
Management Science

Publication details, including instructions for authors and subscription information:
http://pubsonline.informs.org

# Profit Estimation Error in the Newsvendor Model Under a Parametric Demand Distribution

Andrew F. Siegel , Michael R. Wagner

http://pubsonline.informs.org/journal/mnsc

MANAGEMENT SCIENCE
Vol. 67, No. 8, August 2021, pp. 4863–4879
ISSN 0025-1909 (print), ISSN 1526-5501 (online)

# Profit Estimation Error in the Newsvendor Model Under a Parametric Demand Distribution

**Andrew F. Siegel,**[a,b] **Michael R. Wagner**[a]

[a] Information Systems and Operations Management, Michael G. Foster School of Business, University of Washington, Seattle, Washington 98195; [b] Finance and Business Economics, Department of Statistics, University of Washington, Seattle, Washington 98195

**Contact:** asiegel@uw.edu, https://orcid.org/0000-0002-2144-2591 (AFS); mrwagner@uw.edu, https://orcid.org/0000-0003-2077-3564 (MRW)

**Received:** June 16, 2019
**Revised:** December 6, 2019; April 7, 2020; April 29, 2020
**Accepted:** May 17, 2020
**Published Online in Articles in Advance:** December 28, 2020

https://doi.org/10.1287/mnsc.2020.3766



**Abstract.** We consider the newsvendor model in which uncertain demand is assumed to follow a probabilistic distribution with known functional form but unknown parameters. These parameters are estimated, unbiasedly and consistently, from data. We show that the classic maximized expected profit expression exhibits a systematic expected estimation error. We provide an asymptotic adjustment so that the estimate of maximized expected profit is unbiased. We also study expected estimation error in the optimal order quantity, which depends on the distribution: (1) if demand is exponentially or normally distributed, the order quantity has zero expected estimation error; (2) if demand is log-normally distributed, there is a nonzero expected estimation error in the order quantity that can be corrected. Numerical experiments, for light- and heavy-tailed distributions, confirm our theoretical results.

**History:** Accepted by Vishal Gaur, operations management.
**Funding:** The first author gratefully acknowledges the support of the Grant I. Butterbaugh Professorship. The second author gratefully acknowledges the support of a Neal and Jan Dempsey Faculty Fellowship.
**Supplemental Material:** The online electronic companion is available at https://doi.org/10.1287/mnsc.2020.3766.

**Keywords:** newsvendor • estimation error • statistics • Fisher information

## 1. Introduction

In this paper, we show that the classic newsvendor model can systematically overestimate its maximized expected profit; in other words, the expected profit *perceived* by a decision maker can be larger on average than what actually occurs. In practice, the expected profit serves as a *forecast* for the future realized profit, and we effectively show that this forecast is *biased*. The American Production and Inventory Control Society (APICS), a worldwide professional organization founded in 1957 with more than 45,000 members, specifies that "a normal property of a good forecast is that it is not biased" (APICS 2019). Similarly, the Institute of Business Forecasting and Planning provides resources for measuring and correcting forecasting bias (Singh 2017). IBM has also provided guidance on correcting forecasting bias (Parks 2018), which is summarized as "error isn't always addressable but bias usually is. Consider focusing on bias first when trying to improve your forecasts." An article discussing the correction of forecasting bias has also appeared on LinkedIn (Bentzley 2017), and it concludes by stating, "By…driving relentlessly until the forecast has had the bias addressed…the organization can make the most of its efforts and will continue to improve the quality of its forecasts and the supply chain overall." Academic research has also identified the costs of positive and negative bias: Cassar and Gibson (2008) report, "Higher costs of obsolescence and inventory holding (Lee and Adam 1986, Watson 1987) and lower returns on capital investments may result from optimistic forecasts, while stock-out costs and reputation damage are likely consequences of pessimistic forecasts (Ittner and Larcker 1998, Durand 2003)." Furthermore, Lipson (2019) is a University of Virginia case study that teaches financial statement forecasting and states that forecasts should be unbiased (neither optimistic nor pessimistic) because biased forecasts lead to biased decisions. More informally, if a manager chronically overestimates profit, eventually it will be noticed that the errors cannot be solely attributed to forecasting error but rather to mistaken calculations (i.e., forecasting bias). In this paper, we provide an approach to correct these mistaken calculations in the context of the newsvendor model and produce bias-adjusted forecasts of profit.

As a concrete example of the need for unbiased forecasts, we summarize the case of a large Californian consumer electronics firm, disguised as "Leitax," as reported in Oliva and Watson (2009). This paper categorizes forecast biases as either intentional (driven by misaligned incentives, politics, etc.) or unintentional (resulting from procedural and informational blind spots). The authors discuss the implementation at Leitax of a coordination system to address these biases, which results in an increase of forecasting accuracy, defined as one minus the ratio of the absolute deviation of sales from forecast to the forecast itself from 58% to 88%. These forecasting improvements directly enhance operational outcomes: average on-hand inventory decreases from \$55 million to \$23 million, and excess and obsolescence costs decrease from an average of \$3 million per year to practically zero. However, this paper is silent on *statistical* bias in forecasts (a form of unintentional procedural bias), which is the focus of our research and can supplement the coordination system presented in Oliva and Watson (2009).

We study a scenario in which the demand distribution is specified by a finite number of parameters (e.g., mean, standard deviation) that are estimated from data. Even though we utilize unbiased and consistent estimators, when fed through the optimization operator of the newsvendor model, an incorrect estimation of the optimized expected profit materializes. We derive this asymptotic expected estimation error in closed form for a generic continuous demand distribution with $k$ parameters that are estimated from data. We also design an adjustment term, which accounts for estimation error distortions, that provides an unbiased estimation of the maximized profit (asymptotically to second order) that more accurately reflects actual realized profit. We also summarize experiments that show that (1) the expected estimation error is statistically significant, and (2) our adjusted profit has no statistically significant expected estimation error. Finally, we also consider expected estimation error in the order quantity and derive its general asymptotic functional form; in particular, we show that (1) when demand is exponentially or normally distributed, the order quantity has no expected estimation error, and (2) when demand is log-normally distributed, the order quantity has a expected estimation error that can be corrected.

We provide the following numerical example, motivated by valuing reactive capacity in a newsvendor context, which illustrates the practical benefits of our research.

**Example 1.** Consider a newsvendor situation in which the unit selling price $p = 100$, the unit procurement cost $c = 40$, and demand is exponentially distributed with a mean $\theta = 200$. Suppose that we have the following 10 sample demands from this distribution $\{217, 444, 148, 219, 251, 126, 28, 32, 210, 147\}$, which lead to an estimated mean $\hat{\theta} = 182.15$. Using an *estimated* exponential distribution with mean $\hat{\theta} = 182.15$, the optimal newsvendor order quantity is 167, which results in an estimated expected profit of \$4,253. Reactive capacity, as described in section 13.4 of Cachon and Terwiesch (2012), results in the opportunity to place a second order at a higher cost once demand is realized to satisfy any unmet demand from the first order; we assume that utilizing reactive capacity requires a fixed cost of \$5,500 and a per-unit premium of 20% so that the unit cost of reactive capacity is $(1 + 20\%)c = 48$. Applying the analysis from section 13.4 of Cachon and Terwiesch (2012) with the *estimated* exponential distribution with mean $\hat{\theta} = 182.15$, we calculate that we should order 33 units in the original order and utilize reactive capacity for any demand in excess of 33, which results in an expected profit of \$4,100. Therefore, we conclude that reactive capacity is not worth it (\$4,100 < \$4,253). But this would be wrong. Using the exact exponential distribution with mean $\theta = 200$, the nonreactive and reactive expected profits are \$4,670 and \$5,041, respectively, and it is beneficial to utilize the reactive capacity. However, these profit values are typically not available because the true mean $\theta$ is typically not available and is only estimated (with error) as $\hat{\theta} = 182.15$. In our paper, we show that the original profit estimates based on $\hat{\theta} = 182.15$ are biased, and we provide adjustments to correct for the biases that can be used by the newsvendor, who does not have access to the true mean. Using these adjustments (that only require $\hat{\theta}$ and not the true unknown value $\theta$), our bias-adjusted estimates for the nonreactive and reactive expected profits are \$3,947 and \$4,088, respectively, and we make the correct conclusion that reactive capacity is beneficial. Finally, note that the bias adjustment of the (nonreactive) newsvendor profit is 7.2% of its biased estimate and that the true expected profit (with respect to $\theta$) of using reactive capacity is 8.0% higher than the profit of not using reactive capacity.

Our contributions for identifying the newsvendor profit's expected estimation error are derived within the theoretical framework developed by Siegel and Woodgate (2007) for quantifying the expected estimation error in the performance of financial portfolios optimized using estimates of true parameters; the newsvendor model shares the basic structure of optimizing while pretending that estimates are correct. Achieving our closed-form profit adjustment is complex because the order quantity and the exact functional form of the realized profit are both nonlinear functions of the estimation error of the parameter(s). We use the method of statistical differentials to find the second order Taylor series approximation to the expected future profit (averaged over the random future demand realization) and then take its expectation (over the sampling error of parameter estimation). Our result is asymptotically correct when the sample size used for demand-function estimation is large (and simulations demonstrate that our results are useful even for moderate sample sizes) and remains statistically consistent when estimated values are substituted for the unknown parameter(s) of the demand distribution. In effect, we use perturbation analysis to discover how estimation errors are misused by the newsvendor's optimization process, which wrongly "believes" that the estimated parameters are correct. All proofs are provided in the online electronic companion.

### 1.1. Contributions

Our research provides the following main contributions:

- For a generic demand distribution, under appropriate assumptions, we derive a second order approximation of the newsvendor profit's expected estimation error for an arbitrary order quantity that is a smooth function of the estimated parameters. Using these expected estimation error results, we provide an adjusted expected profit expression that can be computed without knowledge of the true parameter(s), which is an asymptotically unbiased estimator of the true expected profit when the estimated order quantity is used. Hypothesis testing experiments confirm our theoretical results in that we show at the 5% test level the unadjusted estimation error is significantly positive (the null is rejected), whereas the adjusted estimation error is not significantly different from zero (the null is accepted).
- We similarly derive expected estimation error expressions for the order quantity, which allows us to also provide adjusted order quantities. We show that some distributions (e.g., exponential and normal) already exhibit order quantities with zero expected estimation error, but other distributions (e.g., log-normal) result in a nonzero expected estimation error. We provide an adjusted order quantity that is an asymptotically unbiased estimate of the true optimal order quantity. Again, hypothesis testing experiments confirm our theoretical results.
- We characterize the relative importance of profit and order quantity expected estimation errors. In particular, the adjustment to the order quantity results in a negligible change in expected profit: for a log-normal distribution of demand, the adjustment results in a 0.001% increase in profit in an example. In contrast, correcting the profit expected estimation error is not negligible: for the same log-normal distribution of demand, the adjustment results in a 1.1% increase in profit. Furthermore, when demand is exponentially distributed, the order quantity is unbiased, but the profit estimation error in an example is 3.4% of the true expected profit. Thus, the profit adjustments can be managerially relevant.

### 1.2. Literature Review

The newsvendor is a fundamental model of operations management, appearing frequently in both practice and academic research. Consequently, there is a vast literature on various aspects of the newsvendor; here, we discuss the most relevant papers to properly position our contributions. To the best of our knowledge, we are the first to document expected estimation error in the maximized expected profit perceived by the newsvendor, who estimates, using available data, the parameters of a demand distribution. Much attention has been focused on various consequences of distributional uncertainty; in particular, we discuss the following streams of literature: (1) the data-driven newsvendor, (2) approximation algorithms for stochastic inventory-management problems, (3) algorithmic inventory management, (4) distributionally robust inventory management, and (5) the newsvendor with operational statistics in which estimation and optimization are performed jointly. However, as we discuss in the following, these streams focus on identifying good ordering strategies but have not studied the expected profit estimate that a decision maker would also require in practice.

