Optimising forecasting models for inventory planning

Nikolaos Kourentzes[a]

Devon K. Barrow[b]

Juan R. Trapero[a,c]

[a] Lancaster Centre for Marketing Analytics and Forecasting

[b] Birmingham University; [c] University Castilla-La Mancha

1st July 2019

LUMS | Marketing Analytics and Forecasting

The problem

- Inaccurate forecasts are costly for operations in terms of stock-outs, over-stocking and poor service level.
- There are two elements to this:
  - Appropriate forecasting model specification;
  - Translation of forecast into inventory decisions.
- Connecting forecast uncertainty to inventory decisions makes the latter adaptive to the quality of forecasts.
- Consider the simple case of an Order Up To policy. The safety stock (SS) is calculated as:

  $SS = k\sigma_L$, where $k = \Phi^{-1}(CSL)$ and σ_L is the standard deviation of the forecasts over lead time.

- The σ_L is often approximated as $\sqrt{L}\sigma_1$, which has been strongly criticised (Chatfield, 2000) or using other empirical approaches, such as KDE (Trapero et al., 2018).
- When a model is used to produce the forecast, then we can derive exact expressions.

CMAF

The problem

- For example, for Single Exponential Smoothing [SES; equivalent to ARIMA(0,1,1)] we have (Johnston & Harisson, 1986; Hyndman et al., 2008):

$$\sigma_L = \sigma_1 \sqrt{L} \sqrt{1 + \alpha(L-1) + \alpha^2 (L-1)(2L-1)/6},$$ where α is the smoothing parameter.

- It is clear that the optimisation of the model parameters (α, σ_1 and the initial level) affect σ_L, by extension SS and eventually the inventory performance, even for forecasts of very similar accuracy.

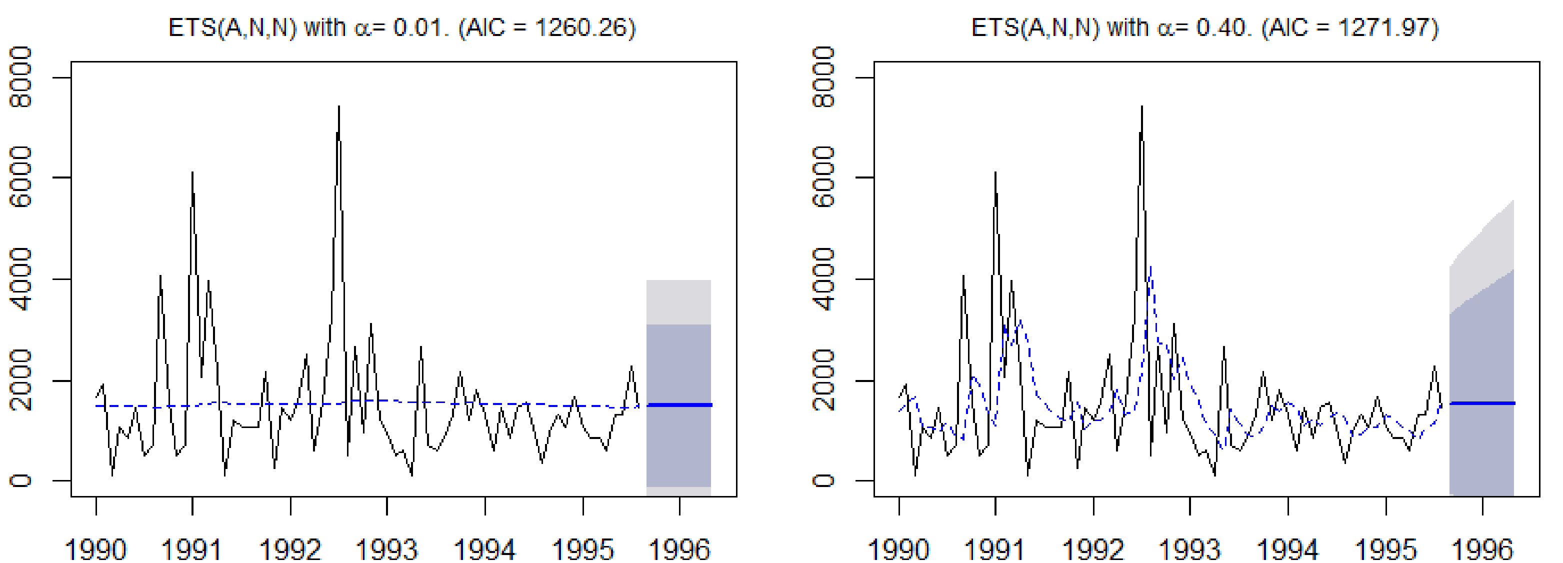

The problem

- The model parameters are typically optimised by minimising the mean squared error (or the negative likelihood), so as to achieve the best fit in the historical demand
- **However, this is connected with the inventory decisions only implicitly and depending on the underlying process and errors in the approximations this connection can vanish entirely** (Fildes and Kingsman, 2010; Kourentzes, 2013; Kourentzes, 2014).
- Or more intuitively, the objective of the optimisation differs from the objective of the inventory decisions.
- This raises the questions:
  - **(i) can we optimise forecasting models in a way that the objectives are aligned?**
  - **(ii) what is the benefit, if any?**

CMAF

Estimating model parameters

- There is long standing research in parameter estimation for forecasting models (Chatfield, 2000; Gardner, 2006).
- What we know:
  - Quadratic errors result in optimal forecasts for the mean (needed for inventory decisions; Gneiting, 2011a).
  - Absolute errors result in optimal forecasts for the median, this provides a connection for quantile forecasting, hence the estimation of SS (Gneiting, 2011b).
  - One-step ahead errors do not represent well multi-step errors, unless the true process is modelled (Xia & Tong, 2011; Barrow & Kourentzes, 2016).
  - Multi-step ahead forecast errors can be useful alternatives, but their performance is inconsistent.
  - Multi-step ahead forecast errors result in shrinkage of parameters of univariate models, thus reducing over-fitting when the model is misspecified, but can overshrink as this is proportional to the forecast horizon.

