Unbiased estimation of maximum expected profits in the Newsvendor Model: a case study analysis

Halkos, George and Kevork, Ilias

University of Thessaly, Department of Economics

17 August 2012

Unbiased Estimation of Maximum Expected Profits in the Newsvendor Model: A Case Study Analysis

By

George E. Halkos and Ilias S. Kevork
Department of Economics, University of Thessaly

ABSTRACT

In the current paper we study a real life inventory problem whose operating conditions match to the principles of the classical newsvendor model. Applying appropriate tests to the available sample of historical demand data, we get the sufficient statistical evidences to support that daily demand is stationary, uncorrelated, and normally distributed. Given that at the start of each day, the selling price, the purchasing cost per unit, and the salvage value are known, and do not change through the whole period under investigation, we derive exact and asymptotic prediction intervals for the daily maximum expected profit. To evaluate their performance, we derive the analytic form of three accuracy information metrics. The first metric measures the deviation of the estimated probability of no stock-outs during the day from the critical fractile. The other two metrics relate the validity and precision of the two types of prediction interval to the variability of estimates for the ordered quantity. Both theoretical and empirical analysis demonstrates the importance of implications of the loss of goodwill to the adopted inventory policy. Operating the system at the optimal situation, this intangible cost element determines the probability of no stock-outs during the day, and assesses the precision of prediction intervals. The rising of the loss of goodwill leads to smaller estimates for the daily maximum expected profit and to wider prediction intervals. Finally, in the setting of the real life newsvendor problem, we recommend the asymptotic prediction interval since with samples over 25 observations this type of interval has higher precision and probability to include the daily maximum expected profit almost equal to the nominal confidence level.

Keywords: Newsvendor model, Loss of goodwill, Target inventory measures, Prediction interval, Accuracy information metric.

JEL Codes: C13: Estimation; C53: Forecasting models; M11: Production Management; M21: Business Economics.

1. Introduction

In the majority of papers in stock control, optimal policies have been developed with hypothetical conditions under which inventory systems operate. Without the existence of real life inventory problems and the relevant data, the practicality of derived optimal inventory policies and the validity of the theoretical results cannot be assessed. Even though the stochastic law of generating demand was known, the application of the theoretical formulae which ensure the optimal situation by replacing population parameters by their estimates might lead to unexpected results for stock control. The choice of the best estimation procedure might not imply the desired target inventory measures from the management point of view. So, the variability of any type of estimates into the theoretical formulae which give the optimal target inventory measures should be evaluated regarding managerial aspects of inventory.

For the current work, we have managed to collect data for a real life inventory problem whose operating conditions match completely to the principles of the classical newsvendor model. The data included the daily demand of a perishable product for 30 days, and the values for the selling price, the purchase cost, and the salvage value. Starting from the 27^{th} day and sequentially reaching the 31^{st} day, for the demand of each one of the five days under consideration, we applied traditional tests for stationarity, autocorrelation, and normality to the available sample of all the previous days. The results gave the sufficient statistical evidences to accept stationarity, normality, and independence for the daily demand. Besides, the sizes of the estimated coefficients of variation justify the choice of the normal distribution against the truncated normal at point zero, following the recommendations of Law (1997) and Halkos & Kevork (2012a). The latter distribution excludes the occurrence of negative values which is very likely to happen when coefficients of variation are large.

The second important characteristic of the real life problem is the fact that the selling price, the purchase cost per unit of product, and the salvage value were known at the start of any day and remained constant through the whole period of 30 days. Besides, at the start of any day, the inventory system was starting with stock equal to the received order quantity. Under these conditions, and assuming that the inventory system operates at the optimal situation, at the start of each one of the five days under consideration we illustrate the process of determining the order quantity and forecasting the daily expected economic profit at the end of the day. At the optimal situation, the probability of no stock-outs during the day equals to a critical fractile which depends upon the overage and the underage cost. An important intangible cost element which determines the underage cost is the loss of goodwill. When the loss of goodwill is zero then the economic profit equals to the accounting profit which is determined from the tangible revenue and cost elements, namely, the price, the purchase cost and the salvage value.

At the optimal situation, two target inventory measures will be under consideration in the current work; the optimal order quantity and the maximum expected profit. Since the mean and the variance of the daily demand are unknown, these measures should be estimated. In the current work we use the maximum likelihood estimators of the mean and the variance of the daily demand. An appropriate adjustment is made to the estimator of the standard deviation in order to eliminate its biasedness. By replacing the sample mean and the unbiased estimator of the standard deviation into the theoretical formula which gives the optimal order quantity, for the latter we obtain an unbiased estimator. Then, an unbiased estimator for the maximum expected profit is derived by replacing the unbiased estimator of the optimal order quantity into the theoretical expression which gives the expected profit of the classical newsvendor model. Finally, the exact and the asymptotic distributions of the maximum expected profit are derived and the corresponding prediction intervals are constructed.

Syntetos et al. (2010) pointed out the usefulness and the necessity of developing appropriate accuracy information metrics to relate efficiency of stock control systems to demand estimation (or forecasting) performance. One of the oldest metric is the expected total operating cost (ETOC) introduced by Hayes (1969) who investigated the inaccuracy in the estimation of target inventory measures in the newsvendor model. Katircioglou (1996) used the ETOC to study the effect of biasing in estimating target inventory measures in three different inventory models when demand follows the normal and gamma distributions. Modeling demand in the newsvendor model by the Johnson Translation System and using the ETOC, Akcay et al. (2011) quantified the inaccuracy in estimating the optimal order quantity as a function of the historical data, the critical fractile, and the shape parameters of the demand distribution.

For the (Q,R) continuous inventory system, Silver & Rahnama (1987) provided a methodology to ascertain the amount of biasing and the resulting expected percentage cost penalty, when demand parameters are replaced by their estimates. Liyanage & Shanthikumar (2005) introduced the a-priori expected profit (an equivalent function to the ETOC) to quantify the bias in the estimated target inventory measures for the newsvendor model under an exponential demand. Ali et al. (2011) explored the relationship between accuracy of demand forecast and inventories holdings in a supply chain with one retailer and one manufacturer when demand at the retailer is stationary and is generated by one of the three models AR(1), MA(1), and ARMA(1,1).

In the current work, we derive the analytic form of three accuracy information metrics. These metrics relate the sampling distributions of the sample mean and the unbiased estimator of the standard deviation to the variability of estimates for the two target inventory measures. The first metric measures the deviation of the estimated probability of not experiencing stock-outs during the day from the critical fractile. The other two information metrics are related to the validity and precision of the two types of prediction interval for the maximum expected profit. Particularly, we evaluate the probability the asymptotic prediction interval to include the maximum expected profit in finite samples, and for both the exact and the asymptotic prediction intervals we examine their relative expected half-length. The latter metric is defined as the expected half-length divided by the maximum expected profit.

The necessity of the first accuracy information metric arises since by using estimates of demand parameters in the theoretical formula which gives the optimal order quantity, the latter is incorrectly computed and the critical fractile does not give the requested probability of no stock-outs during the period. To resolve this problem under normal demand with unknown mean and unknown variance, Ritcken & Sankar (1984) made an adjustment to the safety stock coefficient and developed the appropriate estimator for the optimal order quantity which ensures that the estimated probability of no stock-outs during the period is equal to the requested critical fractile. Janssen et al. (2009) handled the same problem and modified Ritcken and Sankar's estimator when demand follows the normal distribution with unknown mean but known variance. Halkos & Kevork (2012b) resolved the same problem and derived the estimator which ensures the requested probability of no stock-outs during the period when demand is exponential. However, the main drawback of such estimators is their biasedness. In the current paper, and in the context of the real life newsvendor problem, we illustrate that using the suggested unbiased estimator for the optimal order quantity this problem is resolved with historical demand data of length over 25 observations.

The other two accuracy information metrics were adopted in the paper of Kevork (2010), who examined the implications of estimating demand parameters on the variability of estimates for the optimal order quantity and the maximum expected profit when demand is normal. The author derived the asymptotic distribution of a biased estimator for the maximum expected profit, and evaluated the metrics through Monte Carlo simulations in a hypothetical experimental framework regarding the values of the price, the purchasing cost, the salvage value, and the loss of goodwill. Halkos & Kevork (2012c) also used these two metrics to show that with an exponential or Rayleigh demand, the validity and precision of estimates for the maximum expected profit do not depend upon the values assigned to the revenue and cost parameters of the newsvendor model.

Summarizing, therefore, in the area of the classical newsvendor model, the contribution of this paper is two-fold. At the theoretical level, the paper establishes an unbiased estimator for the maximum expected profit when demand follows the normal distribution with unknown mean and unknown variance. Then it derives for the first time its exact and asymptotic distribution when the revenue parameters and the purchasing cost are known at the start of the period and remain constant through a number of successive periods. Under such conditions, first we show that the loss of goodwill is an increasing function of the critical fractile. Second, we prove that the estimated maximum expected profit is minimized and the precision of the corresponding prediction intervals is maximized when the loss of goodwill is zero. At the practical level, we illustrate the verification of the theoretical findings in the context of forming the inventory policy in a real life newsvendor problem. To the extent of our knowledge, the principles of this inventory system operation and the trading of the newsvendor product is met in the literature for the first time.