**1.2.1. Data-Driven Newsvendor.** The newsvendor model can be solved using only data without any assumptions on the form of the demand distribution: Kleywegt et al. (2002) design the sample average approximation (SAA) method using available data to create an empirical demand distribution, and the optimal order quantity is found at the critical fractile of this distribution (in the newsvendor context); these authors also identify a bias in a natural estimation of the optimal objective value (for more general stochastic discrete optimization problems) but do not provide an adjustment term to correct the bias as we do in our paper. Levi et al. (2007a) study the sampling-based SAA method for the newsvendor as well as a multiperiod extension of it and show that their policy has a solution that is provably near optimal with high probability; see the references therein for a more comprehensive literature review related to the data-driven newsvendor. Levi et al. (2015) refine the analysis in Levi et al. (2007a) to produce tighter performance bounds by introducing an additive bias into the order quantity; in addition, they utilize a second order Taylor series in their analysis as we do in ours. Ban and Rudin (2019) extend the data-driven approach of Levi et al. (2007a, 2015) to include features (explanatory variables) that influence the demand distribution, using techniques from machine learning; as in Levi et al. (2007a), performance bounds are derived that hold with a prescribed probability.

He et al. (2012) consider a similar features-based newsvendor model for staffing hospital operating rooms. However, in these references, the benchmark optimal expected profit depends on the true distribution, which is unknown, and hence, the optimal expected profit is not calculable. The only calculable profit expression is with respect to the data-driven empirical distribution, which is the realization of a random variable (because of sampling), and the main results also calculate the expected value of this random variable with respect to the true unknown distribution. Thus, although the performance guarantees of these papers are powerful, they are implicit results, and no calculable expressions exist for the expected profits. In contrast, we focus on natural calculable profit expressions based on distributional parameters estimated from data, which turn out to have nonzero expected estimation errors, and we provide adjustments to the biased expressions, thereby giving the decision maker a calculable unbiased estimate of expected profit. Ban and Rudin (2019) also interpret their results in terms of "generalization error," in which the decisions are based on in-sample data, but the quality of the decisions is evaluated out-of-sample (in terms of the unknown true distribution), and their analysis utilizes results from machine learning on generalization error (see references therein); in this interpretation, we provide explicit closed-form corrections for evaluating the out-of-sample expected profits, a result that is new to the literature to the best of our knowledge. Ban et al. (2019) also study features of demand, but the focus is on algorithmic solutions to multiperiod inventory-management models, whereas we focus on a single-period model in which closed-form solutions for the order quantity, profit expressions, and corrections are available.

**1.2.2. Approximation Algorithms for Stochastic Inventory Management Problems.** A related stream focuses on approximation algorithms for various stochastic inventory management models with provable performance guarantees that do not require a probabilistic qualifier (i.e., the performance guarantees hold with probability one). Levi et al. (2007b) provide two- and three-approximation algorithms for the multiperiod periodic-review stochastic inventory-control and the stochastic lot-sizing problems, respectively. Levi et al. (2008b) design a two-approximation algorithm for a multiperiod capacitated inventory control model. Levi et al. (2008a) derive two-approximation algorithms for stochastic inventory control models with lost sales. Levi and Shi (2013) provide a three-approximation algorithm for a multiperiod stochastic lot-sizing problem with order lead times, and Shi et al. (2014) design a four-approximation algorithm for a similar problem with capacities and setup costs. These papers focus on different aspects of a single location dealing with a single product; Levi et al. (2017) develop various approximation algorithms with performance guarantees between two and three for multiechelon systems with multiple products. However, all these approximation algorithm results require complete knowledge of the true demand distribution to calculate the expected profits—an assumption we relax.

**1.2.3. Algorithmic Inventory Management.** There are a number of other papers that take an algorithmic approach to solving data-driven inventory-management problems. Kunnumkal and Topaloglu (2008) apply a stochastic gradient approach to a multiperiod newsvendor model. Burnetas and Smith (2000) and Huh and Rusmevichientong (2009) also apply stochastic gradient algorithms except, to a scenario with censored demand in which only sales data—not true demand data—are available; Godfrey and Powell (2001), Huh et al. (2011), and Besbes and Muharremoglu (2013) also study the impact of censored demand for the newsvendor via various algorithms. Hannah et al. (2010) study a more general stochastic optimization problem algorithmically, which can be applied to an inventory-management context. Our work differs from these papers in that we focus on a simpler, single-period model so that we may obtain closed-form solutions, and we assume that demand, not just sales, data are available.

**1.2.4. Distributionally Robust Inventory Management.** The sensitivity of the newsvendor solution to the demand distribution has been studied in the distributionally robust optimization literature. Scarf (1958) analyzes the so-called "distribution-free" newsvendor model, in which the mean and standard deviation of demand are given, but the distribution is not; the objective is then to maximize the worst-case expected profit, which is minimized over all distributions with the given mean and standard deviation. Gallego and Moon (1993) provide extensions and a review of similar work. More recently, Perakis and Roels (2008) and Natarajan et al. (2018) consider similar but more sophisticated distribution-free settings in which, again, the focus is to derive order quantities that serve well in the worst case. These papers assume that their information is correct and do not consider the impact of estimation error, which we identify as an assumption with consequences to the newsvendor's anticipated expected profit.

**1.2.5. Operational Statistics.** In the paper perhaps most related to our research, Liyanage and Shanthikumar (2005) consider the single-period newsvendor model under a parametric demand distribution in which data are available to estimate parameters. Assuming that demand is exponentially distributed, the authors intentionally bias the order quantity to obtain higher expected profit via a joint estimation–optimization operation. We also consider the special case in which demand is exponentially distributed, and although the operational-statistics order quantity improves profit, we show that it still exhibits a nonzero expected estimation error. Furthermore, we show that the expected estimation error is considerably larger, asymptotically, than the profit improvement. Chu et al. (2008) study the operational-statistics approach using Bayesian analysis for parametric distributions characterized by location and scale parameters, deriving closed-form solutions for an exponential distribution of demand (matching that of Liyanage and Shanthikumar 2005) as well as for a uniform distribution of demand. In contrast, our results are not limited to location-scale families of distributions and are applicable to a broad array of distributional families for which we derive closed-form asymptotic expressions for the profit's expected estimation error, which can be used to easily correct the error.

## 2. The Newsvendor Model

We let $X$ denote random demand, which we assume has a continuous distribution with cumulative distribution function $F(x,\theta)$, where $\theta$ is a parameter that specifies the demand distribution from a family of distributions. For this introduction, we assume that $\theta$ is a scalar though we extend all our results to the case in which $\theta$ is a finite-dimensional vector. The probability density function of demand for a given $\theta$ is $f(x,\theta)=\frac{\partial}{\partial x}F(x,\theta)$ and is assumed to exist and be positive on a contiguous interval of support. The economics of the newsvendor model are as follows: $p>0$ is the sales price per unit, $c>0$ is the cost per unit, and $q\geq 0$ is the order quantity; we assume $p>c$ throughout. The newsvendor profit is the random variable $p\min(q,X)-cq$, and we focus on the expected profit

$$\pi(q,\theta)=pE_{X\sim F(\cdot,\theta)}[\min(q,X)]-cq.$$

This expected profit is known to be maximized when $q$ is chosen so that $1-F(q,\theta)=\frac{c}{p}$, where the left-hand side of this equation is the probability of selling the $q$th unit; we may write this optimality condition as

$$q(\theta)=[F(\cdot,\theta)]^{-1}\left(1-\frac{c}{p}\right), \tag{1}$$

where the inverse function is taken with respect to the first argument of $F$ at a given value of $\theta$. The expected profit for an arbitrary ordering quantity $q$, when the truth $\theta$ is known, is as follows:

$$\begin{aligned}\pi(q,\theta)&=pE_{X\sim F(\cdot,\theta)}[\min(q,X)]-cq\\&=p\int_0^q xf(x,\theta)dx+pq[1-F(q,\theta)]-cq\\&=(p-c)q-p\int_0^q F(x,\theta)dx,\end{aligned} \tag{2}$$

where the last equality is by integration of parts.

Note that all of the results in our paper continue to hold for the case of a fixed salvage value $s$. Instead of $p\min(q,X)-cq$, the profit becomes $p\min(q,X)-cq+s[q-\min(q,X)]$. Rearranging, the profit is $(p-s)\min(q,X)-(c-s)q$, and so the newsvendor problem with salvage value is entirely equivalent to the newsvendor problem without salvage value provided that the cost and price are both reduced by the salvage value. For ease of exposition, we assume $s=0$ in our paper.

### 2.1. Motivation for a Parametric Model

There are a number of reasons a manager should consider using a parametric model over the purely data-driven approaches reviewed in Section 1.2.1 (e.g., the SAA method). Researchers at Amazon, from their Forecasting Data Science and Supply Chain Optimization Technologies groups, discuss how parametric distributions fit within their broader forecasting activities (Madeka et al. 2018). In particular, they discuss, in a newsvendor context, how (shifted) gamma and log-normal distributions may be utilized to forecast future demand. In particular, they point out that "...the lognormal distribution is suitable for evergreen products with non-negligible demands."

Furthermore, the expected profit is larger (closer to optimal) with a parametric model when such a model can reasonably be assumed (e.g., by examining histograms of data and/or more rigorous distributional goodness-of-fit hypothesis testing). An efficient estimator of the order quantity, for instance, using the maximum likelihood estimator (MLE) for a parametric model, has minimum asymptotic variance among all estimators (see sections 6.2–6.4 of Lehman 1983 for details and regularity conditions). The decrease from optimal expected profit is proportional (asymptotically) to the variance of this estimator as may be seen using the expectation of a second order Taylor series expansion of profit as a function of the estimated order quantity. For example, measuring asymptotic statistical efficiency as the ratio of variances, if the SAA estimate has efficiency 65% relative to a (correct) parametric model, then the drop in expected profit (as compared with using the optimal but unknown order quantity) using the MLE is only

65% as large as the drop in expected profit incurred when using the SAA estimator. In fact, under exponentially distributed demand, the asymptotic statistical efficiency of the SAA order quantity with respect to the parametric order quantity is never larger than 65%; we show this formally in the next section.

Finally, in a biological context, Jabot (2015) argues that parametric forecasting is preferred to nonparametric forecasts because of (1) the ability to diagnose parametric forecasting failure via Bayesian model checking procedures, and (2) forecasting uncertainty can be estimated using synthetic data generated from the fitted parametric model.

## 3. The Impact of Parameter Estimation on Expected Profit

In practice, the parameter $\theta$ can be unknown, in which case it is estimated as $\hat{\theta}$ (i.e., using data). The newsvendor, using the estimated distribution $F(\cdot, \hat{\theta})$, chooses the stock quantity $\hat{q} \triangleq q(\hat{\theta})$. The parameter estimate is assumed to be unbiased and consistent so that $E[\hat{\theta}] = \theta$ and $E[(\hat{\theta}-\theta)^2] = O(\frac{1}{n})$, where $n$ is the size of an independent and identically distributed sample of demand quantities $X_i \sim F(\cdot, \theta)$, $i = 1, \ldots, n$. Detailed assumptions are given in Section 3.2.

*Actual* expected profit is then $\pi(\hat{q}, \theta)$, which is computed using the true parameter value that determines demand; that is, the order quantity is determined using the estimated demand parameter $\hat{\theta}$, but profit is averaged over actual demand $X$ with the correct distribution $F(\cdot, \theta)$. However, the decision maker does not have access to $\theta$, and therefore, the decision maker does not have access to the expected profit $\pi(\hat{q}, \theta)$. The only calculation a decision maker can reasonably do, initially, to have an idea what expected profit will actually be is $\pi(\hat{q}, \hat{\theta})$, which is calculated using the estimated distribution $F(\cdot, \hat{\theta})$; we call $\pi(\hat{q}, \hat{\theta})$ the *perceived* expected profit. Note that both of these expected profits (actual and perceived) are random variables, depending on the estimate $\hat{\theta}$ and its random estimation error, generated when the truth is $\theta$.

