Misspecification Loss in Inventory Systems with Correlated Demand

Guokai Li
Smith School of Business, Queen's University, Kingston, Ontario, Canada, guokai.li@queensu.ca

Jinzhi Bu
Department of Logistics and Maritime Studies, Faculty of Business, The Hong Kong Polytechnic University, Hung Hom, Kowloon, Hong Kong, jinzhi.bu@polyu.edu.hk

Boxiao Chen
Department of Information & Decision Sciences, College of Business Administration, University of Illinois Chicago, Chicago, IL, USA, bbchen@uic.edu

Stefanus Jasin
Department of Technology and Operations, Stephen M. Ross School of Business, University of Michigan, Ann Arbor, MI, USA, sjasin@umich.edu

We study a periodic-review inventory system with positive lead time and correlated demand. While most inventory models assume independent demand for tractability, this simplification may lead to model misspecification in practice. We consider a decision maker who implements a myopic base-stock policy based on a misspecified i.i.d. model, while the true demand follows a stationary Gaussian AR(1) process. We derive closed-form expressions for the optimal policy under the true model and for the cost of the misspecified policy, and quantify misspecification loss via a competitive ratio. Our results show that longer lead times do not necessarily amplify misspecification loss. In particular, when demand exhibits positive correlation, increasing the lead time can attenuate the impact of misspecification. At the same time, this misspecification becomes increasingly consequential when backlog penalties are high. Under positive correlation, underestimation of uncertainty leads to persistent understocking and substantial losses, whereas under negative correlation, misspecification induces overstocking and the resulting loss is less sensitive to stockout penalties. Overall, the cost of using a simplified demand model depends not only on demand variability, but also on whether demand exhibits persistence or reversal over time and the relative costs of backlog versus excess inventory.

1. Introduction

In the classical literature on stochastic inventory control, intertemporally independent demand is often assumed for analytical tractability (see Goldberg et al. 2021 for a review). However, in many practical settings, demand exhibits significant serial dependence and is more appropriately modeled by a time-series process. For instance, demand stimulated by advertising or promotions today may cannibalize future demand, inducing negative correlation over time. Accounting for such dependence typically requires tracking the entire demand history in the associated dynamic program, leading to the well-known curse of dimensionality. As a result, much of the literature focuses on models with independent demand. To implement these models in practice, however, one must approximate the true demand process using a parsimonious forecasting model (e.g., i.i.d. demand), which introduces model misspecification. This raises the question of how such misspecification affects inventory decisions. In particular, we investigate the following question:

How does misspecification loss vary with lead time, correlation level, and backlog penalty?

This paper provides an exact, closed-form answer to this question in an analytically tractable setting. Specifically, we consider a periodic-review backlogged inventory system with deterministic lead time $L > 0$ and an autoregressive demand process AR(1) with correlation parameter $\theta \in (-1, 1)$. We compare two myopic (newsvendor) base-stock implementations for the cumulative lead-time demand. In the first, the decision maker uses the *true* AR(1) model to form the conditional distribution of the $(L{+}1)$-period demand and selects the corresponding conditional quantile. In the second, the decision maker *misspecifies* demand as an AR(0) (i.i.d. Gaussian) process, selects the unconditional quantile, and applies it in the true AR(1) environment. We measure misspecification loss using a *competitive ratio*, defined as the long-run average cost under the misspecified policy divided by that under the policy based on the true model.

Main results. We show that the misspecification loss can be decomposed into two components: *information factor* and *correlation factor*. Information factor, measured by the ratio of unconditional to conditional lead-time demand variance, reflects the relative uncertainty increment over the lead-time window as the seller ignores the last-period demand information. Correlation factor, measured by the ratio of approximated to true lead-time demand variance, captures the distortion arising from approximating correlated demands as independent.

Our first set of results concerns the impact of lead time on misspecification loss. We show that, for sufficiently large L, the competitive ratio is monotone in L, with the direction depending on θ (and a threshold $\tilde{\theta}$): when $\theta < \tilde{\theta}$ (e.g., demand is sufficiently negatively correlated), a longer lead time *amplifies* misspecification loss, whereas when $\theta > \tilde{\theta}$ (e.g., demand is sufficiently positively correlated), a longer lead time *mitigates* it. This finding is surprising and shows that lead time does not simply make misspecification worse. Instead, its effect depends on how demand evolves over time and on which of the two components of the misspecification loss has a dominant effect. On the one hand, the information factor decreases in L because the influence of the last observed demand on future demand decays geometrically. On the other hand, the correlation factor increases in L because a longer lead-time involves more correlated demand realizations, making the independence approximation less accurate.

When demand is positively correlated, the impact of the last-period demand information accumulates over time, resulting in a relatively large information factor. Meanwhile, since the demand process mainly reinforces movement in the same direction, the independent-demand approximation is relatively accurate and the induced correlation factor is relatively small. As a result, the information factor dominates, so that the misspecification loss decreases in L. In contrast, when demand is negatively correlated, the impact of the last-period demand information offsets over time, resulting in a relatively small information factor. In this case, since demand tends to alternate over time, an independent model fails to capture this oscillating dependence structure and thus incurs a substantially larger correlation factor. Consequently, the correlation factor dominates and the misspecification loss increases in L.

Our second set of results characterizes limiting regimes that further clarify the role of correlation. We derive the limiting competitive ratio as $L \to \infty$ and show that it increases in the absolute correlation level $|\theta|$. Furthermore, for a fixed $|\theta|$, positive (negative, respectively) correlation leads to higher misspecification loss when $|\theta|$ is small (large, respectively). We also study the effect of the backlog penalty, identifying regimes in which the competitive ratio diverges under positive correlation as $b \to \infty$, while remaining bounded under negative correlation.

These asymptotic results provide simple diagnostics for when demand correlation must be modeled carefully and when simplified forecasting may be sufficient from a cost perspective.

Related literature. Our work contributes to the inventory management literature; see Goldberg et al. (2021) and Song (2023) for comprehensive reviews. Much of this literature assumes that demand is intertemporally independent, primarily for analytical tractability. As Goldberg et al. (2021) note, "When demand is not independent over time, even simple models can become inherently high dimensional ... this phenomenon arises in many realistic models for demand." As a result, practitioners often rely on simplified models, which introduce misspecification. This raises the question of how costly such misspecification can be.

Our work is closely related to the literature that provides managerial insights under correlated demand. For tractability, most studies (including ours) focus on backlogged systems and an AR(1) demand model with correlation coefficient $\theta \in (-1, 1)$, which admits explicit formulas. For example, Lee et al. (1997) study the bullwhip effect in a backlogged inventory system and show that under positive correlation (i.e., $\theta \in (0, 1)$), the order variance at upstream stages exceeds that of end-customer demand, and this distortion increases with lead time. The case of negative correlation is not considered. Chen et al. (2000) incorporate demand forecasting and also find that the bullwhip effect is amplified by lead time and positive correlation. Lee et al. (2000) show that the inventory reduction from information sharing increases with the correlation level, the coefficient of variation, and the backlog penalty.

In contrast to these works, which focus on variances or inventory levels, we directly study the ratio of costs under the true and misspecified models. Our results reveal a more nuanced role of lead time: rather than uniformly worsening performance under misspecification, a longer lead time can mitigate its impact when demand is positively correlated or only mildly negatively correlated.

In addition to analytical results, there is also a stream of work on algorithms for correlated demand models. Many papers consider backlogged systems under autoregressive moving-average (ARMA) demand and allow negative orders. A common approach is a greedy policy that minimizes one-period inventory cost. For example, Johnson and Thompson (1975) prove the optimality of a greedy policy with zero lead time, and Aviv (2003) extends this result to positive lead times. When negative orders are not allowed, Caldentey et al. (2022) provide an upper bound on the gap between a modified greedy policy and the greedy policy. Since the greedy policy provides a lower bound on the optimal cost, their bound also upper-bounds the optimality gap. They further show numerically that this gap is typically small, justifying the use of the greedy policy as a benchmark. In addition, Levi et al. (2007), Chao et al. (2015), and Chao et al. (2018) develop algorithms with worst-case guarantees, while Hihat et al. (2023) and Qi et al. (2024) study learning-based approaches.

Organization. The remainder of the paper is organized as follows. In Section 2, we introduce the model. In Section 3, we analyze the inventory policy and define the misspecification loss. In Section 4, we study the comparative statics of the misspecification loss. In Section A, we conduct numerical experiments to test the robustness of our results. Section 5 concludes.

2. Model

We consider a periodic-review, fully backlogged inventory system over an infinite discrete-time horizon with deterministic lead time $L \geq 0$.

Ground-truth AR(1) **model.** Following the classic supply chain literature (see, e.g., Lee et al. 1997, Chen et al. 2000, Lee et al. 2000), we model intertemporal demand correlation via a stationary first-order autoregressive process. Specifically, demand $\{D_t\}_{t\geq 1}$ evolves as

$$D_t = \alpha + \theta \cdot D_{t-1} + \epsilon_t,$$

where $\alpha \in \mathbb{R}$ is a constant intercept, $\theta \in (-1, 1)$ is the autocorrelation coefficient, and $\{\varepsilon_t\}$ is i.i.d. Gaussian white noise with mean zero and variance σ^2. The restriction $\theta \in (-1, 1)$ is necessary and sufficient for stationarity; outside this range the process is explosive. Let $\mu := \alpha/(1-\theta)$ denote the stationary mean. Then the process can be written in mean-centered form as

$$D_t - \mu = \theta(D_{t-1} - \mu) + \epsilon_t,$$

or equivalently, in moving-average form as $D_t - \mu = \sum_{k=0}^{\infty} \theta^k \varepsilon_{t-k}$. The lag-0 and lag-1 autocovariances are

$$\gamma_0 = \mathbb{E}\left[(D_t - \mu)(D_t - \mu)\right] = \frac{1}{1-\theta^2}\sigma^2, \qquad \gamma_1 = \mathbb{E}\left[(D_t - \mu)(D_{t-1} - \mu)\right] = \frac{\theta}{1-\theta^2}\sigma^2. \tag{1}$$

Misspecified AR(0) **model.** As noted in the introduction, much of the inventory literature focuses on independent demand, which leads practitioners to approximate correlated demand with i.i.d. models. We capture this by assuming the seller ignores the serial dependence in demand and fits a misspecified AR(0) model:

$$\hat{D}_t - \hat{\mu} = \hat{\epsilon}_t,$$

where $\hat{\epsilon}_t \sim \mathcal{N}(0, \hat{\sigma}^2)$. Applying the Yule-Walker equations (see Section 2.5.3 in Brockwell and Davis 2002) to match the sample moments of the true AR(1) process yields

$$\hat{\mu} = \mu, \qquad \hat{\sigma}^2 = \gamma_0 = \frac{1}{1-\theta^2}\sigma^2,$$

where the last equality follows from (1). In other words, the misspecified model correctly captures the unconditional mean and variance of demand, but ignores all serial dependence.

Inventory system. The system can be modeled as a Markov decision process with state $\boldsymbol{s}_t := (I_t, x_{t-1}, \ldots, x_{t-L+1}, D_{t-1})$ at period t, where I_t denotes the on-hand inventory and x_{t-i} denotes the order quantity placed in period $t-i$. Since unmet demand is fully backlogged, inventory evolves as

$$I_t = I_{t-1} - D_{t-1} + x_{t-L}.$$

There are no fixed or variable ordering costs. The per-period holding and backlog cost is

$$C(\boldsymbol{s}_t, D_t) = h \cdot (I_t - D_t)^+ + b \cdot (D_t - I_t)^+,$$

where $h > 0$ is the unit holding cost and $b > 0$ is the unit backlog penalty. Following the classical literature (e.g., Zipkin 2000), we assume $b > h$. An admissible policy $\pi = \{\pi_t\}_{t \geq 0}$ is a sequence of decision rules such that in each period t, the order quantity is given by $x_t = \pi_t(\mathcal{H}_t) \in \mathbb{R}$ where $\mathcal{H}_t = \{\boldsymbol{s}_\ell\}_{\ell=0}^t$. Note that we allow negative order quantities (i.e., $x_t < 0$) for analytical tractability. The seller's objective is to minimize the long-run average cost over all admissible policies π:

$$V^\pi = \limsup_{T \to \infty} \mathbb{E}\left[\frac{1}{T} \sum_{t=0}^{T-1} C(\boldsymbol{s}_t^\pi, D_t)\right],$$

where $\boldsymbol{s}_t^\pi$ denotes the state at period t under policy π.

