

# Using information from the business environment to improve forecasting

**Nikolaos Kourentzes**

Professor in Predictive Analytics

Skövde Artificial Intelligence Lab

Skövde University

Forecasting Forum Scandinavia Workshop

16/09/2020

# A few words about the forum

Enable closer interaction between industry and academia, in questions of business forecasting and predictive analytics. Catalyst in:

- providing **innovative solutions** to real business problems;
- shorten the path to implementing **innovative** and **impactful research to practice**;
- create **consortia** between and within **industry** and **academia** to facilitate ambitious research by sharing know-how, resources, and risk.

Our principle: is to keep this forum an open community. Members can contribute with talks, questions, and ideas to progress the level and quality of business forecasting and predictive analytics.

We aim to host forum workshops bi-annually, with speakers from practice and academia to present their views and challenges on predictive problems, followed by an open discussion with the audience.

Join our LinkedIn group  
<https://bit.ly/2ZJseVk>



# Workshop schedule

Using information from the business environment to improve forecasting

13:00-13:15	Introduction to Forecasting Forum Scandinavia.
13:15-14:15	The academic perspective, Nikolaos Kourentzes.
14:15-14.30	Break
14:30-15:15	The practice perspective, Alexander Norén.
15:15-16:00	Open discussion. Moderator: Anette Tänneryd.

# A few words about the speaker

- BA in Strategic Management
- MSc in Operational Research
- PhD in AI & Forecasting
- Professor in Predictive Analytics

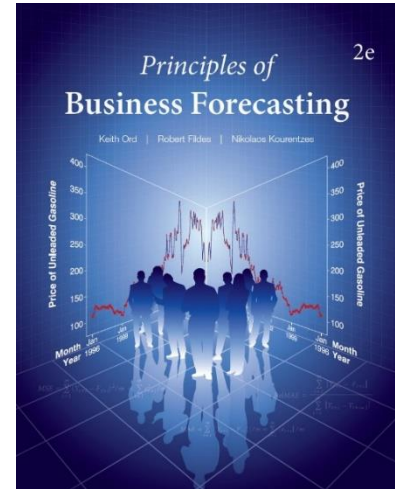
Research interests and consulting experience in various applications of forecasting, including:

- Business Forecasting and Demand Planning
- Promotional/Pricing Modelling
- Supply Chain Forecasting

*Core question:  
How to handle  
model uncertainty*

Long experience in applied research projects with industry in various sectors, including: retailing, FMCG manufacturers, pharmaceuticals, tech, cosmetics, media among others.

Research blog: <http://nikolaos.kourentzes.com>



Business  
forecasting  
textbook

# A few words about the Skövde AI Lab

## Skövde Artificial Intelligence Lab

- Long standing research group in AI & data science
- 8 full time great colleagues and a thriving doctoral community!
- Track record on diverse AI research projects with industry
- Main areas of interest
  - deep learning;
  - predictive analytics;
  - reasoning under uncertainty;
  - visual analytics;
  - transparent decision support;
  - data privacy;
  - recommender systems.



Visit our group at



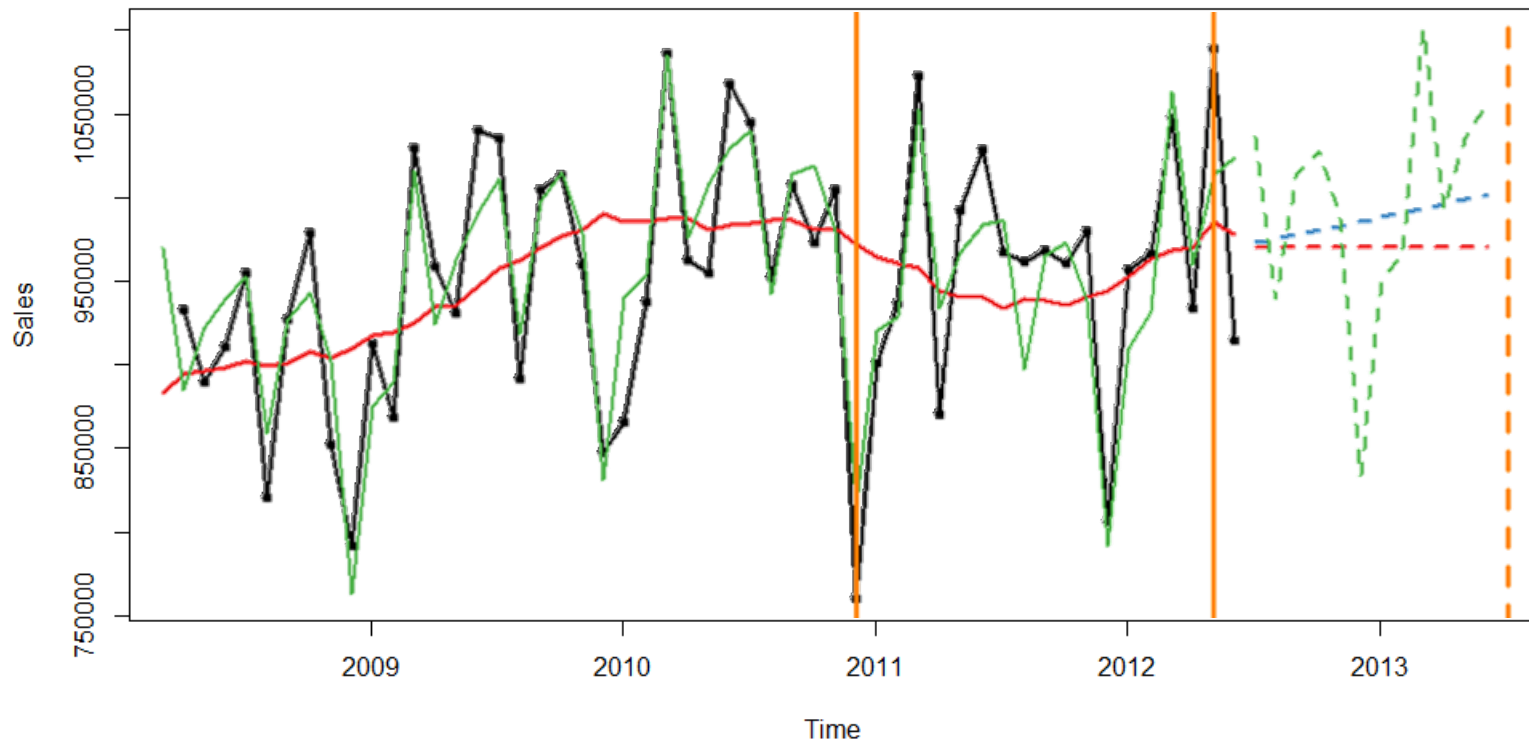
<https://bit.ly/35oVXXe>

# Agenda

1. What makes a good forecast?
2. What is the standard approach to business forecasting?
3. What are its limitations?
4. Enriching forecasts with external information
  - From the economy
  - From the market and our customers
5. A holistic forecasting approach: the way forward!

