Estimating Demand Uncertainty Over Multi-Period Lead Times

ISIR 2016

Patrick Saoud & Nikolaos Kourentzes & John Boylan

Department of Management Science - Lancaster University

August 23, 2016

Table of Contents

Main Formula for Safety Stocks

In an Order-Up-To policy, the O.U.T level is estimated as:

$$OUT = \widehat{CF} + z_{\alpha} \times \sqrt{Var(e)} \tag{1}$$

where:

- OUT is the OUT level.
- CF is the cumulative forecast.
- α is the prescribed service level.
- z is the inverse of the Cumulative Distribution function of the forecast errors.
- $Var(e)$ is the variance of the forecast error.

Main Formula for Safety Stocks

In an Order-Up-To policy, the O.U.T level is estimated as:

$$OUT = \widehat{CF} + z_\alpha \times \sqrt{Var(e)} \tag{2}$$

Where do the possible discrepancies in the O.UT. level stem from?

Main Formula for Safety Stocks

In an Order-Up-To policy, the O.U.T level is estimated as:

$$OUT = \widehat{CF} + z_\alpha \times \sqrt{Var(e)} \tag{2}$$

Where do the possible discrepancies in the O.UT. level stem from?

- The cumulative forecasts are not accurate.

Main Formula for Safety Stocks

In an Order-Up-To policy, the O.U.T level is estimated as:

$$OUT = \widehat{CF} + z_\alpha \times \sqrt{Var(e)} \tag{2}$$

Where do the possible discrepancies in the O.UT. level stem from?

- The cumulative forecasts are not accurate.
- The Cumulative Distribution Function chosen is not the adequate one.

Main Formula for Safety Stocks

In an Order-Up-To policy, the O.U.T level is estimated as:

$$OUT = \widehat{CF} + z_\alpha \times \sqrt{Var(e)} \tag{2}$$

Where do the possible discrepancies in the O.UT. level stem from?

- The cumulative forecasts are not accurate.
- The Cumulative Distribution Function chosen is not the adequate one.
- The variance of the errors is wrongly estimated.

Types of Uncertainty

Uncertainty

In modeling demand, there exists three types of uncertainties:

Types of Uncertainty

Uncertainty

In modeling demand, there exists three types of uncertainties:

1. Sample Size Uncertainty (for example Phillips [1979]).

Types of Uncertainty

Uncertainty

In modeling demand, there exists three types of uncertainties:

1. Sample Size Uncertainty (for example Phillips [1979]).
2. Parameter Uncertainty (for example [Ansley and Newbold, 1981])

Types of Uncertainty

Uncertainty

In modeling demand, there exists three types of uncertainties:

1. Sample Size Uncertainty (for example Phillips [1979]).
2. Parameter Uncertainty (for example [Ansley and Newbold, 1981])
3. Model Uncertainty (An excellent discussion is found in [Chatfield, 1993] and [Chatfield, 1995])

Variance of errors

MSE

The forecasts are assumed to be unbiased, and so the Mean Squared Error is equated with the variance

Variance of errors

MSE

The forecasts are assumed to be unbiased, and so the Mean Squared Error is equated with the variance

From the Bias-Variance Decomposition, we have that:

$$MSE(Y_t - \hat{Y}_t) = E[(Y_t - \hat{Y}_t)^2] = Bias^2 + Var(e_t) \qquad (3)$$

Variance of errors

MSE

The forecasts are assumed to be unbiased, and so the Mean Squared Error is equated with the variance

$$MSE(Y_t - \hat{Y}_t) = E[(Y_t - \hat{Y}_t)^2] = Bias^2 + Var(e_t) \qquad (4)$$

When Bias is present, $MSE(e_t) > Var(e_t)$. Its impact is rarely studied in the safety stocks literature [Manary and Willems, 2008].