3. Misspecification Loss

In this section, we derive closed-form expressions for the average cost under each policy and characterize the resulting misspecification loss. Let $\phi(\cdot)$ and $\Phi(\cdot)$ denote the standard normal PDF and CDF, respectively. We also define $\beta := b/(b+h)$ and $z_\beta := \Phi^{-1}(\beta)$. We begin by recalling a standard result for the newsvendor problem with normally distributed demand.

Lemma 1 (**Newsvendor with normal demand; Zipkin 2000**). *Let $Y \sim \mathcal{N}(\mu_0, \sigma_0^2)$ with $\mu_0 \in \mathbb{R}$ and $\sigma_0 > 0$. Fix holding cost $h > 0$ and backlog penalty $b > 0$. For a given order-up-to level $q \in \mathbb{R}$, define cost*

$$J(q) := \mathbb{E}\big[h\,(q-Y)^+ + b\,(Y-q)^+\big].$$

Then the following statements hold:

(a) For any $q \in \mathbb{R}$, $J(q) = \sigma_0 \cdot [(h+b)\phi(\eta) + \eta\big((h+b)\Phi(\eta) - b\big)]$, where $\eta := (q-\mu_0)/\sigma_0$.

(b) $J(q)$ is convex in q and admits a unique minimizer $q^\star$.

(c) The unique minimizer is $q^\star = \mu_0 + z_\beta \cdot \sigma_0$.

(d) The minimal expected cost is $J(q^\star) = (h+b)\phi(z_\beta) \cdot \sigma_0$.

3.1. Myopic Policy Under the True Model

We first introduce the inventory policy under the true AR(1) model. We focus on the *myopic policy*, which in each period solves a newsvendor problem for the cumulative lead-time demand. Specifically, given state $\boldsymbol{s}_t$, the myopic policy sets the target inventory level q_t^* by solving

$$q_t^* = \operatorname*{arg\,min}_{q_t \in \mathbb{R}} \mathbb{E}\left[h \cdot (q_t - \Delta_{t,L})^+ + b \cdot (\Delta_{t,L} - q_t)^+ \middle| \boldsymbol{s}_t\right],$$

where $\Delta_{t,L} := \sum_{k=0}^{L} D_{t+k}$ denotes the cumulative demand over periods t through $t+L$. The corresponding order quantity is then $x_t^* = q_t^* - I_t - \sum_{k=1}^{L-1} x_{t-k}$.

Given state $\boldsymbol{s}_t$ and the true AR(1) model, we can expand the cumulative lead-time demand as

$$\begin{aligned}
\Delta_{t,L} = \sum_{k=0}^{L} D_{t+k} &= \sum_{k=0}^{L}\left(\mu + \theta^{k+1}(D_{t-1} - \mu) + \sum_{j=0}^{k} \theta^{k-j}\epsilon_{t+j}\right) \\
&= (L+1)\mu + \left(\sum_{k=0}^{L} \theta^{k+1}\right)(D_{t-1} - \mu) + \sum_{k=0}^{L}\left(\sum_{i=0}^{L-k} \theta^i\right)\epsilon_{t+k} \\
&= (L+1)\mu + \frac{\theta - \theta^{L+2}}{1-\theta}(D_{t-1} - \mu) + \sum_{k=0}^{L} \frac{1-\theta^{L-k+1}}{1-\theta}\epsilon_{t+k},
\end{aligned} \tag{2}$$

where the last step applies the geometric series formula. Since $\Delta_{t,L}$ is a linear combination of the Gaussian noises $\{\epsilon_{t+k}\}$ and the state $\boldsymbol{s}_t$, it follows that $\Delta_{t,L} \mid \boldsymbol{s}_t$ is Gaussian with conditional mean and variance

$$\begin{aligned}
\mathbb{E}[\Delta_{t,L} \mid \boldsymbol{s}_t] &= (L+1)\mu + \frac{\theta - \theta^{L+2}}{1-\theta}(D_{t-1} - \mu), \\
\mathbb{V}[\Delta_{t,L} \mid \boldsymbol{s}_t] &= \sigma^2 \sum_{j=0}^{L}\left(\frac{1-\theta^{L-j+1}}{1-\theta}\right)^2 = \frac{\sigma^2}{(1-\theta)^2}\left(L + 1 - \frac{2\theta(1-\theta^{L+1})}{1-\theta} + \frac{\theta^2(1-\theta^{2L+2})}{1-\theta^2}\right).
\end{aligned}$$

By Lemma 1(c), the optimal target inventory level is

$$\begin{aligned}
q_t^* &= \mathbb{E}\left[\Delta_{t,L} \mid \boldsymbol{s}_t\right] + z_\beta \cdot \sqrt{\mathbb{V}\left[\Delta_{t,L} \mid \boldsymbol{s}_t\right]} \\
&= (L+1)\mu + \frac{\theta - \theta^{L+2}}{1-\theta}(D_{t-1} - \mu) + z_\beta \cdot \frac{\sigma}{1-\theta}\sqrt{L + 1 - \frac{2\theta(1-\theta^{L+1})}{1-\theta} + \frac{\theta^2(1-\theta^{2L+2})}{1-\theta^2}}.
\end{aligned} \tag{3}$$

The corresponding long-run average cost under the myopic policy is, by Lemma 1(d),

$$\begin{aligned}
V^* &= \limsup_{T\to\infty} \mathbb{E}\left[\frac{1}{T}\sum_{t=0}^{T-1}(h+b)\cdot\phi(z_\beta)\cdot\sqrt{\mathbb{V}\left[\Delta_{t,L} \mid \boldsymbol{s}_t\right]}\right] \\
&= (h+b)\cdot\phi(z_\beta)\cdot\sqrt{\mathbb{V}\left[\Delta_{t,L} \mid \boldsymbol{s}_t\right]} \\
&= (h+b)\cdot\phi\left(z_\beta\right)\cdot\frac{\sigma}{1-\theta}\sqrt{L + 1 - \frac{2\theta(1-\theta^{L+1})}{1-\theta} + \frac{\theta^2(1-\theta^{2L+2})}{1-\theta^2}},
\end{aligned}$$

where the second equality uses the fact that $\mathbb{V}[\Delta_{t,L} \mid \boldsymbol{s}_t]$ is constant across periods, i.e., $\mathbb{V}[\Delta_{t,L} \mid \boldsymbol{s}_t] = \mathbb{V}[\Delta_{t',L} \mid \boldsymbol{s}_{t'}]$ for all t, t'. Since the myopic policy is known to be optimal for this class of systems (see, e.g., Lee et al. 1997, Caldentey et al. 2022), V^* is indeed the optimal long-run average cost.

3.2. Myopic Policy Under the Misspecified Model

We now discuss the myopic policy when the seller uses the misspecified AR(0) model. Given state $\boldsymbol{s}_t$, the target inventory level $\hat{q}_t^*$ solves

$$\hat{q}_t^* = \operatorname*{arg\,min}_{\hat{q}_t \in \mathbb{R}} \mathbb{E}\left[h \cdot \left(\hat{q}_t - \hat{\Delta}_{t,L}\right)^+ + b \cdot \left(\hat{\Delta}_{t,L} - \hat{q}_t\right)^+ \middle| \boldsymbol{s}_t \right],$$

where $\hat{\Delta}_{t,L} := \sum_{k=0}^{L} \hat{D}_{t+k}$ is the cumulative demand forecast under the misspecified model. Since the misspecified model posits $\hat{D}_t - \mu = \hat{\epsilon}_t$ with $\hat{\epsilon}_t \sim \mathcal{N}(0, \gamma_0)$ independently across periods, the estimated cumulative demand is

$$\hat{\Delta}_{t,L} = \sum_{k=0}^{L} \hat{D}_{t+k} = (L+1)\mu + \sum_{\ell=t}^{t+L} \hat{\epsilon}_\ell,$$

which is Gaussian with mean and variance

$$\mathbb{E}[\hat{\Delta}_{t,L}] = (L+1)\mu, \quad \mathbb{V}[\hat{\Delta}_{t,L}] = (L+1)\gamma_0 = \frac{(L+1)\sigma^2}{1-\theta^2}.$$

By Lemma 1(c), the target stock level under the misspecified model is therefore

$$\hat{q}_t^* = \mathbb{E}[\hat{\Delta}_{t,L}] + z_\beta \cdot \sqrt{\mathbb{V}[\hat{\Delta}_{t,L}]} = (L+1)\mu + z_\beta \sqrt{\frac{L+1}{1-\theta^2}}\sigma, \tag{4}$$

and the corresponding order quantity is $\hat{x}_t^* = \hat{q}_t^* - I_t - \sum_{k=1}^{L-1} x_{t-k}$.

We now compute the long-run average cost incurred when the misspecified policy $\hat{q}_t^*$ is deployed in the true AR(1) environment:

$$\hat{V} = \limsup_{T\to\infty} \mathbb{E}\left[\frac{1}{T} \sum_{t=0}^{T-1} \left(h \cdot (\hat{q}_t^* - \Delta_{t,L})^+ + b \cdot (\Delta_{t,L} - \hat{q}_t^*)^+ \right) \right].$$

Crucially, the cost is evaluated against the true cumulative demand $\Delta_{t,L}$, not the misspecified forecast $\hat{\Delta}_{t,L}$. Since $\hat{q}_t^*$ is a constant (see (4)) and $\Delta_{t,L}$ is stationary, the induced one-period cost process $h(\hat{q}_t^* - \Delta_{t,L})^+ + b(\Delta_{t,L} - \hat{q}_t^*)^+$ is stationary. Therefore, its expectation is the same in every period, and

$$\hat{V} = \mathbb{E}\left[h \cdot (\hat{q}_t^* - \Delta_{t,L})^+ + b \cdot (\Delta_{t,L} - \hat{q}_t^*)^+ \right].$$

From (2), $\Delta_{t,L}$ is Gaussian with mean $\mathbb{E}[\Delta_{t,L}] = (L+1)\mu$. Its variance follows from the law of total variance:

$$\begin{aligned} \mathbb{V}[\Delta_{t,L}] &= \mathbb{V}\left[\sum_{\ell=1}^{L+1} \theta^\ell (D_{t-1} - \mu) \right] + \mathbb{V}[\Delta_{t,L} \mid \boldsymbol{s}_t] \\ &= \left(\frac{\theta - \theta^{L+2}}{1-\theta} \right)^2 \gamma_0 + \frac{\sigma^2}{(1-\theta)^2} \left(L + 1 - \frac{2\theta(1-\theta^{L+1})}{1-\theta} + \frac{\theta^2(1-\theta^{2L+2})}{1-\theta^2} \right) \\ &= \frac{\sigma^2}{(1-\theta)^2} \left(L + 1 - \frac{2\theta(1-\theta^{L+1})}{1-\theta^2} \right). \end{aligned} \tag{5}$$

Note that the unconditional variance $\mathbb{V}[\Delta_{t,L}]$ is different from the estimation $\mathbb{V}[\hat{\Delta}_{t,L}]$ based on the misspecified model. In the following lemma, we establish their relationship.