Alternative objective functions

- All the typical cost functions follow the same logic:
  - Mean squared error, 1-step ahead (likelihood). (i) Standard objective function; (ii) Assumes model to be true, otherwise approximates only short term behaviour of demand; (iii) very easy to use.

$$MSE_{t+1} = \frac{1}{n}\sum_{t=1}^{n}\left(y_t - \hat{y}_{t|t-1}\right)^2$$

  - Mean squared error, L-steps ahead. (i) Attempts to recognise that long-term forecasting is more difficult and focuses there; (ii) all else equal, results in lower smoothing parameters; (iii) lessens training sample.

$$MSE_{t+L} = \frac{1}{n-L+1}\sum_{t=L+1}^{n}\left(y_t - \hat{y}_{t|t-L}\right)^2$$

  - Mean squared error, 1 to L-steps ahead. (i) Inventory decisions happen over lead time, so minimise trace forecast error.

$$MSE_{t+1\,-t+L} = \frac{1}{L}\sum_{h=1}^{L} MSE_{t+h}$$

CMAF

Alternative objective functions

- Instead of minimising the forecast trace error, one can minimise the cumulative error over lead time: we are interested in meeting the total demand over lead time, not per period, which is a more difficult problem.
  - Cumulative mean squared error over L. (i) Implies smoothing of data, easier to minimise (removes timing complication); (ii) equivalent to overlapping temporal aggregation; (iii) lessens training sample; (iv) shrinkage implications (data less volatile).

$$cMSE_{t+L} = \frac{1}{n-L+1}\sum_{t=L+1}^{n}\left(\sum_{j=t-L+1}^{t} y_j - \sum_{j=t-L+1}^{t} \hat{y}_{j|t-L}\right)^2$$

