10 mistakes that will ruin your demand forecast - and how to fix them

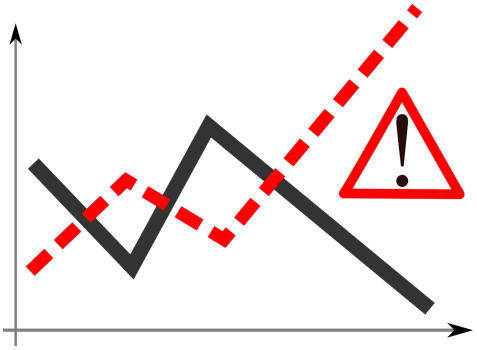

Demand planning is extremely important for any supply chain manager. And the heart of demand planning is demand forecasting. If you get your forecasts wrong, costs go up, revenue goes down, and so does customer satisfaction.

So here are 10 typical mistakes that ruin your demand forecasts. And since I'm here to help, I'll outline how to fix each one.

Mistake #1: Relying on sales forecasts

Let's get this one out of the way right off the bat. Salespeople are indispensable and deserving of our love and gratitude etc. Nevertheless: Unless your sales organization uses solid, data-driven methods to forecast sales, these forecasts are as a rule not very reliable. They are often based on gut feelings and rough levels of closing probabilities, and they can be motivated by various factors that are not related to forecasting accuracy - such as bonus schemes.

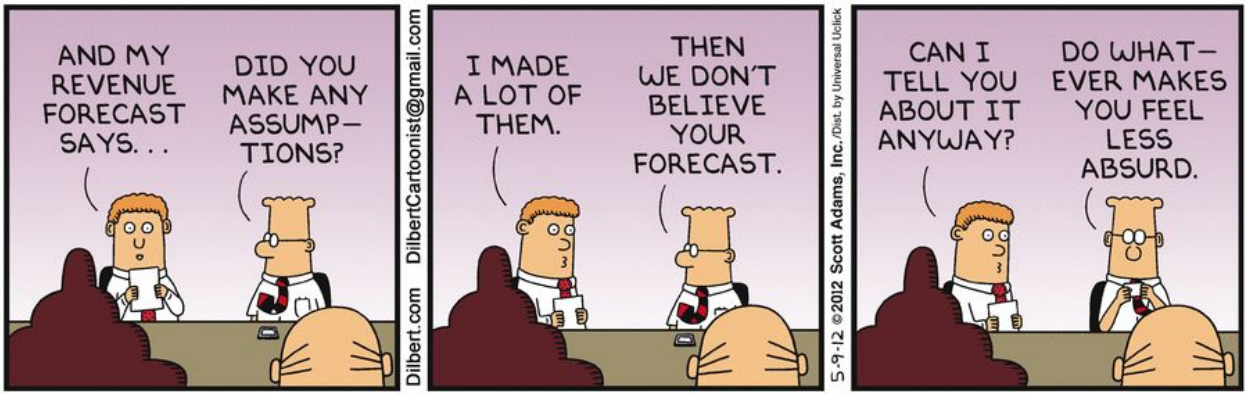

Credit: Scott Adams. https://dilbert.com/strip/2012-05-09

More often than not, sales forecasts are too optimistic. Worse, the level of optimism relative to the truth can vary independently of reality as the causes of said optimism are not necessarily correlated with what is really going on in the market.

Fix: Stop relying on sales forecasts as your demand forecasts. Instead, implement a data driven forecasting process.

Mistake #2: Using a biased model

From here on out, we'll assume that our demand forecasting uses a machine learning or statistical model to predict demand. The model parameters are learned from your historical data.

Any forecasting model must be validated using historical data. This can be done by keeping the most recent data hidden from the model. If you intend to predict one month ahead, then keeping the last month makes sense. Call this a "validation set."