The aforementioned arguments and remarks lead the rest of the paper to be structured as follows. The next section gives a literature review consisting of indicative papers in the area of the newsvendor model, which assessed the proposed methodologies on real data. In section 3, we give the description of the real life newsvendor problem and establish the theoretical framework which matches to the system operating principles. In section 4, we develop the estimators for the optimal order quantity and the maximum expected profit and comment on their properties. We also evaluate analytically the deviations of the estimated probability of no stock-outs from the critical fractile. In section 5, we derive the exact and the asymptotic prediction intervals for the maximum expected profit and obtain the analytic forms of the accuracy information metrics regarding their precision and validity. The verification of the theoretical findings is illustrated in section 6 using the demand data of the real problem and the values for the revenue and cost parameters. Finally, the last section concludes the paper summarizing the most important findings.

2. Literature Review

In the current section we review a number of recent relevant papers from the area of the newsvendor model. These papers are classified into two main streams of research. In the first stream, information about the real operating conditions of inventory systems has been made available to the authors, and the derived optimal inventory policies have been assessed on real data from such systems. The second steam of research includes papers which study the behavior of decision makers who participate in computerized laboratory experiments. These experiments are organized according to the newsvendor settings, and *"artificially"* real data for the order quantity are recorded from each participant for a number of decision periods.

Regarding the first stream of research, Choi et al. (2011) derived the optimal stocking policy of a fast fashion retailer by maximizing an objective function called safety-first objective. This function depends on the expected profit, the standard deviation of profit, and a pre-determined target profit threshold. The optimal policy was applied to real data for one particular style from a fast fashion brand in Hong-Kong. The data included periodic forecast for the quantity of the product, cost and revenue parameters, and profit targets. Demand distribution was estimated from the periodic forecasts and appeared to be normal.

Su & Pearn (2011) developed a statistical test of hypothesis to select, among two newsvendor products, that one which has a higher probability of achieving a target profit under the optimal ordering policy. The proposed statistical methodology was applied to select among basic, intermediate and high level of an English teaching magazine which are supplied in the beginning of each month by a magazine publisher in Taipei, Taiwan. The data included demand for 100 months plus values for the cost and revenue parameters which remained fixed for the whole period under investigation.

Mostard et al. (2011) performed a case study to compare new proposed forecasting methods for demand to existing ones which are based on advance demand information and on expert judgments. Based on a data set of around 700 stock keeping units of a mail order apparel company in Netherlands, and for two selling seasons, the authors compared the accuracy of the proposed and the existing forecasting methods based on advance demand information. Also for a smaller number of stock keeping units, the authors compared expert judgments to methods based on advance demand information. Data included sales and lost demand which is registered through the call center and the voice response system.

Beutel & Minner (2011) compared three approaches to estimate the optimal order quantity when for a sample of selling seasons data are available for demand and several explanatory variables of demand, and price changes between successive selling seasons. The authors assumed that demand follows the normal and gamma distributions, and its average size is a linear function of price (with negative slope). The three approaches were assessed by using real data from 64 stores of a large European retail chain for a certain salad which has the characteristics of a newsvendor product. The data included daily sales, prices, and weather information. From the sample, the authors excluded those selling seasons for which stock outs were observed.

Olivares et al. (2008) developed a general structural model to impute the overage and underage costs of the classical newsvendor model based on observed decisions for the order quantity. For two specified models, the authors obtained consistent estimates of the parameters and derived the asymptotic distribution of the corresponding estimators. The proposed methodology was applied to a decision concerning the length of operating room time which should be reserved to a specific surgical case using real data from cardiac surgery. The dataset included for 258 cardiac surgery cases the reserved and the actual operating room time, as well as, patient and procedure characteristics.

Mostard & Teunter (2006) derived an exact formula to compute the optimal order quantity when the newsvendor product is returned with a certain probability and can be resold any number of times. The application of this formula was illustrated using real data of a mail order retailer. The data were available for 427 products but for only one selling season, and included gross and net demand (gross demand minus number of resalable returns), return rates, salvage values, purchase costs, and selling prices. Unfortunately, no indication was given about the right value of the loss of goodwill. To obtain the results, the author assumed normality for the gross demand and then he showed that net demand is also approximately normally distributed.

To the extent of our knowledge, the earliest experimental study in the second stream of research was conducted by Schweitzer & Cachon (2000). For 15 hypothetical successive selling seasons, 34 MBA students made decisions for the order quantity knowing that demand was uniformly distributed. The authors concluded that there were several behavioral factors which forced in this laboratory experiment the participants to order quantities which were deviating from the optimal ones.

The laboratory experiment of Benzion et al. (2008) included 60 sophomores and junior management students who were divided into four groups. The participants in the first group knew a-priori that demand was uniformly distributed, and the participants of the second group that demand was normally distributed. Decisions for the order quantity were made for 100 successive selling seasons. For all the possible scenarios which were adopted in this experiment, the authors concluded that the purchase order converged to a size between the average demand and the optimal order quantity.

The number of participants increased to 121 students in the experiment of Benzion et al. (2010). Participants were divided into two clusters of four groups. For the one cluster, demand was generated from the uniform distribution, while for the other cluster demand was normally distributed. In each cluster, only two groups knew a-priori the type and the parameters of demand distribution. The experimental findings indicated that although participants who knew the demand distribution behaved differently, the knowledge of demand distribution did not lead to order quantities closer to the optimal ones.

More about this second stream of research and its relation to the growing area of behavioral operations management can be found in Feng et al. (2011) who conducted a similar type of laboratory experiment to examine cross-national differences in ordering policies between Chinese and American decision makers.

Our crucial remark for the majority of the aforementioned works is the very little importance which has been given to the loss of goodwill. Authors ignored this intangible cost element or indirectly they assumed that it is zero, and for the newsvendor settings such scenarios are not realistic. In the current work, we demonstrate how to handle this cost element and examine its implications on the optimal inventory policies. Unfortunately, when the other cost and revenue elements are known and remain constant through successive periods, at the optimal situation a positive loss of goodwill (a) determines the value of the critical fractile and hence the requested probability of no stock-outs during the period, and (b) assesses the precision of prediction intervals for the maximum expected economic profit at the end of each period.

3. The real-life inventory problem and a newsvendor model specification

The city of Volos is the capital of the county of Magnesia located in Thessaly which is one of the largest in size and population regions in Greece. Volos has one of the biggest commercial ports in Greece, and the port of Volos is one of the very few intermediate stations to reach by sea the Sporades Islands. Northern Sporades are located to the east of Volos and include the islands of Skiathos, Skopelos, Alonnissos, as well as, other smaller islands. Along the coastline of the islands, there are beautiful beaches, and with a nice continental climate, the islands are very attractive to tourists.

Every morning during the summer, a number of passenger and car-ferry ships are leaving the port of Volos and sailing for a one-day voyage for Sporades. The ships carry natives and tourists, vehicles (cars and Lorries) and goods from Volos to the islands, from island to island, and finally from the islands to Volos. At the night of the same day, the ships return to the port of Volos.

During the summer of 2011, the authors managed to collect real data for the daily demand of a certain type of product which was being sold in the canteen of one of the ships. This product was a traditional type of cheese-pie which is produced at the local market of Volos and its name is «sfoliata cheese-pie». Every morning, before the ship sailed for the islands, the canteen was receiving the amount of sfoliata cheese-pies which had been ordered from the previous day. This amount should be enough to satisfy the demand which would occur during the whole voyage, since no other delivery of pies reached the ship from the ports of the islands where it sailed into. At night, when the ship returned to Volos, any pie which had not been sold during the voyage was thrown away. The canteen was purchasing each pie at €1.20 and was selling it at €2.95. No fixed costs were being incurred with the delivery of the ordered quantity, and *the purchase cost and the selling price remained the same for all the summer of 2011*.

The description of the supply and trading processes of the sfoliata cheese-pie, as well as, the physical characteristics of the specific product, satisfy the conditions which are required for considering the classical newsvendor model. Demand of the product lasts one day, and so the day is considered as the period, or the inventory cycle. Since any pie which is not sold during the voyage of the previous day is thrown away at night when the ship returns to the port (salvage value is zero), at the start of the next day the canteen has as stock only the quantity which receives in the morning before the ship sails for the islands. No fixed costs are incurred with the delivery of the ordered quantity at the start of the day, and no other delivery of the product is made during the period, namely, during the day.

Our analysis will be focused on the period from 21/7/2011 up to 22/8/2011 which constituted the peak-time for tourism. Within this period, data for demand were available only for 30 days, since the ship was not in service for three successive days. Figure 1 displays the time-series plot of demand. For confidentiality reasons, the values of demand are missing on the vertical axis, and upon the staff of the canteen request, we cannot disclose the name of the ship and cannot give detailed information for the process of collecting the demand data. Also, it is easily realized that no data about daily sales and daily ordered quantities for the product were made available to us from the canteen for obvious reasons.

Being at the start of the $(n+1)^{th}$ day, with $n = 27,28,...,31$, and having available demand data for a sample of the previous n successive days, we applied in E-Views the classical statistical tests for the existence of stationarity, autocorrelation, and normality. The results of the tests are presented in table 1.