In this paper, we show that $\pi(\hat{q}, \hat{\theta})$ is a biased estimator of the true expected profit $\pi(\hat{q}, \theta)$. A main contribution of our paper is to derive an adjustment to $\pi(\hat{q}, \hat{\theta})$ that results in an asymptotically unbiased estimate of the true expected profit despite the inaccessibility of $\theta$ itself:

$$E\left[\pi(\hat{q}, \hat{\theta}) - \text{Adjustment}\right] = E\left[\pi(\hat{q}, \theta)\right] + o(1/n),$$

where the expectation is taken with respect to the sampling distribution. This adjustment is calculated via Taylor series expansions and can be used without knowledge of the true value of $\theta$. In addition, our analysis is not limited to the order quantity $\hat{q}$; we provide an asymptotically unbiased estimate of $E[\pi(\tilde{q}, \theta)]$ that does not require knowledge of the true value of $\theta$, where $\tilde{q}$ is an arbitrary order quantity defined as a smooth function of $\hat{\theta}$. We are abusing notation here for simplicity of explanation: in fact, $\tilde{q} = \tilde{q}_n$, a sequence of functions that we assume converges to the constant $q(\theta)$ in the limit for a large sample size $n$ for which $\hat{\theta}$ is converging to $\theta$.

In addition, our results have implications for improving the estimation of the expected mismatch cost, which is defined as $cE[\max\{q - X, 0\}] + (p - c)E[\max\{X - q, 0\}]$. In particular, it is straightforward to show that the expected mismatch cost is equivalent to $(p - c)E[X] - \pi(q, \theta)$; see, for instance, section 13.1 of Cachon and Terwiesch (2012). Because expected demand $E[X]$ is estimated unbiasedly (assuming the mean is included in the parameter vector $\theta$), any bias in forecasting expected profit flows directly and dollar-for-dollar as a bias of the expected mismatch cost. It, therefore, follows that, for the naive newsvendor, the expected mismatch cost is estimated with bias and that, by using our profit bias adjustment, one can obtain an (asymptotically) unbiased estimate of the expected mismatch cost.

Finally, we apply our Taylor series technique to derive asymptotically unbiased estimates of the true optimal order quantity that also depend on the unknown value of $\theta$ to achieve the following result:

$$E\left[\hat{q} - \text{Adjustment}\right] = q(\theta) + o(1/n).$$

In the next two Sections 3.1 and 3.2, we focus on the expected profit adjustment and consider the cases of an exponential distribution of demand and a general distribution of demand, respectively. In the exponential case, we are able to derive an exact adjustment (i.e., no $o(1/n)$ error term). In the general case, under appropriate assumptions, we derive an asymptotic adjustment via Taylor series analysis.

### 3.1. Closed-Form Solutions for an Exponential Distribution of Demand

In this section, we consider the case in which demand is exponentially distributed with unknown mean $\theta$ that is estimated as $\hat{\theta} = \frac{1}{n}\sum_{i=1}^{n} X_i$. In this case, the order quantity is $\hat{q} = \ln(p/c)\hat{\theta}$. The order quantity $\hat{q}$ is of the form $\tilde{q}(\hat{\theta}) = a\hat{\theta}$, where $a$ is a constant; many of the following results are presented in terms of the more general $\tilde{q}(\hat{\theta}) = a\hat{\theta}$. In our literature review, we describe the joint estimation–optimization procedure studied in Liyanage and Shanthikumar (2005), which derives $q_{OS} \triangleq n[(p/c)^{1/(n+1)} - 1]\hat{\theta}$ as the optimal order quantity for an exponential distribution of demand (the subscript $OS$ standing for operational statistics). Thus, the general form of ordering quantity, $\tilde{q}(\hat{\theta}) = a\hat{\theta}$, unifies the classical and the operational statistics methods because $\hat{q} = a_C\hat{\theta}$ with $a_C = \ln(p/c)$ and $q_{OS} = a_{OS}\hat{\theta}$, where $a_{OS} = n[(p/c)^{1/(n+1)} - 1]$. Our first result, from Liyanage and Shanthikumar (2005), provides the *actual* expected profit with respect to the true mean $\theta$ averaged over the sampling distribution for the general order quantity $\tilde{q}(\hat{\theta}) = a\hat{\theta}$; we provide a new proof of this result. Note that we are unaware of any similar result in the literature for a distribution other than the exponential distribution.

**Lemma 1** (Liyanage and Shanthikumar 2005)**.** *The overall expectation of actual expected profit for exponentially distributed demand with order quantity $\tilde{q}(\hat{\theta}) = a\hat{\theta}$ and constant a at a fixed sample size n is*

$$E\left[\pi\left(\tilde{q}(\hat{\theta}), \theta\right)\right] = \left[p - ac - p\left(\frac{n}{n+a}\right)^n\right]\theta.$$

Our next result provides the perceived expected profit for the general order quantity $\tilde{q}(\hat{\theta}) = a\hat{\theta}$, where the estimated distribution $F(\cdot, \hat{\theta})$ is used instead of the true unknown distribution $F(\cdot, \theta)$.

**Lemma 2.** *The overall expectation of perceived expected profit for exponentially distributed demand with order quantity $\tilde{q}(\hat{\theta}) = a\hat{\theta}$ and constant a at a fixed sample size n is*

$$E\left[\pi\left(\tilde{q}(\hat{\theta}), \hat{\theta}\right)\right] = \left[p - ac - pe^{-a}\right]\theta.$$

Comparing the overall expectation of actual and perceived expected profit expressions of Lemmas 1 and 2, we observe that they are clearly different. Their difference suggests an adjustment for the perceived expected profit when the true mean $\theta$ is not known, and we propose an adjusted expected profit

$$\hat{\pi}_a\left(\tilde{q}(\hat{\theta}), \hat{\theta}\right) \triangleq \pi\left(\tilde{q}(\hat{\theta}), \hat{\theta}\right) - p\hat{\theta}\left[\left(\frac{n}{n+a}\right)^n - e^{-a}\right]. \quad (3)$$

Note that this expression does not require knowledge of the true value $\theta$ (which contrasts with Lemmas 1 and 2) and only requires its unbiased and consistent estimator $\hat{\theta}$. This adjusted expected profit is an unbiased estimator for the actual expected profit as presented in the next proposition. In addition, with the classical order quantity, this adjusted expected profit has lower variance than the unadjusted expected profit, thereby reducing the mean squared error (the sum of variance and squared bias) by reducing both contributions.

**Proposition 1.** *If demand is exponentially distributed with fixed sample size n and order quantity $\tilde{q}(\hat{\theta}) = a\hat{\theta}$, then the adjusted perceived expected profit, defined in Equation (3), is an unbiased estimator for the actual expected profit, $E[\hat{\pi}_a(\tilde{q}(\hat{\theta}), \hat{\theta})] = E[\pi(\tilde{q}(\hat{\theta}), \theta)]$, and the adjustment is negative, decreasing the perceived expected profit and thereby correcting for its over-optimism. In addition, the adjusted perceived expected profit has smaller variance than the unadjusted perceived expected profit when the classical order quantity ($a = a_C$) is used.*

#### 3.1.1. Asymptotic Properties for the Exponential Distribution.

We next develop the asymptotic properties of the adjustment's improvement and of the expected profit improvement from using the OS order quantity $q_{OS} = a_{OS}\hat{\theta} = n[(p/c)^{1/(n+1)} - 1]\hat{\theta}$ instead of the classical choice $\hat{q} = a_C\hat{\theta} = \ln(p/c)\hat{\theta}$. This asymptotic analysis allows us to show that, although the OS expected profit improvement becomes negligible at moderate sample sizes, the adjustment's improvement remains economically meaningful in this range.

**Proposition 2.** *For exponentially distributed demand, the asymptotic expected profit improvement from using the OS order quantity in place of the classical choice is*

$$E\left[\pi\left(q_{OS}, \theta\right)\right] - E\left[\pi\left(\hat{q}, \theta\right)\right] = \frac{c\theta\left(-2\ln(p/c) + \left[\ln(p/c)\right]^2\right)^2}{8n^2} + O\left(\frac{1}{n^3}\right).$$

From Proposition 2, we see that the actual expected profit improvement from using $q_{OS}$ instead of $\hat{q}$ is $O(1/n^2)$. The next two results provide asymptotic expressions for the profit expected estimation errors when the *unadjusted* perceived expected profits are used to estimate the actual expected profits for both $\hat{q}$ and $q_{OS}$ ordering quantities, respectively. These errors can also be interpreted as improvements when the *adjusted* expected profits are used instead, which we contrast with the profit improvement of Proposition 2.

**Proposition 3.** *For exponentially distributed demand, the asymptotic expected estimation error for the unadjusted perceived expected profit when the classical order quantity is used is*

$$E\left[\pi\left(\hat{q}, \hat{\theta}\right) - \pi\left(\hat{q}, \theta\right)\right] = c\theta\left[\frac{\left[\ln(p/c)\right]^2}{2n} + \frac{-8\left[\ln(p/c)\right]^3 + 3\left[\ln(p/c)\right]^4}{24n^2}\right] + O\left(\frac{1}{n^3}\right),$$

*and when the OS order quantity is used, it is*

$$E\left[\pi\left(q_{OS}, \hat{\theta}\right) - \pi\left(q_{OS}, \theta\right)\right] = c\theta\left[\frac{\left[\ln(p/c)\right]^2}{2n} + \frac{-24\left[\ln(p/c)\right]^2 + 16\left[\ln(p/c)\right]^3 - 3\left[\ln(p/c)\right]^4}{24n^2}\right] + O\left(\frac{1}{n^3}\right).$$

From Proposition 3, we observe that the asymptotic expected estimation error for the OS order quantity is identical to that of the classical order quantity, up to $O(1/n)$, and differs only in the $O(1/n^2)$ term. Furthermore, these results show that the expected profit estimation errors for $\hat{q}$ and $q_{OS}$ are both $O(1/n)$, which is asymptotically greater than the actual profit improvement of $O(1/n^2)$. Thus, if one were to use $q_{OS}$ instead of $\hat{q}$, a small amount of additional profit would be realized; however, if a naive estimation of expected profit is used (i.e., using $F(\cdot, \hat{\theta})$), the expected estimation error dwarfs the improvement. Therefore, we recommend decision makers adopt our unbiased adjusted expected profit expression to eliminate the expected estimation error whether or not they use the OS order quantity.

**3.1.2. Numerical Verification for the Exponential Distribution.** In this section, we (numerically) compare the magnitude of the operational statistics approach's profit improvement over the classic order quantity (solid line) with that of its expected estimation error (dashed line) for demand that is exponentially distributed with mean $\theta = 200$. In the left pane of Figure 1, for $c = 0.4$ and $p = 1$, we see the improvement and expected estimation error as a function of sample size $n$, and we observe that, although the profit improvement of the operational statistics approach is indeed positive, the profit expected estimation error still exists and is much larger than the improvement. Furthermore, the profit improvement becomes negligible for moderate values of $n$, yet the expected estimation error remains economically meaningful for much larger values of $n$. In the right pane, we provide the profit improvement and expected estimation error as a function of $c/p$ for $n = 20$ and we observe similar results.