CMAF

Alternative objective functions

$$MSE_{t+1}$$

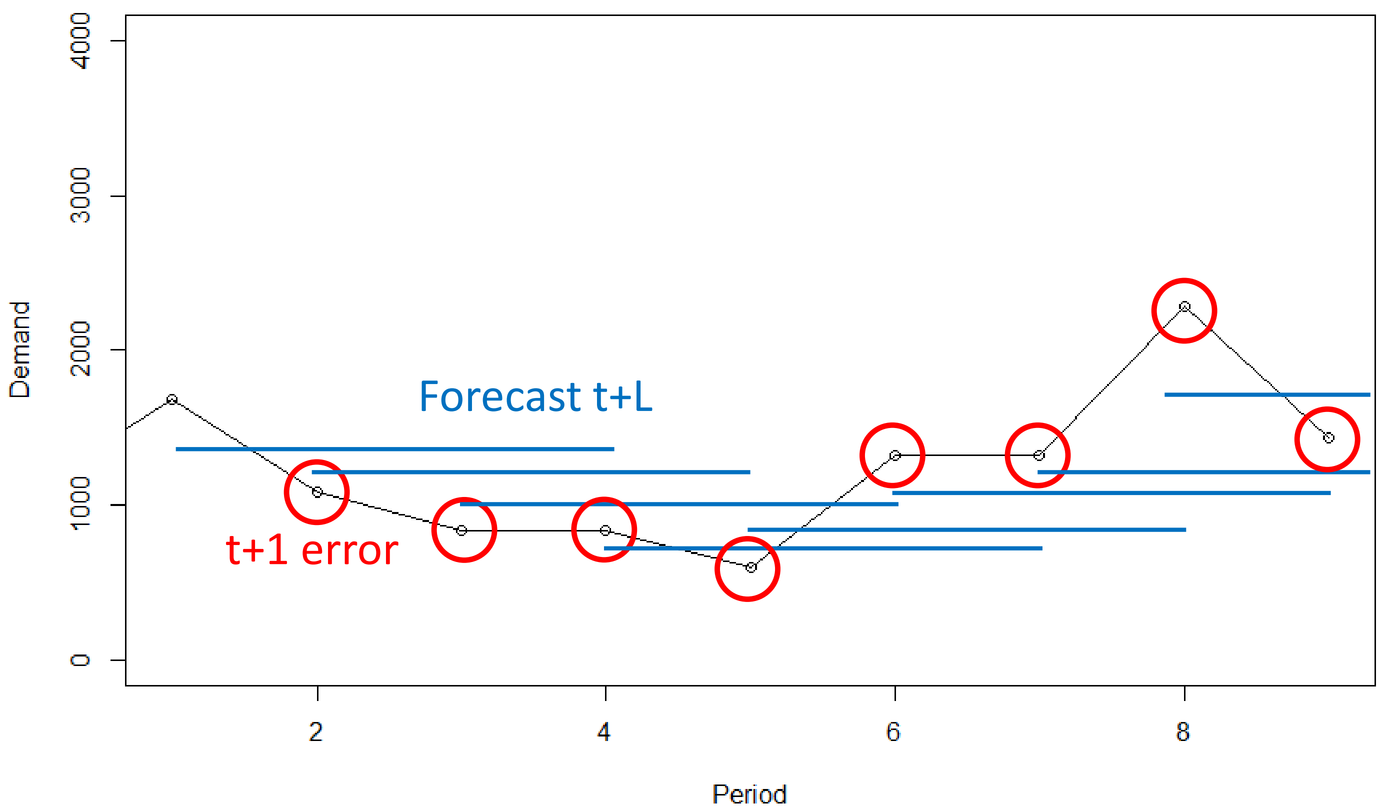

Alternative objective functions

$$MSE_{t+L}$$

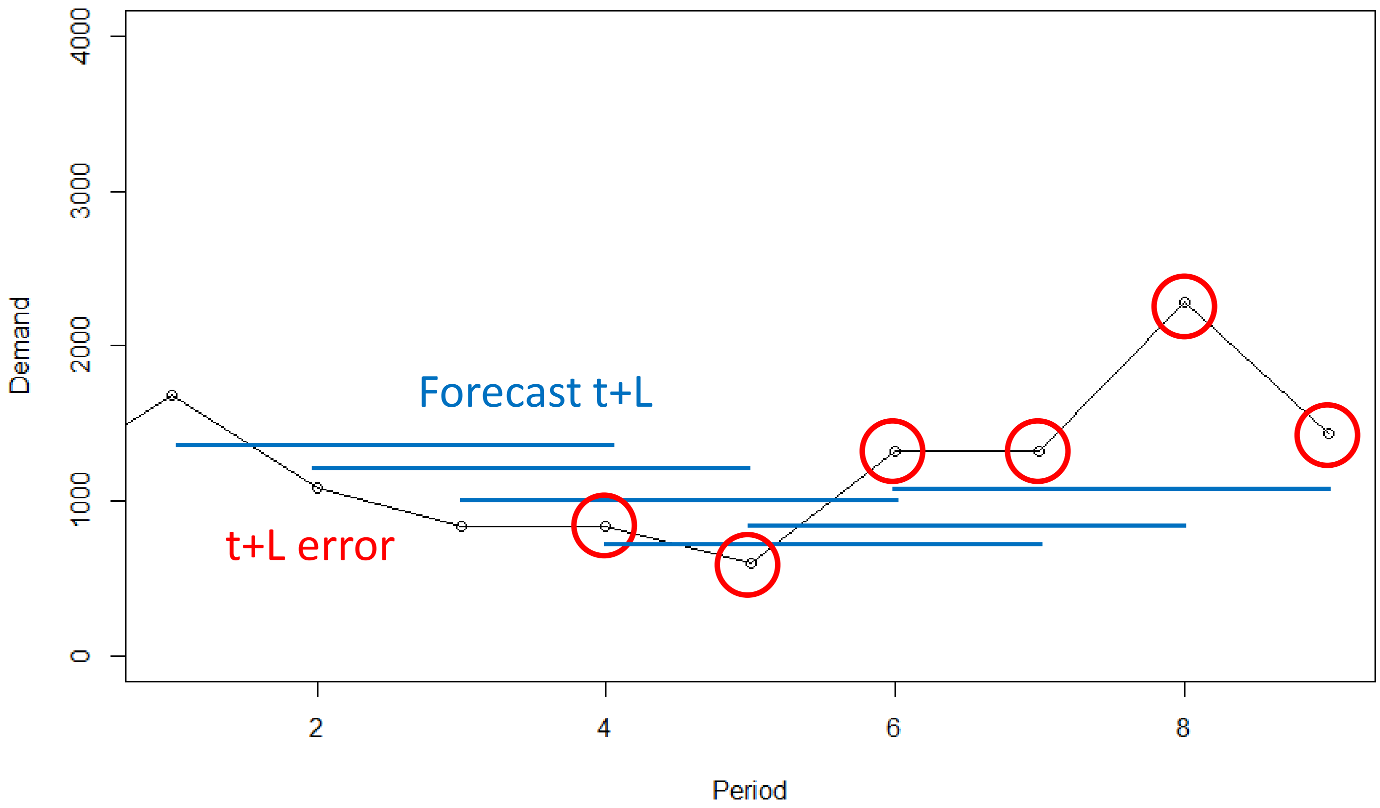

Alternative objective functions

$$MSE_{t+1\ -t+L}$$

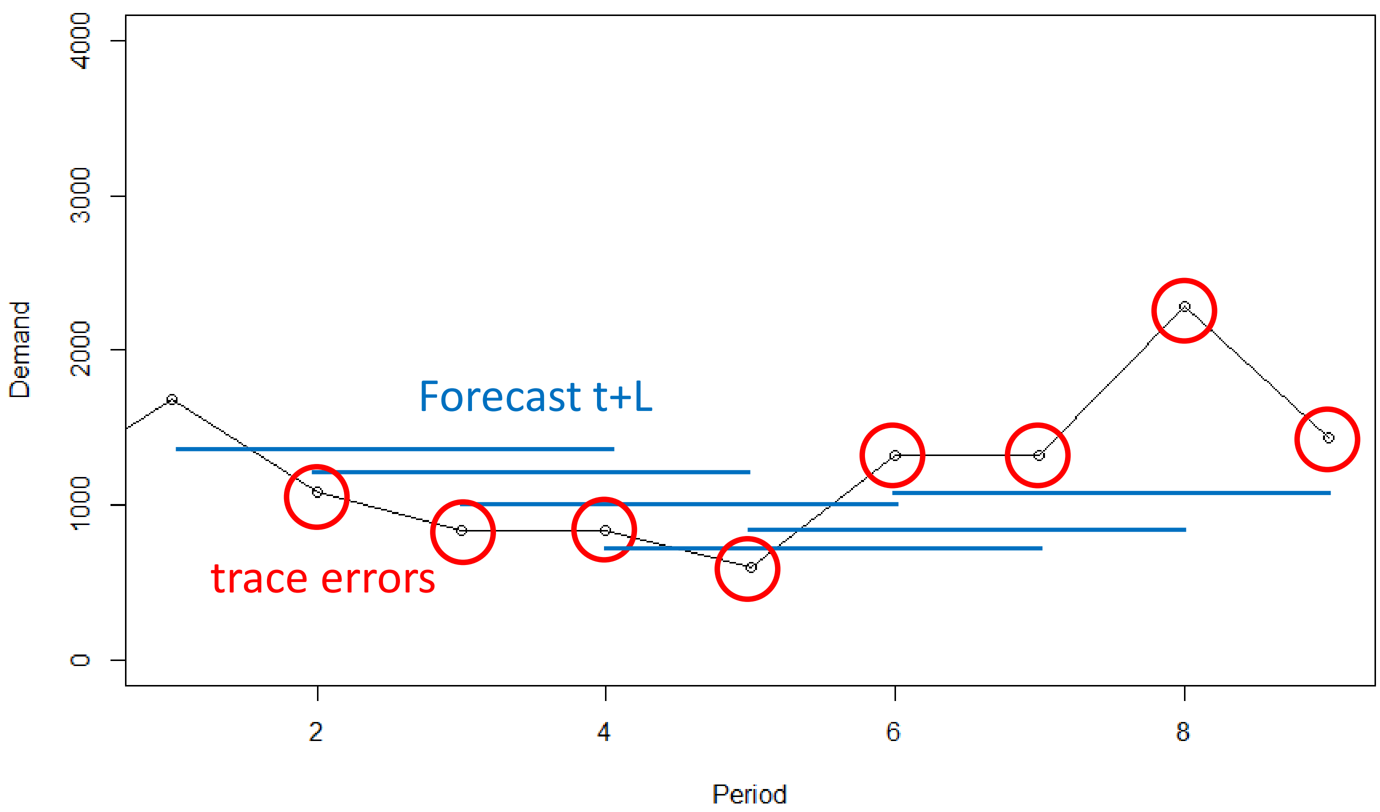