Variance of errors

MSE

The forecasts are assumed to be unbiased, and so the Mean Squared Error is equated with the variance

- Demand parameters are usually optimized to yield a minimum MSE.
- Reaching MMSE might not be optimal in terms of setting safety stocks [Strijbosch et al., 2011].
- An unbiased estimator is not optimal in an inventory context [Janssen et al., 2011].
- Some authors recommend adding Bias for better safety stocks performance (for e.g. Silver and Rahnama [1986] and Silver and Rahnama [1987])

Estimating variance over lead-time

- Demand is forecasted over a lead-time.
- $\widehat{CY}_{t+L} = \sum_{i=1}^{L} \hat{Y}_{t+i}$

Estimating variance over lead-time

- Demand is forecasted over a lead-time.
- $\widehat{CY}_{t+L} = \sum_{i=1}^{L} \hat{Y}_{t+i}$
- The conventional approximation used to estimate error variance is: $Var(\hat{e}_{t+1} + \hat{e}_{t+2} + ...) = L \times Var(\hat{e}_{t+1})$

Estimating variance over lead-time

- The conventional approximation used to estimate error variance is: $Var(e_{t+1} + e_{t+2} + ...) = L \times Var(\hat{e}_{t+1})$

Estimating variance over lead-time

- The conventional approximation used to estimate error variance is: $Var(e_{t+1} + e_{t+2} + ...) = L \times Var(\hat{e}_{t+1})$
- This is the L-steps-ahead variance for a Random Walk ([Koehler, 1990] and [Chatfield and Koehler, 1991]).

Estimating variance over lead-time

- The conventional approximation used to estimate error variance is: $Var(e_{t+1} + e_{t+2} + ...) = L \times Var(\hat{e}_{t+1})$
- This is the L-steps-ahead variance for a Random Walk ([Koehler, 1990] and [Chatfield and Koehler, 1991]).
- One assumption it makes is that $Var(e_{t+1}) = Var(e_{t+2}) = ...Var(e_{t+L})$

Estimating variance over lead-time

- The conventional approximation used to estimate error variance is: $Var(e_{t+1} + e_{t+2} + ...) = L \times Var(\hat{e}_{t+1})$
- This is the L-steps-ahead variance for a Random Walk ([Koehler, 1990] and [Chatfield and Koehler, 1991]).
- One assumption it makes is that $Var(e_{t+1}) = Var(e_{t+2}) = ...Var(e_{t+L})$
- That is typically not the case. It holds for some processes, but not for all.

Multiple-Steps-Ahead Forecasts

- Suppose demand follows an AR(1) model, given by: $Y_t = \phi Y_{t-1} + \epsilon_t$ with ϵ i.i.d and $\epsilon \sim N(0, \sigma^2)$.
- At time t, its $t+1$ forecast is: $\hat{Y}_{t+1|t,t-1\ldots} = \phi Y_t + \epsilon_{t+1}$
- Its variance is σ^2
- Its $t+2$ forecast is $\hat{Y}_{t+2|t,t-1\ldots} = \phi \hat{Y}_{t+1|t,t-1\ldots} + \epsilon_{t+2} = \phi^2 Y_t + \phi \epsilon_{t+1} + \epsilon_{t+2}$
- Its variance is $(1+\phi^2)\sigma^2$
- $Var(Y_{t+2|t,t-1\ldots}) > Var(Y_{t+1|t,t-1\ldots})$

Estimating variance over lead-time

- The conventional approximation used to estimate error variance is: $Var(e_{t+1} + e_{t+2} + ...) = L \times Var(\hat{e}_{t+1})$
- Another approach would consist of summing the variances across horizons [Barrow and Kourentzes, 2016].

Estimating variance over lead-time

- The conventional approximation used to estimate error variance is: $Var(e_{t+1} + e_{t+2} + ...) = L \times Var(\hat{e}_{t+1})$
- Another approach would consist of summing the variances across horizons [Barrow and Kourentzes, 2016].
- Retrieve $t + 1$ in-sample errors, $t + 2$ in-sample errors,... and then sum calculate their respective variance.

Estimating variance over lead-time

- The conventional approximation used to estimate error variance is: $Var(e_{t+1} + e_{t+2} + ...) = L \times Var(\hat{e}_{t+1})$
- Another approach would consist of summing the variances across horizons [Barrow and Kourentzes, 2016].
- Retrieve $t + 1$ in-sample errors, $t + 2$ in-sample errors,... and then sum calculate their respective variance.
- This approach is independent of any assumptions on the forecasting model or method.
- $Var(e_{t+1} + e_{t+2} + ...) = \sum_{i=1}^{L}(Var(e_{t+i}))$

Estimating variance over lead-time

- The conventional approximation used to estimate error variance is: $Var(e_{t+1} + e_{t+2} + ...) = L \times Var(\hat{e}_{t+1})$
- Another approach would consist of summing the variances across horizons [Barrow and Kourentzes, 2016].
- Retrieve $t + 1$ in-sample errors, $t + 2$ in-sample errors,... and then sum calculate their respective variance.
- This approach is independent of any assumptions on the forecasting model or method.
- $Var(e_{t+1} + e_{t+2} + ...) = \sum_{i=1}^{L}(Var(e_{t+i}))$
- Again this is flawed as it overlooks correlations between the forecasting errors.