Figure 1: Time-series plot of daily demand

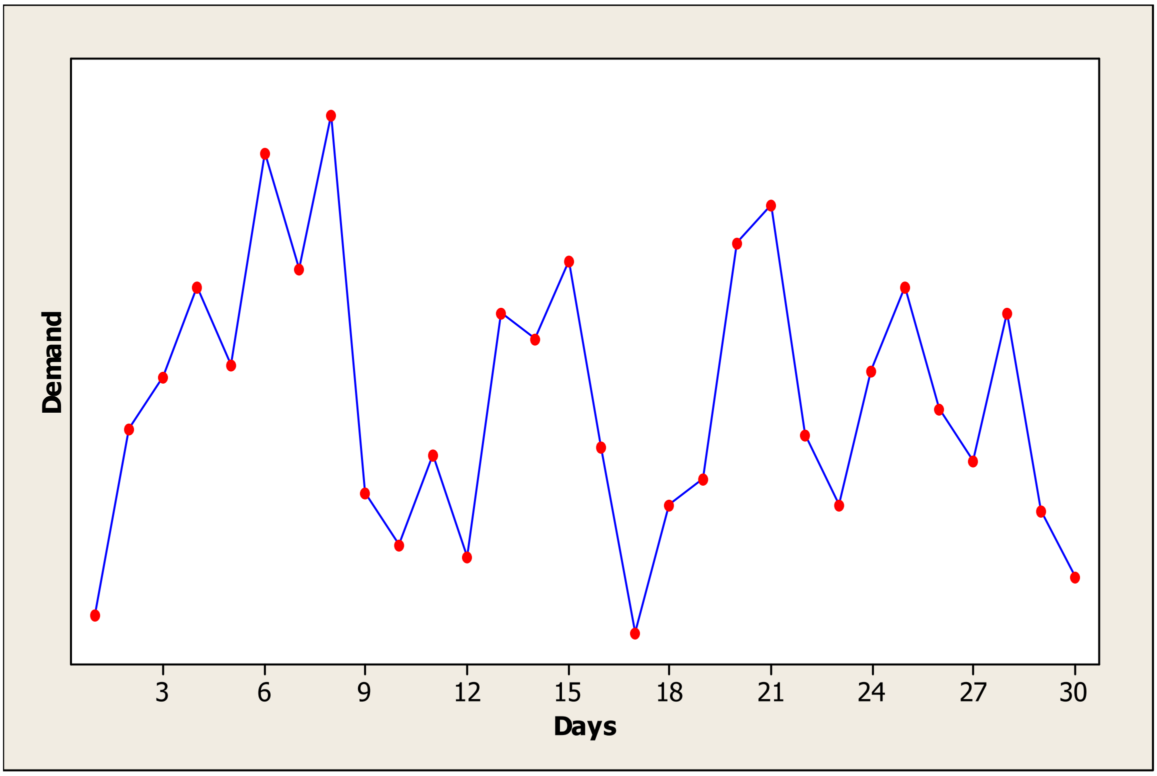

Table 1: Results of the statistical tests

n	Testing for a unit root			Testing for autocorrelation				Testing for normality
	p-value of (D-F) test	p-value for variables		p-values for the Ljung-Box test				p-value of Jarque-Bera test
		Intercept	$Demand_{t-1}$	Lag 3	Lag 6	Lag 9	Lag 12	
26	0.0091	0.0011	0.0010	0.065	0.103	0.247	0.347	0.728
27	0.0081	0.0010	0.0009	0.059	0.107	0.269	0.398	0.721
28	0.0057	0.0006	0.0006	0.074	0.124	0.265	0.337	0.753
29	0.0041	0.0005	0.0004	0.103	0.141	0.289	0.438	0.701
30	0.0063	0.0008	0.0006	0.118	0.169	0.253	0.390	0.635

To explore if at the start of the $(n+1)^{th}$ day, the realization of demand for the previous n days comes from a stationary time series, we applied the classical Dickey-Fuller test for the existence of a unit root (e.g. Halkos & Kevork, 2005). We used two test equations, with the first to include only an intercept term, and the second to include both trend and intercept terms. Since the trend coefficient appeared as non-statistically significant, in table 1 we present the results for the test equation including only the intercept term. For testing autocorrelation we used the Box-Ljung test statistic reporting the p-values at different lags, and finally for testing normality, the Jarque-Bera test was applied each time to the available sample of demand data. For the latter two tests see for example Harvey (1992).

From the reported p-values, being at the start of any of the five days under consideration, and using the sample of demand data from the previous n successive days, (a) at level of significance 1% we reject the hypotheses of a unit root, (b) at level of significance 5% we cannot reject the hypothesis of a white noise series, and (c) at level of significance 10% we cannot reject the null hypothesis of a normal distribution for the daily demand. Consequently, at the start of the day 27th, 28th, 29th, 30th, and 31st, we have sufficient statistical evidences to support that daily demand, X, is stationary, uncorrelated, and normally distributed.

Consequently, to specify the ordering policy of the canteen regarding the newsvendor product «sfoliata cheese-pie», we shall accept that each day demand, X, is formed independently and follows the normal distribution with mean μ and variance σ^2. Denoting also by Q the order quantity, p the selling price, c the purchase cost per unit of product, v the salvage value, and s the loss of goodwill per unit of product, the expected profit for the classical newsvendor model is derived in Silver et al. (1998) and is given by,

$$E(\pi) = (p-v)\mu - (c-v)Q - (p-v+s)\{\phi_z - z(1-\Phi_z)\}\sigma =$$

$$= (p-c)Q + s(Q-\mu) - (p-v+s)\{(Q-\mu)\Phi_z + \sigma\phi_z\}, \qquad (1)$$

where $z = (Q-\mu)/\sigma$, and ϕ_z, Φ_z are respectively the probability density function and the cumulative distribution function of the standard normal evaluated at z.

Providing that the coefficient of variation of daily demand is not large (e.g. Halkos & Kevork, 2012a), the optimal order quantity, Q^*, which maximizes (1) satisfies the sufficient optimality condition

$$\Phi_{z_R} = \Pr(X \leq Q^*) = \Pr(Z \leq z_R) = \frac{p-c+s}{p-v+s} = R ,$$

and is computed from $Q^* = \mu + z_R \cdot \sigma$. Lapin (1994) defines the loss of goodwill, s, as the present value of future profits which are expected to be lost from present unsatisfied customers who will not come back to the business. So, setting $s = \delta(p - c)$ in the critical fractile equation $R = (p - c + s)/(p - v + s)$, we obtain

$$C_u + C_o = p - v + s = \frac{(1+\delta)(p-c)}{R}, \tag{2}$$

where $C_u = p - c + s$ is the underage cost and $C_o = c - v$ is the overage cost. Then replacing Q in (1) with $Q^* = \mu + z_R \cdot \sigma$, and using (2), the maximum expected profit per monetary unit profit (e.g. per euro's profit) is given by

$$\xi^* = \frac{E(\pi)^*}{p-c} = \mu - \frac{(1+\delta)}{R}\phi_{z_R}\sigma \tag{3}$$

Further, solving the critical fractile equation with respect to s, the parameter δ is determined from

$$\delta = \frac{p-v}{p-c} \cdot \frac{R}{1-R} - \frac{1}{1-R}. \tag{4}$$

When the optimality condition holds, the critical fractile R gives the probability of no stock-outs during the period, and is equivalent to the cycle service level (denoted as P_1) which is used in other types of inventory models (e.g. Syntetos et al., 2010). In the context of the classical newsvendor model, R can take any value between zero and one. Particularly, following Schweitzer & Cachon (2000), when the specified values for p, v, c, s, give an $R < 0{,}5$ (or $R > 0{,}5$), then the product is classified as low-profit (or high-profit respectively).

From (4), it is easily deduced that δ is increasing in R since $d\delta/dR = (c - v)/\{(p - c)(1 - R)^2\} > 0$ with limits $\lim_{R \to 0} \delta = -1$ and $\lim_{R \to 1} \delta = \infty$. To ensure, therefore, that $\delta \geq 0$, R should satisfy the condition $R \geq (p - c)/(p - v)$. For the real life problem with the «sfoliata cheese-pie» to be the newsvendor product under consideration, this condition holds for any $R \geq (2.95-1.20)/(2.95-0) \approx 0.593$. So the product is considered as high-profit no matter which values the loss of goodwill takes on.

4. Unbiased estimators for Q^* and ξ^*

Suppose that $X_1, X_2, \ldots, X_n$ is a sequence of independent normal random variables with mean μ and variance σ^2 representing demand for a sample of the most recent n successive periods. Since in practice μ and σ^2 are unknown, the optimal order quantity, Q^*, cannot be computed precisely and should be estimated. Being at the start of period $n+1$, the most traditional approach for estimating Q^* is to replace μ and σ^2 by their estimates in the theoretical formula $Q^* = \mu + z_R \sigma$.

In the current work we shall use for μ and σ^2 their corresponding maximum likelihood (ML) estimators $\overline{X}_n = \sum_{t=1}^{n} X_t \Big/ n$ and $\hat{\sigma}_n^2 = \sum_{t=1}^{n} (X_t - \overline{X}_n)^2 \Big/ n$. Replacing, however, the ML estimators of μ and σ^2 into the theoretical formula of Q^*, the resulting estimator is biased, and this is due to the fact that $\hat{\sigma}_n$ is a biased estimator for σ. This is verified as $E(\hat{\sigma}_n) = g_n \sigma$, where $g_n = (2/n)^{1/2} \Gamma(n/2) / \Gamma((n-1)/2)$ and $\Gamma(x)$ is the gamma function evaluated at x (e.g. Lindgren, 1976, pp. 341). It is easily deduced that such an estimator for Q^* would be also biased even if we used, instead of $\hat{\sigma}_n^2$, the unbiased estimator of the population variance.