Alternative objective functions

$$cMSE_{t+L}$$

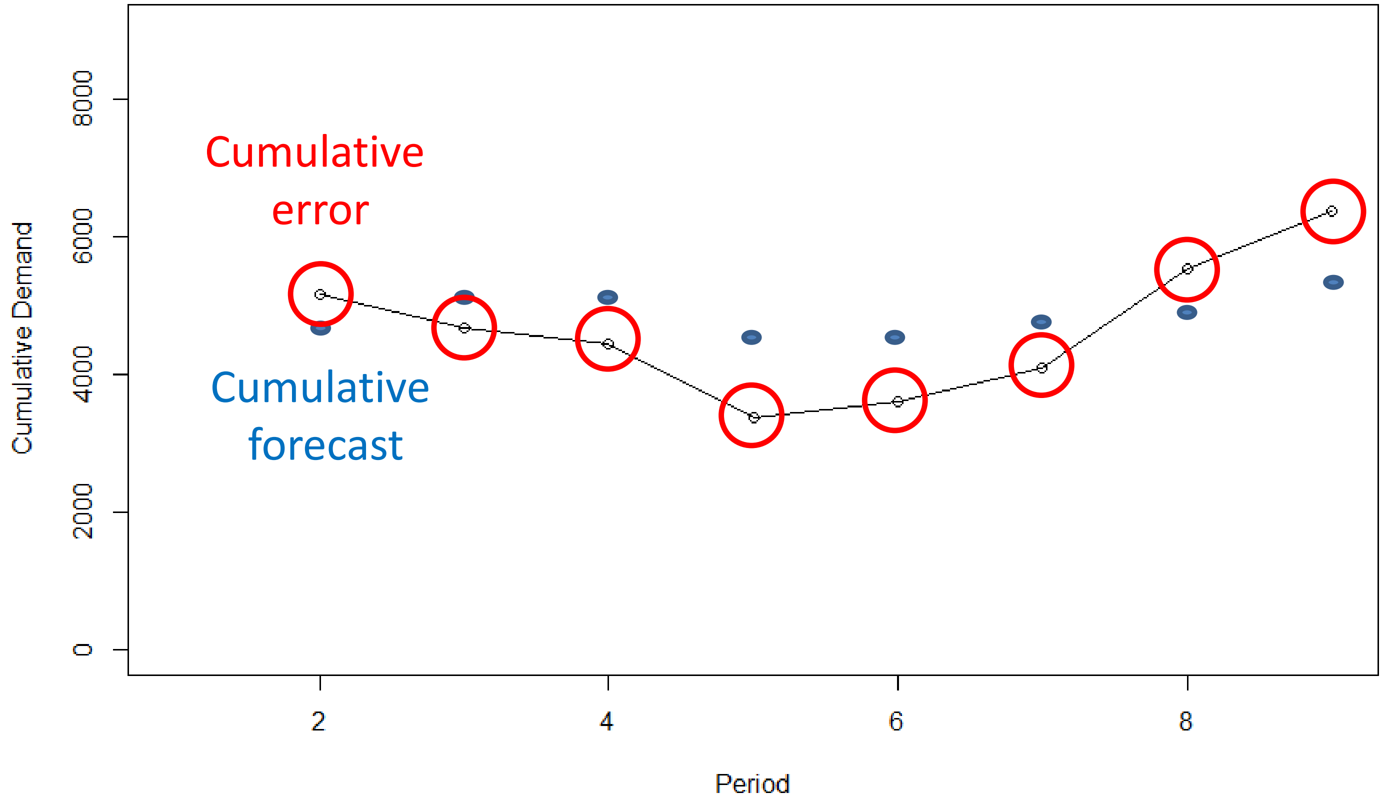

Proposed objective function

- All the previous objective functions measure how closely we follow the observed demand. Instead focus directly on the inventory decision.