Correlations between forecast errors

Correlations between forecast errors

- Multi-steps-ahead forecast errors are correlated with each other ([Johnston and Harrison, 1986] ,[Box et al., 1994], [Barrow and Kourentzes, 2016] and [Prak et al., 2016]).

Correlations between forecast errors

- Multi-steps-ahead forecast errors are correlated with each other ([Johnston and Harrison, 1986] ,[Box et al., 1994], [Barrow and Kourentzes, 2016] and [Prak et al., 2016]).
- Even in the absence of autocorrelation within demand, Prak et al. [2016] showed the existence of this correlation

Correlations between forecast errors

- Multi-steps-ahead forecast errors are correlated with each other ([Johnston and Harrison, 1986] ,[Box et al., 1994], [Barrow and Kourentzes, 2016] and [Prak et al., 2016]).
- Even in the absence of autocorrelation within demand, Prak et al. [2016] showed the existence of this correlation
- This is prevalent in real-life modeling due to Model Uncertainty and Parameter Uncertainty.

Cumulative Errors

Cumulative Errors

Since demand is forecasted over a lead-time, it would seem natural to use the errors over lead-time as well.

$$CE_{t+h} = \sum_{i=1}^{h}(e_{t+i}) = \sum_{i=1}^{h}(Y_{t+i} - \hat{Y}_{t+i})$$

Cumulative Errors

Cumulative Errors

Since demand is forecasted over a lead-time, it would seem natural to use the errors over lead-time as well.

$$CE_{t+h} = \sum_{i=1}^{h}(e_{t+i}) = \sum_{i=1}^{h}(Y_{t+i} - \hat{Y}_{t+i})$$

- This has been used in the literature (see for e.g. Eppen and Martin [1988], Lee [2014].)

Cumulative Errors

Cumulative Errors

Since demand is forecasted over a lead-time, it would seem natural to use the errors over lead-time as well.

$$CE_{t+h} = \sum_{i=1}^{h}(e_{t+i}) = \sum_{i=1}^{h}(Y_{t+i} - \hat{Y}_{t+i})$$

- This has been used in the literature (see for e.g. Eppen and Martin [1988], Lee [2014].)
- No motivation is provided nevertheless in the literature.

Cumulative Errors

Cumulative Errors

Since demand is forecasted over a lead-time, it would seem natural to use the errors over lead-time as well.

$$CE_{t+h} = \sum_{i=1}^{h}(e_{t+i}) = \sum_{i=1}^{h}(Y_{t+i} - \hat{Y}_{t+i})$$

Cumulative Errors

Cumulative Errors

Since demand is forecasted over a lead-time, it would seem natural to use the errors over lead-time as well.

$$CE_{t+h} = \sum_{i=1}^{h}(e_{t+i}) = \sum_{i=1}^{h}(Y_{t+i} - \hat{Y}_{t+i})$$

- This captures all the lead-time uncertainties, and contains the aggregate properties of the errors.

Cumulative Errors

Cumulative Errors

Since demand is forecasted over a lead-time, it would seem natural to use the errors over lead-time as well.

$$CE_{t+h} = \sum_{i=1}^{h}(e_{t+i}) = \sum_{i=1}^{h}(Y_{t+i} - \hat{Y}_{t+i})$$

- This captures all the lead-time uncertainties, and contains the aggregate properties of the errors.
- It circumvents the need to model uncertainties at each horizon and reconstruct them.