Lemma 2. *$\mathbb{V}[\Delta_{t,L}] \le \mathbb{V}[\hat{\Delta}_{t,L}]$ for $\theta \in (-1,0)$ and $\mathbb{V}[\Delta_{t,L}] \ge \mathbb{V}[\hat{\Delta}_{t,L}]$ for $\theta \in (0,1)$.*

By Lemma 2, the misspecified model overestimates (underestimates, resp.) demand variance over the lead-time window when the demand process is negatively (positively, resp.) correlated.

Proof of Lemma 2. We analyze the following difference:

$$\mathbb{V}[\Delta_{t,L}] - \mathbb{V}[\hat{\Delta}_{t,L}] = \frac{2\sigma^2\theta}{(1+\theta)(1-\theta)^2}\left(L - \sum_{k=1}^{L}\theta^k\right).$$

Since $\frac{2\sigma^2}{(1+\theta)(1-\theta)^2}\left(L-\sum_{k=1}^{L}\theta^k\right) \ge 0$ for $\theta \in (-1,1)$, the comparison between $\mathbb{V}[\Delta_{t,L}]$ and $\mathbb{V}[\hat{\Delta}_{t,L}]$ only depends on the sign of θ: $\mathbb{V}[\Delta_{t,L}] \ge \mathbb{V}[\hat{\Delta}_{t,L}]$ for $\theta \in (-1,0)$ and $\mathbb{V}[\Delta_{t,L}] \le \mathbb{V}[\hat{\Delta}_{t,L}]$ for $\theta \in (0,1)$. □

Define the normalized (true and estimated) standard deviation ratios

$$\delta_T := \sqrt{\frac{\mathbb{V}[\Delta_{t,L}]}{\sigma^2}} = \frac{1}{1-\theta}\sqrt{L+1-\frac{2\theta(1-\theta^{L+1})}{1-\theta^2}}, \qquad \delta_E := \sqrt{\frac{\mathbb{V}[\hat{\Delta}_{t,L}]}{\sigma^2}} = \sqrt{\frac{L+1}{1-\theta^2}}.$$

Applying Lemma 1(a) to evaluate $\hat{V}$ gives

$$\begin{aligned}\hat{V} &= \sqrt{\mathbb{V}[\Delta_{t,L}]}\left[(h+b)\cdot\phi\left(\frac{\hat{q}_t^* - \mathbb{E}[\Delta_{t,L}]}{\sqrt{\mathbb{V}[\Delta_{t,L}]}}\right) + \frac{\hat{q}_t^* - \mathbb{E}[\Delta_{t,L}]}{\sqrt{\mathbb{V}[\Delta_{t,L}]}}\left((h+b)\cdot\Phi\left(\frac{\hat{q}_t^* - \mathbb{E}[\Delta_{t,L}]}{\sqrt{\mathbb{V}[\Delta_{t,L}]}}\right) - b\right)\right]\\ &= (h+b)\phi\left(\frac{\delta_E}{\delta_T}z_\beta\right)\delta_T\sigma + \left((h+b)\Phi\left(\frac{\delta_E}{\delta_T}z_\beta\right) - b\right)\delta_E z_\beta \sigma.\end{aligned}$$

Due to the optimality of V^*, we have $\hat{V} \ge V^*$.

3.3. An Auxiliary Policy for Decomposing Misspecification Loss

We also introduce an auxiliary myopic policy that that does not use the last-period demand information D_{t-1}, but knows the accurate value of the unconditional demand variance over the lead time window, i.e., $\mathbb{V}[\Delta_{t,L}]$. In particular, given state $\boldsymbol{s}_t$, the target inventory level $\tilde{q}_t^*$ solves

$$\tilde{q}_t^* = \arg\min_{\tilde{q}_t \in \mathbb{R}} \mathbb{E}\left[h\cdot(\tilde{q}_t - \tilde{\Delta}_{t,L})^+ + b\cdot(\tilde{\Delta}_{t,L} - \tilde{q}_t)^+ \middle| \boldsymbol{s}_t\right],$$

where $\tilde{\Delta}_{t,L}$ is Gaussian variable with mean $(L+1)\mu$ and variance $\mathbb{V}[\Delta_{t,L}]$. Thus, the auxiliary policy sets the target inventory level as

$$\tilde{q}_t^* = (L+1)\mu + z_\beta \cdot \sqrt{\mathbb{V}[\Delta_{t,L}]}, \tag{6}$$

and the corresponding long-run average cost is denoted is

$$\tilde{V} = (h+b)\phi(z_\beta)\frac{\sigma}{1-\theta}\sqrt{L+1-\frac{2\theta(1-\theta^{L+1})}{1-\theta^2}}$$

It can be verified that the auxiliary policy is optimal among all policies ignoring the last-period information. In the following, we will use the auxiliary policy as a bridge between the true-model policy and the misspecified-model policy to explain the misspcefication loss.

3.4. Competitive Ratio

To measure the misspecification loss, we define the competitive ratio

$$\rho(\theta,L,\beta) := \frac{\hat{V}}{V^*} = \frac{\delta_T}{\frac{1}{1-\theta}\sqrt{L+1-\frac{2\theta(1-\theta^{L+1})}{1-\theta}+\frac{\theta^2(1-\theta^{2L+2})}{1-\theta^2}}} \cdot \frac{\phi\left(\frac{\delta_E}{\delta_T}z_\beta\right)+\frac{\delta_E}{\delta_T}z_\beta\left(\Phi\left(\frac{\delta_E}{\delta_T}z_\beta\right)-\Phi(z_\beta)\right)}{\phi\left(z_\beta\right)}, \quad (7)$$

which satisfies $\rho(\theta,L,\beta)\geq 1$ since $\hat{V}\geq V^*$. Two immediate observations follow from (7). First, $\rho(\theta,L,\beta)$ is independent of both the demand mean μ and the noise variance σ^2. Second, when $L=0$, we have $\delta_E/\delta_T=1$, which yields $\rho(\theta,0,\beta)=1/\sqrt{1-\theta^2}$, a quantity that depends only on θ and is independent of the cost parameters h and b. In what follows, we focus on the practically relevant case of positive lead time, $L\geq 1$.

Since the misspecified model ignores serial correlation, both the conditional mean and variance it uses may differ substantially from their true counterparts. To disentangle the two sources of loss, we decompose the competitive ratio as $\rho(\theta,L,\beta)=f(\theta,L)\cdot g(\theta,L,\beta)$, where

$$f(\theta,L) := \frac{\tilde{V}}{V^*} = \sqrt{\frac{\mathbb{V}[\Delta_{t,L}]}{\mathbb{V}[\Delta_{t,L}\mid \boldsymbol{s}_t]}} = \frac{\delta_T}{\frac{1}{1-\theta}\sqrt{L+1-\frac{2\theta(1-\theta^{L+1})}{1-\theta}+\frac{\theta^2(1-\theta^{2L+2})}{1-\theta^2}}}, \quad (8)$$

$$g(\theta,L,\beta) := \frac{\hat{V}}{\tilde{V}} = \frac{G\left(\frac{\delta_E}{\delta_T}z_\beta\right)}{G(z_\beta)} = \frac{\phi\left(\frac{\delta_E}{\delta_T}z_\beta\right)+\frac{\delta_E}{\delta_T}z_\beta\left(\Phi\left(\frac{\delta_E}{\delta_T}z_\beta\right)-\Phi(z_\beta)\right)}{\phi\left(z_\beta\right)}, \quad (9)$$

with $G(x) := \phi(x)+x(\Phi(x)-\Phi(z_\beta))$.

In the following, we explain the two components of the misspecification loss.

1. The *information factor* $f(\theta,L)$ measures the loss due to the ignorance of the last-period demand information D_{t-1}. Since $f(\theta,L)=\tilde{V}/V^*$, it captures the difference between the auxiliary policy and the true-demand policy. According to (3) and (6), due to the ignorance of D_{t-1}, the auxiliary policy deals with a more volatile demand, i.e., $\mathbb{V}[\Delta_{t,L}]\geq\mathbb{V}[\Delta_{t,L}\mid \boldsymbol{s}_t]$, and hence prepares more safety stocks. Furthermore, by (5), we can further quantify the difference $\mathbb{V}[\Delta_{t,L}]-\mathbb{V}[\Delta_{t,L}\mid \boldsymbol{s}_t]=\mathbb{V}\left[\sum_{\ell=1}^{L+1}\theta^\ell(D_{t-1}-\mu)\right]$.
2. The *correlation factor* $g(\theta,L,\beta)$ measures the loss due to the ignorance of the correlation between future demands. Since $g(\theta,L,\beta)=\hat{V}/\tilde{V}$, it captures the difference between the misspecified-model policy and the auxiliary policy. According to (4) and (6), the misspecified model treats future demands as independent random variables and derives wrong variance estimation $\mathbb{V}[\hat{\Delta}_{t,L}]$, resulting in an inappropriate amount of safety stock quantity. In particular, according to Lemma 2, we have $\mathbb{V}[\hat{\Delta}_{t,L}]\geq\mathbb{V}[\Delta_{t,L}]$ for $\theta\in(-1,0)$ and $\mathbb{V}[\hat{\Delta}_{t,L}]\leq\mathbb{V}[\Delta_{t,L}]$ for $\theta\in(0,1)$, implying that the misspecified-model policy overstocks (i.e., $\tilde{q}_t^*\leq\hat{q}_t^*$) under negative correlation while understocks (i.e., $\tilde{q}_t^*\geq\hat{q}_t^*$) under positive correlation.

In the following lemma, we show that both factors are greater than one.

Lemma 3. *$f(\theta,L)\geq 1$ and $g(\theta,L,\beta)\geq 1$.*

By Lemma 3, both factors induce negative effects and hence $\rho(\theta, L, \beta) \geq 1$.

Proof of Lemma 3. Since $\mathbb{V}[\Delta_{t,L}] \geq \mathbb{V}[\Delta_{t,L} \mid \boldsymbol{s}_t]$, we can deduce that $f(\theta, L) \geq 1$. Since $G'(x) = \Phi(x) - \Phi(z_\beta)$, we can deduce that $G(x)$ decreases when $x \leq z_\beta$ and increases when $x > z_\beta$, implying $G(x) \geq G(z_\beta)$ and hence $g(\theta, L, \beta) \geq 1$. □

4. Comparative Statics

In this section, we analyze how the competitive ratio $\rho(\theta, L, \beta)$ varies with the lead time L, the correlation coefficient θ, and the backlog penalty b. We also define $z_\infty := z_\beta\sqrt{(1-\theta)/(1+\theta)}$, which will serve as a key quantity in the asymptotic analysis that follows.

4.1. Impact of Lead Time L

We begin by establishing the asymptotic behavior of each factor as L grows large.

Proposition 1 **(Asymptotic Behavior of the Two Factors)**.

1. *The information factor admits the expansion*

$$f(\theta, L) = 1 + \frac{\theta^2}{2(1-\theta^2)(L+1)} + \frac{\theta^3(3\theta+8)}{8(1-\theta^2)^2(L+1)^2} + O\left(\frac{1}{(L+1)^3}\right),$$

and is decreasing in L for all sufficiently large L, with $\lim_{L\to\infty} f(\theta, L) = 1$.

2. *The correlation factor admits the expansion*

$$g(\theta, L, \beta) = \frac{G(z_\infty)}{\phi(z_\beta)} + \frac{\theta z_\infty\left(\Phi(z_\infty) - \Phi(z_\beta)\right)}{\phi(z_\beta)(L+1)(1-\theta^2)} + \frac{\theta^2 z_\infty\left(3\Phi(z_\infty) - 3\Phi(z_\beta) + z_\infty\phi(z_\infty)\right)}{2\phi(z_\beta)(L+1)^2(1-\theta^2)^2} + O\left(\frac{1}{(L+1)^3}\right),$$

and is increasing in L for all sufficiently large L, with $\lim_{L\to\infty} g(\theta, L, \beta) = G(z_\infty)/\phi(z_\beta)$.

