ANALYZE - SIMULATE - AUTOMATE

Abels & Kemmner

WHITE PAPER

Dr. Götz-Andreas Kemmner

The market is picking up and delivery capacity takes a nosedive! Could there be a clearer sign of unrealistic sales forecasts in many companies?

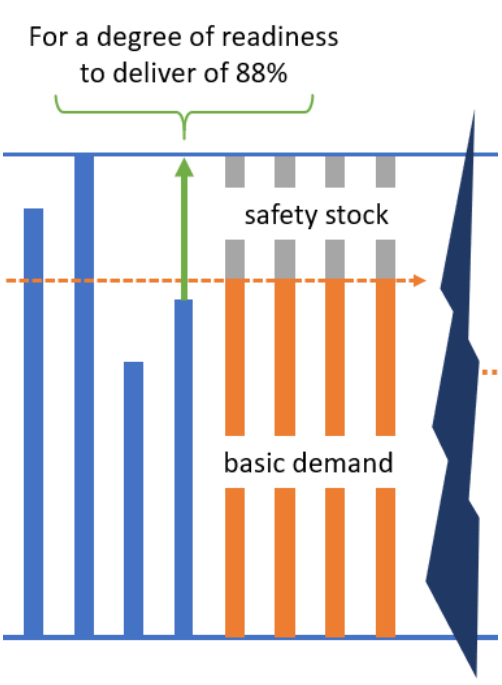

BEST PRACTICE RULES for effective DEMAND FORECASTS

Inhalt

ANALYZE – SIMULATE – AUTOMATE A QUANTUM LEAP AHEAD	**Abels & Kemmner** brings to light the 0ptimization potential of supply chains and replaces gut feeling with facts.
Thanks to our **unique consulting approach**, we help companies achieve sustainable concepts that we validate and optimize and implement in a secure and agile manner.	**As a pioneer** in logistics simulation and automation in supply chain management, we combine strategic and operational consulting with powerful digital methods.

© Abels & Kemmner
Gesellschaft für
Unternehmensberatung GmbH
Kaiserstr. 100
52134 Aachen / Herzogenrath

AACHEN - EINDHOVEN - LONDON

www.ak-online.com

ak@ak-online.com
+49 (0) 2407 9565-0

Best Practice Rules for effective Sales Forecasts

Dr. Götz-Andreas Kemmner
LinkedIn

The market is picking up and delivery capacity takes a nosedive! Could there be a clearer sign of unrealistic sales forecasts in many companies? Companies with sales forecasting that works - we know a whole range of them and have helped many of them ourselves to achieve a reliable sales forecast - responded faster at both the start of the recession and when the market bounced back again.

Despite everything, the importance of the sales forecast for corporate success is underestimated by many people: On the one hand, we come across companies who believe so faithfully in their own expertise and current ERP system that they do not even notice how bad they are doing at the moment, compared to their actual potential. On the other hand, some companies are of the opinion that a sales forecast is no longer required in the age of production that is in sync with the market, and that poor readiness for delivery or high stocks are attributed to the poor performance of their logistics. But how are you going to know how to trim the sails (= plan and schedule) if you don't know where the wind is coming from (= future demand)?!

How are you going to know how to trim the sails if you don't know where the wind is coming from?!

Using the following 12 basic principles and best practice components for efficient and effective sales forecasts, we want to correct any fallacies and point out the important success factors for good sales forecasting.

You are not likely to be able to implement all the best practice rules. Either because you lack the tools for this or because your company's structure or your product portfolio does not allow for it. But the more these best practice rules are observed and implemented, the more effective your sales forecast will be.

Basic principle 1: Market demand does not stick to corporate plans

Anyone responsible for supply chain, inventories and readiness for delivery, sees something completely different in sales planning or sales forecasting compared to Sales or Category Management.

Imagine you are responsible for sales and preparing for the planning meeting for the next financial year. The markets are uncertain, the global economy is sputtering to a halt and you assume that successfully achieving the previous year's sales might prove difficult.

Gesellschaft für
Unternehmensberatung GmbH
Kaiserstr. 100
52134 Aachen / Herzogenrath

www.ak-online.com
ak@ak-online.com
+49 (0) 2407 9565-0

During the meeting there is the classic exchange of opinions. Meaning, you enter the meeting with your own opinion and leave it with the management board's opinion - there has to be another 10% in it. The head of logistics now crosses your path and inquires about your expectations for the next few months. Do you tell him to plan and schedule cautiously or will you get him in the mood for the great 10% increase?

Anyone who is responsible for a supply chain's economic efficiency and inventories has a completely different grasp of sales planning than the company's divisions which are focused on sales.