Cumulative Errors

Cumulative Errors

Since demand is forecasted over a lead-time, it would seem natural to use the errors over lead-time as well.

$$CE_{t+h} = \sum_{i=1}^{h}(e_{t+i}) = \sum_{i=1}^{h}(Y_{t+i} - \hat{Y}_{t+i})$$

- This captures all the lead-time uncertainties, and contains the aggregate properties of the errors.
- It circumvents the need to model uncertainties at each horizon and reconstruct them.
- We know from (overlapping) temporal aggregation that this will smooth the values.

Cumulative Errors

Cumulative Errors

Since demand is forecasted over a lead-time, it would seem natural to use the errors over lead-time as well.

Cumulative Errors

Cumulative Errors

Since demand is forecasted over a lead-time, it would seem natural to use the errors over lead-time as well.

- $Var(CE_{t+h}) = Var(\sum_{i=1}^{h}(e_{t+i})) = \sum_{i=1}^{h}(Var(e_{t+i})) + 2\sum_{i=1}^{h}\sum_{j\neq i}^{h} 2Cov(\epsilon_{t+i}, \epsilon_{t+j})$

Cumulative Errors

Cumulative Errors

Since demand is forecasted over a lead-time, it would seem natural to use the errors over lead-time as well.

- $Var(CE_{t+h}) = Var(\sum_{i=1}^{h}(e_{t+i})) = \sum_{i=1}^{h}(Var(e_{t+i})) + 2\sum_{i=1}^{h}\sum_{j\neq i}^{h} 2Cov(\epsilon_{t+i}, \epsilon_{t+j})$
- The correlations are captured in the variance of cumulative errors!

Experimental Setup

- O.U.T Policy
- Deterministic Lead times $L = \{0, 2, 5\}$
- Review Period = 1
- Horizon = Lead Time + Review = {1,3,6 }

Experimental Setup

- O.U.T Policy
- Deterministic Lead times $L = \{0, 2, 5\}$
- Review Period = 1
- Horizon = Lead Time + Review = {1,3,6 }
- Three types of uncertainties are explored.

Experimental Setup

- O.U.T Policy
- Deterministic Lead times $L = \{0, 2, 5\}$
- Review Period = 1
- Horizon = Lead Time + Review = {1,3,6 }
- Three types of uncertainties are explored.
- Three methods of estimating variance are tested.

Experimental Setup

- O.U.T Policy
- Deterministic Lead times $L = \{0, 2, 5\}$
- Review Period = 1
- Horizon = Lead Time + Review = {1,3,6 }
- Three types of uncertainties are explored.
- Three methods of estimating variance are tested.
- Threefold data split $\{100; 200; 100\}$

Experimental Setup

5 demand processes are generated

- I(1): $Y_t = Y_{t-1} + \epsilon_t$
- AR(1): $Y_t = \phi Y_{t-1} + \epsilon_t$
- MA(1):$Y_t = \epsilon_t - \theta\epsilon_{t-1}$
- IMA(1,1):$Y_t = Y_{t-1} + \epsilon_t - \theta\epsilon_{t-1}$
- ARMA(1,1): $Y_t = \phi Y_{t-1} + \epsilon_t - \theta\epsilon_{t-1}$
- For all processes, ϵ is i.i.d and $\epsilon \sim N(0, \sigma^2)$
- 500 replications are produced

Results

- In order to contrast the methods, the deviation from service level is measured.
- The trade-off curves are plotted in parallel.

Results

- In order to contrast the methods, the deviation from service level is measured.
- The trade-off curves are plotted in parallel.
- The results reported are for $L = 3$ and a service level of 90%.

Results

- In order to contrast the methods, the deviation from service level is measured.
- The trade-off curves are plotted in parallel.
- The results reported are for $L = 3$ and a service level of 90%.
- The results for $L = 6$ and other service levels are proportional.

AR(1) Results

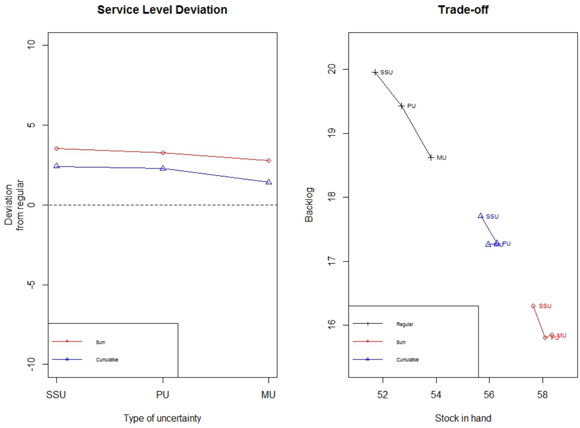
Service Level Deviation
Deviation from regular
10
5
0
-5
-10
SSU
PU
MU
Type of uncertainty
Sum
Cumulative
Trade-off
Backlog
20
19
18
17
16
52
54
56
58
Stock in hand
SSU
PU
MU
Regular
Sum
Cumulative