Proposition 1 reveals that the two factors exert opposing forces as L grows: the information factor decreases toward one, while the correlation factor increases toward a limit that depends on θ and β. The intuition is as follows. First, as mentioned, the information factor is related to $\mathbb{V}[\Delta_{t,L}] - \mathbb{V}[\Delta_{t,L} \mid \boldsymbol{s}_t] = \mathbb{V}\left[\sum_{\ell=1}^{L+1} \theta^\ell (D_{t-1} - \mu)\right]$. The influence of the historical demand D_{t-1} on future periods decays geometrically (i.e., at rate θ^ℓ), so its cumulative effect over the lead-time window remains bounded. As L increases, this bounded cumulative effect becomes negligible relative to the growing lead-time demand $\Delta_{t,L}$, and the information factor consequently converges to one. For the correlation factor, the key observation is that the misspecified model treats the demands $\{D_{t+k}\}$ as independent, so each additional period in the lead-time window contributes an independent variance of γ_0 to $\mathbb{V}[\hat{\Delta}_{t,L}] = (L+1)\gamma_0$. In contrast, the true unconditional variance $\mathbb{V}[\Delta_{t,L}]$ grows more slowly with L under negative correlation (and faster under positive correlation) due to the offsetting (reinforcing, resp.) effect of the serial dependence. As the lead time increases, the involvement of more correlated demand variables makes the independent estimation less accurate, causing the correlation factor to grow.

Proof of Proposition 1. We first rewrite $f(\theta, L)$ in (8) as follows:

$$f(\theta, L) = \sqrt{\frac{\mathbb{V}[\Delta_{t,L}]}{\mathbb{V}[\Delta_{t,L} \mid \boldsymbol{s}_t]}} = \sqrt{\frac{\mathbb{V}[\Delta_{t,L} \mid \boldsymbol{s}_t] + \mathbb{E}[\Delta_{t,L} \mid \boldsymbol{s}_t]}{\mathbb{V}[\Delta_{t,L} \mid \boldsymbol{s}_t]}} = \sqrt{1 + \frac{\theta^2(1-\theta^{L+1})^2}{(L+1)(1-\theta^2) - \theta(1-\theta^{L+1})(2+\theta-\theta^{L+2})}}.$$

where the second identity follows from the law of total variance and the fact that $\mathbb{V}[\Delta_{t,L} \mid \boldsymbol{s}_t]$ and $\mathbb{E}[\Delta_{t,L} \mid \boldsymbol{s}_t]$ are both constants independent of s_t. By Taylor expansion that $\sqrt{1+x} = 1 + \frac{x}{2} - \frac{x^2}{8} + O(x^3)$ and the fact that $\frac{\theta^2(1-\theta^{L+1})^2}{(L+1)(1-\theta^2)-\theta(1-\theta^{L+1})(2+\theta-\theta^{L+2})} = \frac{\theta^2}{(1-\theta^2)(L+1)} + \frac{\theta^3(\theta+2)}{(1-\theta^2)^2(L+1)^2} + O\left(\frac{1}{(L+1)^3}\right)$, we can derive that

$$f(\theta, L) = 1 + \frac{\theta^2}{2(1-\theta^2)(L+1)} + \frac{\theta^3(3\theta+8)}{8(1-\theta^2)^2(L+1)^2} + O\left(\frac{1}{(L+1)^3}\right),$$

which also implies that $\lim_{L\to\infty} f(\theta, L) = 1$. Then, we have

$$f(\theta, L+1) - f(\theta, L) = -\frac{\theta^2}{2(1-\theta^2)(L+1)(L+2)} + O\left(\frac{1}{(L+1)^3}\right) \le 0,$$

which implies that $f(\theta, L)$ decreases in L when L is sufficiently large.

In the following, we analyze the correlation factor $g(\theta, L, \beta)$ in (9). By Taylor expansion $(1-x)^{-1/2} = 1 + \frac{x}{2} + \frac{3}{8}x^2 + O(x^3)$, we have

$$\frac{\delta_E}{\delta_T} z_\beta = z_\beta \sqrt{\frac{1-\theta}{1+\theta}} \cdot \sqrt{\frac{1}{1 - \frac{2\theta(1-\theta^{L+1})}{(L+1)(1-\theta^2)}}} = z_\infty \cdot \left(1 + \frac{\theta}{(L+1)(1-\theta^2)} + \frac{3\theta^2}{2(L+1)^2(1-\theta^2)^2} + O\left(\frac{1}{(L+1)^3}\right)\right).$$

Then, by Taylor expansion, we have

$$\begin{aligned} G\left(\frac{\delta_E}{\delta_T} z_\beta\right) &= G(z_\infty) + G'(z_\infty)\left(\frac{\delta_E}{\delta_T} z_\beta - z_\infty\right) + \frac{1}{2} G''(z_\infty)\left(\frac{\delta_E}{\delta_T} z_\beta - z_\infty\right)^2 + O\left(\left(\frac{\delta_E}{\delta_T} z_\beta - z_\infty\right)^3\right) \\ &= G(z_\infty) + \frac{\theta z_\infty\left(\Phi(z_\infty) - \Phi(z_\beta)\right)}{(L+1)(1-\theta^2)} + \frac{\theta^2 z_\infty\left(3\Phi(z_\infty) - 3\Phi(z_\beta) + z_\infty \phi(z_\infty)\right)}{2(L+1)^2(1-\theta^2)^2} + O\left(\frac{1}{(L+1)^3}\right), \end{aligned}$$

which implies that

$$g(\theta, L, \beta) = \frac{G(z_\infty)}{\phi(z_\beta)} + \frac{\theta z_\infty\left(\Phi(z_\infty) - \Phi(z_\beta)\right)}{\phi(z_\beta)(L+1)(1-\theta^2)} + \frac{\theta^2 z_\infty\left(3\Phi(z_\infty) - 3\Phi(z_\beta) + z_\infty \phi(z_\infty)\right)}{2\phi(z_\beta)(L+1)^2(1-\theta^2)^2} + O\left(\frac{1}{(L+1)^3}\right).$$

Similar to the proof for $f(\theta, L)$, we have $\lim_{L\to\infty} g(\theta, L, \beta) = G(z_\infty)/\phi(z_\beta)$ and the monotonicity of $g(\theta, L, \beta)$ regarding L depends on the sign of $\left(\Phi(z_\infty) - \Phi(z_\beta)\right)\theta$. If $\theta \in (-1, 0)$, we have $\Phi(z_\infty) - \Phi(z_\beta) \ge 0$ and $\theta \le 0$, implying $\left(\Phi(z_\infty) - \Phi(z_\beta)\right)\theta \le 0$; if $\theta \in (0, 1)$, we have $\Phi(z_\infty) - \Phi(z_\beta) \le 0$ and $\theta \ge 0$, implying $\left(\Phi(z_\infty) - \Phi(z_\beta)\right)\theta \le 0$. Since $\left(\Phi(z_\infty) - \Phi(z_\beta)\right)\theta \le 0$ for any $\theta \in (-1, 1)$, we can deduce that $g(\theta, L, \beta)$ increases in L when L is sufficiently large. $\square$

The following theorem establishes the monotonicity of the competitive ratio with respect to L; Figure 1 provides an illustration. We say that a function $h(x)$ eventually increases in x (eventually decreases, resp.) if there exists a constant $x_0 \in \mathbb{R}$ such that $h(\cdot)$ is increasing (decreasing, resp.) on $[x_0, \infty)$.

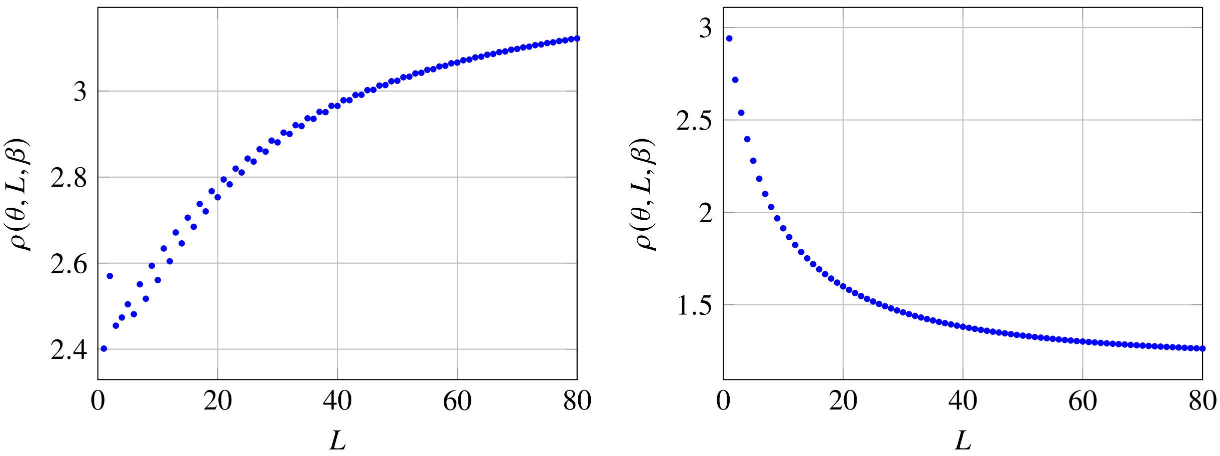

(a) $\theta = -0.95$ (b) $\theta = 0.95$

Figure 1 Competitive ratio $\rho(\theta, L, \beta)$ as functions of L when $h = 1$, $b = 3$, and $\theta \in \{-0.95, 0.95\}$.

THEOREM 1 **(Impact of Lead Time** L**).** *There exists a constant* $\tilde{\theta} \in (-1, 1)$ *such that:*

(i) For any given $\theta \in (-1, \tilde{\theta})$, $\rho(\theta, L, \beta)$ *eventually increases in* L*;*

(ii) For any given $\theta \in (\tilde{\theta}, 1)$, $\rho(\theta, L, \beta)$ *eventually decreases in* L*.*

Moreover, the threshold satisfies $\tilde{\theta} < 0$ *if* $z_\beta < 1/\sqrt{2}$, $\tilde{\theta} = 0$ *if* $z_\beta = 1/\sqrt{2}$, *and* $\tilde{\theta} > 0$ *if* $z_\beta > 1/\sqrt{2}$.

Theorem 1 shows that as L increases, the misspecification loss tends to be mitigated when $\theta > \tilde{\theta}$ and amplified when $\theta < \tilde{\theta}$. The direction is determined by which of the two countervailing factors dominates: the information factor decreases in L while the correlation factor increases in L.

When correlation is sufficiently negative, the impacts of D_{t-1} on different periods offsets and hence $\mathbb{V}[\Delta_{t,L}] - \mathbb{V}[\Delta_{t,L} \mid \boldsymbol{s}_t] = \mathbb{V}\left[\sum_{\ell=1}^{L+1} \theta^\ell (D_{t-1} - \mu)\right]$ is relatively small, resulting in a minuscule information factor. At the same time, the misspecified model ignores the strong mean-reversion pattern in demand and hence the corresponding variance estimation $\mathbb{V}[\hat{\Delta}_{t,L}]$ is far from the real value $\mathbb{V}[\Delta_{t,L}]$, leading to a significant correlation factor. The correlation factor therefore dominates, and $\rho(\theta, L, \beta)$ increases in L. When correlation is sufficiently positive, the impacts of D_{t-1} on different periods accumulates and hence $\mathbb{V}[\Delta_{t,L}] - \mathbb{V}[\Delta_{t,L} \mid \boldsymbol{s}_t]$ is relatively large, making the information factor significant. Meanwhile, the demand process reinforces in the same direction and hence the variance estimation $\mathbb{V}[\hat{\Delta}_{t,L}]$ does not deviate from $\mathbb{V}[\Delta_{t,L}]$ too much, so the correlation factor is negligible. The information factor therefore dominates, and $\rho(\theta, L, \beta)$ eventually decreases in L.