To develop, therefore, an unbiased estimator for Q^*, appropriate adjustments should be made to $\hat{\sigma}_n$ so that to build an unbiased estimator for the standard deviation of demand per period. Denoting the unbiased estimator of σ by $\hat{V}_n$ and setting $\hat{V}_n = g_n^{-1} \cdot \hat{\sigma}_n$, it is easily verified that the following estimator for Q^*

$$\hat{Q}_{n+1} = \overline{X}_n + z_R \hat{V}_n = \overline{X}_n + z_R g_n^{-1} \hat{\sigma}_n \tag{5}$$

is unbiased.

However, the above process of developing the estimator $\hat{Q}_{n+1}$ has a major drawback from a theoretical point of view. Having replaced μ with $\overline{X}_n$ and σ with $\hat{V}_n$, we shall take order quantities which differ from the optimal one, and so the requested critical fractile R might not be attained. This remark makes necessary the evaluation of the actual critical fractile, which is defined as the probability at the start of period $n+1$ the estimator $\hat{Q}_{n+1}$ to meet the realized demand of this period. We shall denote this probability by R_{act}.

Syntetos et al. (2010) pointed out the usefulness of developing appropriate accuracy information metrics to relate efficiency of stock control systems to demand estimation (or forecasting) performance. Under this framework, we suggest as a first metric the size of discrepancies $R - R_{act}$. Evaluating these discrepancies at different values of R, we relate the actual critical fractile to the unbiased estimators for demand distribution, $\overline{X}_n$, $\hat{V}_n$, as well as, to the statistical properties of $\hat{Q}_{n+1}$.

To evaluate $R - R_{act}$, first we should obtain the values for R_{act}. By considering $\hat{Q}_{n+1}$ as an estimator for the R^{th} percentile of the demand distribution, in the appendix we show that R_{act} is computed analytically from

$$R_{act} = \Pr\left(t_{n-1} \leq \sqrt{\frac{n-1}{n+1}} g_n^{-1} z_R \right), \tag{6}$$

where t_{n-1} is the central student-t distribution with $n-1$ degrees of freedom. Table 2 displays the values of R_{act} for different combinations of n and R. To compute precisely the ratio of the two gamma functions which enter into the formula which determines g_n, we used the transformation $\exp[\ln(\Gamma(n/2)) - \ln(\Gamma((n-1)/2))]$.

Table 2: Values of R_{act} which are attained by using $\hat{Q}_{n+1}$

	Desired critical fractile							
n	R=0.2	R=0.3	R=0.4	R=0.6	R=0.8	R=0.9	R=0.95	R=0.99
5	0.230	0.319	0.409	0.591	0.770	0.859	0.907	0.957
10	0.215	0.310	0.405	0.595	0.785	0.880	0.929	0.976
15	0.210	0.307	0.403	0.597	0.790	0.886	0.936	0.981
20	0.208	0.305	0.402	0.598	0.792	0.890	0.940	0.984
25	0.206	0.304	0.402	0.598	0.794	0.892	0.942	0.985
30	0.205	0.303	0.402	0.598	0.795	0.893	0.943	0.986
40	0.204	0o.303	0.401	0.599	0.796	0.895	0.945	0.987
50	0.203	0.302	0.401	0.599	0.797	0.896	0.946	0.988
100	0.202	0.301	0.401	0.599	0.798	0.898	0.948	0.989

From the data of table 2, R_{act} is greater than R when the product is low-profit (R<0,5) and smaller than R for high-profit products (R>0,5). Besides, the discrepancies $R - R_{act}$ are getting larger as R is taking on values closer either to zero or to one. In spite of that, for a range of R between 0.2 and 0.99 (such a range is expected from the practitioners' point of view) the discrepancies lie below 5% even with samples of just five observations. For the real-life newsvendor problem analyzed in this paper, the samples range between 26 and 30 observations, and R should be greater than 0.593 in order to ensure non-negative values for the loss of goodwill. Under such conditions, we see from Table 2 that the discrepancies $R - R_{act}$ are negligible. For example, for R=0.99, and for samples over 25 observations, the discrepancies are below 0.5%.

The previous arguments indicate that with samples over 25 observations and R greater or equal to 0.593, the use of $\hat{Q}_{n+1}$ results in estimated probabilities not to have stock-outs during the period which are almost identical to the requested critical fractiles. Thus, it is legitimate to use $\hat{Q}_{n+1}$ to proceed to estimate the maximum expected profit. Replacing in (1), μ with $\overline{X}_n$, σ with $\hat{V}_n$, and Q with $\hat{Q}_{n+1}$, and using also (2), we obtain the corresponding estimator for the maximum expected profit per monetary unit profit,

$$\hat{\xi}_{n+1} = \frac{(p-c)\hat{Q}_{n+1} + s\left(\hat{Q}_{n+1} - \overline{X}_n\right) - (p-v+s)\left\{\left(\hat{Q}_{n+1} - \overline{X}_n\right)\Phi_{\hat{z}} + \hat{V}_n\phi_{\hat{z}}\right\}}{p-c} =$$

$$= \frac{(p-c)\overline{X}_n + (p-c+s)z_R\hat{V}_n - (p-v+s)\left(z_R R + \phi_{z_R}\right)\hat{V}_n}{p-c} =$$

$$= \overline{X}_n - \frac{(1+\delta)}{R}\phi_{z_R}\hat{V}_n\,, \qquad (7)$$

as $\hat{z} = \frac{\hat{Q}_{n+1} - \overline{X}_n}{\hat{V}_n} = z_R$, and hence $\Phi_{\hat{z}} = R = (p-c+s)/(p-v+s)$.

From (7), it is easily deduce that $\hat{\xi}_{n+1}$ is an unbiased estimator for ξ^*. Besides, we find out that $\hat{\xi}_{n+1}$ can be concluded directly from (3) by replacing μ with $\overline{X}_n$ and σ with $\hat{V}_n$.

To describe $\hat{\xi}_{n+1}$ (or ξ^*) for different combinations of δ and R, we need to study the properties of function $(1+\delta)\phi_{z_R}/R$. Given a realization of historical demand data $x_1, x_2, ..., x_n$, and the corresponding estimates $\overline{x}_n$, $\hat{v}_n$, (or given μ and σ regarding ξ^*), $\hat{\xi}_{n+1}$ and ξ^* attain their maximum values for those values of δ and R where $(1+\delta)\phi_{z_R}/R$ becomes minimum.

As we have shown in the previous section, δ is an increasing function of R, and hence we cannot deduce without further analysis that the minimum of this function at any R is attained by setting $\delta = 0$. To resolve this problem, we replace (4) into $(1+\delta)\phi_{z_R}/R$ and we take $[(c-v)/(p-c)]\phi_{z_R}/R$. So the problem of studying $(1+\delta)\phi_{z_R}/R$ for different combinations of δ and R is equivalent to the problem of finding the properties of the function $\psi(R) = \phi_{z_R}/(1-R)$. In Proposition 1 we study the monotonicity and range of $\psi(R)$ when R takes on values on the interval $(0,1)$.

Proposition 1: For $0 < R < 1$, $\psi(R)$ is strictly increasing in R with $\lim_{R \to 0} \psi(R) = 0$ and

$$\lim_{R \to 1} \psi(R) = \infty$$

Proof: See in the Appendix.

In the previous section we also showed that the condition $R \geq (p-c)/(p-v)$ should hold in order the loss of goodwill to be non-negative. Hence, the domain of function $\psi(R)$ should be restricted to the interval $\frac{p-c}{p-v} \leq R < 1$. From proposition 1 we deduced that $\psi(R)$, and the equivalent function $(1+\delta)\phi_{z_R}/R$, have their minimum values at $R = (p-c)/(p-v)$. We conclude, therefore, that for the real-life newsvendor problem $\hat{\xi}_{n+1}$ and ξ^* attain their maximum values when $R = 1.75/2.95$, and for this value of R the loss of goodwill is zero.

5. Validity and precision of $\hat{\xi}_{n+1}$

In this section, for any finite n, as well as, for the case where n tends to infinity, we derive the sampling distribution of $\hat{\xi}_{n+1}$. Based on the exact and asymptotic distributions of $\hat{\xi}_{n+1}$, we give the corresponding prediction intervals for ξ^*. To relate the sampling variability of $\overline{X}_n$ and $\hat{V}_n$ to the precision and validity of the two types of prediction interval we use two additional accuracy information metrics. The first metric is the relative expected half-length (REHL) which is defined as the expected half-length of the interval divided by ξ^*. Dividing by ξ^*, we resolve the problem of comparability of precision, since different combinations of δ and R lead to different sizes of ξ^*. The second metric is the effective probability in finite samples the asymptotic prediction interval to include ξ^* when the population standard deviation of demand is replaced by its unbiased estimator $\hat{V}_n$. We call this metric as actual confidence level (ACL).

In the appendix we show that for any finite n, an $(1-\alpha)100\%$ prediction interval for ξ^* is constructed from

$$\hat{\xi}_{n+1} + \left[\frac{\lambda g_n^{-1}}{\sqrt{n}} - \frac{t'_{n-1,1-\alpha/2}(\lambda)}{\sqrt{n-1}}\right]\hat{\sigma}_n \le \xi^* \le \hat{\xi}_{n+1} + \left[\frac{\lambda g_n^{-1}}{\sqrt{n}} - \frac{t'_{n-1,\alpha/2}(\lambda)}{\sqrt{n-1}}\right]\hat{\sigma}_n \qquad (8)$$

where $t'(\lambda)$ is the non-central student-t distribution with non-centrality parameter $\lambda = \sqrt{n}\,\frac{1+\delta}{R}\phi_{z_R}$. The REHL of the interval (8) is computed from

$$\text{REHL} = \frac{E(HL)}{\xi^*} = \frac{t'_{n-1,1-\alpha/2}(\lambda) - t'_{n-1,\alpha/2}(\lambda)}{2\sqrt{n-1}\left(\mu - \frac{1+\delta}{R}\phi_{z_R}\sigma\right)} E(\hat{\sigma}_n) =$$

$$= \frac{t'_{n-1,1-\alpha/2}(\lambda) - t'_{n-1,\alpha/2}(\lambda)}{\sqrt{2n(n-1)}\left(CV^{-1} - \frac{1+\delta}{R}\phi_{z_R}\right)} \frac{\Gamma(n/2)}{\Gamma((n-1)/2)}. \qquad (9)$$

where $CV = \sigma/\mu$ is the coefficient variation of demand.