# What makes a good forecast?

- We start from a fairly uncontested statement: A good **forecast** should capture the key features of the demand series we are trying to predict.



Do we  
ever  
expect  
the  
forecast  
to be  
spot on?

Forecast = ( local level (of sales) + slope x seasonal profile ) + special events

and a **random unforecastable** element that is the unexplained variability of the data.

# What makes a good forecast?

- We are not producing forecasts for the sake of forecasting, but to support decisions.
  - More **accurate forecasts** → **less uncertainty** about the future → **better decisions**.
  - In producing good forecasts, our objective is to model the uncertainty that relates to our decision, which itself is unknown.
- First, we need understand a bit more about this uncertainty.
  - Suppose we sell some product, and we expect on average to sell 100 units per day. We understand that some days we may sell less or more than a 100. Why?
  - To make our life easy let's suppose our product has no seasonality, trend, or other special factors. Where does this randomness comes from?
  - **Having perfect information** about the complete economic system, all decisions of our customers & competitors **is infeasible**, nor do we understand these interactions fully. We capture key features of our demand, but not all details → these are part of the randomness.

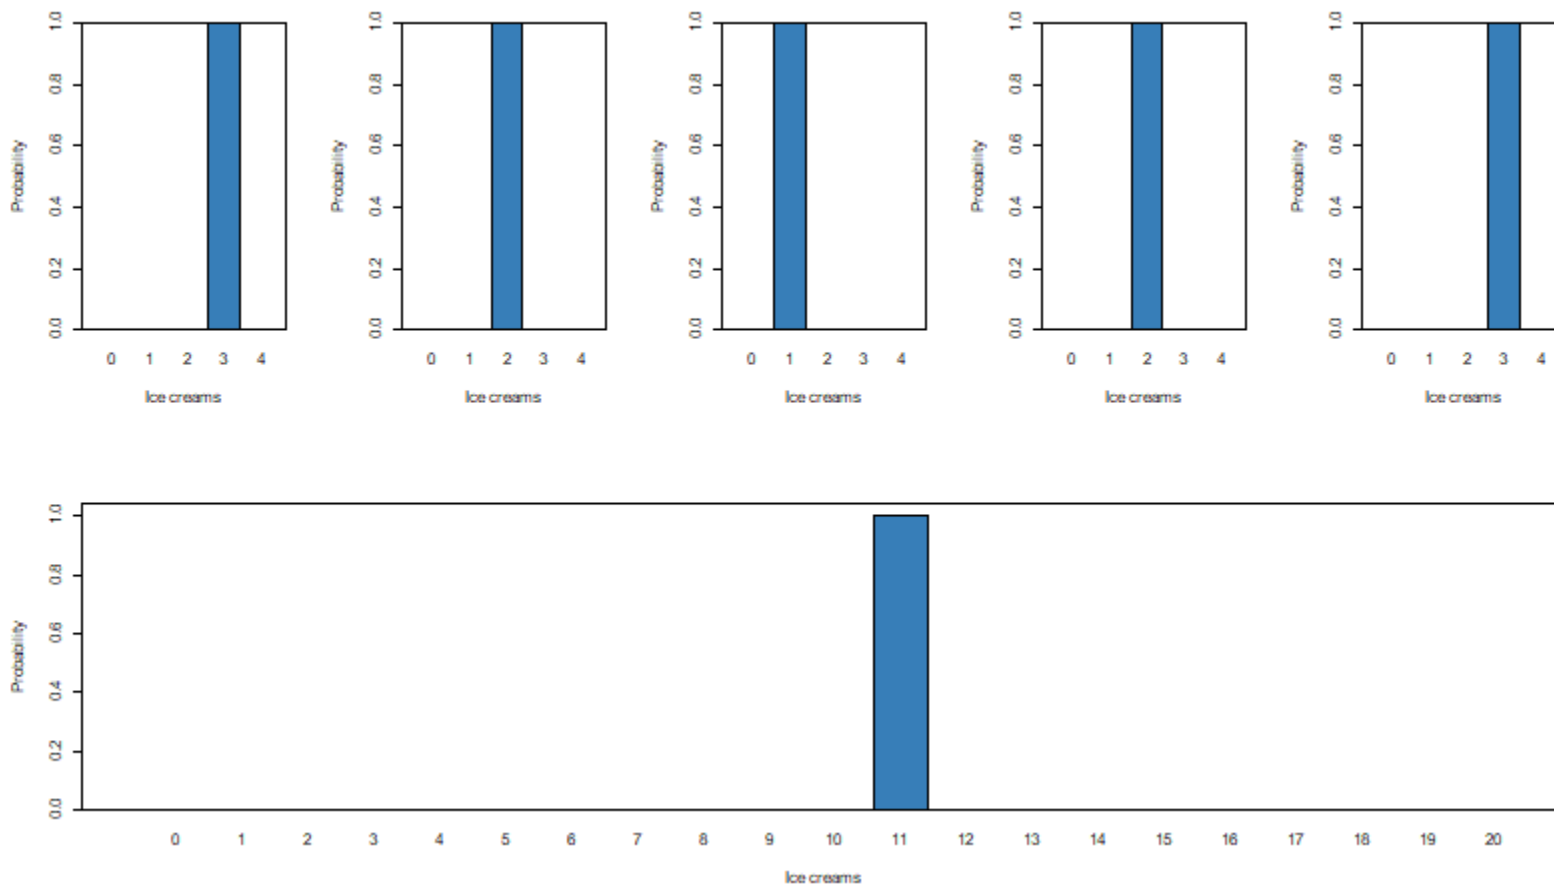
# What can we tell about this randomness?