**3.1.3. Efficiency of Parametric and SAA Estimates of Order Quantity for the Exponential Distribution.** In this section, we formalize our argument from Section 2.1 that the parametric estimate of the optimal order quantity is asymptotically more efficient than the SAA estimate, which orders the $1 - c/p$ percentile of the *empirical* distribution. The statistical efficiency of the estimator $\hat{q}_{SAA}$ with respect to the MLE $\hat{q}$ is defined as $\frac{Var(\hat{q})}{Var(\hat{q}_{SAA})}$.

**Lemma 3.** *The asymptotic statistical efficiency of the SAA order quantity $\hat{q}_{SAA}$ as compared with the MLE order quantity $\hat{q} = \hat{\theta}\ln(p/c)$ under the exponential demand parametric model tends asymptotically (as sample size $n$ grows) to*

$$\frac{Var(\hat{q})}{Var(\hat{q}_{SAA})} \rightarrow \frac{[\ln(p/c)]^2}{p/c - 1},$$

*which can be no larger than 64.8%.*

## 3.2. Asymptotic Solutions for a General Distribution of Demand

In this section, we extend our results from the exponential distribution to a larger class of smooth distributions with smoothness conditions given as follows. Furthermore, we consider an arbitrary order quantity $\tilde{q}(\hat{\theta})$, where $\tilde{q}$ is a smooth function of $\hat{\theta}$ in the sense that $\tilde{q}$ is differentiable as a function of $\hat{\theta}$, and the derivative is locally Lipschitz continuous in $\hat{\theta}$. We also assume that $\tilde{q}(\hat{\theta})$ is a consistent estimator of $\tilde{q}(\theta)$ in the sense that $\tilde{q}(\hat{\theta})$ converges in probability to $\tilde{q}(\theta)$ as sample size $n$ grows—an example of which appears as $q_{OS}$ in Proposition 2. For such a distribution and order quantity, using a Taylor series approximation, we expand the expression for $\pi(\tilde{q}(\hat{\theta}), \hat{\theta}) - \pi(\tilde{q}(\hat{\theta}), \theta)$ about $\theta$ to second order in $\hat{\theta} - \theta$ and then take the expected value; the result is the *second order expected estimation error*, which is $O(1/n)$. We let $\nabla$ denote the gradient

**Figure 1.** (Color online) Demand Has an Exponential Distribution with $\theta = 200$

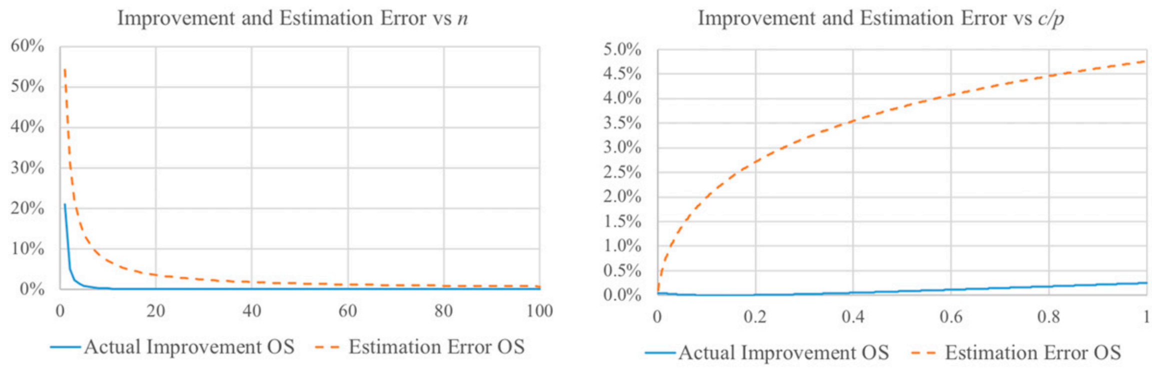


*Notes.* (Left) Operational statistics profit improvement (over classic order quantity) and profit expected estimation error as a function of $n$. (Right) Operational statistics profit improvement (over classic order quantity) and profit expected estimation error as a function of $c/p$.

operator with respect to $\theta$ and let $\nabla^2$ denote the Hessian matrix operator with respect to $\theta$.

Our smoothness conditions (which, henceforth, are to be assumed) on the parametric family of distributions are adapted from Lehman (1983, chapter 6, corollary 2.3, theorems 2.3 and 1.1) and are sufficient to ensure that the maximum likelihood estimator exists and is efficient. In particular, these assumptions are satisfied by any one-parameter exponential family of distributions (see, e.g., Lehman 1983).

#### 3.2.1. Assumptions for General Demand Distribution Family.

1. The parameter space for $\theta$ is an open interval (possibly unbounded), and the distribution $F(\cdot,\theta)$ has density $f(\cdot,\theta)$. The observations are independent and identically distributed with distribution $F(\cdot,\theta)$ for some true value of $\theta$ that is in the parameter space.
2. The distributions $F(\cdot,\theta)$ have common support $A=\{x: f(x,\theta)>0\}$ and are distinct to ensure estimation is possible.
3. For every $x\in A$, the density $f(\cdot,\theta)$ is three times differentiable with respect to $\theta$, and the third derivative is continuous in $\theta$.
4. The integral $\int f(x,\theta)dx$ can be twice differentiated under the integral sign.
5. The Fisher information $I(\theta)$ is positive and finite (see Lehman 1983, section 6.4 for additional assumptions for the vector parameter case).
6. For any given $\theta_0$ in the parameter space, there exists $\nu>0$ and a function $M(x)$ (both of which may depend on $\theta_0$) such that $E_{\theta_0}[M(X)]<\infty$ and also $|\partial^3 \ln f(x,\theta)/\partial\theta^3|\le M(x)$ for all $x\in A$ and $\theta_0-\nu<\theta<\theta_0+\nu$.
7. The probability of multiple roots to the likelihood equation $\sum \partial \ln[f(x_i,\theta)]/\partial\theta=0$ tends to zero as the sample size $n\to\infty$.

In the following proposition, for a multidimensional parameter $\theta=(\theta_1,\ldots,\theta_k)'$, where $k$ is a finite integer, we show that the newsvendor's perceived expected profit $\pi(\tilde{q}(\hat{\theta}),\hat{\theta})$ is systematically biased in that it differs, on average, from the actual expected profit $\pi(\tilde{q}(\hat{\theta}),\theta)$ that is realized based on the true, but unknown, demand parameter $\theta$. Subsequently, we show how these expected estimation errors can be used to create adjusted expected profits that more accurately reflect reality (i.e., they are unbiased). In addition, note that the sign of the expected estimation error is not yet evident; we explore the direction of the error in Section 3.3 by applying our generic expressions to specific distributions.

**Proposition 4.** *If $\tilde{q}(\hat{\theta})$ is smooth as defined at the start of the section, then the profit expected estimation error $E[\pi(\tilde{q}(\hat{\theta}),\hat{\theta})-\pi(\tilde{q}(\hat{\theta}),\theta)]$ for $\tilde{q}(\hat{\theta})$ for the case of a k-dimensional parameter $\theta=(\theta_1,\ldots,\theta_k)'$ with unbiased estimator $\hat{\theta}$ equals*

$$-p\cdot tr\left[\mathrm{Cov}(\hat{\theta})\left((\nabla\tilde{q}(\theta))(\nabla F)'+\frac{1}{2}\int_0^{\tilde{q}(\theta)}\nabla^2F(x,\theta)dx\right)\right]+o\left(\frac{1}{n}\right),$$

*where* tr *denotes the matrix trace operator, $\nabla\tilde{q}(\theta)$ is the gradient of $\tilde{q}(\theta)$ with respect to $\theta$, $\nabla F$ is the gradient of F with respect to $\theta$ and evaluated at $(\tilde{q}(\theta),\theta)$, $\nabla^2F(x,\theta)$ is the $k\times k$ Hessian matrix of F with respect to $\theta$, and $Cov(\hat{\theta})$ is the covariance matrix of $\hat{\theta}$.*

The expression identified in Proposition 4 provides evidence that a nonzero expected estimation error indeed exists; we provide further evidence via numerical experiments, focused on identifying statistically significant nonzero expected estimation errors, in Section 3.5.

We can apply Proposition 4 to the special case in which $\tilde{q}(\hat{\theta})=\hat{q}$ for which $\nabla\hat{q}(\theta)=-\nabla F/f$ from Lemma EC.4, which results in the following corollary.

**Corollary 1.** *The profit expected estimation error $E[\pi(\hat{q},\hat{\theta})-\pi(\hat{q},\theta)]$ for $\hat{q}$ for the case of a k-dimensional parameter $\theta=(\theta_1,\ldots,\theta_k)'$ with unbiased estimator $\hat{\theta}$ equals*

$$p\cdot tr\left[\mathrm{Cov}(\hat{\theta})\left(\frac{(\nabla F)(\nabla F)'}{f}-\frac{1}{2}\int_0^q\nabla^2F(x,\theta)dx\right)\right]+o\left(\frac{1}{n}\right),$$

*where $q=q(\theta)$, tr denotes the matrix trace operator, $f=f(q,\theta)$, $\nabla F$ is the gradient of F with respect to $\theta$ and evaluated at $(q,\theta)$, $\nabla^2F(x,\theta)$ is the $k\times k$ Hessian matrix of F with respect to $\theta$, and $\mathrm{Cov}(\hat{\theta})$ is the covariance matrix of $\hat{\theta}$.*

The next corollary considers the scalar $\theta$ case of Proposition 4 for the general ordering quantity $\tilde{q}(\hat{\theta})$.

**Corollary 2.** *If $\tilde{q}(\hat{\theta})$ is smooth as defined at the start of the section, then the profit expected estimation error $E[\pi(\tilde{q}(\hat{\theta}),\hat{\theta})-\pi(\tilde{q}(\hat{\theta}),\theta)]$ for $\tilde{q}(\hat{\theta})$ for the case of a one-dimensional parameter $\theta$ with unbiased estimator $\hat{\theta}$ equals*

$$-p\cdot\mathrm{Var}(\hat{\theta})\cdot\left(\tilde{q}'(\theta)F_\theta+\frac{1}{2}\int_0^{\tilde{q}(\theta)}F_{\theta\theta}(x,\theta)dx\right)+o\left(\frac{1}{n}\right),$$

*where $\tilde{q}'(\theta)=\frac{\partial\tilde{q}(\theta)}{\partial\theta}$, $F_\theta=\frac{\partial F}{\partial\theta}(\tilde{q}(\theta),\theta)$, $F_{\theta\theta}(x,\theta)=\frac{\partial^2F(x,\theta)}{\partial\theta^2}$, and $Var(\hat{\theta})$ is the variance of $\hat{\theta}$ under its sampling distribution.*

Note that Corollary 2 is applicable to both the naive order quantity $\hat{q}=[F(\cdot,\hat{\theta})]^{-1}(1-c/p)$ as well as the operational statistics order quantity $q_{OS}=n[(p/c)^{1/(n+1)}-1]\hat{\theta}$ for an exponential distribution of demand. Furthermore, the following corollary results when assigning $\tilde{q}(\hat{\theta})=\hat{q}$ because $\tilde{q}'(\theta)=-F_\theta/f$ for the naive order quantity as follows by differentiating the equation $1 - c/p = F(\hat{q}(\theta), \theta)$ with respect to $\theta$ to find $0 = f(\hat{q}(\theta), \theta)\hat{q}'(\theta) + F_\theta(\hat{q}(\theta), \theta) = f\hat{q}'(\theta) + F_\theta$.