Sales is compelled to align itself to the sales plan. In some companies there are also interests of other corporate areas which need to be considered. Controlling might look at keeping the scheduling, quantity and price deviations as low as possible and production tries to make sure capacity utilization corresponds to the targets. As a result, the sales plan figures are often bent into shape anywhere that sales and operations planning takes place.

However, logistics and the entire supply chain need the **demand reality** expected in future.

Anyone indicating sales quantities that do not turn up later, and do not show up in any statistics, raises the supply chain's costs. Shifting sales planning figures between product groups, on the other hand, in contrary to realistic expectations, for low planning deviation's sake, robs the company of market opportunities, revenue and often even increases costs, too.

As a result, Best Practice Rule 1 is:

A responsible sales forecast ensures that reality, and not hope, is the starting point of supply chain planning.

It helps to identify the actual expected demand volume, so it can be incorporated into and established in the further planning process if there is sufficient liquidity, production and supplier capacity to do so.

Basic principle 2: Forecasts at product group level

Forecasts at product group level are comfortable for sales and useless for production and purchasing.

If sales forecasts are produced by sales, this is usually done at product group level, i.e. for a condensed group of similar but not identical articles. However, purchasing has to order specific individual parts and production orders have to be produced for specific parts as well.

Split factors are often applied to get from the sales' blanket statement to individual material numbers. A product group's forecasted demand, e.g. the product group television sets, is therefore broken down into the individual sets' demands: On average 22% of the demand is attributed to the HD-TS 15 set, 8% to the HD-TS 16 set etc. In case you like working with split factors, have you ever tried calculating the monthly fluctuation of split factors? Probably not, otherwise you would have a more critical perception of split factors...

Split factors almost always fluctuate so drastically that they become useless for deducing the demand of an individual material number from a product group forecast in many cases. This is known as the "split factor problem".

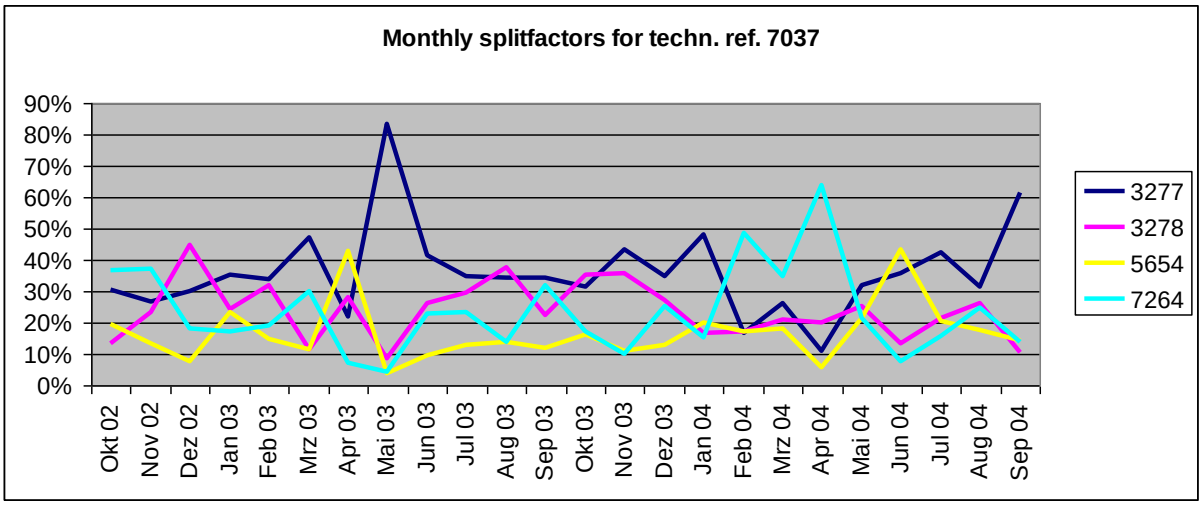

Figure 1: The breaking down of a product group related forecast into individual articles usually fails due to the split factor problem

If you need statements about the future demand of specific articles, materials numbers, SKUs (Stock Keeping Units), inventory items, etc. you cannot avoid an article-specific forecast.

An additional consideration at product group level, plus a look at monetary units instead of quantities, may make sense if you are comparing statistical forecasts and sales expectations with each other; but exclusively for that.

Your Best Practice Rule 2 is therefore:

Forecast at individual part /SKU level and avoid split factors.

Basic principle 3: Sales is completely out of its depth …

Sales is completely out of its depth with a forecast at article/material number level.

This might sound a bit trivial at first, and it is one reason why sales usually submit its sales forecasts at condensed levels, e.g. at product group level.