One might expect $\tilde{\theta} = 0$ to always hold, i.e., that zero correlation is the natural dividing line between the two regimes. Theorem 1 reveals, however, that this boundary is shifted by z_β. When the backlog penalty is high (i.e., $z_\beta > 1/\sqrt{2}$), since positive correlation induces systematic understocking (i.e., $\tilde{q}_t^* \geq \hat{q}_t^*$ for $\theta > 0$), which incurs large backlog costs. This additional penalty weakens the relative importance of the information factor, so a strictly positive threshold $\tilde{\theta} > 0$ is required before increasing L begins to mitigate misspecification loss.

Proof of Theorem 1. According to Proposition 1, we have

$$\rho(\theta,L,\beta)=f(\theta,L)\cdot g(\theta,L,\beta)=\frac{G(z_\infty)}{\phi(z_\beta)}+\frac{\theta^2G(z_\infty)+2\theta z_\infty\left(\Phi(z_\infty)-\Phi(z_\beta)\right)}{2\phi(z_\beta)(1-\theta^2)(L+1)}+\frac{c_1(\theta,\beta)}{(L+1)^2}+O\left(\frac{1}{(L+1)^3}\right),$$

where $c_1(\theta,\beta)$ is a term independent of L and its explicit expression is omitted. Then, we have

$$\rho(\theta,L+1,\beta)-\rho(\theta,L,\beta)=-\frac{\theta^2G(z_\infty)+2\theta z_\infty\left(\Phi(z_\infty)-\Phi(z_\beta)\right)}{2\phi(z_\beta)(1-\theta^2)L(L+1)}+O\left(\frac{1}{(L+1)^3}\right).$$

In the following, we analyze the sign of the following function

$$H(\theta):=\theta^2G(z_\infty)+2\theta z_\infty\left(\Phi(z_\infty)-\Phi(z_\beta)\right)=\theta^2\phi(z_\infty)+(\theta+2)\theta z_\infty\left(\Phi\left(z_\infty\right)-\Phi(z_\beta)\right).$$

If $H(\theta)>0$, then $\rho(\theta,L,\beta)$ decreases in L when L is sufficiently large; if $H(\theta)<0$, then $\rho(\theta,L,\beta)$ increases in L when L is sufficiently large.

We first analyze the zero roots of the function $H(\theta)$. First, we reparametrize $H(\theta)$ by defining $u=z_\infty$ and replacing θ with $\frac{z_\beta^2-u^2}{z_\beta^2+u^2}$. There exists a one-to-one mapping between u and θ, and $u=z_\beta$ corresponds to $\theta=0$. Thus, we can focus on the solutions to the following equation:

$$\frac{z_\beta^2-u^2}{z_\beta^2+u^2}\phi(u)+\frac{3z_\beta^2+u^2}{z_\beta^2+u^2}u(\Phi(u)-\Phi(z_\beta))=0,$$

which is equivalent to

$$\frac{\Phi(u)-\Phi(z_\beta)}{\phi(u)}=\frac{u^2-z_\beta^2}{u(3z_\beta^2+u^2)}.$$

We define

$$L(u):=\frac{\Phi(u)-\Phi(z_\beta)}{\phi(u)}\qquad R(u):=\frac{u^2-z_\beta^2}{u(3z_\beta^2+u^2)},$$

and we have $L(z_\beta)=R(z_\beta)$. We will prove that if $u\neq z_\beta$, then $L(u)=R(u)$ has at most one zero root.

Then, we compute the gradient of $T(u):=e^{-u^2/2}(L(u)-R(u))$ as follows:

$$\begin{aligned}T'(u)&=e^{-u^2/2}\left(L'(u)-R'(u)-uL(u)+uR(u)\right)\\&=e^{-u^2/2}\left(1+uR(u)-R'(u)\right)\\&=\frac{2u^6+(8z_\beta^2+1)u^4+6z_\beta^2(z_\beta^2-1)u^2-3z_\beta^4}{e^{u^2/2}u^2(u^2+3z_\beta^2)^2}.\end{aligned}$$

Since the denominator $e^{u^2/2}u^2(u^2+3z_\beta^2)^2$ is positive, we focus on the numerator and define the following function:

$$P(s)=2s^3+(8z_\beta^2+1)s^2+6z_\beta^2(z_\beta^2-1)s-3z_\beta^4.$$

Since $P''(s) = 12s + 16z_\beta^2 + 2 > 0$, $P(0) = -3z_\beta^4 < 0$, and $\lim_{s\to\infty} P(s) = +\infty$, we can deduce that $P(s)$ is strictly convex in s and changes sign exactly once on $(0,\infty)$. Therefore, there exists a constant c_2, such that $T'(u) \le 0$ when $u \in (0, c_2)$ and $T'(u) \ge 0$ when $u \in (c_2, \infty)$. Therefore, $T(u)$ has at most two zero roots (denoted by (u_1, u_2)), and one of them is $u = z_\beta$.

For the root $u_1 = z_\beta$, we have $P(z_\beta^2) = 8z_\beta^4(2z_\beta^2 - 1)$ and hence $\mathrm{sign}(T'(z_\beta)) = \mathrm{sign}(P(z_\beta^2)) = \mathrm{sign}(2z_\beta^2 - 1)$. Then, we analyze the zero root $u_2 \ne z_\beta$ under different z_β:

1. If $z_\beta = 1/\sqrt{2}$, we have $T'(z_\beta) = 0$ and hence $u_2 = u_1 = z_\beta$.
2. If $z_\beta < 1/\sqrt{2}$, we have $T'(z_\beta) < 0$, which implies that $u_2 > z_\beta$ and the corresponding θ is $\theta_2 < 0$.
3. If $z_\beta > 1/\sqrt{2}$, we have $T'(z_\beta) > 0$, which implies that $u_2 < z_\beta$ and the corresponding θ is $\theta_2 > 0$.

Lastly, we analyze the sign of $H(\theta)$. Before proceeding, we first clarify several terms:

$$\lim_{\theta\to-1} H(\theta) = -\infty < 0, \quad H(1) = \frac{1}{2} > 0, \quad H(0) = 0, \quad H'(0) = 0, \quad H''(0) = 2(1 - 2z_\beta^2)\phi(z_\beta).$$

Then, we consider three cases based on the value of z_β:

1. If $z_\beta = 1/\sqrt{2}$, then $H(\theta) = 0$ only has one zero root $\theta = 0$. Since $H(-1) < 0$ and $H(1) > 0$, we can deduce that $H(\theta) < 0$ when $\theta \in (-1, 0)$ and $H(\theta) > 0$ when $\theta \in (0, 1)$.
2. If $z_\beta < 1/\sqrt{2}$, then $H(\theta) = 0$ has two zero roots $\theta_1 = 0$ and $\theta_2 < 0$. Since $H(0) = 0$, $H'(0) = 0$, and $H''(0) > 0$, we can deduce that $H(\theta)$ is positive around $\theta = 0$. Thus, $H(\theta)$ changes the sign if and only if $\theta = \theta_2$. Therefore, $H(\theta) < 0$ for $\theta \in (-1, \theta_2)$ and $H(\theta) > 0$ for $\theta \in (\theta_2, 1)$.
3. If $z_\beta > 1/\sqrt{2}$, then $H(\theta) = 0$ has two zero roots $\theta_1 = 0$ and $\theta_2 > 0$. Since $H(0) = 0$, $H'(0) = 0$, and $H''(0) < 0$, we can deduce that $H(\theta)$ is negative around $\theta = 0$. Similarly, we have $H(\theta) < 0$ for $\theta \in (-1, \theta_2)$ and $H(\theta) > 0$ for $\theta \in (\theta_2, 1)$.

Above all, there exists a constant $\tilde{\theta}$ such that $H(\theta) < 0$ for $\theta \in (-1, \tilde{\theta})$ and $H(\theta) > 0$ for $\theta \in (\tilde{\theta}, 1)$. Therefore, if $\theta \in (-1, \tilde{\theta})$, then $\rho(\theta, L, \beta)$ increases in L when L is sufficiently large; if $\theta \in (\tilde{\theta}, 1)$, then $\rho(\theta, L, \beta)$ decreases in L when L is sufficiently large. Moreover, as for the threshold $\tilde{\theta}$, we have $\tilde{\theta} < 0$ if $z_\beta < 1/\sqrt{2}$, $\tilde{\theta} > 0$ if $z_\beta > 1/\sqrt{2}$, and $\tilde{\theta} = 0$ if $z_\beta = 1/\sqrt{2}$. □

4.2. Impact of Correlation Coefficient θ

We now analyze how the competitive ratio varies with θ. To obtain clean analytical results, we focus on the limiting regime $L \to \infty$, in which Proposition 1 gives $f(\theta, L) \to 1$ and the competitive ratio converges to

$$\Gamma(\theta, \beta) := \lim_{L\to\infty} \rho(\theta, L, \beta) = \frac{\phi(z_\infty) + z_\infty\left(\Phi(z_\infty) - \Phi(z_\beta)\right)}{\phi(z_\beta)}.$$

The following proposition characterizes how $\Gamma(\theta, \beta)$ varies with θ; Figure 2a illustrates the result, and Figure 2b shows that the same qualitative behavior holds under finite lead times.

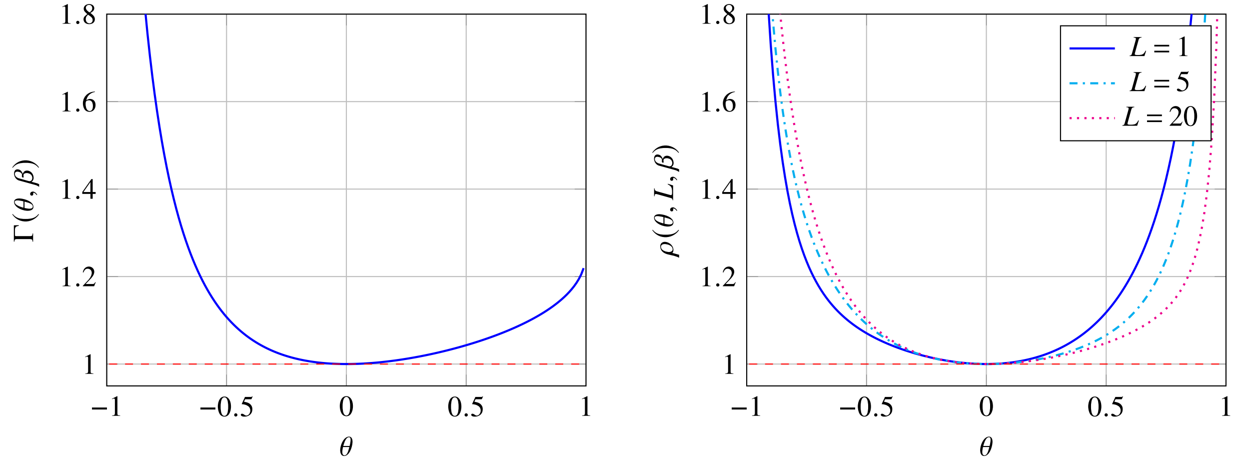

(a) $L \to \infty$ (b) $L \in \{1,5,20\}$

Figure 2 Competitive ratios $\Gamma(\theta,\beta)$ and $\rho(\theta,L,\beta)$ as functions of θ when $h=1$ and $b=3$.

Proposition 2. *$\Gamma(\theta,\beta)$ is decreasing in θ for $\theta<0$ and increasing in θ for $\theta\geq 0$, attaining its minimum at $\theta=0$. Moreover,*

$$\lim_{\theta\to-1}\Gamma(\theta,\beta)=+\infty, \qquad \lim_{\theta\to 1}\Gamma(\theta,\beta)=\phi(0)/\phi(z_\beta).$$

Proposition 2 shows that the limiting competitive ratio $\Gamma(\theta,\beta)$ is U-shaped in θ, increasing in $|\theta|$: stronger correlation in either direction leads to greater misspecification loss. Figure 2b confirms that this pattern persists under finite lead times.

