z\left(s\right)\right) dB\left(s\right)$ is a martingale and $v^{\star\prime}\left(Z\left(s\right)\right) dI\left(s\right) = 0$ for all $s \geq 0$ a.s. by the condition $v^{\star\prime}\left(0\right) = 0$ and I only increases when $Z\left(s\right) = 0$. As $v^\star$ and $\kappa^\star$ jointly solve (44), we have

$$\frac{1}{t}\mathbb{E}\left[\int_0^t \left\{ n^{1/3}\psi\Delta\left(Z\left(s\right)\right) + \frac{\phi}{\left(\lambda f\left(\bar{p}\right)\right)^2}\Delta\left(Z\left(s\right)\right)^2 + hZ\left(s\right) \right\} ds\right] \geq \kappa^\star - \mathbb{E}\left[\frac{v^\star\left(Z\left(t\right)\right) - v^\star\left(Z\left(0\right)\right)}{t}\right]$$

for any $t \geq 0$ and Δ such that $\inf_{z\geq 0}\Delta\left(z\right) = \delta < 0$. Because $v^\star\left(q\right) \leq Cq^{3/2}$, we have

$$\mathbb{E}\left[\frac{v^\star\left(Z\left(t\right)\right)}{t}\right] \leq \mathbb{E}\left[\frac{Z\left(t\right)^{3/2}}{t}\right] \leq \left(\frac{1}{t^{1/3}}\mathbb{E}\left[\frac{Z\left(t\right)^2}{t}\right]\right)^{3/4} \tag{46}$$

by Jensen's inequality. By Ito's lemma, we have

$$\begin{aligned}
\mathbb{E}\left[Z\left(t\right)^2\right] &= \mathbb{E}\left[Z\left(0\right)^2\right] + 2t + \mathbb{E}\left[\int_0^t 2\Delta\left(Z\left(s\right)\right) Z\left(s\right) ds\right] \\
&= \mathbb{E}\left[Z\left(0\right)^2\right] + 2t + \mathbb{E}\left[\int_0^t 2\Delta\left(Z\left(s\right)\right) Z\left(s\right) 1_{\{Z(s)\leq\hat{z}\}} ds\right] + \mathbb{E}\left[\int_0^t 2\Delta\left(Z\left(s\right)\right) Z\left(s\right) 1_{\{Z(s)>\hat{z}\}} ds\right] \\
&\leq \mathbb{E}\left[Z\left(0\right)^2\right] + 2t + \mathbb{E}\left[\int_0^t 2\Delta\left(Z\left(s\right)\right) Z\left(s\right) 1_{\{Z(s)\leq\hat{z}\}} ds\right] \\
&\leq \mathbb{E}\left[Z\left(0\right)^2\right] + 2\left(1 + \Delta\left(\hat{z}\left(0\right)\right)\right)t,
\end{aligned}$$

implying that

$$\limsup_{t\to\infty} \mathbb{E}\left[\frac{Z\left(t\right)^2}{t}\right] < \infty.$$

Therefore, by (46), we have

$$\limsup_{t\to\infty} \mathbb{E}\left[\frac{v^\star(Z(t))}{t}\right] = 0,$$

(it is without loss of generality to assume $\mathbb{E}\left[Z(0)^2\right] < \infty$)so that (43) follows.

To prove that $\hat{R}_n^\star = \kappa^\star$, let $Z^\star$ be the (weak) solution of

$$Z^\star(t) = \int_0^t \hat{\Delta}_n^\star(Z^\star(s))\,ds + \sqrt{2}B^\star(t) + I^\star(t),$$

where $B^\star$ is a standard Brownian motion that is independent of B and $I^\star$ is an increasing process such that $Z^\star(s)\,dI^\star(s) = 0$ for all $s \geq 0$ a.s.

Observe that because $g^\star(z)$ is increasing in z and $\lim_{z\to\infty} g^\star(z) = \infty$, $\hat{\Delta}_n^\star(z)$ is decreasing in z and $\lim_{z\to\infty} \hat{\Delta}_n^\star(z) = -\infty$. Therefore, we can repeat the previous arguments to obtain

$$\frac{1}{t}\mathbb{E}\left[\int_0^t \left\{n^{1/3}\psi\hat{\Delta}_n^\star(Z^\star(s)) + \frac{\phi}{(\lambda f(\bar{p}))^2}\hat{\Delta}_n^\star(Z^\star(s))^2 + hZ^\star(s)\right\} ds\right] = \kappa^\star - \mathbb{E}\left[\frac{v^\star(Z^\star(t)) - v^\star(Z^\star(0))}{t}\right]$$

and

$$\limsup_{t\to\infty} \mathbb{E}\left[\frac{v^\star(Z(t)) - v^\star(0)}{t}\right] = 0.$$

Therefore, we have

$$\liminf_{t\to\infty} \frac{1}{t}\mathbb{E}\left[\int_0^t \left\{n^{1/3}\psi\hat{\Delta}_n^\star(Z^\star(s)) + \frac{\phi}{(\lambda f(\bar{p}))^2}\Delta^\star(Z^\star(s))^2 + hZ^\star(s)\right\} ds\right] = \kappa^\star,$$

proving the optimality of $\hat{\Delta}_n^\star$. □

C Proofs of Lemmas

This section is omitted for brevity. The complete version is available from authors upon request.