**Corollary 3.** *The profit expected estimation error* $E[\pi(\hat{q}, \hat{\theta}) - \pi(\hat{q}, \theta)]$ *for* $\hat{q}$ *for the case of a one-dimensional parameter* $\theta$ *with unbiased estimator* $\hat{\theta}$ *equals*

$$p \cdot \mathrm{Var}(\hat{\theta}) \cdot \left(\frac{F_\theta^2}{f} - \frac{1}{2}\int_0^q F_{\theta\theta}(x, \theta)dx\right) + o\left(\frac{1}{n}\right),$$

*where* $q = q(\theta)$, $F_\theta = \frac{\partial F(q,\theta)}{\partial \theta}$, $f = f(q, \theta)$, $F_{\theta\theta}(x, \theta) = \frac{\partial^2 F(x,\theta)}{\partial \theta^2}$, *and* $Var(\hat{\theta})$ *is the variance of* $\hat{\theta}$ *under its sampling distribution.*

**3.2.2. Estimating the Covariance Matrix.** Our framework assumes that there is a single vector of data, $(X_1, \ldots, X_n)$, that is available to approximate $\theta = (\theta_1, \ldots, \theta_k)'$, which leads to a single vector estimate $\hat{\theta} = (\hat{\theta}_1, \ldots, \hat{\theta}_k)'$. However, this setup is not sufficient to directly estimate the covariance matrix $\mathrm{Cov}(\hat{\theta})$ required for Proposition 4 and Corollary 1 because we would need multiple samples of the vector estimate $\hat{\theta}$. Consequently, we approximate $\mathrm{Cov}(\hat{\theta})$ by using the Fisher information matrix

$$I(\theta) = -E_{X\sim F(\cdot,\theta)}\left[\nabla^2(\ln f(X, \theta))\right],$$

where the expectation is taken over the random demand $X$ with distribution $F(\cdot, \theta)$, parameterized by the true value $\theta$, and $\nabla^2$ is the Hessian second derivative operator with respect to $\theta$. Using the asymptotic theory of maximum likelihood estimation (see, for instance, Lehman 1983, sections 6.2–6.4, in particular theorems 2.3 and 4.1) together with our assumptions for general demand distribution family near the start of Section 3.2, we may approximate the covariance matrix of $\hat{\theta}$ as follows:

$$\mathrm{Cov}(\hat{\theta}) = \frac{1}{n}I^{-1}(\theta) + o\left(\frac{1}{n}\right).$$

Similarly, for the univariate parameter case, the Fisher information is the scalar

$$I(\theta) = -E_{X\sim F(\cdot,\theta)}\left[\frac{\partial^2}{\partial\theta^2}\ln f(X, \theta)\right],$$

which allows the variance of $\hat{\theta}$ to be approximated as $\mathrm{Var}(\hat{\theta}) = \frac{1}{nI(\theta)} + o(\frac{1}{n})$. The following corollary uses these expressions to rewrite the expected estimation errors.

**Corollary 4.** *If* $\tilde{q}(\hat{\theta})$ *is smooth as defined at the start of section* 3.2, *then the profit expected estimation error* $E[\pi(\tilde{q}(\hat{\theta}), \hat{\theta}) - \pi(\tilde{q}(\hat{\theta}), \theta)]$ *for* $\tilde{q}(\theta)$ *for the case of a* $k$*-dimensional parameter* $\theta = (\theta_1, \ldots, \theta_k)'$ *with unbiased estimator* $\hat{\theta}$ *equals*

$$\frac{p}{n}tr\left[\left(E_{X\sim F(\cdot,\theta)}\left[\nabla^2(\ln f(X, \theta))\right]\right)^{-1} \times \left((\nabla\tilde{q}(\theta))(\nabla F)' + \frac{1}{2}\int_0^{\tilde{q}(\theta)} \nabla^2 F(x, \theta)dx\right)\right] + o\left(\frac{1}{n}\right),$$

*and the case of a one-dimensional parameter* $\theta$ *with unbiased estimator* $\hat{\theta}$ *equals*

$$\frac{p}{nE_{X\sim F(\cdot,\theta)}\left[\frac{\partial^2}{\partial\theta^2}\ln f(X, \theta)\right]} \times \left(\tilde{q}'(\theta)F_\theta + \frac{1}{2}\int_0^{\tilde{q}(\theta)} F_{\theta\theta}(x, \theta)dx\right) + o\left(\frac{1}{n}\right).$$

Corollary 4 allows us to derive adjustments for the newsvendor's perceived expected profit that are easily computable to better reflect reality; we explore this topic in the next section.

**3.2.3. Unbiased Adjusted Expected Profit Estimation.** In this section, we provide an adjusted expected profit estimator, denoted as $\hat{\pi}_a$, which better reflects the actual expected profit observed by the newsvendor than its naive estimate, in which the hat over the $\pi$ indicates that this is an estimated value that is available to the newsvendor. To construct this adjusted expected profit, we modify the newsvendor's perceived expected profit $\pi(\tilde{q}(\hat{\theta}), \hat{\theta})$ by subtracting a consistent estimate of the second-order expected estimation error that is obtained by replacing unknown quantities in the error formulas with their estimates and ignoring the $o(1/n)$ error term. The result is that this adjusted expected profit is unbiased (to the second order) with respect to the actual expected profit $\pi(\tilde{q}(\hat{\theta}), \theta)$ in the sense that their expectations are equal to the second order.

For the case of a one-dimensional parameter $\theta$ with unbiased estimator $\hat{\theta}$, we propose the following adjusted expected profit

$$\hat{\pi}_a(\tilde{q}(\hat{\theta}), \hat{\theta}) \triangleq \pi(\tilde{q}(\hat{\theta}), \hat{\theta}) - \frac{p}{nE_{X\sim F(\cdot,\hat{\theta})}\left[\frac{\partial^2}{\partial\theta^2}\ln f(X, \hat{\theta})\right]} \times \left(\tilde{q}'(\hat{\theta})\hat{F}_\theta + \frac{1}{2}\int_0^{\tilde{q}(\hat{\theta})} F_{\theta\theta}(x, \hat{\theta})dx\right),$$

where $\hat{F}_\theta = \frac{\partial F}{\partial\theta}(\tilde{q}(\hat{\theta}), \hat{\theta})$ and $F_{\theta\theta}(x, \hat{\theta}) = \frac{\partial^2 F}{\partial\theta^2}(x, \hat{\theta})$ are both available without knowledge of the true value of $\theta$.

Similarly, for the case of a $k$-dimensional vector parameter $\theta$ with unbiased estimator $\hat{\theta}$, we propose

$$\hat{\pi}_a(\tilde{q}(\hat{\theta}),\hat{\theta}) \triangleq \pi(\tilde{q}(\hat{\theta}),\hat{\theta}) - \frac{p}{n} tr\left[\left(E_{X\sim F(\cdot,\hat{\theta})}\left[\nabla^2\left(\ln f(X,\hat{\theta})\right)\right]\right)^{-1} \times \left(\nabla\tilde{q}(\hat{\theta})(\nabla\hat{F})' + \frac{1}{2}\int_0^{\tilde{q}(\hat{\theta})} \nabla^2 F(x,\hat{\theta})dx\right)\right],$$

where $\nabla\hat{F} = \nabla F(\tilde{q}(\hat{\theta}),\hat{\theta})$ and $\nabla^2 F(x,\hat{\theta})$ is the Hessian of $F$ evaluated at $(x,\hat{\theta})$, which are also available without knowing the true value of $\theta$. This leads to the following theorem, one of the main results of our paper.

**Theorem 1.** *If $\tilde{q}(\hat{\theta})$ is smooth as defined at the start of Section 3.2, then, in the case of a k-dimensional, $k \geq 1$, vector parameter $\theta$ with unbiased estimator $\hat{\theta}$, the adjusted expected profit eliminates the $O(1/n)$ bias and is unbiased up to $o(1/n)$. That is,*

$$E\left[\hat{\pi}_a(\tilde{q}(\hat{\theta}),\hat{\theta})\right] = E\left[\pi(\tilde{q}(\hat{\theta}),\theta)\right] + o\left(\frac{1}{n}\right).$$

### 3.3. Illustrative Examples

In this section, we apply the general results derived in the previous sections to representative distributions and the classic order quantity $\hat{q}$ to learn how the expected estimation error depends on specific situations. We consider the exponential and normal distributions. In particular, we show that the second-order expected estimation error is positive in all cases.

**3.3.1. Exponential Distribution.** We first consider the exponential distribution that depends on a single parameter: its mean $\theta$. We may use the unbiased estimator $\hat{\theta} = \frac{1}{n}\sum_{i=1}^n X_i$, where $E[\hat{\theta}] = \theta$ and $Var(\hat{\theta}) = \theta^2/n$. If $\theta$ were known, the optimal order quantity would be $q = \theta\ln(p/c)$. However, the newsvendor naively uses $\hat{q} = \hat{\theta}\ln(p/c)$ instead. Applying Equation (2), we may find the expected profit with exponentially distributed demand for an arbitrary order quantity $\tilde{q}$ and mean $\theta$: $\pi(\tilde{q},\theta) = p\theta(1 - e^{-\tilde{q}/\theta}) - c\tilde{q}$. This allows us to determine the estimation error for $\hat{q}$, $\pi(\hat{q},\hat{\theta}) - \pi(\hat{q},\theta) = p\hat{\theta}(1 - e^{-\hat{q}/\hat{\theta}}) - p\theta(1 - e^{-\hat{q}/\theta})$, whose expected value with respect to the sampling distribution of $\hat{\theta}$ is provided in the following proposition.

**Proposition 5.** *The (positive) profit expected estimation error for $\hat{q}$ for the case of an exponential distribution with mean $\theta$ and computed using the general distribution result of Corollary 2 equals*

$$E\left[\pi(\hat{q},\hat{\theta}) - \pi(\hat{q},\theta)\right] = \frac{c\theta\left[\ln(p/c)\right]^2}{2n} + o\left(\frac{1}{n}\right).$$

We may contrast the result of Proposition 5, obtained by applying the general result of Corollary 3 to the exponential distribution, with that of Proposition 3, which is derived using the form of the exponential distribution. Note that the $\frac{1}{n}$ term is identical in both expressions, and the value of using Proposition 3 is that the coefficient of the $1/n^2$ term is known exactly, whereas it is not known exactly in Proposition 5, using the more general result.

Note that this second order expected estimation error $c\theta[\ln(p/c)]^2/(2n)$ for the exponential distribution is positive. We also note that, all else equal, the second order expected estimation error is monotonically increasing in $p$, is proportional to expected demand $\theta$, and is inversely proportional to $n$. As $c$ changes, all else equal, the behavior depends on the sign of the partial derivative

$$\frac{\partial}{\partial c}\left(\frac{c\theta\left[\ln(p/c)\right]^2}{2n}\right) = \frac{\theta\ln(p/c)\left[\ln(p/c) - 2\right]}{2n},$$

which is positive when $c < pe^{-2} \approx p/7.389$. Thus, the second order expected estimation error initially increases with $c$, reaching its maximum value of $2p\theta/(ne^2)$ when $c = pe^{-2}$, and then decreases to zero when $c = p$ with $q = \hat{q} = 0$. In particular, whenever $c > p/7.389$, the second order expected estimation error is decreasing with $c$. Finally, replacing the unknown $\theta$ value in the second order expected estimation error expression with its known unbiased estimator $\hat{\theta}$, we can determine the adjusted expected profit.

**Corollary 5.** *The adjusted expected profit for an exponential distribution is*

$$\hat{\pi}_a(\hat{q},\hat{\theta}) = p\hat{\theta}(1 - c/p) - c\hat{q} - \frac{c\hat{\theta}\left[\ln(p/c)\right]^2}{2n},$$

*where $E[\hat{\pi}_a(\hat{q},\hat{\theta})] = E[\pi(\hat{q},\theta)] + o(\frac{1}{n})$.*

**3.3.2. Normal Distribution.** We now consider the normal distribution, which is parameterized by the two-vector $\theta = (\mu,\sigma)$, where $\mu$ is the mean and $\sigma$ is the standard deviation of demand. If $\mu$ and $\sigma$ are known, the optimal order quantity is $q = q(\theta) = \mu + \sigma\Phi^{-1}(1 - \frac{c}{p})$. The estimated order quantity, based on an estimate $\hat{\theta} = (\hat{\mu},\hat{\sigma})$ for $\theta = (\mu,\sigma)$, is $\hat{q} = q(\hat{\theta}) = \hat{\mu} + \hat{\sigma}\Phi^{-1}(1 - \frac{c}{p})$.