- Randomness may be unforecastable, but we can still characterise it.
- Suppose we have 5 customers with equal chance of buying 0 to 4 ice creams per day, i.e. on they average buy  $(0+4)/2 = 2$  ice creams per day (our forecast). For a single day:



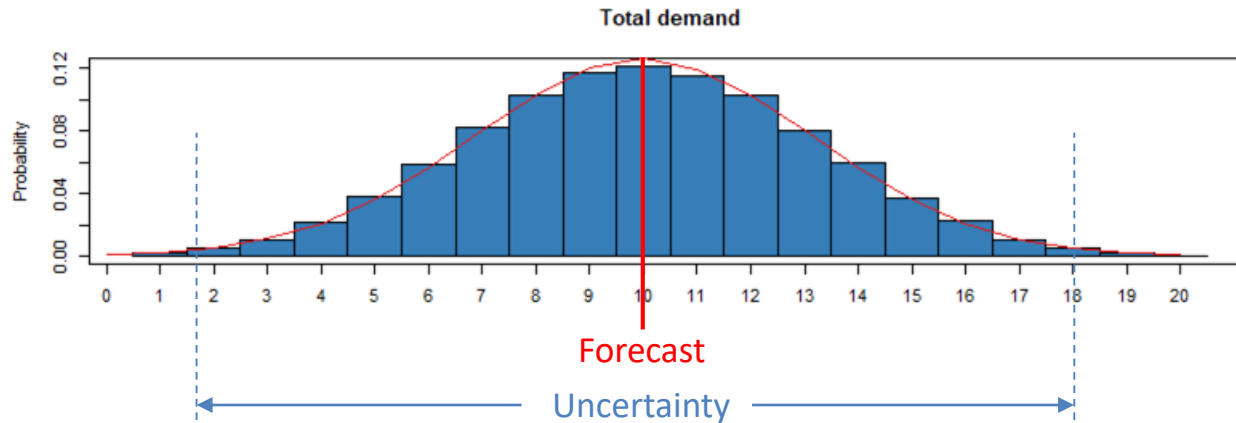
# What can we tell about this randomness?

- Across multiple days:



# What can we tell about this randomness?

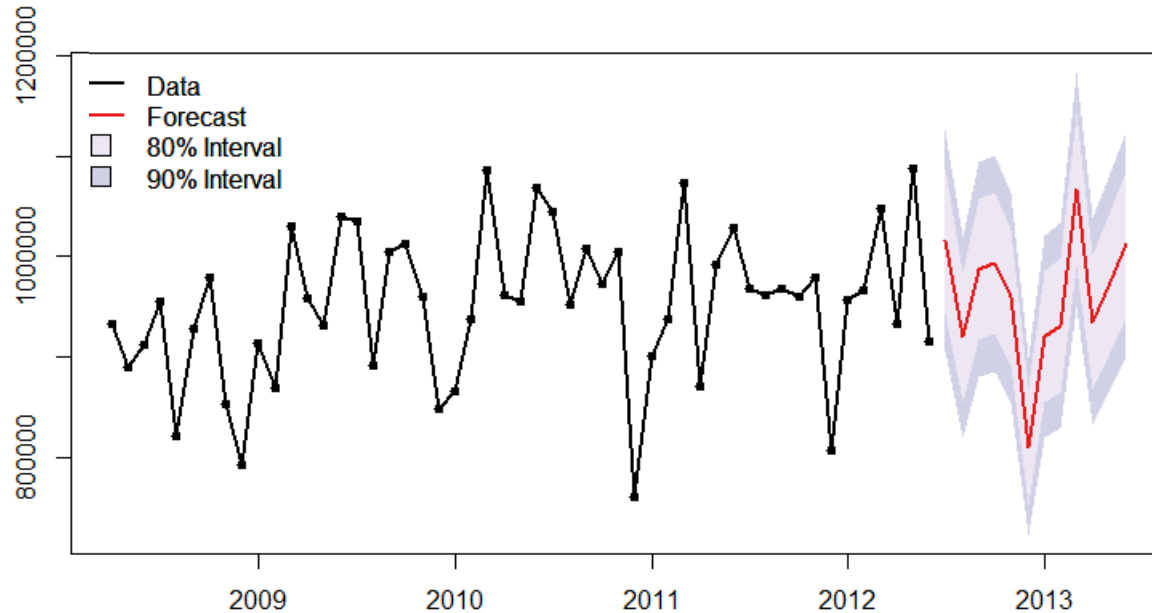
- If we observe the demand across many periods (or customers, or stores, ...) we can see that we have a clear “**on average**” sale and **bounds for the uncertainty** → **this is a complete forecast**.



- Reality is a lot messier, but a good forecast needs to have these two elements! Good forecasts try to capture:
  - Any **structure of the demand**;
  - Characterise the remaining **randomness**.

# A good forecast

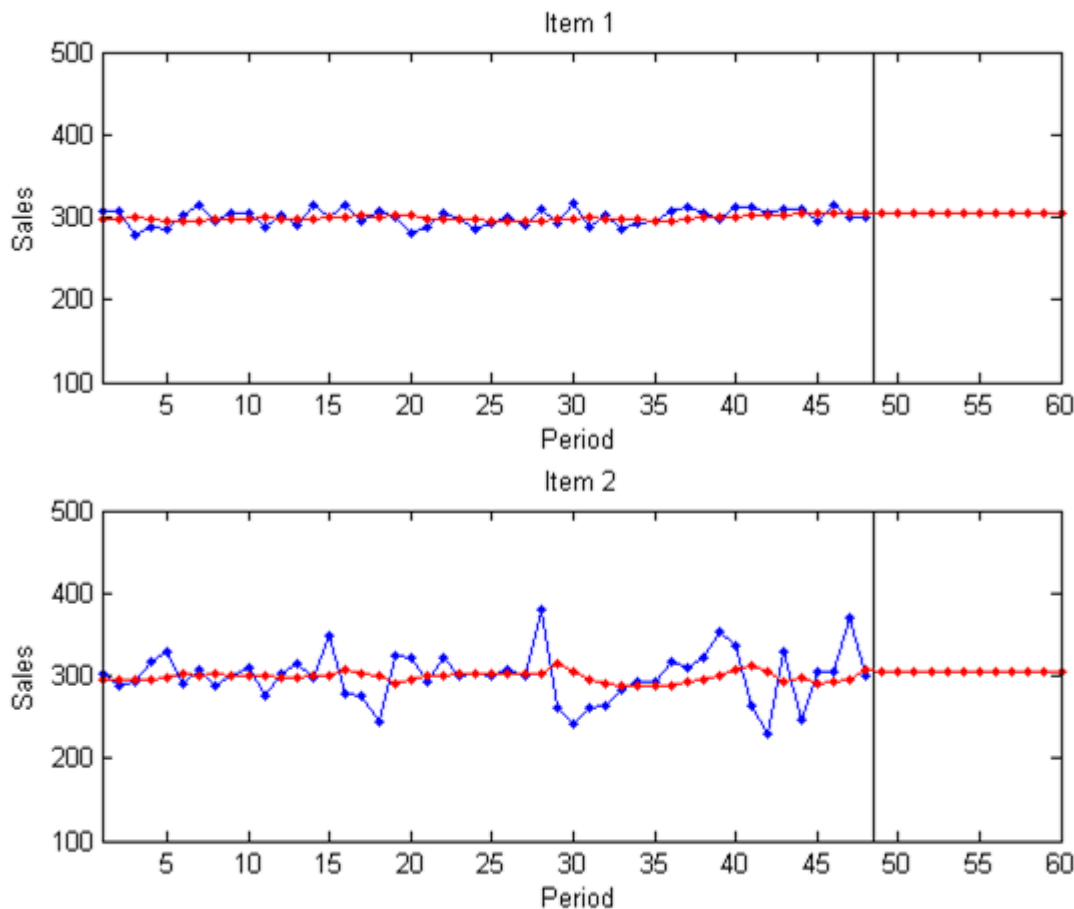
- Everything contains some randomness, and that cannot be forecasted.
- We need to remember to resist our urge to over-explain the data.
- Instead, we need to characterise the risk associated with a prediction.