To derive the asymptotic distribution of $\hat{\xi}_{n+1}$, we need two prerequisite results which are stated in Propositions 2 and 3.

Proposition 2: Consider the random variables $Z_1 = \frac{\sqrt{n}(\overline{X}_n - \mu)}{\sigma}$ and $Y_n = \frac{\sqrt{n}(\hat{\sigma}_n - g_n\sigma)}{\sigma\sqrt{n(1-g_n^2)-1}}$.

As $n \to \infty$, then Z_1 and Y_n are asymptotically independent.

Proof: See in the Appendix.

Proposition 3: Let $\beta_1, \beta_2, \ldots$ be a sequence of real numbers with $\beta_n = \frac{\sqrt{n(1-g_n^2)-1}}{g_n}$.

As $n \to \infty$, then $\lim_{n\to\infty} \beta_n = \frac{1}{\sqrt{2}}$.

Proof: See in the Appendix.

Let Z_1 and Z_2 be independent standard normal random variables and $\omega_1, \omega_2, \ldots$ be a sequence of real numbers. If a random variable Y_n converges in distribution to Z_2 and $\omega_n \to \omega$ as $n \to \infty$, where ω is also a real number, then Severini (2005, pp. 337) states that any statistic of the form $T_n = Z_1 + \omega_n Y_n$ is asymptotically normally distributed with mean zero and variance $1 + \omega^2$.

In proposition 2, we specified the random variables Z_1 and Y_n, from (A10) of the appendix Y_n converges in distribution to the standard normal, and Z_1, Y_n are asymptotically independent. Setting also $\omega_n = (1+\delta)\phi_{z_R}\beta_n / R$, then, by the result of proposition 3, we have

$$\lim_{n\to\infty} \omega_n = \omega = \frac{1}{\sqrt{2}} \frac{1+\delta}{R} \phi_{z_R}.$$

Under the aforementioned specifications, the statistic T_n takes the form

$$T_n = \frac{\sqrt{n}\left(\overline{X}_n - \mu\right)}{\sigma} - \frac{1+\delta}{R} \phi_{z_R} g_n^{-1} \frac{\sqrt{n}\left(\hat{\sigma}_n - g_n \sigma\right)}{\sigma} = \frac{\sqrt{n}\left(\hat{\xi}_{n+1} - \xi^*\right)}{\sigma},$$

and

$$\frac{\sqrt{n}\left(\hat{\xi}_{n+1} - \xi^*\right)}{\sigma} \xrightarrow{D} N\left(0, 1 + \frac{1}{2}\left(\frac{1+\delta}{R}\phi_{z_R}\right)^2\right).$$

Hence the asymptotic $(1-\alpha)100\%$ prediction intervals for ξ^* has the form

$$\hat{\xi}_{n+1} - z_{\alpha/2} \frac{\sigma}{\sqrt{n}} \sqrt{1 + \frac{1}{2}\left(\frac{1+\delta}{R}\phi_{z_R}\right)^2} \le \xi^* \le \hat{\xi}_{n+1} + z_{\alpha/2} \frac{\sigma}{\sqrt{n}} \sqrt{1 + \frac{1}{2}\left(\frac{1+\delta}{R}\phi_{z_R}\right)^2}.$$

To use the asymptotic prediction interval in finite samples, we replace the unknown σ with its unbiased estimator $\hat{V}_n$, and so we obtain the following approximate $(1-\alpha)100\%$ prediction interval

$$\hat{\xi}_{n+1} \pm z_{\alpha/2} \frac{g_n^{-1}\hat{\sigma}_n}{\sqrt{n}} \sqrt{1 + \frac{1}{2}\left(\frac{1+\delta}{R}\phi_{z_R}\right)^2}. \tag{10}$$

The effective probability the prediction interval (10) to include ξ is not equal to the nominal confidence level $1-\alpha$. We shall call this probability as Actual Confidence Level (ACL). In the appendix we show that ACL is analytically computed from

$$\mathrm{ACL}=\Pr\left\{-z_{\alpha/2}\frac{\sqrt{\frac{n-1}{n}\left(1+\frac{\lambda^2}{2n}\right)}-\lambda}{g_n}\le t'_{n-1}(\lambda)\le z_{\alpha/2}\frac{\sqrt{\frac{n-1}{n}\frac{n-1}{n}\left(1+\frac{\lambda^2}{2n}\right)}+\lambda}{g_n}\right\}. \tag{11}$$

Finally the REHL of the approximate interval (10) is easily derived and is computed from

$$\mathrm{REHL}_\infty=z_{\alpha/2}\frac{\sqrt{1+\frac{1}{2}\left(\frac{1+\delta}{R}\phi_{z_R}\right)^2}}{\sqrt{n}\left(CV^{-1}-\frac{1+\delta}{R}\phi_{z_R}\right)}. \tag{12}$$

From (9) and (12) we establish that the size of the REHL of both the exact and the approximate prediction intervals depends upon the values of the function $(1+\delta)\phi_{z_R}/R$. In the previous section we showed that this function attains its minimum at $R=[(p-c)/(p-v)]$ where $\delta=0$. Hence, we conclude that both the exact and the approximate prediction intervals attain their maximum precision when the loss of goodwill is eliminated.

6. Results for the real-life newsvendor problem

The analysis which preceded in sections 4 and 5 showed that in order to proceed to predict ξ^* when the order quantity is determined from (5), we should make up our mind about two considerations. First we have to decide for the size of the loss of goodwill using any information provided from the business, and hence to determine the value of δ. Second, we should choose among the exact and the approximate prediction interval which one eventually we shall use to give the limits within which ξ^* lies with a pre-specified probability. These two issues are addressed in the current section and in the setting of the real-life newsvendor problem.

Regarding the first issue, the staff of the canteen did not have any knowledge at all about the size of the loss of goodwill. This was expected because the impersonal environment which dominates in selling the cheese-pie at the canteen does not allow records to be kept for the customers' identity (Foreign and Greek tourists, natives who are making this voyage regularly etc.). To resolve this problem we decided to give results for three alternative cases according to which the size of the loss of goodwill is coded as small, moderate, and large. In section 3, we found out that, given the price, the purchase cost, and the salvage value of the cheese-pie, to avoid a negative loss of goodwill, the critical fractile R should be greater than 0.593. Further, we established a positive relation between δ and R. So, we decided to set R=0.6, 0.8, and 0.95, and these values to indicate respectively a small, moderate, and large size for the loss of goodwill. Replacing into (4), we obtain the corresponding values of δ, which are 0.029, 1.743, and 12.03 respectively.

For the three selected values of R and for each one of the five days under consideration, Table 3 displays the estimated coefficient of variation, $\hat{V}_n / \overline{X}_n$, the estimates for Q^* and the predictions for ξ^*. The latter two quantities are obtained respectively from (5) and (7). From the data of the table, we verify the theoretic result from proposition 1 that $\hat{\xi}_{n+1}$ is decreasing in δ as R is getting larger. So, the increase of R leads to larger order quantities, and to the rise of total revenue. But the increase of cost due to the rise of δ, which is implied by increasing R, eventually dominates over the increase of the revenue, and eventually the expected profit is getting smaller. From table 3, we also observe than the estimated coefficients of variation are below 0.25. These sizes of CV justify the use of the normal distribution as a good approximation to the truncated normal at point zero. The latter distribution does not allow the occurrence of negative demand, which is likely to occur if the normal distribution had a large coefficient of variation. For a further discussion about this topic, see Law (1997) and Halkos & Kevork (2012a).

Table 3: Estimates for the optimal order quantity and predictions for the maximum expected profit per monetary unit profit for the real-life newsvendor model

n	Estimated CV	R=0.6 , δ=0.029		R=0.8 , δ=1.743		R=0.95 , δ=12.03	
		$\hat{Q}_{n+1}$	$\hat{\xi}_{n+1}$	$\hat{Q}_{n+1}$	$\hat{\xi}_{n+1}$	$\hat{Q}_{n+1}$	$\hat{\xi}_{n+1}$
26	0.239	96	76.03€	108	69.61€	126	59.81€
27	0.236	95	75.82€	108	69.49€	125	59.84€
28	0.232	96	76.43€	108	70.19€	125	60.66€
29	0.233	95	75.81€	107	69.59€	124	60.09€
30	0.239	94	74.66€	107	68.35€	124	58.71€

Regarding the second issue for choosing among the exact and the asymptotic form of the prediction interval for ξ^*, we constructed Table 4. The REHL's and the ACL's have been computed for CV=0.25 and at nominal confidence level equal to 0.95 using (9) and (12). The critical values of the non-central student-t distribution were obtained through the statistical package MINITAB.