Proof of Proposition 2. We compute the gradient of $\Gamma(\theta,\beta)$ as follows:

$$\frac{\partial}{\partial\theta}\Gamma(\theta,\beta)=\frac{z_\beta}{\phi(z_\beta)}\left(\Phi(z_\infty)-\Phi(z_\beta)\right)\left(-\frac{1}{\sqrt{(1+\theta)^3(1-\theta)}}\right).$$

When $\theta<0$, we have $\frac{\partial}{\partial\theta}\Gamma(\theta,\beta)\leq 0$ and hence $\Gamma(\theta,\beta)$ decreases in θ; when $\theta>0$, we have $\frac{\partial}{\partial\theta}\Gamma(\theta,\beta)\geq 0$ and hence $\Gamma(\theta,\beta)$ increases in θ.

Then, we analyze the limits. As θ tends to -1, we have $z_\infty\to\infty$, and hence

$$\lim_{\theta\to-1}\Gamma(\theta,\beta)=\frac{z_\infty(1-\Phi(z_\beta))}{\phi(z_\beta)}=\infty.$$

As θ tends to 1, we have $z_\infty\to 0$, and hence $\lim_{\theta\to 1}\Gamma(\theta,\beta)=\frac{\phi(0)}{\phi(z_\beta)}$.

In the following theorem, we compare the misspecification loss under positive and negative correlations. To this end, we define $\xi(|\theta|,L,\beta):=\rho(|\theta|,L,\beta)/\rho(-|\theta|,L,\beta)$. Figure 3 illustrates how ξ varies with $|\theta|$.

Theorem 2 **(Impact of Correlation Sign)**. *There exists a constant $\kappa\in[0,1)$ such that*

(i) For any $\theta\in(0,\kappa)$, there exists $\ell(\theta)$ satisfying $\rho(\theta,L,\beta)\geq\rho(-\theta,L,\beta)$ for $L\geq\ell(\theta)$.

(ii) For any $\theta\in(\kappa,1)$, there exists $\ell(\theta)$ satisfying $\rho(\theta,L,\beta)\leq\rho(-\theta,L,\beta)$ for $L\geq\ell(\theta)$.

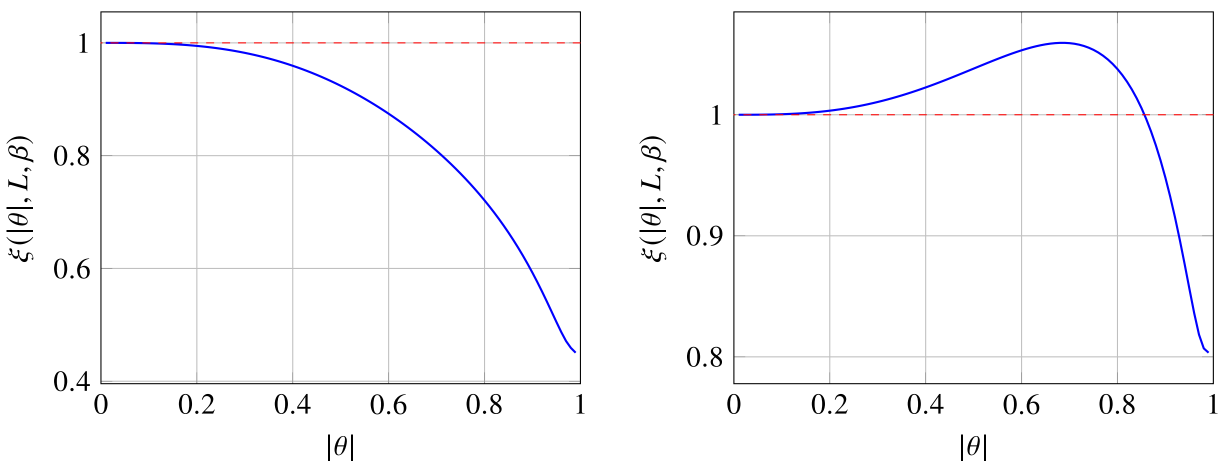

(a) $b = 10$ (b) $b = 30$

Figure 3 $\xi(|\theta|, L, \beta)$ as a function of $|\theta|$ when $h = 1$, $b \in \{10, 30\}$ and $L = 50$.

Moreover, $\kappa = 0$ *when* $z_\beta \leq \sqrt{3}$ *and* $\kappa > 0$ *when* $z_\beta > \sqrt{3}$.

Theorem 2 and Figure 3 show that positive correlation is more costly than negative correlation when $|\theta|$ is small, while this ordering reverses when $|\theta|$ is large. Moreover, the first region disappears when the backlog penalty is low. This has a practical implication: when the correlation level is small (large, resp.), sellers should be more cautious about model misspecification if demand is positively (negatively, resp.) correlated. Furthermore, when the backlog penalty b is low (i.e., $z_\beta \leq \sqrt{3}$, so that $\kappa = 0$), positive correlation is always more costly than negative correlation, regardless of $|\theta|$.

To understand the intuition, consider the limiting case $L \to \infty$, where the information factor vanishes and only the correlation factor remains. When $|\theta|$ is large, the correlation factor dominates under negative correlation, because the strong mean-reversion pattern induces a far more inaccurate variance estimator $\mathbb{V}[\hat{\Delta}_{t,L}]$. When $|\theta|$ is small and b is high, however, the correlation factor is more significant under positive correlation: even a modest positive correlation causes systematic understocking (i.e., $\tilde{q}_t^* \geq \hat{q}_t^*$), which is heavily penalized by a large backlog cost.

Proof of Theorem 2. According to Proposition 1, we have $\lim_{L\to\infty} f(\theta, L) = 1$ and $\lim_{L\to\infty} g(\theta, L, \beta) = G(z_\infty)/\phi(z_\beta)$. For any $\theta \in (0,1)$, we can derive the following limit:

$$\lim_{L\to\infty} \frac{\rho(\theta, L, \beta)}{\rho(-\theta, L, \beta)} = \frac{G\left(z_\beta\sqrt{\frac{1-\theta}{1+\theta}}\right)}{G\left(z_\beta\sqrt{\frac{1+\theta}{1-\theta}}\right)}.$$

Therefore, it is sufficient to analyze the sign of $G\left(z_\beta\sqrt{\frac{1-\theta}{1+\theta}}\right) - G\left(z_\beta\sqrt{\frac{1+\theta}{1-\theta}}\right)$.

We define $F(s) := G(e^{-s}z_\beta) - G(e^s z_\beta)$ and will analyze its property for $s \in (0, \infty)$. First, we have

$$F(s) = 0 \quad \Leftrightarrow \quad \phi(e^{-s}z_\beta) + e^{-s}z_\beta\left(\Phi(e^{-s}z_\beta) - \Phi(z_\beta)\right) = \phi(e^s z_\beta) + e^s z_\beta\left(\Phi(e^s z_\beta) - \Phi(z_\beta)\right)$$

$$\Leftrightarrow \quad \Phi(z_\beta) = \frac{\phi(e^{-s}z_\beta) + e^{-s}z_\beta\Phi(e^{-s}z_\beta) - \phi(e^s z_\beta) - e^s z_\beta \Phi(e^s z_\beta)}{e^{-s}z_\beta - e^s z_\beta} =: Q(s).$$

Furthermore, we have $F(s) = (e^s z_\beta - e^{-s}z_\beta)\left(\Phi(z_\beta) - Q(s)\right) = (e^s z_\beta - e^{-s} z_\beta)\,(\beta - Q(s))$. Then, we analyze the property of $Q(s)$.

First, since $\phi'(x) = -x\phi(x)$, we can compute the gradient of $Q(s)$ as follows:

$$Q'(s) = \frac{z_\beta}{\left(e^{-s}z_\beta - e^s z_\beta\right)^2}\left((e^{-s}+e^s)[\phi(e^{-s}z_\beta) - \phi(e^s z_\beta)] - 2z_\beta[\Phi(e^s z_\beta) - \Phi(e^{-s}z_\beta)]\right).$$

Then, we define

$$M(s) := (e^{-s}+e^s)[\phi(e^{-s}z_\beta) - \phi(e^s z_\beta)] - 2z_\beta[\Phi(e^s z_\beta) - \Phi(e^{-s}z_\beta)],$$

whose gradient can be computed as follows:

$$M'(s) = (e^s - e^{-s})\left[\left(1 - e^{-2s}z_\beta^2\right)\phi\left(e^{-s}z_\beta\right) - \left(1 - e^{2s}z_\beta^2\right)\phi\left(e^s z_\beta\right)\right].$$

If $z_\beta \le 1$ or $s > \log z_\beta$, then $M'(s) > 0$. In the following, we analyze the case with $z_\beta > 1$ and $s \in (0, \log z_\beta)$. In this case, we can deduce that $e^s z_\beta > e^{-s} z_\beta > 1$ and

$$\begin{aligned}
M'(s) = 0 \quad &\Leftrightarrow \quad \frac{e^{2s}z_\beta^2 - 1}{e^{-2s}z_\beta^2 - 1} = \frac{\phi(e^{-s}z_\beta)}{\phi(e^s z_\beta)} = e^{\frac{e^{2s}z_\beta^2 - e^{-2s}z_\beta^2}{2}} \\
&\Leftrightarrow \quad \log\left(e^{2s}z_\beta^2 - 1\right) - \log\left(e^{-2s}z_\beta^2 - 1\right) - \frac{e^{2s}z_\beta^2 - e^{-2s}z_\beta^2}{2} = 0 \\
&\Leftrightarrow \quad N(x) := \log(xz_\beta^2 - 1) - \log(z_\beta^2/x - 1) - \frac{z_\beta^2 x - z_\beta^2/x}{2} = 0, \quad \text{with } x = e^{2s} \in (1, z_\beta^2).
\end{aligned}$$

Furthermore, we have $M'(s) = (e^s - e^{-s})(e^{-2s}z_\beta^2 - 1)\phi(e^{-s}z_\beta)\left(e^{N(e^{2s})} - 1\right)$, so $\operatorname{sign}[M'(s)] = \operatorname{sign}[N(e^{2s})]$ for $s \in (0, \log z_\beta)$.

Then, we compute the gradient of $N(x)$ as follows:

$$N'(x) = \frac{z_\beta^2}{2x^2(z_\beta^2 x - 1)(z_\beta^2 - x)}\left(-3x(1+x^2) + (x^4 + 6x^2 + 1)z_\beta^2 - (x^3 + x)z_\beta^4\right).$$

Thus, the sign of $N'(x)$ is the same as $\left(-3x(1+x^2) + (x^4+6x^2+1)z_\beta^2 - (x^3+x)z_\beta^4\right)$. By some algebra, we can derive the following results:

1. If $z_\beta \in (1, \sqrt{3}]$, then $N'(x) \ge 0$ for $x \in (1, z_\beta^2)$. Since $N(1) = 0$, we have $N(x) \ge 0$ for $x \in (1, z_\beta^2)$ and hence $M'(s) \ge 0$ for $s \in (0, \log z_\beta)$.
2. If $z_\beta \in (\sqrt{3}, \infty)$, then there exists a constant $\kappa_1 \in (1, z_\beta^2)$ such that $N'(x) < 0$ for $x \in (1, \kappa_1)$ and $N'(x) \ge 0$ for $x \in [\kappa_1, z_\beta^2)$. Since $N(1) = 0$ and $N(z_\beta^2) = +\infty$, we can deduce that there exists a constant $\kappa_2 \in (1, z_\beta^2)$ such that $N(x) < 0$ for $x \in (1, \kappa_2)$ and $N(x) \ge 0$ for $x \in [\kappa_2, z_\beta^2)$. Thus, $M'(s) < 0$ for $s \in (0, \frac{1}{2}\log \kappa_2)$ and $M'(s) \ge 0$ for $s \in [\frac{1}{2}\log\kappa_2, \log z_\beta)$.