- Wide-spread adoption of these prediction intervals has been the main driver in capitalising on the benefits from improvements in forecasting.

# Why is that “a good forecast”?

- As an example of decision making supported by forecasting we use inventory management.
- We consider two different products (low and high variability) and their forecasts

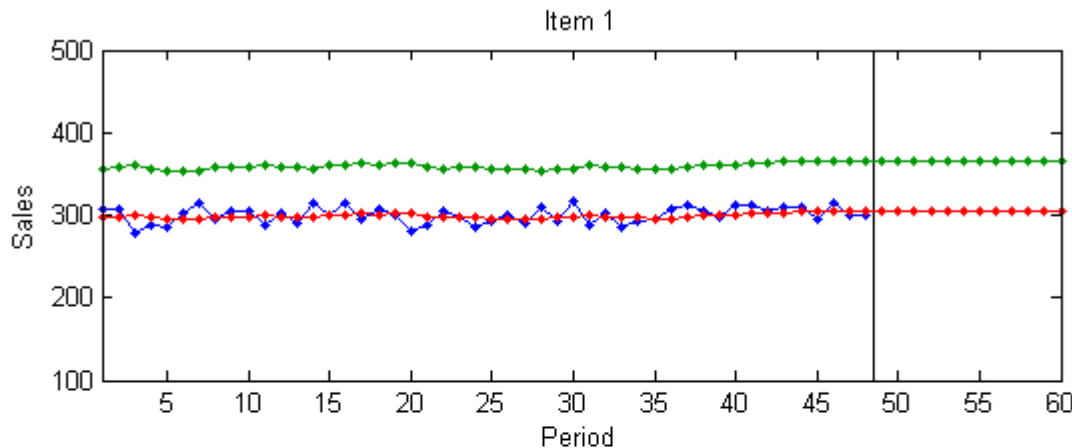


Let's calculate how much stock we should have

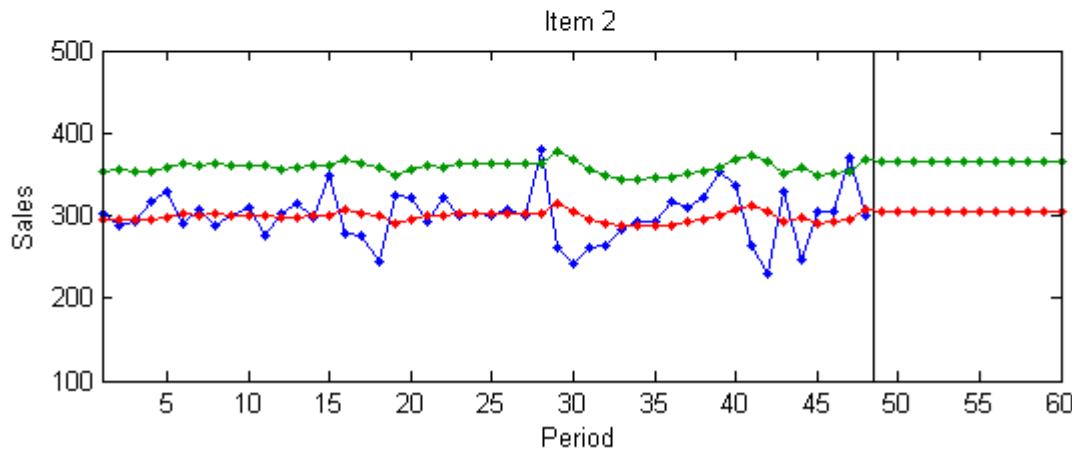
# Why is that “a good forecast”?

For this example let's us always be able to satisfy demand for a given period of time.

Stock = Forecast + Days of Supply (= forecast x some factor)



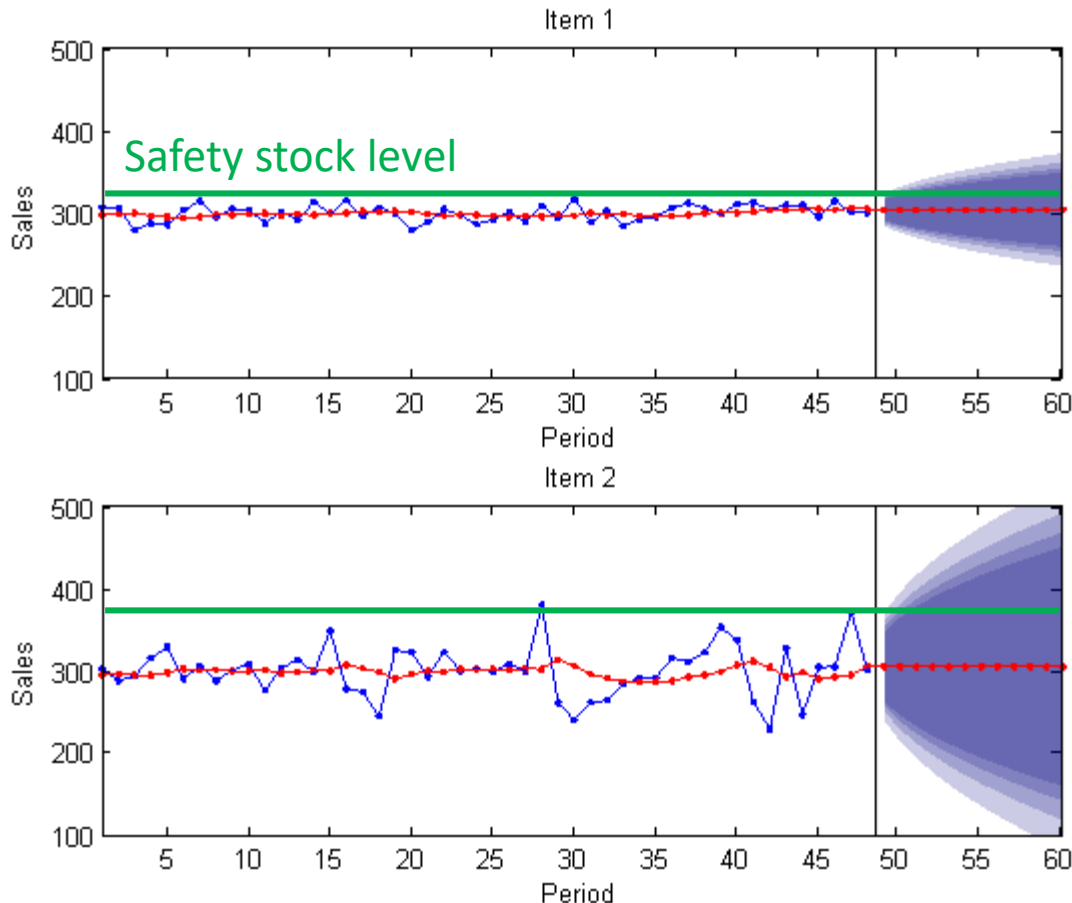
Cover all demand!  
But do I need all  
that stock?