Interesting conclusions are drawn from the examination of the data of Table 4. In any case the approximate form of the prediction interval gives lower REHL's. Given n and increasing R, (a) the REHL's of both forms (exact and approximate) of the prediction interval are getting larger, and (b) the ACL's of the approximate form are marginally declining. Examining the deviations of ACL's from 0.95, we find out that with samples over 25 observations and for the three selected values of R these deviations range below 1.3%. The lower REHL's that the approximate interval attains, in combination with the very small deviations of the ACL's from 0.95 constitute the necessary findings in order to select prediction interval (10) for estimating the daily maximum expected profit from selling the sfoliata cheese-pie.

For the real-life newsvendor problem, Table 5 displays for each one of the five days under consideration the lower and upper limits of the 95% exact and approximate prediction intervals for ξ^*. In the same table, the estimated half-length (HL) is the half-distance between the upper and the lower limit, and the estimates for the REHL's were obtained by replacing in

(9) and (12) the CV's with their estimates (see table 2). The general remark is the verification of the theoretical findings of previous sections. The exact prediction interval presents lower precision in terms of the estimated REHL's and HL's and the increase of R, which is implied by the rising of the loss of goodwill, results in wider prediction intervals.

Table 4: ACL and REHL of confidence intervals at 95% nominal confidence level, CV=0.25

	Actual Confidence Level (ACL)			Relative Expected Half Length (REHL)					
n	Asymptotic form			Exact from			Asymptotic form		
	R=0,6	R=0,8	R=0,95	R=0,6	R=0,8	R=0,95	R=0,6	R=0,8	R=0,95
5	0.8910	0.8868	0.8791	0.4047	0.5043	0.6201	0.2900	0.3485	0.4795
10	0.9228	0.9203	0.9161	0.2347	0.2864	0.3426	0.2051	0.2464	0.3390
15	0.9323	0.9306	0.9277	0.1822	0.2211	0.2624	0.1674	0.2012	0.2768
20	0.9369	0.9356	0.9334	0.1542	0.1866	0.2206	0.1450	0.1742	0.2397
25	0.9396	0.9385	0.9368	0.1361	0.1645	0.1940	0.1297	0.1558	0.2144
30	0.9414	0.9405	0.9390	0.1232	0.1487	0.1752	0.1184	0.1423	0.1957
40	0.9436	0.9429	0.9418	0.1056	0.1273	0.1497	0.1025	0.1232	0.1695
50	0.9449	0.9443	0.9434	0.0939	0.1131	0.1329	0.0917	0.1102	0.1516
100	0.9475	0.9472	0.9467	0.0656	0.0789	0.0925	0.0648	0.0779	0.1072
300	0.9492	0.9491	0.9489	0.0376	0.0452	0.0529	0.0374	0.0450	0.0619

Finally, in the last column of table 5, we give the 95% approximate prediction interval for the maximum expected profit. This is obtained multiplying all the terms of (10) by the profit margin $p-c$. Suppose now that we were at the start of the 31^{st} day, and the loss of goodwill had such a size which would give a critical fractile R=0.8. Having available the demand data for the previous 30 days, making the necessary analysis using the theoretical framework of sections 3, 4, and 5, and finally obtaining the results of tables 2,3,4 and 5, we would make the following recommendation to the canteen: To order 107 sfoliata cheese-pies ensuring in that way a probability of 79.5% not to experience a stock-out during the day. With the specific order quantity, the maximum expected profit at the end of the day will range between 103.75€ έως 135.66€ with probability approximately equal to 94%.

Table 5: 95% confidence intervals for the daily maximum expected profit regarding the real-life newsvendor problem

	n	Confidence Intervals for ξ^*						
		Exact Form			Approximate Form			Approximate confidence intervals for daily maximum expected profit
		Limits	Estimated HL's	Estimated RHL's	Limits	Estimated HL's	Estimated RHL's	
R=0,6	26	65.31€ – 84.48€	9.59€	0.1261	66.88€ – 85.18€	9.15€	0.1204	117.03€ – 149.06€
	27	65.49€ – 83.99€	9.25€	0.1220	66.97€ – 84.67€	8.85€	0.1167	117.19€ – 148.17€
	28	66.46€ – 84.36€	8.95€	0.1171	67.86€ – 85.01€	8.58€	0.1122	118.75€ – 148.77€
	29	66.08€ – 83.59€	8.75€	0.1155	67.41€ – 84.21€	8.40€	0.1108	117.97€ – 147.37€
	30	64.99€ – 82.42€	8.72€	0.1167	66.28€ – 83.04€	8.38€	0.1122	116.00€ – 145.31€
R=0,8	26	57.44€ – 78.54€	10.55€	0.1515	59.60€ – 79.63€	10.02€	0.1439	104.29€ – 139.35€
	27	57.78€ – 78.14€	10.18€	0.1465	59.81€ – 79.18€	9.69€	0.1394	104.66€ – 138.57€
	28	58.89€ – 78.59€	9.85€	0.1403	60.80€ – 79.58€	9.39€	0.1337	106.41€ – 139.26€
	29	58.57€ – 77.83€	9.63€	0.1383	60.40€ – 78.79€	9.19€	0.1321	105.70€ – 137.88€
	30	57.40€ – 76.57€	9.58€	0.1402	59.18€ – 77.52€	9.17€	0.1341	103.57€ – 135.66€
R=0,95	26	45.04€ – 69.92€	12.44€	0.2080	48.09€ – 71.53€	11.72€	0.1960	84.16€ – 125.18€
	27	45.63€ – 69.63€	12.00€	0.2006	48.50€ – 71.17€	11.34€	0.1894	84.88€ – 124.55€
	28	46.96€ – 70.18€	11.61€	0.1913	49.67€ – 71.65€	10.99€	0.1811	86.93€ – 125.38€
	29	46.74€ – 69.43€	11.35€	0.1888	49.33€ – 70.85€	10.76€	0.1791	86.33€ – 123.99€
	30	45.46€ – 68.05€	11.29€	0.1923	47.99€ – 69.44€	10.73€	0.1827	83.98€ – 121.53€

7. Conclusions

In the current paper we studied for a number of successive periods the classical newsvendor model under the following two conditions, when in the beginning of any period the inventory system starts with the optimal stocking level: (a) Demand is formed independently in successive periods and for each period is normally distributed with the same unknown mean and the same unknown variance, and (b) the selling price, the purchase cost per unit, and the salvage value are known quantities at the start of any period and they do not change from period to period.

These conditions are met to a real-life inventory problem which we analyzed in the current work. The supply and trading process of the product, as well as, its physical characteristics agree completely with the principles of the classical newsvendor model with the period to be a single day. Having available historical demand data for 30 days, we assume that the order quantity should be determined at the start of any day, from the 27th up to 31st, using the available sample of all previous days. At first, by applying appropriate tests for stationarity, autocorrelation, and normality, we resulted in sufficient statistical evidences to support that at the start of any of the five days under consideration demand is normally distributed.

To estimate the mean and the standard deviation of the normal distribution we used their maximum likelihood estimators. An appropriate adjustment was made to the estimator of the standard deviation of demand in order the get an unbiased estimator. Then an unbiased estimator of the optimal order quantity was built by replacing the sample mean and the unbiased estimator of the standard deviation of demand into the theoretical formula which gives this optimal quantity. Using this estimator, we evaluated analytically the estimated probability not to have stock-outs during the period. The deviations of this probability from the requested critical fractile were also studied for different sample sizes. We found out that over a reasonable range of the requested critical fractile, and with samples over 25 observations, the estimated probability of no stock-outs is almost identical to the requested critical fractile.

Replacing the estimator of the optimal order quantity into the expression which gives the expected profit of the model, we developed an unbiased estimator for the maximum expected profit. When the price, the purchasing cost and the salvage value are known at the start of the period, at the optimality condition of the model, the value of the critical fractile depends upon the size of the loss of goodwill. We showed that the loss of goodwill is an increasing function of the requested critical fractile, and when the loss of goodwill is zero the estimated maximum expected profit becomes minimum. For the estimator of the maximum expected profit, its exact and asymptotic sampling distributions were derived, and the corresponding prediction intervals were constructed. To evaluate the performance of the two types of prediction interval we derived the analytic forms of two accuracy information metrics: (a) the expected half-length of the interval divided by the maximum expected profit, and (b) the actual probability with which in finite samples the asymptotic interval includes the maximum expected profit.

The estimators for the optimal order quantity and the maximum expected profit were used to form the inventory policy for the real-life newsvendor problem for each one of the five days under consideration. The two main theoretical findings of the paper were verified using the real data for demand. First, the asymptotic prediction interval gave higher precision than the exact and with probability to include the maximum expected profit almost identical to the nominal confidence level. Second, the increase of the requested critical fractile, which was implied by the rising of the loss of goodwill, leaded to lower precision for both the exact and the asymptotic prediction interval. Finally, we closed this work by recommending for similar real-life inventory problems the use of the asymptotic prediction interval for the maximum expected profit when samples over 25 observations are available.

APPENDIX

Proof of (6)

From the definition of R_{act}

$$R_{act} = \Pr\left(X_{n+1} \le \hat{Q}_{n+1}\right) = \Pr\left(X_{n+1} \le \overline{X}_n + z_R g_n^{-1} \hat{\sigma}_n\right) =$$

$$= \Pr\left(\frac{X_{n+1} - \overline{X}_n}{\hat{\sigma}_n} \le z_R g_n^{-1}\right). \tag{A1}$$

Ritchken & Sankar (1984) stated that the statistic

$$\frac{\sqrt{n-1}\left(X_{n+1} - \overline{X}_n\right)}{\sqrt{\sum_{t=1}^{n}\left(X_t - \overline{X}_n\right)^2}\sqrt{\frac{n+1}{n}}} \tag{A2}$$

follows the central student-t distribution with $n-1$ degrees of freedom.