IMA(1,1) Results

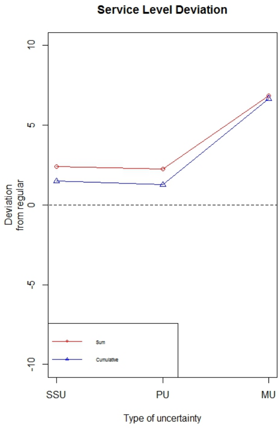

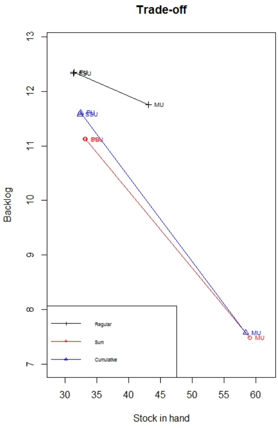

I(1) Results

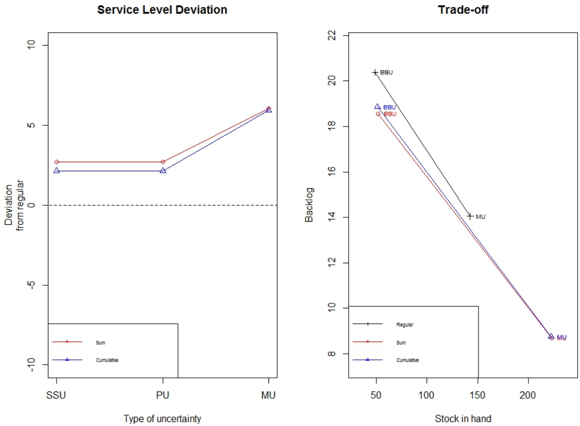
Service Level Deviation
Deviation from regular
10
5
0
-5
-10
SSU
PU
MU
Type of uncertainty
Sum
Cumulative
Trade-off
Backlog
22
20
18
16
14
12
10
8
50
100
150
200
Stock in hand
BBU
MU
Regular
Sum
Cumulative

MA(1) Results

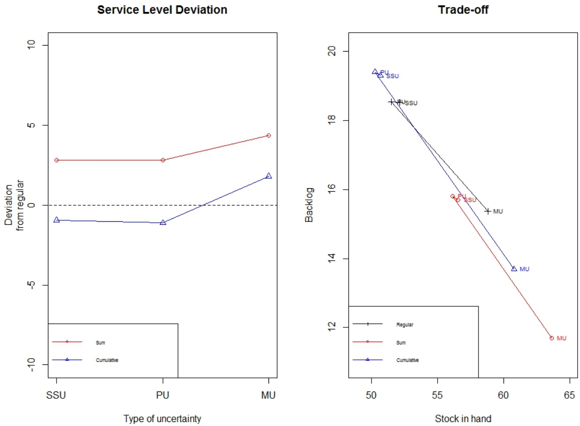
Service Level Deviation
Deviation from regular
10
5
0
-5
-10
Sum
Cumulative
SSU
PU
MU
Type of uncertainty
Trade-off
Backlog
20
18
16
14
12
PU
SSU
PU
SSU
PU
SSU
MU
MU
MU
Regular
Sum
Cumulative
50
55
60
65
Stock in hand

ARMA(1) Results

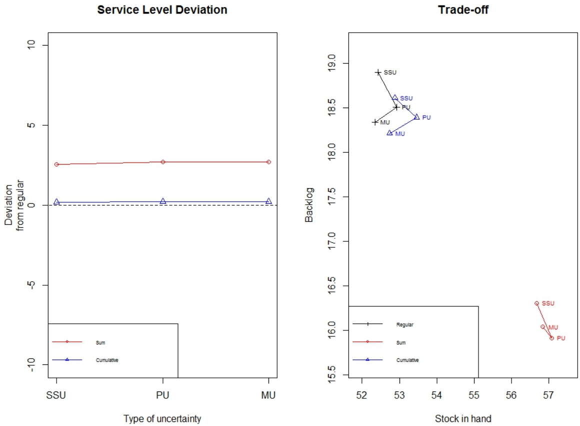

Summary Findings

Summary Findings

- As the level of uncertainty increases, the variance increases which results in higher service levels at the cost of higher inventories.