Cover most demand  
and probably the  
stock is about  
right.

# Why is that “a good forecast”?

Can we do any better? What if we tie the forecast uncertainty to the decision?

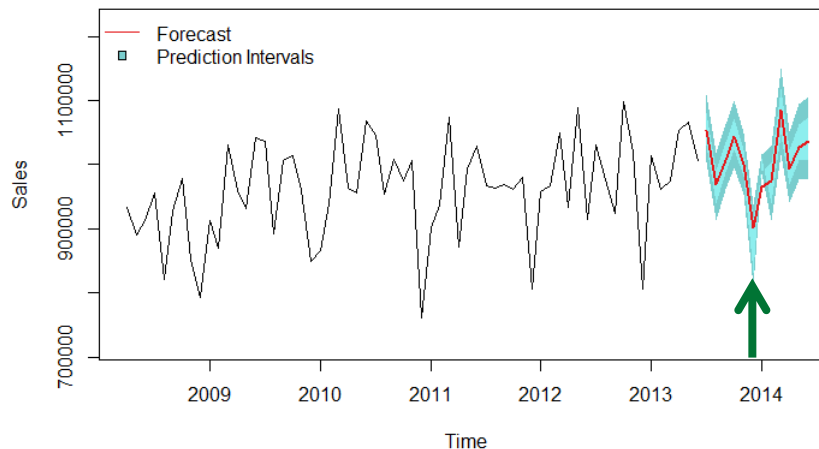


- Stock now adjusts accordingly
- Note that the forecast did not change, but it was incomplete!
- We do understand that the world contains randomness, but we often forget it when we look at forecasts to support our decisions!

# How do we generate these?

Given some historical data, a model based forecast will attempt to capture the key patterns in the data and extrapolate these to the future.

- + Marketing actions
- + Macroeconomic indicators
- + Online behaviour of customers
- + etc.



} Related to the cost of the associated decision

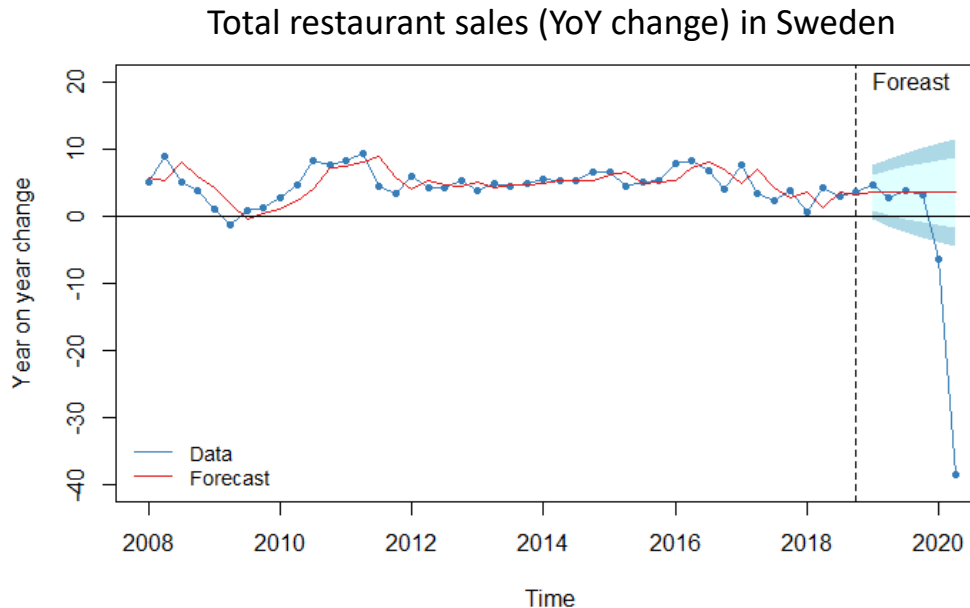
- Better forecasts will capture more parts of the “uncertainty”
  - This is the main benefit of improved forecasts, as the point prediction will not change dramatically.
- The forecast can be enhanced with additional explanatory (causal?) information.
- Superimposed with expert judgment to account for contextual information.

# Extrapolative forecasting

- The most widely used type of forecasts are extrapolative, they extrapolate past patterns in the future (easy to think in terms of seasonality, trend, calendar events, etc.)
- Most software that produce forecasts have an option (or solely) for using extrapolative forecasting (models like exponential smoothing, ARIMA belong here, as well as many AI and ML methods can be extrapolative), because:
  - They are easy to use and build.
  - They are self-contained, past observations are the only information we need.
  - They are easy to automate reliably.
  - They have a proven track record (in practice and academia).
- But they rely on the strong assumption that this structure is immutable in time. Instead of modelling the rules underlying our demand, they model the shape, which will change in periods of disruption.

# Extrapolative forecasting

- But this comes with caveats.



The extrapolative forecast is oblivious to Covid-19. The impact is so significant that sales drop below the 95% lower probability bound.

- The relevant question (now) is not if we could forecast the onset of Covid-19, but rather what the future looks like, when are we back to normal? is it the same normal? etc.