Rearranging the terms in (A2), the next statistic follows also the central student-t distribution with $n-1$ degrees of freedom:

$$\sqrt{\frac{n-1}{n+1}}\,\frac{X_{n+1} - \overline{X}_n}{\hat{\sigma}_n}.$$

Hence, the proof is completed after multiplying both sides of the inequality inside the probability of (A1) by $\left[(n-1)/(n+1)\right]^{1/2}$.

Proof of proposition 1

To derive the result of proposition 1, we need the derivative $\frac{d\phi_{z_R}}{dR} = \frac{d\phi_{z_R}}{dz_R} \cdot \frac{dz_R}{dR}$. For the standard normal, z_R can be treated as a quantile function of R, and using expression (2) of Steinbrecher & Shaw (2008), we take

$$\frac{dz_R}{dR} = \frac{1}{\left(\sqrt{2\pi}\right)^{-1} \cdot \exp\left(-z_R^2/2\right)} = \phi_{z_R}^{-1}$$

and

$$\frac{d\phi_{z_R}}{dR} = -z_R \phi_{z_R} \frac{dz_R}{dR} = -z_R \qquad \text{(A3)}$$

Regarding the function $\psi(R)$ it holds

$$\lim_{R\to 0} \psi(R) = \frac{\lim_{R\to 0} \phi_{z_R}}{\lim_{R\to 0}(1-R)} = \phi_\infty = 0,$$

and using (A3)

$$\lim_{R\to 1} \psi(R) = \lim_{R\to 1} \frac{\frac{d}{dR}\phi_{z_R}}{\frac{d}{dR}(1-R)} = \lim_{R\to 1} z_R = +\infty,$$

$$\frac{d}{dR}\psi(R) = (1-R)^{-1}\frac{d\phi_{z_R}}{dR} + \phi_{z_R}\frac{d}{dR}(1-R)^{-1} = \frac{1}{(1-R)^2}\left\{\phi_{z_R} - z_R(1-R)\right\} > 0. \qquad \text{(A4)}$$

The positive sign of (A4) is explained as follows. When $R \to 0$,

$$\lim_{R\to 0}\left(\phi_{z_R} - z_R(1-R)\right) = \lim_{R\to 0}\phi_{z_R} - \left(\lim_{R\to 0} z_R\right)\left(\lim_{R\to 0}(1-R)\right) = 0 - (-\infty)1 = +\infty. \qquad \text{(A5)}$$

For $R \to 1$

$$\lim_{R\to 1} z_R(1-R) = \lim_{R\to 1}\frac{\frac{d}{dR}(1-R)}{\frac{d}{dR}z_R^{-1}} = \lim_{R\to 1}\frac{\frac{d}{dR}z_R^2}{\frac{d}{dR}\phi_{z_R}^{-1}} = \lim_{R\to 1} 2\phi_{z_R},$$

and

$$\lim_{R\to 1}\left(\phi_{z_R} - z_R(1-R)\right) = \lim_{R\to 1}\left(\lim_{R\to 1} - 2\phi_{z_R}\right) = -\lim_{R\to 1}\phi_{z_R} = 0\,. \tag{A6}$$

Further,

$$\frac{d}{dR}\left(\phi_{z_R} - z_R(1-R)\right) = \frac{d\phi_{z_R}}{dR} - (1-R)\frac{dz_R}{dR} - z_R\frac{d}{dR}(1-R) = -\frac{1-R}{\phi_{z_R}} < 0\,. \tag{A7}$$

From (A5), (A6), and (A7), the function $\phi_{z_R} - z_R(1-R)$ is decreasing in R with range $(0,\infty)$.

Hence the derivative in (A4) is positive.

Proof of (8)

Define the statistic

$$\sqrt{n-1}\left[\frac{\hat{\xi}_{n+1} - \xi^*}{\hat{\sigma}_n} + \frac{(1+\delta)}{R}\phi_{z_R} g_n^{-1}\right] = \frac{\sqrt{n-1}\left\{\overline{X}_n - \xi^*\right\}}{\hat{\sigma}_n}, \tag{A8}$$

The expression on the right hand side of (A8) is taken after using (3) and (7). Since

$Z_1 = \dfrac{\sqrt{n}\left(\overline{X}_n - \mu\right)}{\sigma} \sim N(0,1)$ and $\dfrac{n\hat{\sigma}_n^2}{\sigma^2} \sim \chi^2_{n-1}$ (e.g. Lindgren, 1976, pp. 334), the statistic

$$\frac{Z_1 + \lambda}{\sqrt{\dfrac{\chi^2_{n-1}}{n-1}}} = \frac{\dfrac{\sqrt{n}\left(\overline{X}_n - \mu\right)}{\sigma} + \sqrt{n}\,\dfrac{(1+\delta)}{R}\phi_{z_R}}{\sqrt{\dfrac{n\hat{\sigma}_n^2/\sigma^2}{n-1}}} = \frac{\sqrt{n-1}\left\{\overline{X}_n - \xi^*\right\}}{\hat{\sigma}_n} \tag{A9}$$

follows the non-central student-t distribution, $t'_{n-1}(\lambda)$ with non-centrality parameter

$\lambda = \sqrt{n}\,\dfrac{1+\delta}{R}\phi_{z_R}$.

Then from (A8) and (A9) we take

$$\Pr\left\{t'_{n-1,\alpha/2}(\lambda) \le \sqrt{n-1}\left[\frac{\hat{\xi}_{n+1} - \xi^*}{\hat{\sigma}_n} + \frac{(1+\delta)}{R}\phi_{z_R} g_n^{-1}\right] \le t'_{n-1,1-\alpha/2}(\lambda) \le\right\} = 1-\alpha\,,$$

or

$$\Pr\left\{-\lambda g_n^{-1}\frac{\hat{\sigma}_n}{\sqrt{n}} + t'_{n-1,\alpha/2}(\lambda)\,\frac{\hat{\sigma}_n}{\sqrt{n-1}} \le \hat{\xi}_{n+1} - \xi^* \le -\lambda g_n^{-1}\frac{\hat{\sigma}_n}{\sqrt{n}} + t'_{n-1,1-\alpha/2}(\lambda)\,\frac{\hat{\sigma}_n}{\sqrt{n-1}}\right\} = 1-\alpha\,,$$

and finally

$$\Pr\left\{\hat{\xi}_{n+1}+\left[\frac{\lambda g_n^{-1}}{\sqrt{n}}-\frac{t'_{n-1,1-\alpha/2}(\lambda)}{\sqrt{n-1}}\right]\hat{\sigma}_n \le \xi^* \le \hat{\xi}_{n+1}+\left[\frac{\lambda g_n^{-1}}{\sqrt{n}}-\frac{t'_{n-1,\alpha/2}(\lambda)}{\sqrt{n-1}}\right]\hat{\sigma}_n\right\}=1-\alpha$$

from where the prediction intervals is concluded.

Proof of proposition 2

Since $\frac{n\hat{\sigma}_n^2}{\sigma^2}\sim\chi_{n-1}^2$, the random variable $Y=\frac{\sqrt{n}\hat{\sigma}_n}{\sigma}$ follows the chi-distribution with $v=n-1$ degrees of freedom. From Johnson et al. (1994, vol. 1, pp. 421) we take

$$E(Y)=\sqrt{2}\left[\Gamma\left(\frac{n}{2}\right)\Big/\Gamma\left(\frac{n-1}{2}\right)\right],$$

and

$$\mathrm{Var}(Y)=n-1-2\left[\Gamma\left(\frac{n}{2}\right)\Big/\Gamma\left(\frac{n-1}{2}\right)\right]^2.$$

Keiding et al. (1972) report that a chi-distribution with v degrees of freedom approaches quickly the normal distribution as v gets large. So standardizing Y we take the asymptotic distributional result

$$Y_n=\frac{\sqrt{n}(\hat{\sigma}_n-g_n\sigma)}{\sigma\sqrt{n(1-g_n^2)-1}}\xrightarrow{D}N(0,1). \tag{A10}$$

Knight (1999, pp. 258) shows in example 5.14 that the vector $\begin{bmatrix}\sqrt{n}(\overline{X}_n-\mu)\\ \sqrt{n}(\hat{\sigma}_n-\sigma)\end{bmatrix}$ is asymptotically normally distributed with mean vector zero and covariance matrix, $\begin{bmatrix}\sigma^2 & 0\\ 0 & \sigma^2/2\end{bmatrix}$. From the structure of the covariance matrix, we immediately deduce that $\overline{X}_n$ and $\hat{\sigma}_n$ are asymptotically independent (see also Severini, 2005, pp. 403). Hence Z_1 and Y_n

as functions of $\overline{X}_n$ and $\hat{\sigma}_n$ respectively are also asymptotically independent, since functions of independent normal random variables are also independent (Lindgren, 1976, pp. 334) .

Proof of Proposition 3

Taking limits to both sides of the expression which gives β_n, we have

$$\lim_{n\to\infty}\beta_n = \frac{\left[\lim_{n\to\infty} n\left(1-g_n^2\right)-1\right]^{-\frac{1}{2}}}{\lim_{n\to\infty} g_n}$$

Regarding the numerator, Kevork (2010) showed that $\lim_{n\to\infty} n\left(1-g_n^2\right)=\frac{3}{2}$.