Forecast
- For given parameters produce forecasts

Inventory simulation
- For given CSL simulate inventory given forecasts

Objective function
- Target CSL – Realised CSL
- Minimise inventory costs

Iterate until objective function cannot be improved further

CMAF

Proposed objective function

- Challenges:
  - Sample size: same as all t+L objective functions;
  - Initialisation of inventory simulation: these are hyper-parameters, can be set using heuristics or cross-validated. We use heuristics here.
- Advantages:
  - Match inventory decision making costs:
    - This can be gap in service level

$$Cost = (Target\ CSL - Realised\ CSL)^2$$

    - or cost based

$$Cost = (Lost\ sales) + \lambda(Stock\ on\ hand)$$

    - the latter can be solved for given cost ratio (λ) or the complete Pareto frontier for different cost ratios, in a multi-objective context.
  - Account for different inventory policies directly.

CMAF

Empirical evaluation

- Use a real dataset from an FMCG UK manufacturer: 229 SKUs over 173 weeks.
- Absence of seasonality or any strong trends → use SES to predict demand.
- Test on the last 52 weeks, using rolling origin evaluation (with re-optimisation).
- Order-up-to policy.
- L = {3, 5}, CSL = {90%, 95%, 99%}.
- σ_L is model based. We also evaluation the tick loss (directly estimate a model that predicts the desired quantile that matches target CSL → performed poorly).
- Evaluate on forecast accuracy and inventory performance.

CMAF

Accuracy results

- **RMSE**: Root Mean Squared Error over L.
- **RMSCE**: Cumulative RMSE (i.e. first aggregate over L and then calculate error).
- **AME**: Absolute Mean Error over L.

We see that inventory based optimisation performs poorly in terms of accuracy (RMSE, cRMSE), but best in terms of bias size (AME).

	Accuracy		Bias
Cost function	RelRMSE	RelRMSCE	RelAME
	Horizon 3		
MSE	1.000	1.000	1.000
MSE_{t+h}	0.995	0.990	1.016
$\text{MSE}_{t+1 - t+h}$	**0.994**	**0.987**	1.060
MSE_{t+L}	0.995	0.990	1.067
Inventory (90%)	1.033	1.086	0.678
Inventory (95%)	1.033	1.086	0.524
Inventory (99%)	1.062	1.157	**0.376**
	Horizon 5		
MSE	1.000	1.000	1.000
MSE_{t+h}	0.995	0.987	0.968
$\text{MSE}_{t+1 - t+h}$	**0.991**	**0.974**	1.043
MSE_{t+L}	0.994	0.980	1.108
Inventory (90%)	1.009	1.028	0.745
Inventory (95%)	1.008	1.031	0.632
Inventory (99%)	1.023	1.089	**0.473**