We use the arithmetic mean $\hat{\mu} = \frac{1}{n}\sum_{i=1}^n X_i$ as an unbiased estimator for $\mu$. However, the usual estimators of the standard deviation $\sigma$ are known to be biased (these usual estimators include the sample standard deviation that divides the sum of squared residuals by $n-1$ before the square root is taken and the MLE that divides by $n$). Note that the sample *variance* is an unbiased estimator of $\sigma^2$ but that the nonlinearity of the square root function introduces bias into the sample standard deviation as an estimator of $\sigma$, which is used to find the order quantity.

Using properties of the chi distribution (the square root of a chi-squared distribution) with $n-1$ degrees of freedom from (18.14) of Johnson et al. (1994), we may construct an unbiased estimator $\hat{\sigma}$ of $\sigma$ with moments as follows:

$$\hat{\sigma} \triangleq k_n \sqrt{\frac{1}{n-1} \sum_{i=1}^{n} (X_i - \bar{X})^2},$$

where

$$k_n \triangleq \sqrt{n-1}\left(\frac{\Gamma[(n-1)/2]}{\sqrt{2}\Gamma(n/2)}\right) = \frac{1}{1 - \frac{1}{4(n-1)} + O\left(\frac{1}{n^2}\right)} = 1 + \frac{1}{4n} + O\left(\frac{1}{n^2}\right),$$

which implies that $E[\hat{\sigma}] = \sigma$ and $Var(\hat{\sigma}) = \sigma^2(k_n^2 - 1) = \frac{\sigma^2}{2n} + O(\frac{1}{n^2})$. The expansion of $k_n$ (and, therefore, of the variance) are obtained from (18.15) of Johnson et al. (1994), who credit Johnson and Welch (1939).

The expected profit with normally distributed demand for an arbitrary order quantity $\tilde{q}$ and parameters $\theta = (\mu, \sigma)$, where $\Phi$ and $\varphi$ denote the standard normal distribution and density, respectively, is given by Cachon and Terwiesch (2012) as

$$\pi(\tilde{q}, \theta) = (p-c)\tilde{q} - p(\tilde{q} - \mu)\Phi\left(\frac{\tilde{q}-\mu}{\sigma}\right) - p\sigma\varphi\left(\frac{\tilde{q}-\mu}{\sigma}\right). \quad (4)$$

Using this, along with the facts that $\hat{q} = \hat{\mu} + \hat{\sigma}\Phi^{-1}(1 - c/p)$ and $q = \mu + \sigma\Phi^{-1}(1 - c/p)$, the estimation error (given the estimate $\hat{\theta}$) is

$$\begin{aligned}
&\pi(\hat{q}, \hat{\theta}) - \pi(\hat{q}, \theta) \\
&\quad = p\left[(\hat{q} - \mu)\Phi\left(\frac{\hat{q}-\mu}{\sigma}\right) - (\hat{q} - \hat{\mu})\Phi\left(\frac{\hat{q}-\hat{\mu}}{\hat{\sigma}}\right) + \sigma\varphi\left(\frac{\hat{q}-\mu}{\sigma}\right) - \hat{\sigma}\varphi\left(\frac{\hat{q}-\hat{\mu}}{\hat{\sigma}}\right)\right] \\
&\quad = p\left\{(\hat{q} - \mu)\Phi\left(\frac{\hat{q}-\mu}{\sigma}\right) - \hat{\sigma}\left(1 - \frac{c}{p}\right)\Phi^{-1}\left(1 - \frac{c}{p}\right) + \sigma\varphi\left(\frac{\hat{q}-\mu}{\sigma}\right) - \hat{\sigma}\varphi\left[\Phi^{-1}\left(1 - \frac{c}{p}\right)\right]\right\}.
\end{aligned}$$

The expected value of the estimation error with respect to the sampling distribution of $\hat{\theta}$ is provided in the following proposition.

**Proposition 6.** *The profit expected estimation error for $\hat{q}$ for the case of a normal distribution with mean $\mu$ and standard deviation $\sigma$ and unbiased estimators $\hat{\mu}$ and $\hat{\sigma}$, respectively, equals*

$$E[\pi(\hat{q}, \hat{\theta}) - \pi(\hat{q}, \theta)] = \frac{p\sigma\left(2 + [\Phi^{-1}(1 - c/p)]^2\right)}{4n}\varphi(\Phi^{-1}(1 - c/p)) + o\left(\frac{1}{n}\right).$$

Note that the second order expected estimation error for the normal distribution is positive as was the case for the exponential distribution. In contrast to the exponential distribution case, the second order expected estimation error does not depend on the mean $\mu$ (although we note a similar proportionate dependence on the standard deviation for both distributions). All else equal, the second order expected estimation error is proportional to the uncertainty in demand $\sigma$ and is inversely proportional to the sample size $n$. We also find that, for a fixed ratio $c/p$, the second order expected estimation error is proportional to $p$. Note that the second order expected estimation error is symmetric in $\xi \triangleq \Phi^{-1}(1 - c/p)$ so that it remains unchanged if $c$ is changed to $p - c$, replacing $\xi$ with $-\xi$. All else equal, the second order expected estimation error is maximized when the cost is half of the price because $\xi = 0$ at this point, and this is the only root of the derivative

$$\frac{d}{d\xi}(2 + \xi^2)\varphi(\xi) = -\xi^3\varphi(\xi),$$

which changes sign from positive to negative at zero.

The second order expected estimation error approaches zero as the cost $c$ approaches either zero or $p$. As $c$ approaches zero, the order quantity $q$ approaches $\infty$, and the newsvendor's expected profit approaches $p\mu$ for which we have an unbiased estimate $p\hat{\mu}$, and hence, there is no estimation error in this limit. As $c$ approaches $p$, the order quantity $q$ effectively approaches zero, and the newsvendor's profit approaches zero, and hence, there is no estimation error in this limit either as the estimates become irrelevant.

Replacing the unknown $\mu$ and $\sigma$ values with their known unbiased estimators $\hat{\mu}$ and $\hat{\sigma}$, respectively, we can determine the adjusted expected profit.

**Corollary 6.** *The adjusted expected profit for a normal distribution with unbiased estimated mean $\hat{\mu}$ and standard deviation $\hat{\sigma}$ is*

$$\begin{aligned}
\hat{\pi}_a(\hat{q}, \hat{\theta}) &= (p-c)\hat{q} - p(\hat{q} - \hat{\mu})\Phi\left(\frac{\hat{q}-\hat{\mu}}{\hat{\sigma}}\right) - p\hat{\sigma}\varphi\left(\frac{\hat{q}-\hat{\mu}}{\hat{\sigma}}\right) - \frac{p\hat{\sigma}\left(2 + [\Phi^{-1}(1 - c/p)]^2\right)}{4n} \times \varphi(\Phi^{-1}(1 - c/p)) \\
&= (p-c)[\hat{\mu} + \hat{\sigma}\Phi^{-1}(1 - c/p)] - p\hat{\sigma}(1 - c/p) \times \Phi^{-1}(1 - c/p) - p\hat{\sigma}\varphi[\Phi^{-1}(1 - c/p)] \\
&\quad - \frac{p\hat{\sigma}\left\{2 + [\Phi^{-1}(1 - c/p)]^2\right\}}{4n}\varphi[\Phi^{-1}(1 - c/p)],
\end{aligned}$$

*where $E[\hat{\pi}_a(\hat{q}, \hat{\theta})] = E[\pi(\hat{q}, \theta)] + o(\frac{1}{n})$.*

### 3.4. Managerial Insights

In this section, we unify and discuss the observations we made in our analysis of the exponential and normal distributions. In both cases, we observe that the expected profit estimation error is inversely proportional to the sample size; in other words, as the sample size increases, the profit bias decreases. Thus, in practice, one way to reduce the expected profit estimation error is to collect more samples.

We also find that the exponential distribution's profit estimation error is increasing in the mean $\theta$, whereas the normal distribution's error is increasing in $\sigma$ (and does not depend on the mean $\mu$). However, noting that the exponential distribution's mean is also equal to its standard deviation, we conjecture a second commonality: as the demand distribution has more uncertainty, the profit bias increases.

Regarding the economic parameters of the newsvendor, we observe some differences that depend on the distribution: the exponential's profit bias is increasing in $p$, whereas the normal distribution's dependence on $p$ is more subtle. In both cases, we observe that the bias is unimodal in $c$ with a maximum at a distribution-dependent value of $c$. We can argue that this behavior is general: the bias cannot be monotonic in $c$ for fixed $p$ (if there is a bias for any value of $c$) because the bias approaches zero at the endpoints of the interval $c \in (0, p)$. To see this, note that, for a fixed value of $p$, the bias approaches zero as $c \to 0$ because the order quantity $\hat{F}^{-1}(1 - c/p)$ tends to infinity, stock-outs become very rare, and expected profit tends toward $pE[X]$; if the newsvendor has an unbiased estimator of mean demand $E[X]$, then there is no bias (because there is no nonlinearity). The bias also approaches zero as $c \to p$ from below because the order quantity $\hat{F}^{-1}(1 - c/p)$ tends to zero, and the newsvendor's estimated profit tends to agree closely with the true expected profit because they are both small numbers. On the other hand, the bias is monotonic in $c$ when the ratio $c/p$ is held fixed, which releases $c$ from its upper bound. In fact, the bias is proportional to $c$ in this case as it is also with respect to $p$ as may be seen from Corollaries 1 and 3 because $q$ remains unchanged when we fix the ratio.

### 3.5. Numerical Verification: Exponential and Normal Demand Distributions

Simulations were performed for the exponential and normal demand cases, and considerable expected estimation error reduction was observed using our second order approach. Our experimental primitive is as follows: we generate a sample of $n = 25$ observations from the true distribution (as specified by $\theta$) to obtain the estimate $\hat{\theta}$, from which we obtain the order size $\hat{q}$. We then compute two exact expectations (given this $\hat{q}$) of profit: $\pi(\hat{q}, \hat{\theta})$ as perceived by the naive newsvendor and $\pi(\hat{q}, \theta)$ as the newsvendor actually experiences (note that these expressions appear in Sections 3.3.1 and 3.3.2). Their difference, perceived minus actual expected profit, defines the estimation error before adjustment. For the estimation error after adjustment, we subtract the expected estimation error (which is also a function of $\hat{\theta}$). We are interested in exploring the expected values of these two measures using simulation.

To improve the efficiency of the Monte Carlo simulations, for each sample, we create an antithetic sample by replacing each observation in the sample of 25 with its complementary percentile, which preserves distributions and expected values while decreasing the variance of the simulations as a result of negative correlation (Hammersley and Morton 1956). This method uses the fact that, if $F$ is a cumulative distribution function and $U$ is uniformly distributed from zero to one, then both $F^{-1}(U)$ and $F^{-1}(1 - U)$ follow the same distribution $F$ but are negatively correlated. In the normal case, the two values are simple reflections about the mean.

Thus, for each sample of size 25, we compute estimation errors (both before and after adjustment) for both samples (original and antithetic). We then average the two values (original and antithetic) for each of the measures (before and after adjustment). The result is a pair of unbiased estimates for the error (before and after adjustment).

We compute the $t$-statistic (testing against zero estimation error) for each measure (before and after adjustment) by repeating this procedure for 10,000 independent samples (note that, although an individual sample and its antithetic counterpart are negatively correlated, by averaging them, we may focus on independent random variables while preserving the expectation of interest). The result is a pair of $t$-statistics for the expected estimation error: one before and one after adjustment.