For the denominator, applying the transformation $y=\frac{n}{2}$ we obtain

$$\lim_{n\to\infty} g_n = \lim_{n\to\infty}\sqrt{\frac{2}{n}}\frac{\Gamma\left(\frac{n}{2}\right)}{\Gamma\left(\frac{n-1}{2}\right)} = \lim_{y\to\infty}\sqrt{\frac{1}{y}}\frac{\Gamma(y)}{\Gamma\left(y-\frac{1}{2}\right)}.$$

But Abramowitz & Stegun (1972) report that

$$\lim_{y\to\infty} y^{d_2-d_1}\frac{\Gamma(y+d_1)}{\Gamma(y+d_2)}=1$$

Setting therefore $d_1=0$ and $d_2=-\frac{1}{2}$, it is concluded that $\lim_{n\to\infty} g_n=1$, and hence

$$\lim_{n\to\infty}\beta_n = \left(\frac{3}{2}-1\right)^{-\frac{1}{2}} = \frac{1}{\sqrt{2}}.$$

Proof of (9)

$$\text{ACL} = \Pr\left\{ -z_{\alpha/2} \le \frac{\sqrt{n}\left(\hat{\xi}_{n+1} - \xi^*\right)}{g_n^{-1}\hat{\sigma}_n \sqrt{1 + \frac{1}{2}\left(\frac{1+\delta}{R}\phi_{z_R}\right)^2}} \le z_{\alpha/2} \right\} =$$

$$= \Pr\left\{ -z_{\alpha/2}\sqrt{1 + \frac{1}{2}\left(\frac{1+\delta}{R}\phi_{z_R}\right)^2} \le \frac{\sqrt{n}\left(\hat{\xi}_{n+1} - \xi^*\right)}{g_n^{-1}\hat{\sigma}_n} \le z_{\alpha/2}\sqrt{1 + \frac{1}{2}\left(\frac{1+\delta}{R}\phi_{z_R}\right)^2} \right\}. \qquad \text{(A11)}$$

The midterm of the inequality inside (A11) is rewritten as

$$\frac{\sqrt{n}\left(\hat{\xi}_{n+1} - \xi^*\right)}{g_n^{-1}\hat{\sigma}_n} = \frac{\sqrt{n}\left(\overline{X}_n - \frac{1+\delta}{R}\phi_{z_R} g_n^{-1}\hat{\sigma}_n - \xi^*\right)}{g_n^{-1}\hat{\sigma}_n} = \frac{\sqrt{n}\left(\overline{X}_n - \xi^*\right)}{g_n^{-1}\hat{\sigma}_n} - \sqrt{n}\,\frac{1+\delta}{R}\phi_{z_R}$$

Hence

$$\text{ACL} = \Pr\left\{ -z_{\alpha/2}\sqrt{1 + \frac{1}{2}\left(\frac{1+\delta}{R}\phi_{z_R}\right)^2} + \lambda \le \sqrt{\frac{n g_n}{n-1}}\,\frac{\sqrt{n-1}\left(\overline{X}_n - \xi^*\right)}{\hat{\sigma}_n} \le \right.$$

$$\left. \le z_{\alpha/2}\sqrt{1 + \frac{1}{2}\left(\frac{1+\delta}{R}\phi_{z_R}\right)^2} + \lambda \right\}. \qquad \text{(A12)}$$

The proof is completed by noting that $\frac{\sqrt{n-1}\left(\overline{X}_n - \xi^*\right)}{\hat{\sigma}_n}$ in (A12) follows the non-central student-t distribution with $n-1$ degrees of freedom and non-centrality parameter $\lambda = \sqrt{n}\,\frac{1+\delta}{R}\phi_{z_R}$.

REFERENCES

Abramowitz, M., & Stegun, I.E. (1972). Handbook of mathematical functions. National Bureau of Standards, Washington, D.C.

Akcay, A., Biller, B., & Tayur, S. (2011). Improved inventory targets in the presence of limited historical data. *Manufacturing & Service Operations Management*, DOI 10.1287/ msom. 1100.0320.

Ali, M. M., Boylan, J. E., & Syntetos, A. A. (2011). Forecast errors and inventory performance under forecast information sharing. *International Journal of Forecasting*, doi:10.1016/j.ijforecast.2010.08.003.

Benzion, U., Cohen, Y., Peled, R., & Shavit, T. (2008). Decision making and the newsvendor problem: an experimental study. *Journal of the Operational Research Society*, *59*, 1281–1287.

Benzion, U., Cohen, Y., & Shavit, T. (2010). The newsvendor problem with unknown distribution. *Journal of the Operational Research Society*, *61*, 1022–1031.

Beutel, A. L., & Minner, S. (2011). Safety stock planning under causal demand forecasting. *International Journal of Production Economics*, doi:10.1016/j.ijpe.2011.04.017.

Choi, T. M, Chiu, C. H., & Chester, K. M. (2011). A fast fashion safety-first inventory model, *Textile Research Journal*, *81*, 819–826.

Feng, T., Keller, L. R., & Zheng, X. (2011). Decision making in the newsvendor problem: A cross-national laboratory study. *Omega*, *39*, 41–50.

Halkos, G. E., & Kevork, I. S. (2005). A comparison of alternative unit root tests. *Journal of Applied Statistics*, *32*, 45–60.

Halkos, G. E., & Kevork, I. S. (2012a). The classical newsvendor model under normal demand with large coefficients of variation. *MPRA Paper No. 40414*, Online at http://mpra.ub.uni-muenchen.de/40414/.

Halkos, G. E., & Kevork, I. S. (2012b). Evaluating alternative frequentist inferential approaches for optimal order quantities in the newsvendor model under exponential demand. *MPRA Paper No. 39650*, Online at http://mpra.ub.uni-muenchen.de/39650/.

Halkos, G. E., & Kevork, I. S. (2012c). Validity and precision of estimates in the classical newsvendor model with exponential and Rayleigh demand. *MPRA Paper No. 36460*, Online at http://mpra.ub.uni-muenchen.de/36460/.

Harvey, A.C. (1993). Time series models. 2nd Edition, Prentice Hall, An imprint of Pearson Education Limited.

Hayes, R. H. (1969). Statistical estimation problems in inventory control. *Management Science*, *15*, 686–701.

Janssen E, Strijbosch L, & Brekelmans R. (2009). Assessing the effects of using demand parameters estimates in inventory control and improving the performance using a correction function. *International Journal of Production Economics*, *118*, 34–42.

Johnson, N. L., Kotz, S., & Bakakrishnan, N. (1994). Continuous Univariate Distributions. 2nd Edition. Wiley, New York.

Katircioglou, K. (1996). Essays in inventory control. Ph.D thesis, University of British Columbia, Vancouver, BC, Canada.

Keiding, N., Jensen, S. T., & Ranek, L. (1972). Maximum likelihood estimation of the size distribution of liver cell nuclei from the observed distribution in a plane section. *Biometrics*, *28*, 813–829.

Kevork, I. S. (2010). Estimating the optimal order quantity and the maximum expected profit for single-period inventory decisions. *Omega*, *38*, 218–227.

Knight, K. (1999). Mathematical Statistics. Taylor & Francis Ltd.

Lapin, L. L. (1994). Quantitative methods for business decisions with cases, 6th Edition, Duxbury Press, An International Thomson Publishing Company.

Lindgren, B. W. (1976). Statistical theory, 3rd Edition, London: Collier MacMillan International Editions.

Law, H. (1997). Simple formulas for the expected costs in the newsboy problem: an educational note. *European Journal of Operational Research*, *100*, 557–561.

Liyanage, L. H., & Shanthikumar, J.G. (2005). A practical inventory control policy using operational statistics. *Operations Research Letters*, *33*, 341–348.

Mostard, J., & Teunter, R. (2006). The newsboy problem with resalable returns: A single period model and case study. *European Journal of Operational Research*, *169*, 81–96.

Mostard, J., Teunter, R., & de Coster, R. (2011). Forecasting demand for single-period products: A case study in the apparel industry. *European Journal of Operational Research*, *211*, 139–147.

Olivares, M., Terwiesch, C., & Cassorla, L. (2008). Structural estimation of the newsvendor model: An application to reserving operating room time. *Management Science*, *54*, 41–55

Ritchken, P. H., & Sankar, R. (1984). The effect of estimation risk in establishing safety stock levels in an inventory model. *The Journal of the Operational Research Society*, *35*, 1091–1099.

Schweitzer M. E., & Cachon, G. P. (2000). Decision bias in the newsvendor problem with a known demand distribution: experimental evidence. *Management Science*, *46*, 404–420.

Severini, T.A. (2005). Elements of distribution theory. Cambridge University Press.

Silver, E. A., & Rahnama, M. R. (1987). Biased selection of the inventory reorder point when demand parameters are statistically estimated. *Engineering Costs and Production Economics*, *12*, 283–292.

Silver, E. A., Pyke, D. F., & Peterson, R. (1998). Inventory management and production planning and scheduling, 3rd Edition, New York, John Wiley and Sons.

Steinbrecher, G., & Shaw, W. T. (2008). Quantile mechanics. *European Journal of Applied Mathematics*, *19*, 87–112

Su, R. H., & Pearn, W. L. (2011). Product selection for newsvendor-type products with normal demands and unequal costs. *International Journal of Production Economics*, *132*, 214–222.

Syntetos, A. A., Nikolopoulos, K., & Boylan, J. E. (2010). Judging the judges through accuracy-implication metrics: The case of inventory forecasting. *International Journal of Forecasting*, *26*, 134–143.