CMAF

Accuracy results

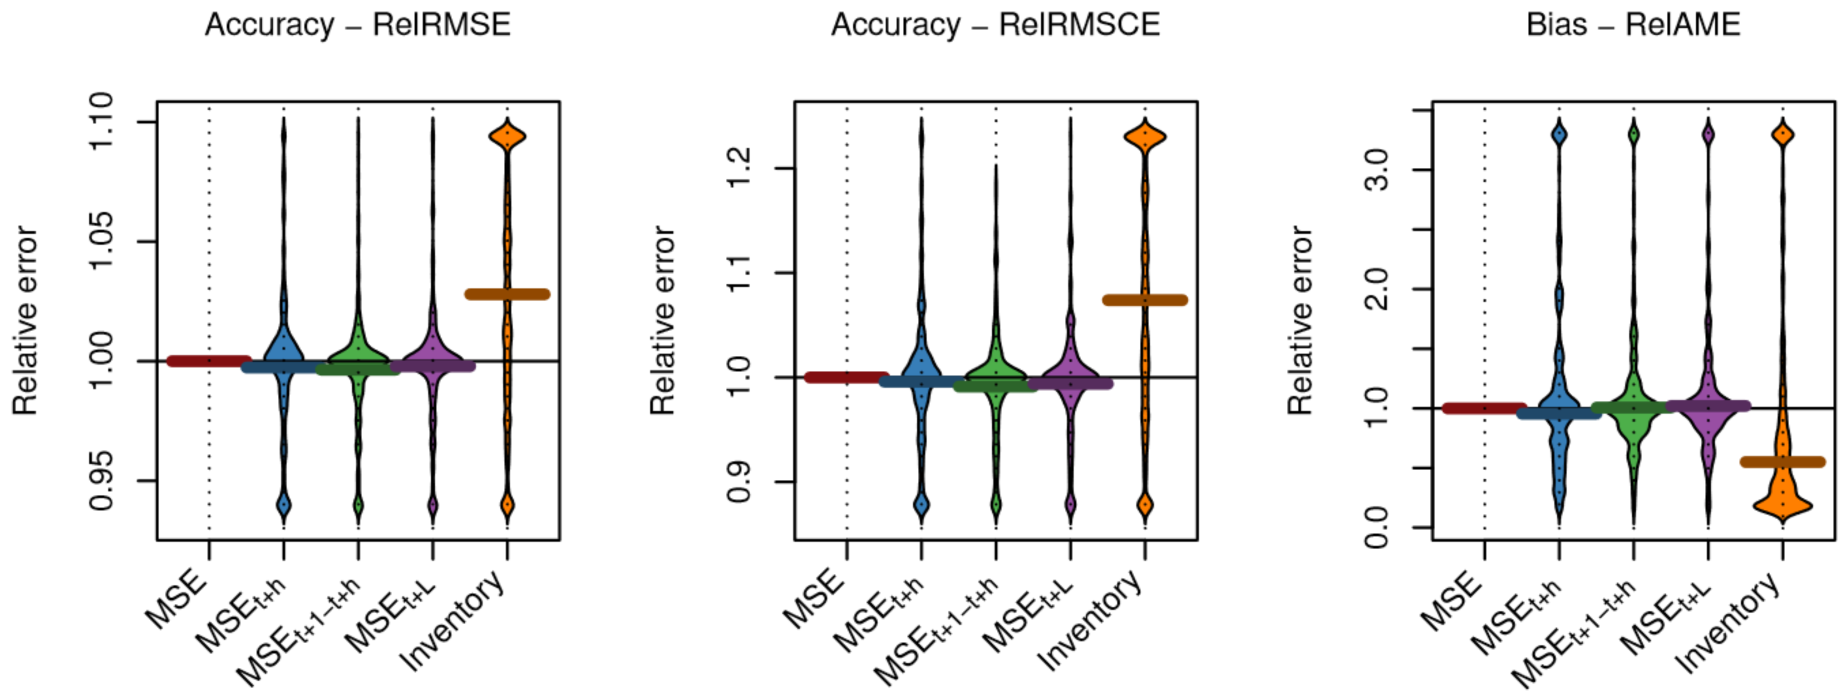

- **RMSE**: Root Mean Squared Error over L.
- **RMSCE**: Cumulative RMSE (i.e. first aggregate over L and then calculate error).
- **AME**: Absolute Mean Error over L.

CMAF

Inventory results (L=3)

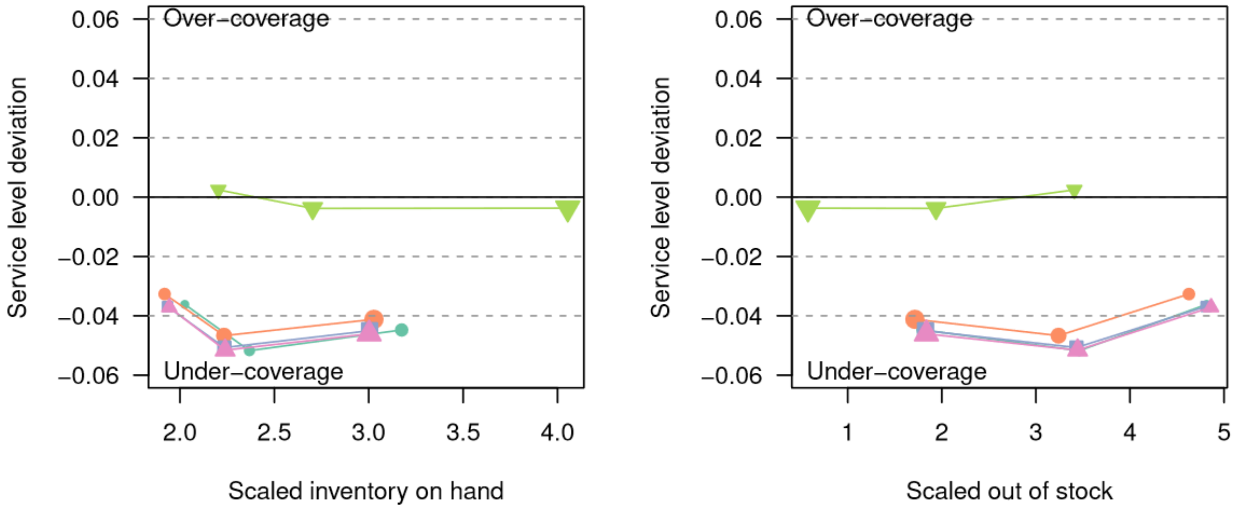

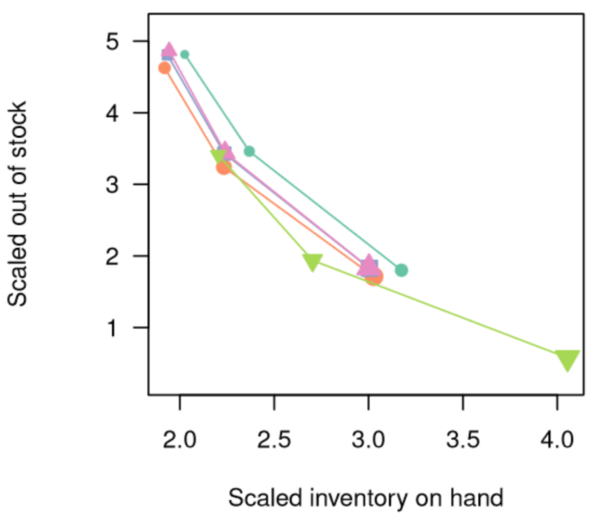

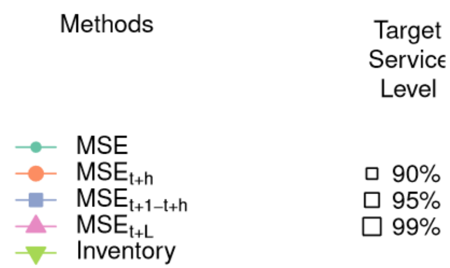

Inventory based results in superior CSL without losing out on lost sales or stock on hand.