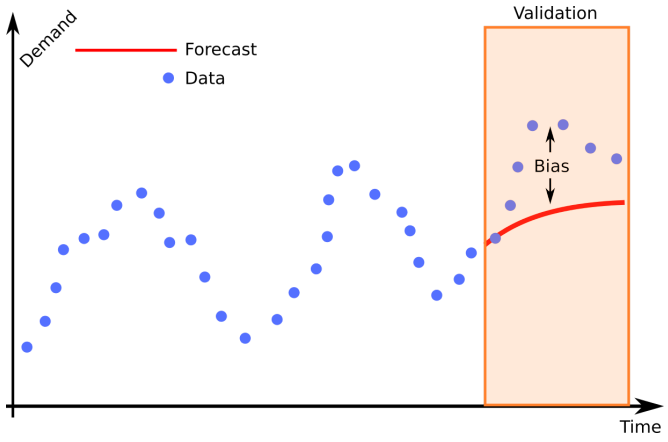

Training the model on the remaining data, run it in prediction mode to see how well it predicts our validation set. If the model is biased, we might see something like this - the forecast is unable to track what happened historically. A biased demand forecast obviously leads to poor demand planning.

Fix: Use a richer model. **Consider increasing model complexity and feeding the model more data, such as external signals. Make sure the model has enough 'memory' to capture longer-term patterns, trends, and seasonality.**

Mistake #3: Using a model with lag

If your model is lagging behind reality, your forecast will always be off. In the worst case, the forecast is so out of sync that it's high when it should be low and vice versa. That can lead to worse demand planning than not having any forecast!

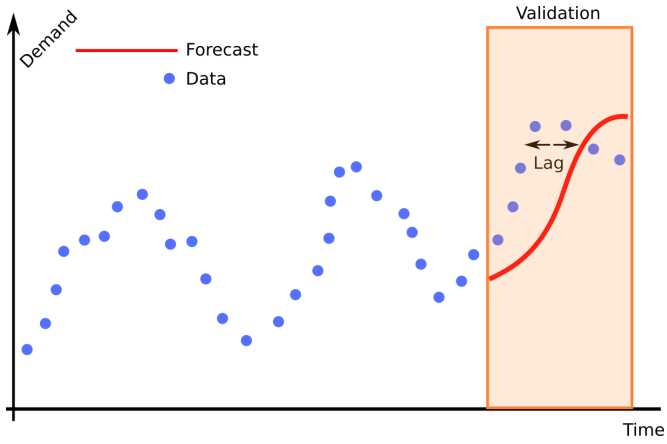

Fix: Implement a more responsive model. Possible options include RNN or LSTM models.

Mistake #4: Forecasting demand for each product separately

So-called 'univariate' models are widely used in demand forecasting. One of the most well-known forecasting tools, Prophet, is fundamentally a univariate tool. The problem is that the demand for any product is probably not independent of the demand for other products. The classic example is the razor and the razor blades. If you aren't capturing this, your forecasts will be much less accurate.

Fix: Implement multivariate forecasting, using for example RNN, LSTM.

Mistake #5: Not sensing demand

Demand sensing is a bit of a buzzword with a sometimes unclear meaning. What does make sense (pardon the pun) is making sure you get the highest possible frequency of data streaming into your pipeline, at the shortest possible lag times.

If something happens to demand at the "front line", but your forecasting relies on weekly - or worse, monthly - batch processing, you'll be caught out. No matter how good your forecasting model is, the forecast will be too late.

Fix: Implement streaming-based forecasting. Get all the important signals flowing through at high frequency and minimal lag.

Mistake #6: Ignoring external signals

A major lesson from years of AI and IoT innovation is that additional data dimensions can improve forecasting accuracy immensely. If demand for your product depends on factors that you aren't even tracking, then obviously your forecast quality will be relatively poor.

Here are a three external signals to consider:

- Weather
- Competitor pricing & locations
- Social media

Fix: Get creative. What could possibly influence demand for your products? Whatever it is and whatever shape or format it takes, there are ways to process and exploit such data with AI.

Mistake #7: Ignoring internal signals

As with external signals, you must exploit any and all internal signals that may carry information about future demand. These may include:

- Price history (not just the present price)
- Promotions
- Demand for other SKUs/products
- Product age, possibly relative to expected life cycle
- Supply chain structure, e.g. a new distribution center

Fix: Review all internal data, including data that could be obtained by investing in sensors. Break down data silos to get a unified data view. Add any internal signal to your forecast model that improves it.