We then repeat this procedure 100 times to obtain 100 pairs of $t$-statistics, each based on 10,000 paired antithetic simulations. In Figure 2, we plot histograms of these 100 paired $t$-statistics for the exponential (left) and normal (right) distributions. For the exponential distribution, we utilize $\theta = 200$; for the normal distribution, we utilize $\mu = 200$ and $\sigma = 65$. In both cases, we set $p = 5$ and $c = 3$. We observe strong evidence that our asymptotic adjustment eliminates the statistically significant estimation error very effectively even in these finite samples. In the left plot of Figure 2, the unadjusted estimation error for the exponential distribution shows high statistical significance (i.e., the $t$-values in the histogram at right are considerably higher than the standard 1.96 critical value), which is successfully eliminated by the asymptotic correction (i.e., the histogram at left is centered at zero).

**Figure 2.** (Color online) Adjustment Eliminates Bias for Exponential and Normal Demand

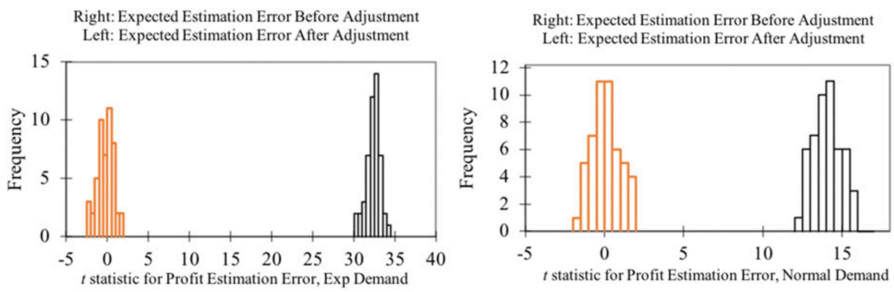


*Notes.* (Left) Exponential distribution with $\theta = 200$, $p = 5$, and $c = 3$ with samples of size $n = 25$. (Right) Normal distribution with $\mu = 200$, $\sigma = 65$, $p = 5$, and $c = 3$ with samples of size $n = 25$.

For exponential demand, the expected profit is 90.4, and the expected estimation error is 3.1, which is a nonnegligible 3.4% of the profit. In the right plot of Figure 2, we observe similar behavior for the normal distribution for which the expected profit is 271.9 and the expected estimation error is 2.6, which is 1% and less than in the exponential case but still nonnegligible. This smaller expected estimation error in the normal case may be due to its smaller standard deviation.

## 4. The Impact of Parameter Estimation on the Order Quantity

In the previous section, we show that, despite utilizing unbiased estimates of the parameters of the probability distribution of demand, the newsvendor profit exhibits a nontrivial estimation error. In this section, we study whether the order quantity itself has an estimation error and show how to correct this error when it exists.

As in the previous section and using Equation (1), we may define the estimation error as the difference between naive and optimal order quantities, $q(\hat{\theta}) - q(\theta) = \hat{q} - q$, which is again a random variable as it depends on the random variable $\hat{\theta}$. Averaging over the sampling distribution of $\hat{\theta}$, we obtain the expected estimation error

$$E[\hat{q} - q].$$

Using similar techniques as in the previous section, we are able to prove the following result.

**Proposition 7.** *The expected estimation error in the order quantity for the case of a k-dimensional parameter* $\theta = (\theta_1, \ldots, \theta_k)'$ *with unbiased estimator* $\hat{\theta}$ *equals*

$$E[\hat{q} - q] = -tr\left[\mathrm{Cov}(\hat{\theta})\left(\frac{f_q(\nabla F)(\nabla F)' - 2f(\nabla F)(\nabla f)' + f^2(\nabla^2 F)}{2f^3}\right)\right] + o\left(\frac{1}{n}\right) = \frac{\gamma}{n} + o\left(\frac{1}{n}\right),$$

*where* $\gamma \triangleq -tr[I^{-1}(\theta)(\frac{f_q(\nabla F)(\nabla F)' - 2f(\nabla F)(\nabla f)' + f^2(\nabla^2 F)}{2f^3})]$. *The expected estimation error in the order quantity for the case of a one-dimensional parameter* $\theta$ *with unbiased estimator* $\hat{\theta}$ *equals*

$$E[\hat{q} - q] = -\mathrm{Var}(\hat{\theta})\frac{f_q F_\theta^2 - 2fF_\theta f_\theta + f^2 F_{\theta\theta}}{2f^3} + o\left(\frac{1}{n}\right).$$

### 4.1. Unbiased Adjusted Order Quantity Estimation

In this section, we define an adjusted order quantity, denoted as $\hat{q}_a$, which provides an unbiased estimate of the true optimal order quantity. To construct this adjusted order quantity, we modify the newsvendor's naive order quantity $q(\hat{\theta})$ by subtracting a consistent estimate of the second order expected estimation error term identified in Proposition 7 (substituting the consistent estimators $\hat{\theta}$ for $\theta$ and $\hat{q}$ for $q$, in particular to estimate the partial derivatives).

**Proposition 8.** *In the case of a k-dimensional parameter* $\theta$ *with unbiased estimator* $\hat{\theta}$*, if we define the adjusted order quantity as*

$$\hat{q}_a \triangleq \hat{q} + \frac{1}{n} tr\left[I^{-1}(\hat{\theta})\left(\frac{\hat{f}_q(\nabla\hat{F})(\nabla\hat{F})' - 2\hat{f}(\nabla\hat{F}) \times (\nabla\hat{f})' + \hat{f}^2(\nabla^2\hat{F})}{2\hat{f}^3}\right)\right] = \hat{q} - \frac{\hat{\gamma}}{n},$$

*then* $E[\hat{q}_a] = q + o(\frac{1}{n})$ *so that the* $O(1/n)$ *term has been removed from the bias of* $\hat{q}$*, where our consistent estimator is*

$$\hat{\gamma} \triangleq -tr\left[I^{-1}(\hat{\theta})\left(\frac{\hat{f}_q(\nabla\hat{F})(\nabla\hat{F})' - 2\hat{f}(\nabla\hat{F})\left(\nabla\hat{f}\right)' + \hat{f}^2(\nabla^2\hat{F})}{2\hat{f}^3}\right)\right].$$

We next show that the order-quantity estimation error does not materially affect the expected-profit estimation error. That is, the naively estimated order quantity can itself have an estimation error, which can be corrected via Proposition 8; however, the second order profit expected estimation error formula remains unchanged (from its form in Corollary 1), and the expected profit is not materially improved by the estimation error correction of the estimated order quantity (i.e., this change is zero to second order). We work with the general case of $k \geq 1$ real parameters in the vector $\theta$. We next show that the second-order profit expected estimation error for the adjusted order quantity is the same as that of the unadjusted order quantity.

**Proposition 9.** *The profit expected estimation error* $E[\pi(\hat{q}_a, \hat{\theta}) - \pi(\hat{q}_a, \theta)]$ *for* $\hat{q}_a$ *for the case of a k-dimensional parameter* $\theta = (\theta_1, \ldots, \theta_k)'$ *with unbiased estimator* $\hat{\theta}$ *equals*

$$p \cdot tr\left[\mathrm{Cov}(\hat{\theta})\left(\frac{(\nabla F)(\nabla F)'}{f} - \frac{1}{2}\int_0^q \nabla^2 F(x,\theta)dx\right)\right] + o\left(\frac{1}{n}\right)$$
$$= \frac{p}{n} tr\left[I^{-1}(\theta)\left(\frac{(\nabla F)(\nabla F)'}{f} - \frac{1}{2}\int_0^q \nabla^2 F(x,\theta)dx\right)\right] + o\left(\frac{1}{n}\right).$$

*Furthermore, the change in actual expected profit (under the true demand parameter* $\theta$*) is negligible; that is,* $E[\pi(\hat{q}_a, \theta) - \pi(\hat{q}, \theta)] = o(\frac{1}{n})$.

We recognize that the first result in Proposition 9 has the same functional form as that in Corollary 1, which is why the second result shows that the order quantity adjustment results in no additional profit to second order.

### 4.2. Illustrative Example: Log-Normal Demand Distribution

We now study the log-normal demand distribution as an example with a biased order quantity. Note that the exponential distribution has an unbiased order quantity because $E[\hat{q}] = E[\hat{\theta}\ln(p/c)] = \theta\ln(p/c) = q$. Similarly, for normally distributed demand with the unbiased estimates $\hat{\mu}$ and $\hat{\sigma}$, we find $E[\hat{q}] = E[\hat{\mu} + \hat{\sigma}\Phi^{-1}(1-c/p)] = \mu + \sigma\Phi^{-1}(1-c/p) = q$. This generalizes to show that the order quantity is unbiased for any location-scale family of demand distributions when unbiased parameter estimates are used.

Log-normal demand may be represented as $X = e^Y$, where $Y$ (the log of demand) has a normal distribution with mean $\mu$ and standard deviation $\sigma$. The unbiased mean and standard deviation estimates ($\hat{\mu}$ and $\hat{\sigma}$ as used in the normal case) may now be computed using the logs of sampled demands. The unadjusted order quantity for the log-normal distribution is, therefore, $\hat{q} = q(\hat{\theta}) = e^{\hat{\mu}+\hat{\sigma}\Phi^{-1}(1-c/p)}$.

**Proposition 10.** *For the log-normal distribution, the asymptotic estimation error-correction term for the estimated order quantity is*

$$\frac{\gamma}{n} = \frac{\sigma^2(2+\xi^2)q}{4n},$$

*so the adjusted estimated order quantity is* $\hat{q}_a = \hat{q} - \hat{\gamma}/n$, *where* $\hat{\gamma} = \frac{\hat{\sigma}^2(2+\xi^2)\hat{q}}{4}$ *and* $\xi = \Phi^{-1}(1 - c/p)$.

The next result provides the profit expected estimation error when demand is log-normally distributed; note that, per Corollary 1 and Proposition 9, the expression is the same when using the unadjusted biased order quantity.