CMAF

Inventory results (L=5)

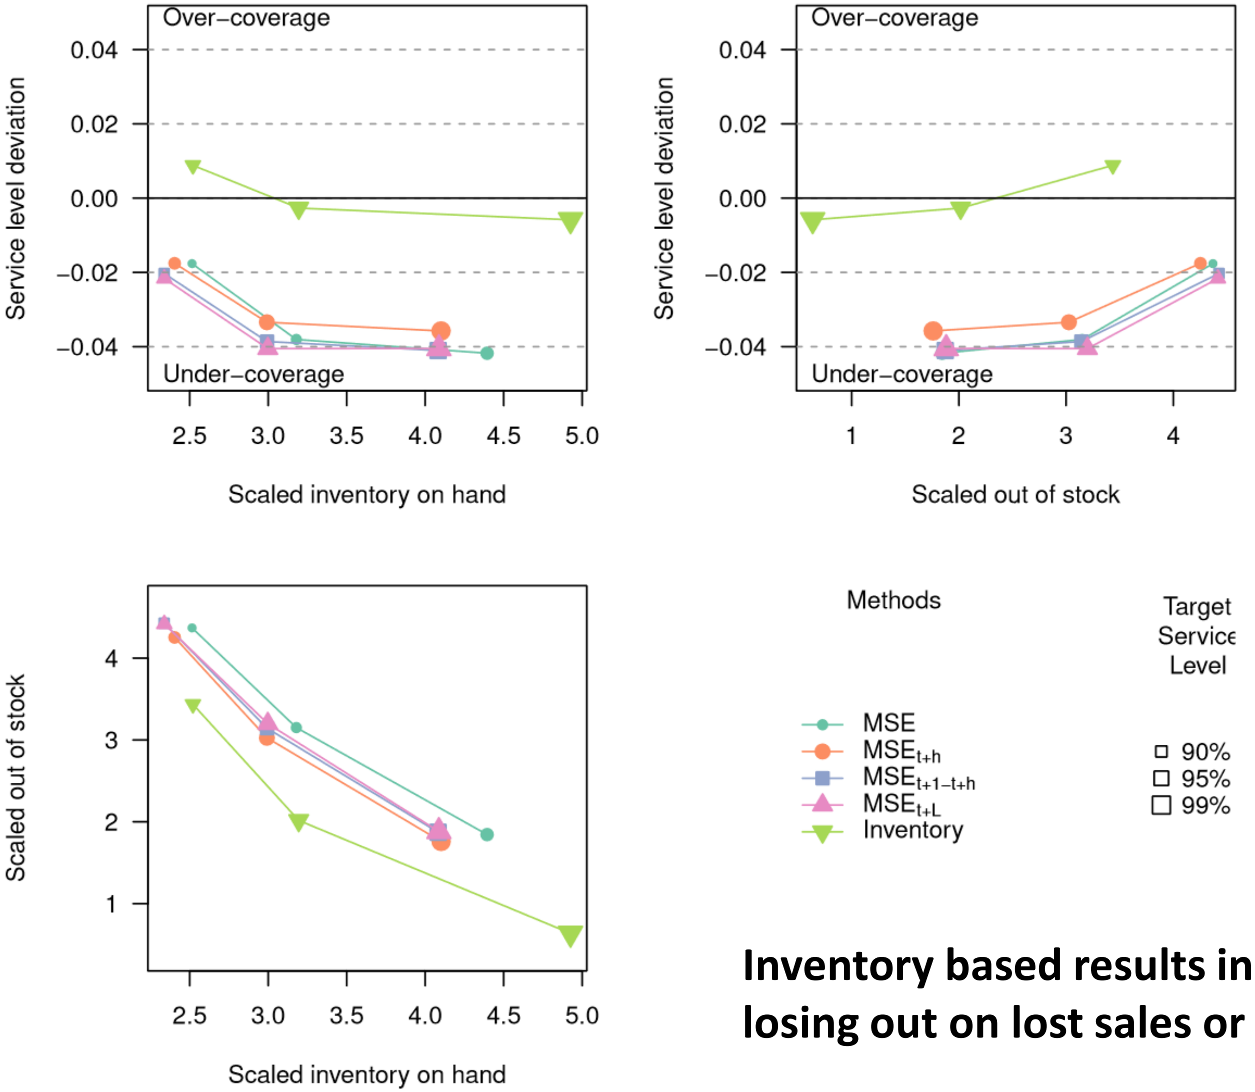

Inventory based results in superior CSL without losing out on lost sales or stock on hand.