Then, we are ready to analyze the property of $F(s)$.

1. If $z_\beta \in (0, \sqrt{3}]$, we have $M'(s) \ge 0$ for $s > 0$, implying that $M(s)$ always increases in s. Since $M(0) = 0$, we have $M(s) \ge 0$ and hence $Q'(s) \ge 0$ for $s > 0$. Therefore, we have $\beta - Q(s) \le \beta - \lim_{t\to 0} Q(t) = 0$ and hence $F(s) \le 0$.
2. If $z_\beta \in (\sqrt{3}, \infty)$, we have $M'(s) < 0$ for $s \in (0, \frac{1}{2}\log \kappa_2)$ and $M'(s) > 0$ for $s \in [\frac{1}{2}\log \kappa_2, \infty)$, implying that $M(s)$ first decreases and then increases. Since $M(0) = 0$ and $M(\infty) = +\infty$, we can deduce that $M(s)$ (and hence $Q'(s)$) is first negative and then positive as s increases, such that $\beta - Q(s)$ first increases and then decreases. Therefore, there exists a threshold $\kappa_3 \in (0, \infty)$ such that $F(s) > 0$ for $s \in (0, \kappa_3)$ and $F(s) \le 0$ for $s \in [\kappa_3, \infty)$.

Since there exists an one-to-one mapping between s and θ such that $e^{-s} = \sqrt{(1-\theta)/(1+\theta)}$, we can deduce that there exists a constant $\kappa \in [0, 1)$ such that

1. For $\theta \in (0, \kappa)$, we have $\lim_{L\to\infty} \frac{\rho(\theta,L,\beta)}{\rho(-\theta,L,\beta)} > 1$.
2. For $\theta \in [\kappa, 1)$, we have $\lim_{L\to\infty} \frac{\rho(\theta,L,\beta)}{\rho(-\theta,L,\beta)} \le 1$.

Moreover, we have $\kappa = 0$ when $z_\beta \le \sqrt{3}$ and $\kappa > 0$ when $z_\beta > \sqrt{3}$. □

4.3. Impact of Backlog Penalty b

We now investigate how the competitive ratio varies with the backlog penalty b. The following theorem establishes monotonicity in b and characterizes the limiting behavior as $b \to \infty$; Figure 4 provides an illustration.

THEOREM 3 **(Impact of Backlog Penalty** b**).** *Recall that* $\beta = b/(h+b)$*. We have:*

(i) $\rho(\theta, L, \beta)$ *is increasing in* b*.*

(ii) As $b \to \infty$ *(equivalently,* $\beta \to 1$*),*

$$\lim_{b\to\infty} \rho(\theta, L, \beta) = \begin{cases} \dfrac{(1-\theta)\sqrt{L+1}}{\sqrt{(L+1)(1-\theta^2) - 2\theta(1-\theta^{L+1}) - (\theta-\theta^{L+2})^2}} & \theta \in (-1, 0], \\ +\infty & \theta \in (0, 1). \end{cases}$$

Theorem 3 and Figure 4 show that the misspecification loss increases in b: as the penalty for unmet demand grows, the cost of ignoring the correlation of future demands becomes more severe. In the limit $b \to \infty$, the competitive ratio remains finite under negative correlation but diverges under positive correlation. Since $f(\theta, L)$ is independent of b, this asymmetry is driven entirely by the correlation factor. Under positive correlation, the misspecified policy systematically understocks ($\tilde{q}_t^* \ge \hat{q}_t^*$), so every stockout incurs the large backlog penalty b, causing the cost ratio to grow without bound. Under negative correlation, the misspecified policy systematically overstocks ($\tilde{q}_t^* \le \hat{q}_t^*$), so stockouts become increasingly rare as $b \to \infty$ and the competitive ratio converges to a finite limit.

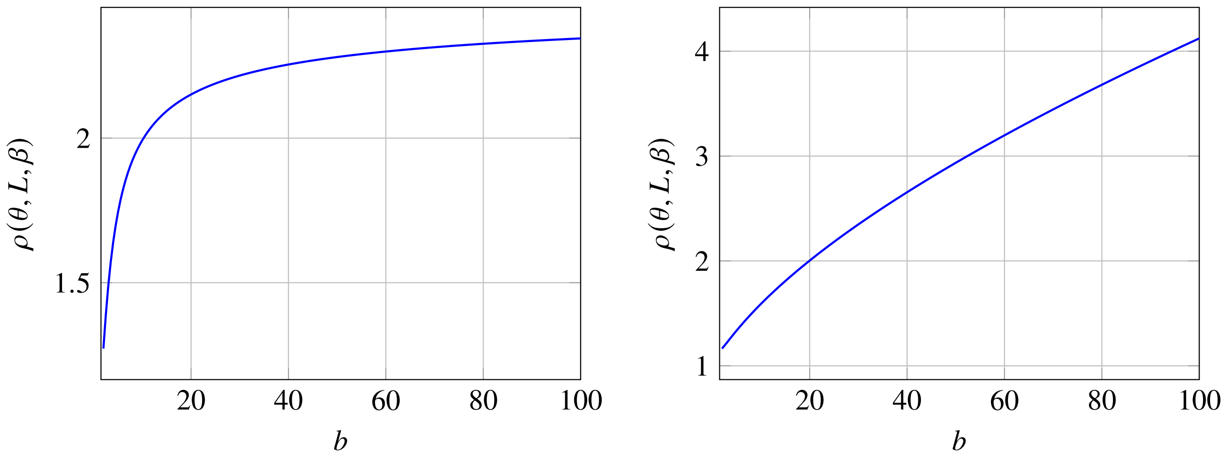

(a) $\theta = -0.8$ (b) $\theta = 0.8$

Figure 4 $\rho(\theta, L, \beta)$ as a function of b when $h = 1$, $\theta \in \{-0.8, 0.8\}$ and $L = 10$.

Proof of Theorem 3. Recall that $\beta = b/(h+b)$ and $\rho(\theta, L, \beta) = f(\theta, L) \cdot g(\theta, L, \beta)$. It suffices to analyze the property regarding β. Since $f(\theta, L)$ is independent of β, we only focus on $g(\theta, L, \beta)$, which can be rewritten as follows:

$$g(\theta, L, \beta) = \frac{\phi(c_3 z_\beta) + c_3 z_\beta(\Phi(c_3 z_\beta) - \Phi(z_\beta))}{\phi(z_\beta)},$$

where $c_3 = \delta_E/\delta_T$ is independent of β. Since z_β increases in β, it suffices to show that $g(\theta, L, \beta)$ increases in z_β. We compute its gradient as follows:

$$\begin{aligned}
\frac{\partial}{\partial z_\beta} g(\theta, L, \beta) &= \frac{1}{\phi(z_\beta)} \left(c_3(1+z_\beta^2)(\Phi(c_3 z_\beta) - \Phi(z_\beta)) + z_\beta \phi(c_3 z_\beta) - c_3 z_\beta \phi(z_\beta) \right) \\
&= \frac{1}{\phi(z_\beta)} \left(c_3(1+z_\beta^2) \int_{z_\beta}^{c_3 z_\beta} \phi(u) du + c_3 z_\beta^2 \int_{z_\beta}^{c_3 z_\beta} \frac{-w^2\phi(w) - \phi(w)}{w^2} dw \right) \\
&= \frac{c_3}{\phi(z_\beta)} \int_{z_\beta}^{c_3 z_\beta} \frac{u^2 - z_\beta^2}{u^2} \phi(u) du.
\end{aligned}$$

When $c_3 \leq 1$, since $c_3 z_\beta \leq z_\beta$ and $u^2 - z_\beta^2 \leq 0$ for $u \in [c_3 z_\beta, z_\beta]$, we can deduce that $\frac{\partial}{\partial z_\beta} g(\theta, L, \beta) \geq 0$; when $c_3 \geq 1$, since $c_3 z_\beta \geq z_\beta$ and $u^2 - z_\beta^2 \geq 0$ for $u \in [z_\beta, c_3 z_\beta]$, we can deduce that $\frac{\partial}{\partial z_\beta} g(\theta, L, \beta) \geq 0$. Therefore, $g(\theta, L, \beta)$ increases in z_β, which implies that $\rho(\theta, L, \beta)$ increases in β.

In the second part, we compute the limiting value of $\rho(\theta, L, \beta)$ as $\beta \to 1$ (i.e., $z_\beta \to \infty$). Define $\bar{\Phi}(x) := 1 - \Phi(x)$. Since the information factor $f(\theta, L)$ is independent of β, we mainly analyze the limit of $g(\theta, L, \beta)$:

$$\begin{aligned}
\lim_{b \to \infty} g(\theta, L, \beta) &= \lim_{z_\beta \to \infty} \frac{1}{\phi(z_\beta)} \left(\phi(c_3 z_\beta) + c_3 z_\beta (\bar{\Phi}(z_\beta) - \bar{\Phi}(c_3 z_\beta)) \right) \\
&= \lim_{z_\beta \to \infty} \frac{1}{\phi(z_\beta)} \left(\phi(c_3 z_\beta) + c_3 \phi(z_\beta) - \frac{\phi(z_\beta)}{z_\beta^2} + o\left(\frac{\phi(z_\beta)}{z_\beta^2} \right) - \phi(c_3 z_\beta) + \frac{\phi(c_3 z_\beta)}{c_3^2 z_\beta^2} + o\left(\frac{\phi(c_3 z_\beta)}{c_3^2 z_\beta^2} \right) \right) \\
&= c_3 + \lim_{z_\beta \to \infty} \left(\frac{\phi(c_3 z_\beta)}{\phi(z_\beta) c_3^2 z_\beta^2} + o\left(\frac{\phi(c_3 z_\beta)}{\phi(z_\beta) c_3^2 z_\beta^2} \right) \right),
\end{aligned}$$

where the second equation is due to the Mills' ratio (see Mills 1926):

$$\frac{\bar{\Phi}(x)}{\phi(x)} = \frac{1}{x} - \frac{1}{x^3} + o\left(\frac{1}{x^3}\right).$$

According to Lemma 2, we can deduce that $c_3 \geq 1$ when $\theta \leq 0$ and $c_3 < 1$ when $\theta > 0$. Then, we compute the limiting value in two cases.

Case I: $\theta \leq 0$**.** In this case, since $c_3 \geq 1$, we have

$$\lim_{z_\beta \to \infty} \frac{\phi(c_3 z_\beta)}{\phi(z_\beta) c_3^2 z_\beta^2} = \lim_{z_\beta \to \infty} \frac{e^{-\frac{1}{2}(c_3^2-1)z_\beta^2}}{c_3^2 z_\beta^2} = 0,$$

which implies $\lim_{\beta \to 1} g(\theta, L, \beta) = c_3$.

Case II: $\theta > 0$**.** In this case, since $c_3 < 1$, we have

$$\lim_{z_\beta \to \infty} \frac{\phi(c_3 z_\beta)}{\phi(z_\beta) c_3^2 z_\beta^2} = \lim_{z_\beta \to \infty} \frac{e^{-\frac{1}{2}(c_3^2-1)z_\beta^2}}{c_3^2 z_\beta^2} = \infty,$$

which implies $\lim_{\beta \to 1} g(\theta, L, \beta) = \infty$.