Mistake #8: Underestimating variance

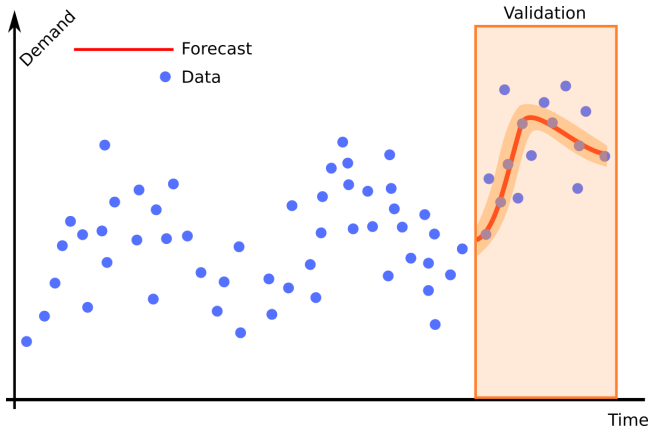

Many forecasting solutions simply provide a single predicted demand quantity for each future time point. We can call that a **bonus mistake: ignoring variance all together.** If you have no idea how uncertain your predictions are, you could be making poor decisions. For example, if your tool predicts a demand of 100, but the prediction could actually be anywhere from 50 to 200, you could quickly end up with dissatisfied customers or wasting cash on inventories that are too large.

If your forecast underestimates variance, you are in for nasty surprises. Suddenly, demand is much higher or lower than expected and you're in trouble.

Fix: Make sure your model captures the variance and uncertainty in the data. Use the right noise model.

Mistake #9: Overestimating variance

This is the opposite of #9. If you overestimate variance, then depending on your overall company strategy you will, to some level, overproduce and stock too much inventory.

Fix: Implement a better model. Investigate model priors and look for signs that training doesn't converge.

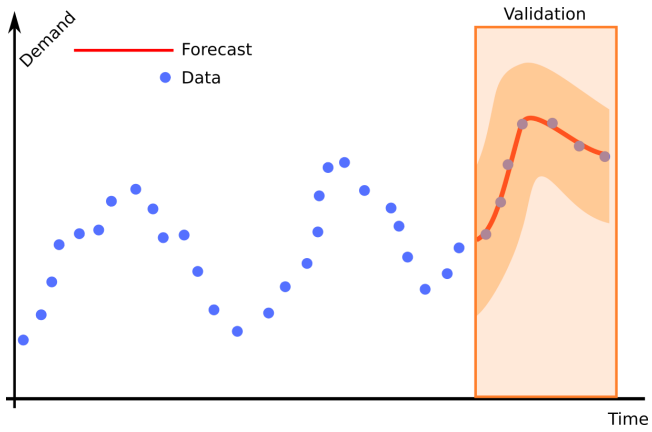

Mistake #10: Not monitoring forecast quality

Once your model is in production, constant monitoring of forecast quality is absolutely crucial. Real demand doesn't know or care about your clever forecasting model. It's up to you to detect a divergence between the forecast and reality.

Fix: Implement constant automated monitoring with alerts and a robust escalation process. Use each case as an opportunity to improve your model further. Always try to figure out what caused the divergence.

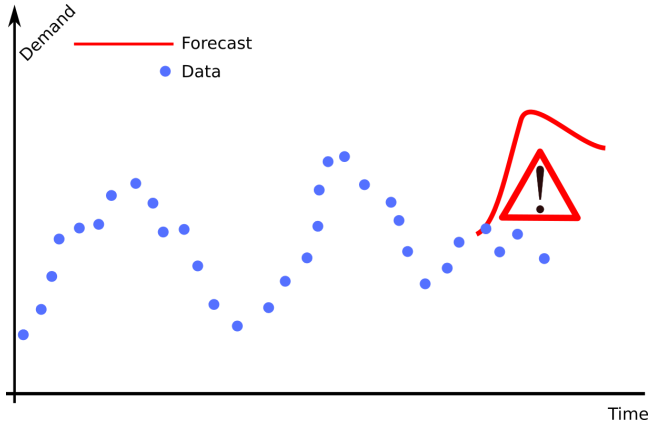

There are other ways to go wrong with forecasting, such as not cleaning your data properly. All of these issues create serious problems for your supply chain. The good news is that they can all be avoided or mitigated by implementing better processes and models and educating your staff about them.

If you want to make sure your forecasting is not misleading you, contact me today - **I'm here to help!** Call me at +45 51 64 66 91 or send me an email at daniel@danieljacobsen.com to set up a free consultation.