Optimisation error surface

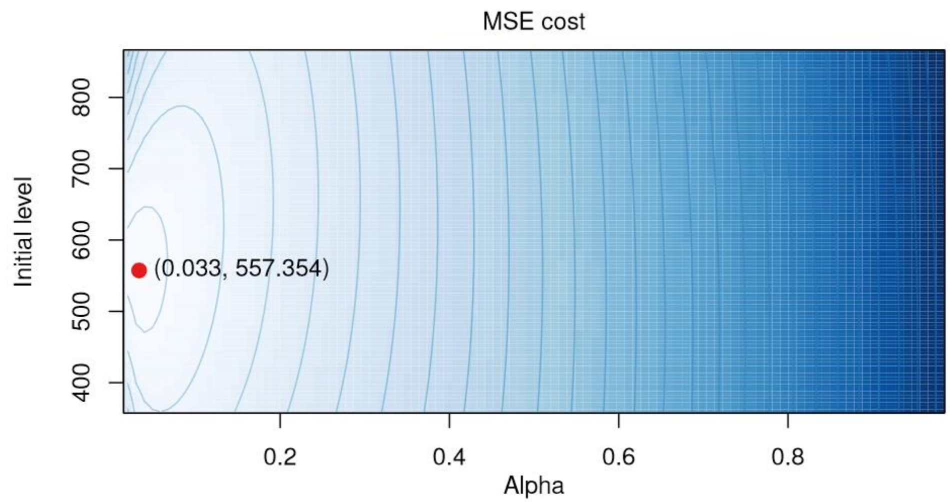

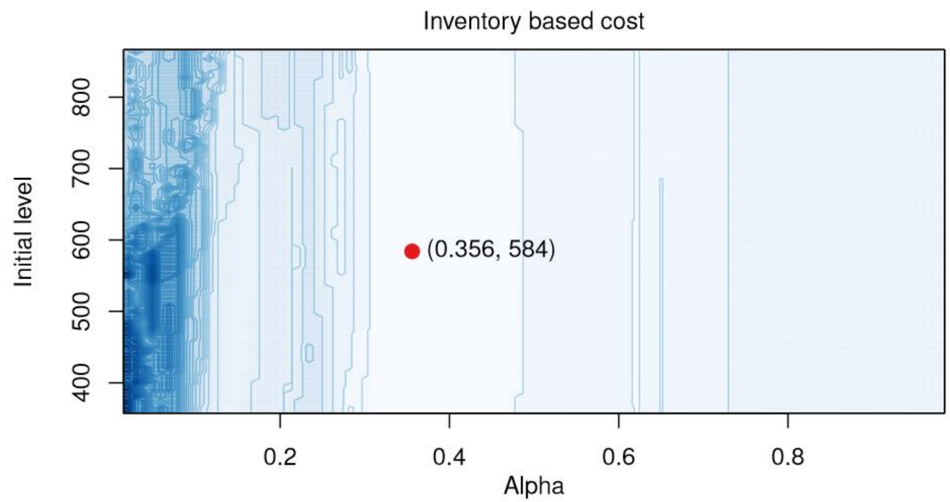

Results for SES (alpha, initial level)

- Surface for MSE well behaved, as expected
- Surface for inventory based cost has multiple local minima → optimization can get stuck.

Optimisation error surface

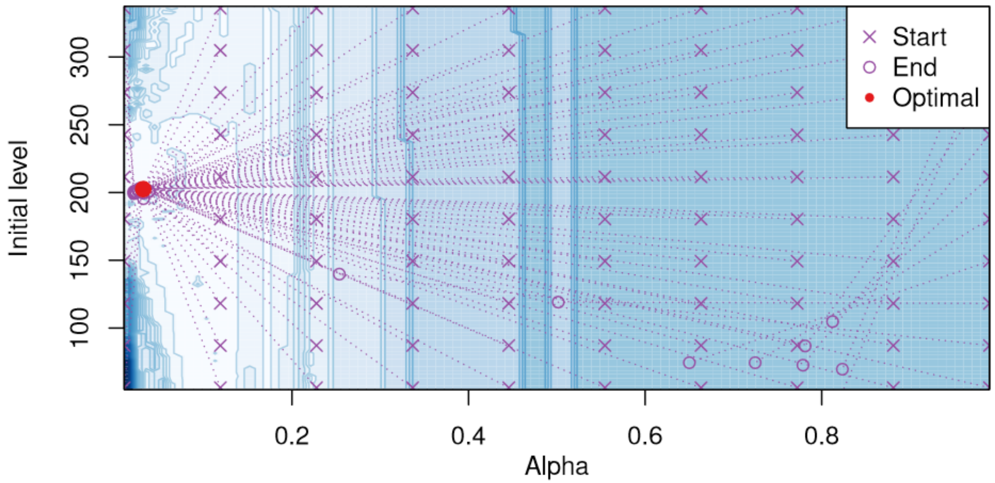

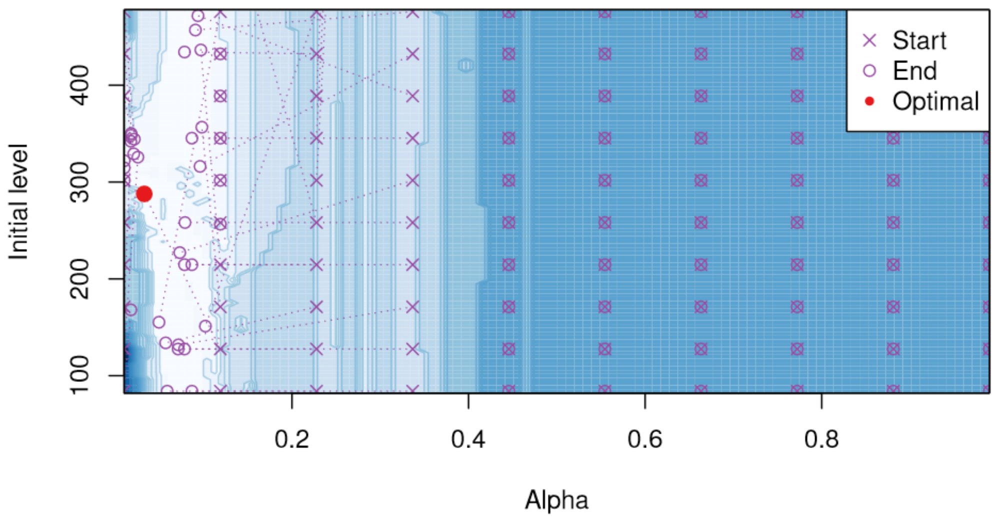

- To resolve the optimization issue we sample the error surface with multiple starts.
- The search is bounded as
  - 0 < alpha < 1
  - Min(y) < initial level <- max(y)

CMAF

Findings

- There is merit into **directly optimising on inventory targets**.
- It is not unreasonably difficult to do so, nor computationally too demanding.
- Can be customised to **match the exact inventory objective** in terms of inventory policy, lead times, service levels, etc.
- Once again: **disconnect between forecast accuracy and inventory performance**.
- Various conventional forecasting objective functions resulted in very similar performance with minor gains when the lead time was considered.
- Idea can be expanded further to **account for any organisational objective** that can be simulated!
- Some issues with optimization, easy to solve, but an elegant solution is future research.
- CSL or Fill rate?

CMAF

Thank you for your attention!
Questions?

Working paper:
https://papers.ssrn.com/sol3/papers.cfm?abstract_id=3363117

Nikolaos Kourentzes (@nkourentz)
email: nikolaos@kourentzes.com
blog: http://nikolaos.kourentzes.com

Marketing Analytics and Forecasting | Lancaster University Management School