# Sourcing information from the business environment

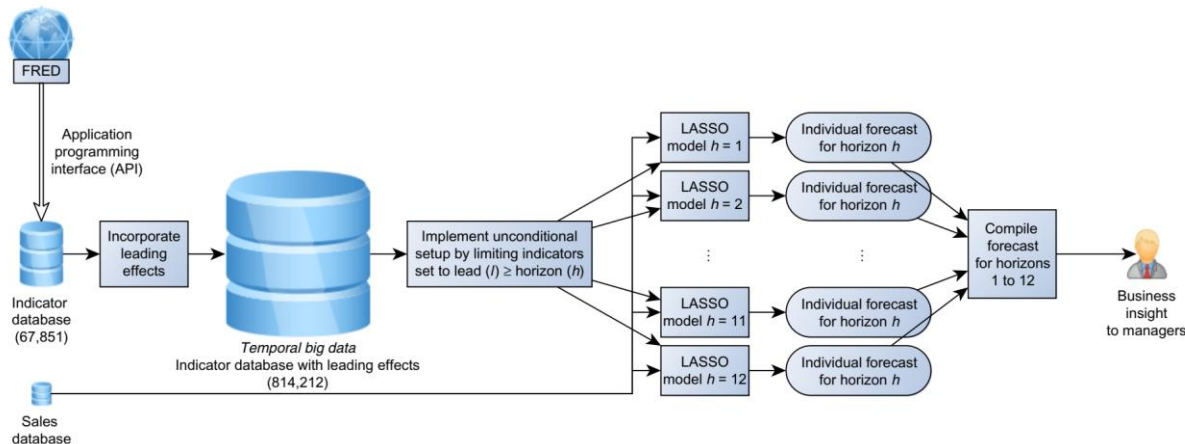
- We can overcome these issues by building “causal” predictive models.

$$\text{Restaurant sales} = \alpha \cdot \text{Covid-19} + \beta \cdot \text{GDP} + \gamma \cdot \text{Unemployment} + \delta \cdot \text{Tourism flows} + \dots$$

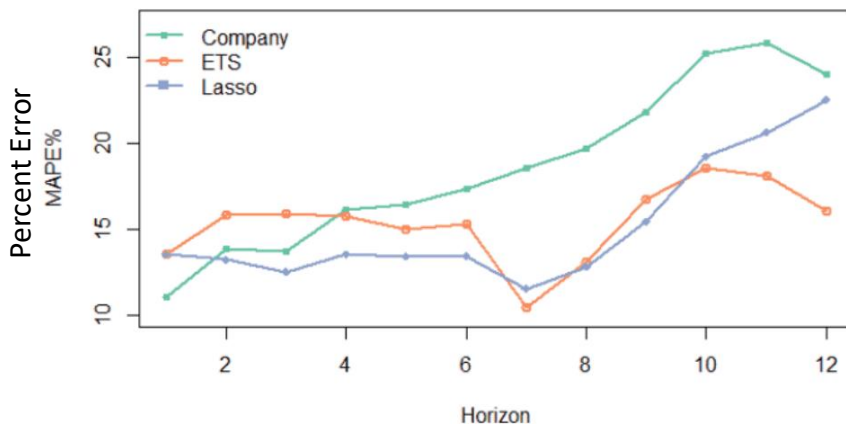
- The model parameters ( $\alpha, \beta, \gamma, \delta, \dots$ ) capture the mechanistic effect of a variable.
- The principle is simple, but the execution can be more involved:
  - Which variables are relevant predictively? If I want to know the sale in the next quarter, I do **not** need just variables that are **correlated now**, but to have **predictive value** for a future period → leading indicators.
  - Select out of the potentially thousand of different indicators the **most informative ones**. This problem is exacerbated by different variables being correlated between them, and us having limited data.

# An example

- We looked at this problem for a manufacturer, for tactical forecasting, and sourced macroeconomic variables (globally).



- A different model for each horizon (how many months ahead)
- Different indicators rise and fall in contribution for different horizons.

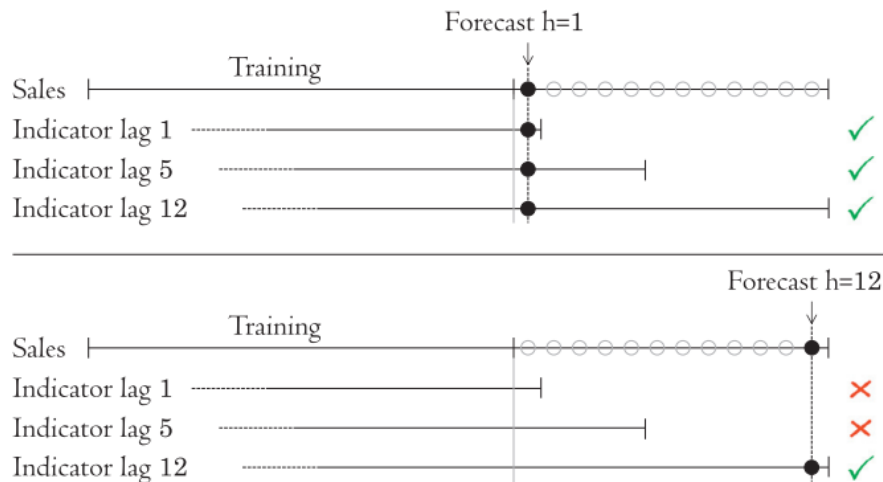


- For very short horizons extrapolative models are well informed.
- For very long horizons the uncertainty makes all forecasts equally bad/good.

# More than just a forecast

- The forecasting model can help experts refine their understanding about the key influencing factors of their market.
- Similarly, experts can help refine the model, and test their understanding of the market.