Summary Findings

- As the level of uncertainty increases, the variance increases which results in higher service levels at the cost of higher inventories.
- We notice a convergence of achieved service levels with Model Uncertainty for the two proposed approaches.

Summary Findings

- As the level of uncertainty increases, the variance increases which results in higher service levels at the cost of higher inventories.
- We notice a convergence of achieved service levels with Model Uncertainty for the two proposed approaches.
- The conventional approach is generally outperformed by the other methods.

Summary Findings

- As the level of uncertainty increases, the variance increases which results in higher service levels at the cost of higher inventories.
- We notice a convergence of achieved service levels with Model Uncertainty for the two proposed approaches.
- The conventional approach is generally outperformed by the other methods.
- The sum of variances returns a superior performance in terms of service levels. The connection between these two has to be explored further.

Conclusion

Any Questions?

Thank you

Craig F. Ansley and Paul Newbold. On the bias in estimates of forecast mean squared error. *Journal of the American Statistical Association*, 76(375):569–578, September 1981.

Devon K. Barrow and Nikolaos Kourentzes. Distributions of forecasting error of forecast combinations: Implications for inventory management. *International Journal of Production Economics*, 177:24–33, 2016.

G. E. P. Box, G. M. Jenkins, and G. C. Reinsel. *Time Series Analysis, Forecasting and Control*. Prentice Hall, 3rd edition, 1994.

Chris Chatfield. Calculating interval forecasts. *Journal of Business & Economic Statistics*, 11(2):121–135, 1993.

Chris Chatfield. Model uncertainty, data mining and statistical inference. *Journal of the Royal Statistical Society. Series A (Statistics in Society)*, 158:419–466, 1995.

Christopher Chatfield and Anne B Koehler. On confusing lead time demand with h-period-ahead forecasts. *International Journal of Forecasting*, 7(2):239–240, 1991.

Gary D. Eppen and R.Kipp Martin. Determining safety stock in the presence of stochastic lead time and demand. *Management Science*, 34(11), 1988.

Elleke Janssen, Leo W.G. Strijbosch, and Ruud Brekelmans. Assessing the effects of using demand parameters estimates in inventory control and improving the performance using a correction function. *International Journal of Production Economics*, pages 34–42, 2011.

F. R. Johnston and P. J. Harrison. The variance of lead time demand. *Journal of Operational Research Society*, 37(3): 303–308, 1986.

Anne B Koehler. An inappropriate prediction interval. *International Journal of Forecasting*, 6(4):557–558, 1990.

Yun Shin Lee. A semi-parametric approach for estimating critical fractiles under autocorrelated demand. *European Journal of Operational Research*, 234:163–173, 2014.

Matthew P. Manary and Sean P. Willems. Setting safety-stocks targets at intel in the presence of forecast bias. *Interfaces*, 38 (2):112–122, 2008.

Peter C.B. Phillips. The sampling distribution of forecasts from a first-order autoregression. *Journal of Econometrics*, 9:241–261, 1979.

Dennis Prak, Ruud Teunter, and Aris Syntetos. On the calculation of safety stocks when demand is forecasted. *European Journal of Operational Research*, 000:1–8, 2016.

Edward A. Silver and Mina Rasty Rahnama. The cost effect of statistical sampling in selecting the reorder point in a common inventory model. *The Journal of the Operational Research Society*, 37(7):705–713, July 1986.

Edward A. Silver and Mina Rasty Rahnama. Biased selection of the inventory reorder point when demand parameters are statistically estimated. *Engineering Costs and Production Economics*, 12:283–292, 1987.

Leo W.G. Strijbosch, Aris A. Syntetos, John E. Boylan, and Elleke Janssen. On the interaction between forecasting and stock control: The case of non-stationary demand. *International Journal of Production Economics*, 133:470–480, 2011.