**Proposition 11.** *For the log-normal distribution, the profit expected estimation error* $E[\pi(\hat{q}_a, \hat{\theta}) - \pi(\hat{q}_a, \theta)]$ *for* $\hat{q}_a$ *equals*

$$\frac{p\sigma}{4n}\left[q(2+\xi^2-\sigma\xi-\sigma^2)\varphi(\xi) + \sigma(3+\sigma^2)\right.$$
$$\left.\times e^{\mu+\sigma^2/2}\Phi(\xi-\sigma)\right] + o\left(\frac{1}{n}\right).$$

*Furthermore, the adjusted expected profit for a log-normal distribution with unbiased estimated parameters* $\hat{\mu}$ *and* $\hat{\sigma}$ *is*

$$\hat{\pi}_a(\hat{q}_a, \hat{\theta}) = (p-c)\hat{q}_a - p\hat{q}_a\Phi\left(\frac{\ln(\hat{q}_a) - \hat{\mu}}{\hat{\sigma}}\right)$$
$$+ pe^{\hat{\mu}+\hat{\sigma}^2/2}\Phi\left(\frac{\ln(\hat{q}_a) - \hat{\mu} - \hat{\sigma}^2}{\hat{\sigma}}\right)$$
$$- \frac{p\hat{\sigma}}{4n}\left[\hat{q}_a(2+\xi^2-\hat{\sigma}\xi-\hat{\sigma}^2)\varphi(\xi)\right.$$
$$\left.+ \hat{\sigma}(3+\hat{\sigma}^2)e^{\hat{\mu}+\hat{\sigma}^2/2}\Phi(\xi-\hat{\sigma})\right].$$

### 4.3. Numerical Verification: Log-Normal Demand Distribution

As in Section 3.5, we perform numerical experiments that provide statistical evidence of nontrivial profit estimation error and now also order-quantity estimation error for the log-normal distribution (a heavy-tailed distribution; see Foss et al. 2013). Demand is simulated from a log-normal distribution with mean 200 and standard deviation 65 so that, on the log scale, we have $\theta = (\mu, \sigma) = (5.248, 0.317)$. The remainder of the experimental setup is as in Section 3.5. We find that the expected profit is 282.1, and the expected estimation error is 3.1. Furthermore, in the left plot of Figure 3 (which uses the bias-adjusted order quantity while exploring profit bias), we again observe strong statistical evidence of the nonzero profit expected estimation error: the unadjusted expected estimation error shows high statistical significance ($t$-values in

**Figure 3.** (Color online) Log-Normal Distribution with Mean 200, Standard Deviation 65, $p = 5$, and $c = 3$ with Samples of Size $n = 25$

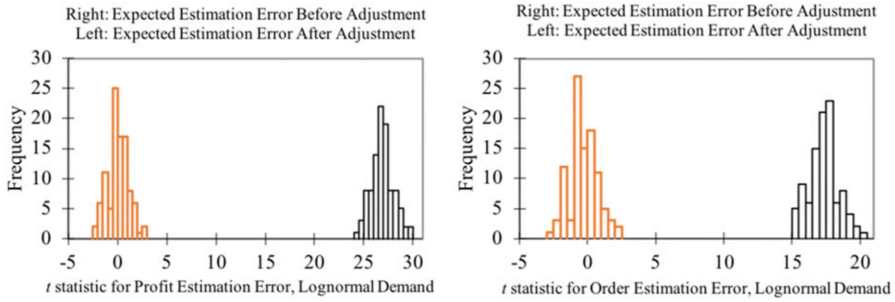


*Note.* The left plot, which presents profit estimation errors, is for the adjusted order quantity $\hat{q}_a$.

the histogram at right are considerably higher than the 1.96 critical value), which is successfully eliminated by the asymptotic correction (histogram at left is centered at zero).

In the right plot of Figure 3, we observe similar evidence for the order quantity: expected estimation error in the unadjusted order quantity $\hat{q}$ shows high statistical significance ($t$-values in the histogram at right are considerably higher than the 1.96 critical value), which is successfully eliminated by the asymptotic correction for the order quantity $\hat{q}_a$ (histogram at left is centered at zero). The naive newsvendor orders, on average, approximately 0.36 to 0.37 more than the true optimal order quantity $q = 175.534$. This does not materially affect the true expected profit as is consistent with Proposition 9: using the unadjusted order quantity $\hat{q}$ reduces true expected profit by only an estimated 0.003 from the expected profit estimated as 282.1 (0.001%) as compared with the material expected estimation error of 3.1 (1.1%) associated with naively computing the profit using $\hat{\theta}$ in place of the true demand parameter $\theta$.

## 5. Conclusion

In this paper, in the context of the newsvendor model, we show that unbiased estimators of distributional parameters, when passed through an optimization operator, result in biased estimates of the optimal objective function value and, depending on the demand distribution, biased order quantities. We derive closed-form second order approximations for the expected estimation errors and use these to form adjusted expected-profit and order-quantity expressions that can be computed by a newsvendor who does not know the true demand distribution. We conduct simulation studies for exponential, normal, and log-normal demand distribution families and find statistically significant estimation errors in each case that are effectively eliminated by our adjustments.

We conclude by commenting on further research. It would be interesting and practically useful to extend this work to the case of sales data (i.e., censored demand data). Finally, Although we focused on analyzing the impact of estimation error for the newsvendor model, we suspect many parametric approaches to optimization under uncertainty that use real data for estimation are significantly affected by estimation error; it might be possible to derive closed-form adjustment terms in other contexts to provide unbiased estimates of meaningful quantities.

## Acknowledgments

The authors thank the department editor, associate editor, and three anonymous referees for very useful feedback that greatly improved the quality of the paper.

## References

APICS (2019) APICS dictionary. Accessed November 1, 2019, http://www.apics.org/apics-for-individuals/publications-and-research/apics-dictionary.

Arnold B, Balakrishnan N, Nagaraja H (2008) *A First Course in Order Statistics* (Society for Industrial and Applied Mathematics, Philadelphia).

Ban G, Rudin C (2019) The big data newsvendor: Practical insights from machine learning. *Oper. Res.* 67(1):90–108.

Ban G, Gallien J, Mersereau A (2019) Dynamic procurement of new products with covariate information: The residual tree method. *Manufacturing Service Oper. Management* 21(4):798–815.

Bentzley J (2017) The overlooked forecasting flaw: Forecast bias and how to tackle it. Accessed November 1, 2019, https://www.linkedin.com/pulse/overlooked-forecasting-flaw-forecast-bias-how-tackle-jim-bentzley.

Besbes O, Muharremoglu A (2013) On implications of demand censoring in the newsvendor problem. *Management Sci.* 59(6):1407–1424.

Burnetas A, Smith C (2000) Adaptive ordering and pricing for perishable products. *Oper. Res.* 48(3):436–443.

Cachon G, Terwiesch C (2012) *Matching Supply with Demand: An Introduction to Operations Management*, 3rd ed. (McGraw-Hill, New York).

Cassar G, Gibson B (2008) Budgets, internal reports, and manager forecast accuracy. *Contemporary Accounting Res.* 25(3):707–737.

Chu L, Shanthikumar J, Shen Z (2008) Solving operational statistics via a Bayesian analysis. *Oper. Res. Lett.* 36(1):110–116.

Durand R (2003) Predicting a firm's forecasting ability: The roles of organizational illusion of control and organizational attention. *Strategic Management J.* 24(9):821–838.

Foss S, Korshunov D, Zachary S (2013) *An Introduction to Heavy-Tailed and Subexponential Distributions*, 2nd ed. (Springer, New York).

Gallego G, Moon I (1993) The distribution free newsboy problem: Review and extensions. *J. Oper. Res. Soc.* 44(8):825–834.

Godfrey G, Powell W (2001) An adaptive, distribution-free algorithm for the newsvendor problem with censored demands, with applications to inventory and distribution. *Management Sci.* 47(8):1101–1112.

Hammersley J, Morton K (1956) A new Monte Carlo technique: Antithetic variables. *Math. Proc. Cambridge Philos. Soc*. 52(3):449–475.

Hannah L, Powell W, Blei D (2010) Nonparametric density estimation for stochastic optimization with an observable state variable. Lafferty JD, Williams CKI, Shawe-Taylor J, Zemel RS, Culotta A, eds. *Advances in Neural Information Processing Systems*, vol. 23 (Curran Associates, Red Hook, NY), 820–828.

He B, Dexter F, Macario A, Zenios S (2012) The timing of staffing decisions in hospital operating rooms: Incorporating workload heterogeneity into the newsvendor problem. *Manufacturing Service Oper. Management* 14(1):99–114.

Huh T, Rusmevichientong P (2009) A nonparametric asymptotic analysis of inventory planning with censored demand. *Math. Oper. Res.* 34(1):103–123.

Huh T, Levi R, Rusmevichientong P, Orlin J (2011) Adaptive data-driven inventory control with censored demand based on Kaplan-Meier estimator. *Oper. Res.* 59(4):929–941.

Ittner C, Larcker D (1998) Are nonfinancial measures leading indicators of financial performance? An analysis of customer satisfaction. *J. Accounting Res.* 36:1–35.

Jabot F (2015) Why preferring parametric forecasting to nonparametric methods? *J. Theoret. Biology* 372:205–210.

Johnson N, Welch B (1939) On the calculation of the cumulants of the chi-distribution. *Biometrika* 31(1/2):216–218.

Johnson N, Kotz S, Balakrishnan N (1994) *Continuous Univariate Distributions*, vol. 1, 2nd ed. (Wiley, New York).

Kim S, Pasupathy R, Henderson S (2015) A guide to sample average approximation. Fu M, ed. *Handbook of Simulation Optimization*, International Series in Operations Research and Management Science, vol. 216 (Springer, New York), 207–243.

Kleywegt A, Shapiro A, Homem-De-Mello T (2002) The sample average approximation method for stochastic discrete optimization. *SIAM J. Optim.* 12(2):479–502.

Kunnumkal S, Topaloglu H (2008) Using stochastic approximation methods to compute optimal base-stock levels in inventory control problems. *Oper. Res.* 56(3):646–664.

Lee T, Adam E (1986) Forecasting error evaluation in material requirements planning (MRP) production-inventory systems. *Management Sci.* 32(9):1186–1205.

Lehman E (1983) *Theory of Point Estimation*, 1st ed. (Wiley, New York).

Levi R, Shi C (2013) Approximation algorithms for the stochastic lot-sizing problem with order lead times. *Oper. Res.* 61(3):593–602.

Levi R, Janakiraman G, Nagarajan M (2008a) A 2-approximation algorithm for stochastic inventory control models with lost sales. *Math. Oper. Res.* 33(2):351–374.

Levi R, Perakis G, Uichanco J (2015) The data-driven newsvendor problem: New bounds and insights. *Oper. Res.* 63(6):1294–1306.

Levi R, Roundy R, Shmoys D (2007a) Provably near-optimal sampling-based policies for stochastic inventory control models. *Math. Oper. Res.* 32(4):821–839.

Levi R, Pál M, Roundy R, Shmoys D (2007b) Approximation algorithms for stochastic inventory control models. *Math. Oper. Res.* 32(2):284–302.

Levi R, Roundy R, Shmoys D, Truong V (2008b) Approximation algorithms for capacitated stochastic inventory control models. *Oper. Res.* 56(5):1184–1199.

Levi R, Roundy R, Truong V, Wang X (2017) Provably near-optimal balancing policies for multi-echelon stochastic inventory control models. *Math. Oper. Res.* 42(1):256–276.

Lipson M (2019) Financial analytics toolkit: Financial statement forecasting. Case Study UV 7915, University of Virginia, Charlottesville.

Liyanage L, Shanthikumar J (2005) A practical inventory control policy using operational statistics. *Oper. Res. Lett.* 33(4):341–348.

Madeka D, Swiniarski L, Foster D, Razoumov L, Wen R, Torkkola K (2018) Sample path generation for probabilistic demand forecasting. Working paper, Amazon, New York.

Natarajan K, Sim M, Uichanco J (2018) Asymmetry and ambiguity in newsvendor models. *Management Sci.* 64(7):3146–3167.

Oliva R, Watson N (2009) Managing functional biases in organizational forecasts: A case study of consensus forecasting in supply chain planning. *Production Oper. Management* 18(2):138–151.

Parks J (2018) Breaking down forecasting: The power of bias. Accessed November 1, 2019, https://www.ibm.com/blogs/think/2018/04/forecasting-and-bias/.

Perakis G, Roels G (2008) Regret in the newsvendor model with partial information. *Oper. Res.* 56(1):188–203.

Scarf H (1958) A min-max solution to an inventory problem. Arrow K, Karlin S, Scarf H, eds. *Studies in Mathematical Theory of Inventory and Production* (Stanford University Press, Palo Alto, CA), 201–209.

Shi C, Zhang H, Chao X, Levi R (2014) Approximation algorithms for capacitated stochastic inventory systems with setup costs. *Naval Res. Logist.* 61(4):304–319.

Siegel A, Woodgate A (2007) Performance of portfolios optimized with estimation error. *Management Sci.* 53(6):1005–1015.

Singh S (2017) A critical look at measuring and calculating forecast bias. Accessed November 1, 2019, https://demand-planning.com/2017/01/13/a-critical-look-at-measuring-and-calculating-forecast-bias/.

Watson R (1987) The effects of demand-forecast fluctuations on customer service and inventory cost when demand is lumpy. *J. Oper. Res. Soc.* 38(1):75–82.