We are not looking just for correlation, but for informative leading variables



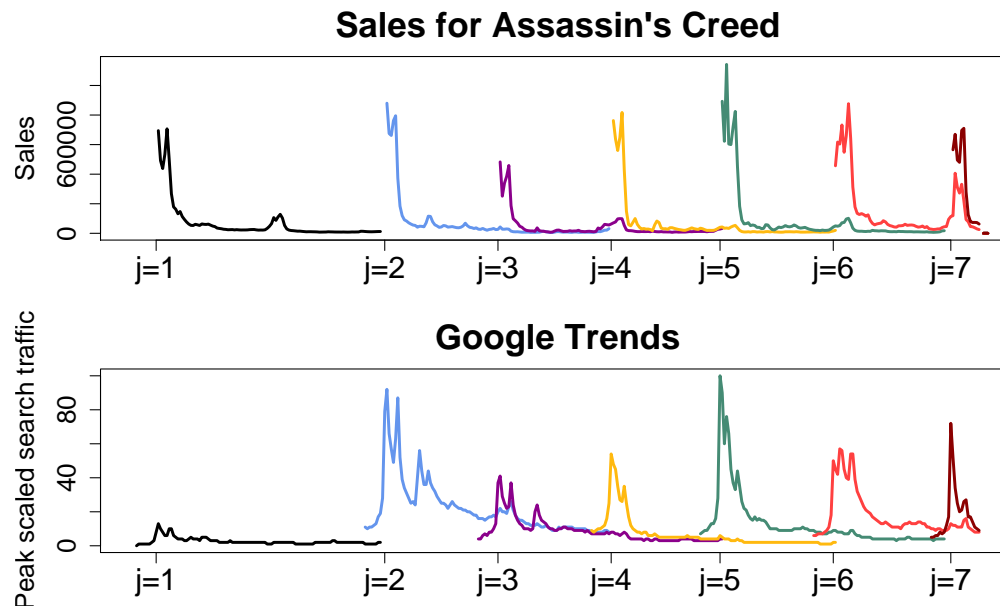
Type of indicator	Percent selected	Indicator units
Labor	18.5	Persons, USD/week, % growth, hours
Consumer price index	12.4	Index, growth rate, % growth
Government services	9.6	Persons, % growth, index
Retail trade	9.3	Persons, index, hours
Financial activities	8.5	Persons, index, % growth, USD
Private services	8.3	Persons, % growth
Transport and automotive	6.4	Persons, USD, index, % growth
Manufacturing	5.8	Hours, USD, net % growth, index
Food manufacturing	5.0	Persons, rate, % growth
Education and health services	4.4	Persons, % growth
Wholesale trade	3.6	Persons, USD, % growth, USD/week
Tourism	2.6	Persons
Recession indicator	1.5	Index
Construction	1.1	Hours, net % growth, index, USD/week
Mining	1.0	Persons, hours, % growth
International trade	1.0	Growth rate, USD, national currency, index
Exchange rate	0.6	Index, % growth

# Some fallacies of causal models

- Causal models are intuitively pleasing: X happens because of A, B, C, ..., etc.
- But they can be rather dangerous!
  - Are the explanatory variables known in the future? With what certainty?
  - A very large model (trying to capture everything) is very sensitive to small issues, and is probably over-fit to the past, suggesting false and spurious connections.
  - Most software and techniques underestimate the uncertainty of the forecast, due to not accounting for the uncertainty of the explanatory variables.
  - They give us a false sense of control, leading to very dubious decisions!
- This has led to the common argument: different models for describing the data and different models for predicting the data → **No! A valid model/theory should describe and predict at the same time.**
- These warnings are the reason why extrapolative forecasts are so common, yet they have well known limitation.

# Beyond macro-data

- Beyond official statistics sources, we can mine variables from the online interactions with or between our customers.
  - The literature has been very optimistic on this. We found that for established products, there is in fact limited value (the decision horizon is the limiting factor).
  - But what about product launch/adoption? On-going work is more promising.



*What about influencers?*

# A classic business problem

- Companies rely on forecasts to support decision making at different levels and functions.

Level	Horizon	Scope	Forecasts	Methods	Information
Operational	Short	Local	Way too many	Statistical	Univariate/Hard
Tactical	Medium	Regional	↕	↕	↕
Strategic	Long	Global	Few expensive	Experts	Multivariate/Soft

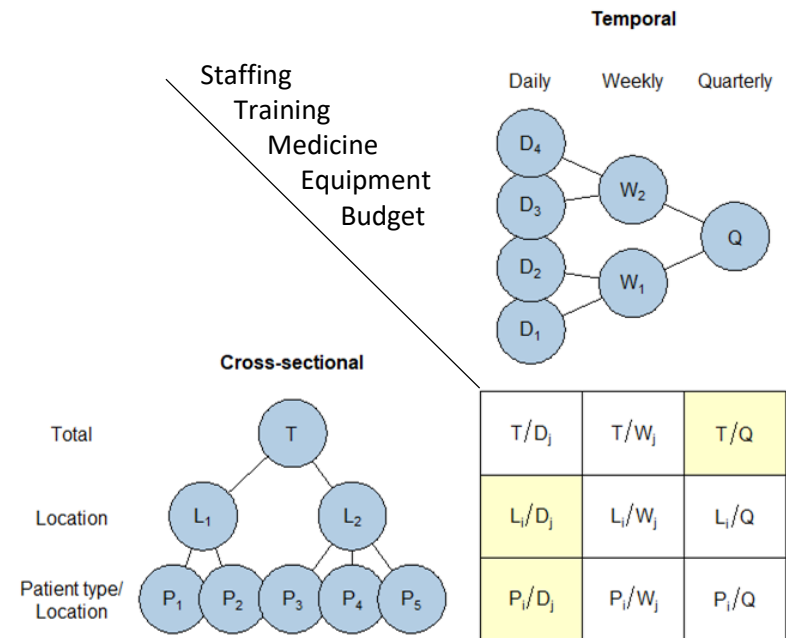
- Operational forecasts are short term and dominated by extrapolative approaches.
- Strategic/tactical can benefit substantially by models that are using leading indicators, macro-economical or otherwise sourced, but are typically done using expert judgement.
- Nonetheless, different functions within an organisation and different planning horizons must point to the same future! Yet, forecasts are often misaligned.
  - Misaligned forecasts → Misaligned decisions → Waste/lost opportunities.