Above all, the limiting value of $\rho(\theta, L, \beta)$ is

$$\lim_{b \to \infty} \rho(\theta, L, \beta) = \begin{cases} \frac{(1-\theta)\sqrt{L+1}}{\sqrt{(L+1)(1-\theta^2) - 2\theta(1-\theta^{L+1}) - (\theta - \theta^{L+2})^2}} & \theta \in (-1, 0), \\ +\infty & \theta \in (0, 1). \end{cases} \qquad \square$$

5. Concluding Remarks

Although the assumption of independent demand is ubiquitous in the inventory management literature, real demand processes are often serially correlated, and practitioners who ignore this dependence face a misspecification cost that is not well understood. This paper takes a step toward filling this gap by deriving exact, closed-form expressions for the misspecification loss in a periodic-review system with AR(1) demand, measured via a competitive ratio. Our analysis reveals that the interplay between lead time and correlation is more subtle than previously recognized: rather than uniformly worsening performance, a longer lead time can actually mitigate misspecification loss when demand is positively correlated. Beyond lead time, we find that misspecification loss is U-shaped in the correlation coefficient, so that stronger dependence in either direction is more costly to ignore, and that the loss always increases in the backlog penalty, diverging to infinity under positive correlation but remaining bounded under negative correlation. Together, these results offer practitioners useful guidance on when the added complexity of explicitly modeling demand correlation is warranted, and when a simpler i.i.d. approximation carries only a modest cost. While we allow negative orders for analytical tractability, Appendix A reports numerical experiments with nonnegative order constraints, which exhibit similar patterns to Section 4 and confirm the robustness of our results.

Appendix A: Numerical Experiments

This section presents numerical experiments to test robustness of the analytical results of Section 4. We consider a practical setting that departs from the base model by imposing a non-negativity constraint on order quantities, and analyze the misspecification cost of a modified myopic policy in which any negative order quantity is replaced by zero. In particular, under the true demand model, the order quantity in period t is $x_t^{*,+} = \max\{x_t^*, 0\}$, where x_t^* is computed in Section 3.1, and the corresponding long-run average cost is denoted by $V^{*,+}$. Similarly, under the misspecified model, the order quantity in period t is $\hat{x}_t^{*,+} = \max\{\hat{x}_t^*, 0\}$, where $\hat{x}_t^*$ is computed in Section 3.2, and the corresponding long-run average cost is denoted by $\hat{V}^+$. We define the competitive ratio $\rho_+(\theta, L, \beta) = \hat{V}^+ / V^{*,+}$. Similarly, we define the corresponding function $\xi_+(|\theta|, L, \beta) = \rho_+(|\theta|, L, \beta)/\rho_+(-|\theta|, L, \beta)$. We set $\mu = 5$, $\sigma = 1$, and $h = 1$. Each data point is computed by averaging over 100 independent sample paths, each consisting of 50,000 periods.

Impact of lead time. Figure 5 plots the simulated competitive ratio $\rho_+(\theta, L, \beta)$ as a function of L for $\theta \in \{-0.95, 0.95\}$, confirming the monotonicity established in Theorem 1: the ratio increases in L under strong negative correlation and decreases in L under strong positive correlation.

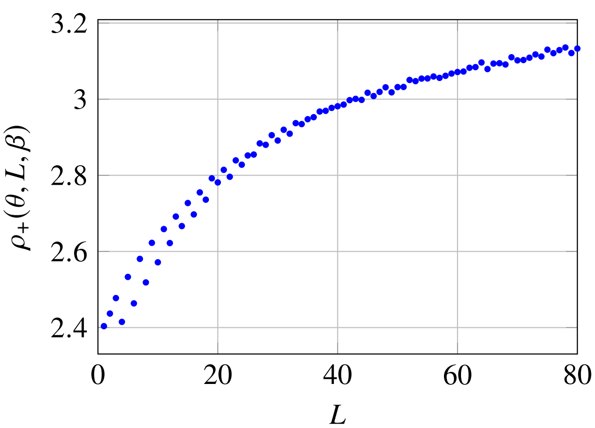

(a) $\theta = -0.95$

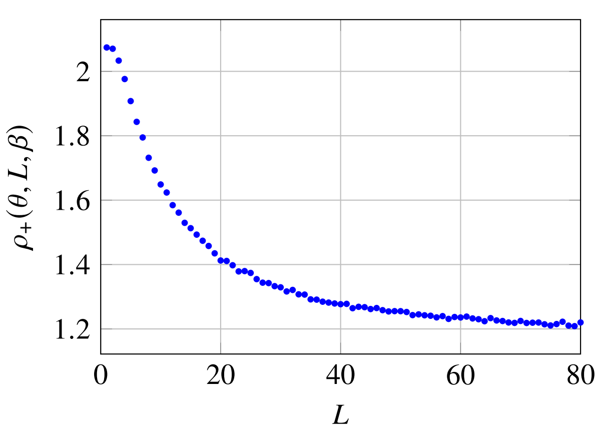

(b) $\theta = 0.95$

Figure 5 Numerical results: Competitive ratio $\rho_+(\theta, L, \beta)$ as functions of L when $\mu = 5$, $\sigma = 1$, $h = 1$, $b = 3$, and $\theta \in \{-0.95, 0.95\}$.

Impact of correlation sign. Figure 6 plots the simulated ratio $\xi_+(|\theta|, L, \beta)$ as a function of $|\theta|$ for $b \in \{10, 30\}$, confirming the threshold structure established in Theorem 2: ξ_+ exceeds one for small $|\theta|$ and falls below one for large $|\theta|$, with the crossover point shifting as b increases.

Impact of backlog penalty. Figure 7 plots the simulated competitive ratio $\rho_+(\theta, L, \beta)$ as a function of b for $\theta \in \{-0.8, 0.8\}$, confirming the monotonicity established in Theorem 3: the ratio increases in b in both cases, and remains bounded under negative correlation while growing rapidly under positive correlation.

According to Figures 5-7, even if negative order quantities are prohibited, numerical experiments still support the results in Theorems 1-3.

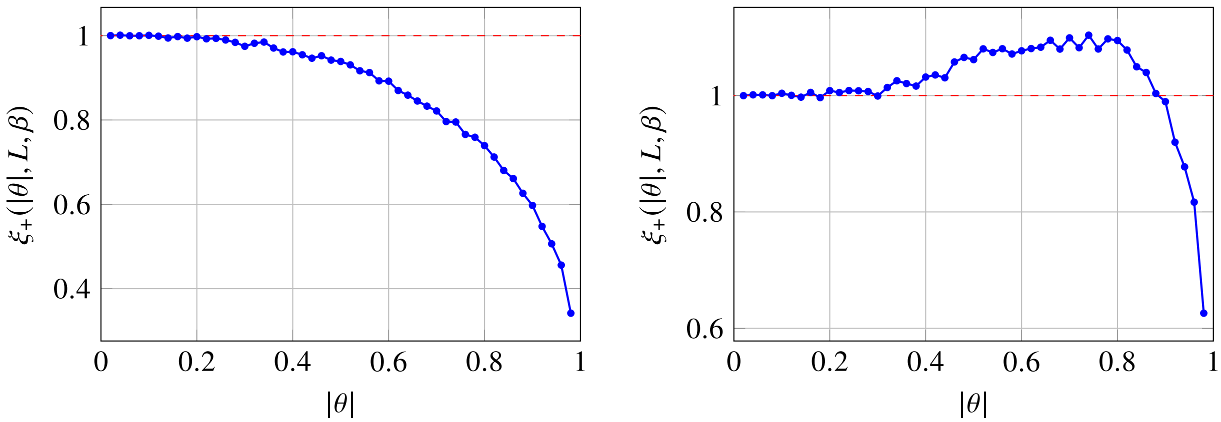

(a) $b = 10$ (b) $b = 30$

Figure 6 Numerical results: $\xi_+(|\theta|, L, \beta)$ as a function of $|\theta|$ when $\mu = 5$, $\sigma = 1$, $h = 1$, $b \in \{10, 30\}$ and $L = 50$.

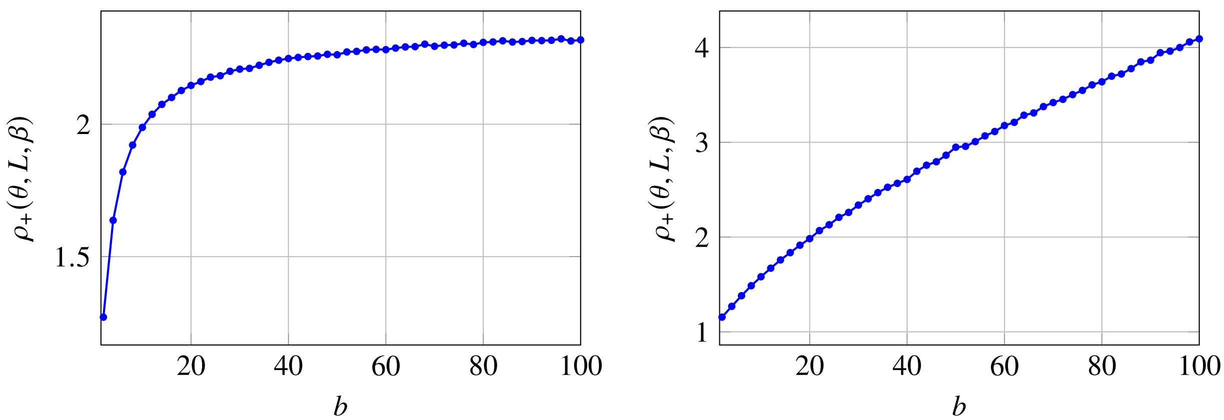

(a) $\theta = -0.8$ (b) $\theta = 0.8$

Figure 7 Numerical results: $\rho_+(\theta, L, \beta)$ as a function of b when $h = 1$, $\theta \in \{-0.8, 0.8\}$ and $L = 10$.

References

Aviv Y (2003) A time-series framework for supply-chain inventory management. *Operations Research* 51(2):210–227.

Brockwell PJ, Davis RA (2002) *Introduction to Time Series and Forecasting* (Springer).

Caldentey R, Giloni A, Hurvich C (2022) Performance bound for myopic order-up-to inventory policies under stationary demand processes. *Operations Research Letters* 50(5):588–595.

Chao X, Gong X, Shi C, Yang C, Zhang H, Zhou SX (2018) Approximation algorithms for capacitated perishable inventory systems with positive lead times. *Management Science* 64(11):5038–5061.

Chao X, Gong X, Shi C, Zhang H (2015) Approximation algorithms for perishable inventory systems. *Operations Research* 63(3):585–601.

Chen F, Drezner Z, Ryan JK, Simchi-Levi D (2000) Quantifying the bullwhip effect in a simple supply chain: The impact of forecasting, lead times, and information. *Management Science* 46(3):436–443.

Goldberg DA, Reiman MI, Wang Q (2021) A survey of recent progress in the asymptotic analysis of inventory systems. *Production and Operations Management* 30(6):1718–1750.

Hihat M, Gaïffas S, Garrigos G, Bussy S (2023) Online inventory problems: Beyond the i.i.d. setting with online convex optimization. *Proceedings of 37th Conference on Neural Information Processing Systems*, 20421–20440.

Johnson GD, Thompson HE (1975) Optimality of myopic inventory policies for certain dependent demand processes. *Management Science* 21(11):1303–1307.

Lee HL, Padmanabhan V, Whang S (1997) Information distortion in a supply chain: The bullwhip effect. *Management Science* 43(4):546–558.

Lee HL, So KC, Tang CS (2000) The value of information sharing in a two-level supply chain. *Management science* 46(5):626–643.

Levi R, Pál M, Roundy RO, Shmoys DB (2007) Approximation algorithms for stochastic inventory control models. *Mathematics of Operations Research* 32(2):284–302.

Mills JP (1926) Table of the ratio: Area to bounding ordinate, for any portion of normal curve. *Biometrika* 18(3/4):395–400.

Qi M, Shen ZJ, Zheng Z (2024) Learning newsvendor problems with intertemporal dependence and moderate non-stationarities. *Production and Operations Management* 33(5):1196–1213.

Song JSJ (2023) *Research Handbook on Inventory Management* (Edward Elgar Publishing).

Zipkin PH (2000) *Foundations of Inventory Management* (McGraw-Hill).