# A hierarchical perceptive

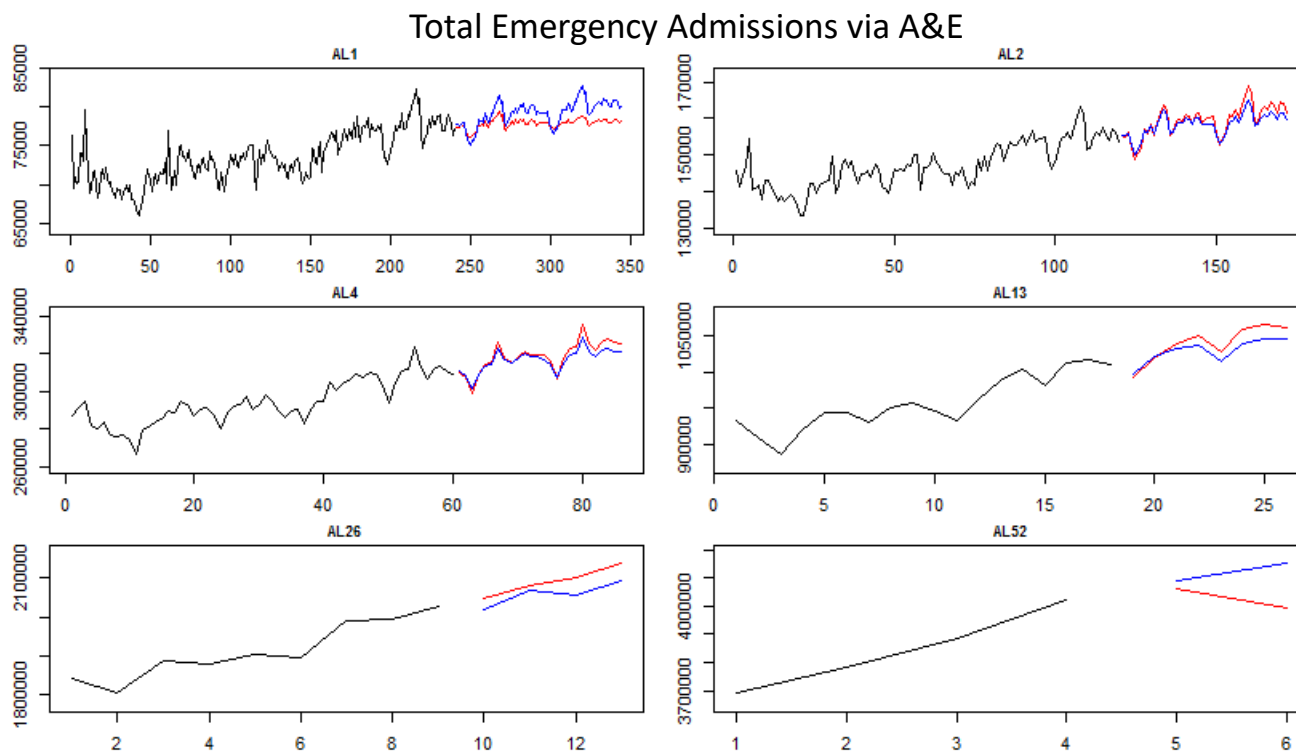
Let us consider the practicalities of health response to an outbreak like Covid-19.

## Multiple problem dimensions

- Planning for local and regional resources.
- Various decisions have different lead times, yet they all need to materialise for the health system to operate.
- Need for aligned/coherent decision making to avoid lack of critical resources/waste.
- The norm is to model the “nodes” of interest independently, but this is a wasted opportunity as there is additional information in neighbouring nodes.



# A visual example on planning horizons



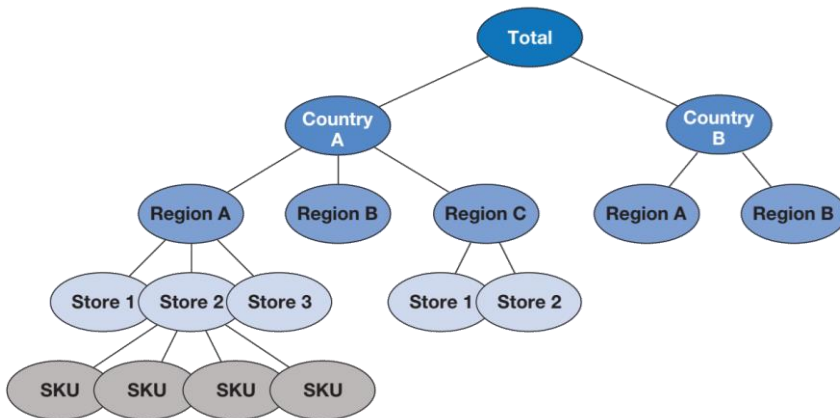
**Red** is the prediction of the base model – at each level separately

**Blue** is the hierarchical forecasts

Observe how information is 'borrowed' between temporal levels. Base models provide very poor weekly and annual forecasts, as the local information is inadequate.

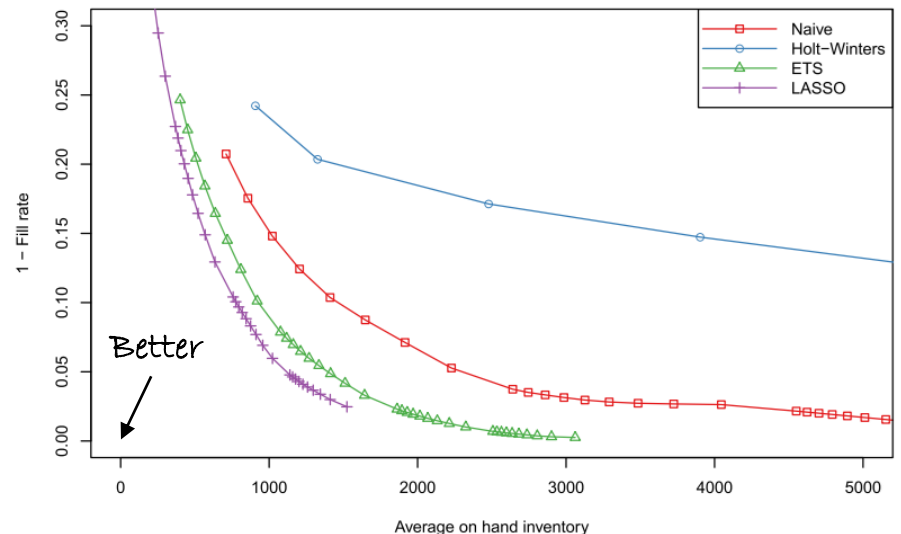
# Putting everything together

- We return to the example with the manufacturer and the tactical forecast based on leading indicators. Can we help operations?
- They often cannot implement such a model, but have to rely on extrapolative forecasts (and expert adjustments, since these people are close to the market!)



... to enrich decisions at all levels, e.g. Inventory decisions at the most disaggregate level.

We rely on the way the company has organised its markets and operations to **connect the independent forecasts, across regions and planning horizons.**



# Conclusions

- Better forecasts → better decisions.
- It is not just about getting the point forecast right, but also about characterising its uncertainty → the difficult part.
- Disruptions in the market increase the uncertainty → we need to enrich our forecasts with additional predictive information.
- Extrapolative forecasts are great, but have their limitations → additional information from the economic and market environments, as well as the behaviour of our customers.
- At an aggregate level, predictions with leading indicators can be very effective → the planning horizon is central!
- Hierarchical approaches can help meld information supporting different decisions and planning horizons to achieve an aligned view of the future throughout organisations.
  - It is no longer about top-down or bottom-up, but rather about diverse information.
- **Join the forum! Help us make this a beneficial community!**

**Thank you for your attention!**  
**Questions?**

Nikolaos Kourentzes

email: [nikolaos@kourentzes.com](mailto:nikolaos@kourentzes.com)

twitter [@nkourentz](https://twitter.com/nkourentz)

Blog: <http://nikolaos.kourentzes.com>