However, when it is supposed to deal with possibly thousands of material numbers, sales is not just out of its depth in terms of quantities, but also in terms of quality. If you have never worked in sales, you cannot imagine how difficult it is to develop an accurate idea of how the demand for one individual article will change in the future. You might have a definite feeling for whether the demand will increase, remain constant or drop for a group of similar articles, but not for an individual article.

Additionally, sales staff usually must consider two sides: On the one hand, they have a certain demand expectation and on the other hand, there is a sales quota to fulfil which often has very little to do with said demand expectation. Which future sales volume should a sales representative use in their forecast for the next few months: The ones they expect, based on their experience, or the ones that are expected "from the top"?

Finally, all of us, including sales, tend to let our expectations regarding the future be heavily influenced by recent experience. Which is called "Recent Past Bias”.

We have no other choice than to reduce the influence of the human factor a bit when it comes to objectively forecasting the future sales volumes of articles.

This brings us to Best Practice Rule 3: The following applies to most industries and markets: Sales forecasts have to be produced in the form of technical (= statistical) forecasts. Information should only be requested from Sales where technical forecasts are insufficient.

Sales can only produce forecasts for sales articles, whether they're standard products or possibly spare parts. It cannot make meaningful statements about individual assembly groups or individual parts that it never sells as such. This is another reason why it is not always possible to use sales' forecasts as...

This brings us to Best Practice Rule 3:

The following applies to most industries and markets: Sales forecasts have to be produced in the form of technical (= statistical) forecasts. Information should only be requested from Sales where technical forecasts are insufficient.

Sales can only produce forecasts for sales articles, whether they're standard products or possibly spare parts. It cannot make meaningful statements about individual assembly groups or individual parts that it never sells as such. This is another reason why it is not always possible to use sales' forecasts as...

Basic principle 4: The finished product level

The finished product level is not always the right forecast level.

We often hear from companies that they cannot provide correct sales forecasts as the variety of their end products is too large and practically every end product is customized. Yet, this kind of company typically builds its end products out of standardized or partially standardized assembly groups and components, and just one part of the components is really customized.

The customized end product is therefore mostly the result of the variable arrangement of "Lego bricks" and not of individually manufactured components. Usually, the acceptable delivery period is at best sufficient for building a customized solution largely made of standard components. The "logistical decoupling point", i.e. the point in the value stream from which customer-order based production begins, is usually right before final assembly; sometimes even further up the value stream.

Above the decoupling point, produced parts have no customer relation. They are manufactured to stock. This is precisely the point in the value stream for which a sales forecast has to be generated.

And therefore Best Practice Rule 4 tells us:

The warehousing stage at which forecasting should be done is at the logistical decoupling point, this is after the last production stage without customer relation.

The three aforementioned basic principles should have highlighted that the technical forecast is of considerable importance in the sales forecasting process. However, in practice technical sales

forecasts are dealt with so cluelessly and sometimes ignorantly, it is not surprising that many companies do not trust statistical forecasts. Statistical forecasts are sometimes off the mark but on average they are much more reliable than human forecasters., Especially as latter can usually only handle a few selected articles whereas the rest of the articles have to be treated rather globally as a result of their sheer numbers.

Statistical forecasts are sometimes off the mark, but on average they are more reliable than human forecasters, provided that the statistical forecasts are compiled properly.

However, the prerequisite for this is that you draw up your statistical forecasts correctly. Which is why we want to deal with important basic principles and best practice rules for effective technical forecasts in the following. The best practice rules are partly organizational requirements for your forecasting process and some have to be implemented using software.

Basic principle 5: Every statistical forecast consists of three components

Every statistical forecast consists of three components and not just a forecast figure.

This situation is too familiar from the weather report: On Saturday, we were supposed to have a bright day with temperatures around the 70ies. On Saturday evening at 8 p.m. we were standing in drizzly weather at 64° with freezing guests cursing the damp charcoal. The obvious allegation is presumably that the weather forecast was wrong. If we hadn't relied on the weather forecast from the weather expert on television, but had taken a look at the weather service online, we might have possibly learned that there was a 45% probability of rain on Saturday evening and a temperature of 72° with a 25% probability of it fluctuating by +/-4° and a probability of 67% that the temperature fluctuation range was within +/-10°. Drizzle and 64° were therefore definitely covered by the forecast's fluctuation range.

With forecasts, regardless of whether they refer to the weather or the future demand of articles, we mainly have two problems: Firstly, we don't have a feeling for probabilities and secondly, we like to pick out what we want to hear from the information provided. That is why the TV weather report capitulated straight away and simply gave up on broadcasting the probability and temperature fluctuation range.

The analogue model makes it clear:
The higher the readiness for delivery should be, the higher the safety stock must be.

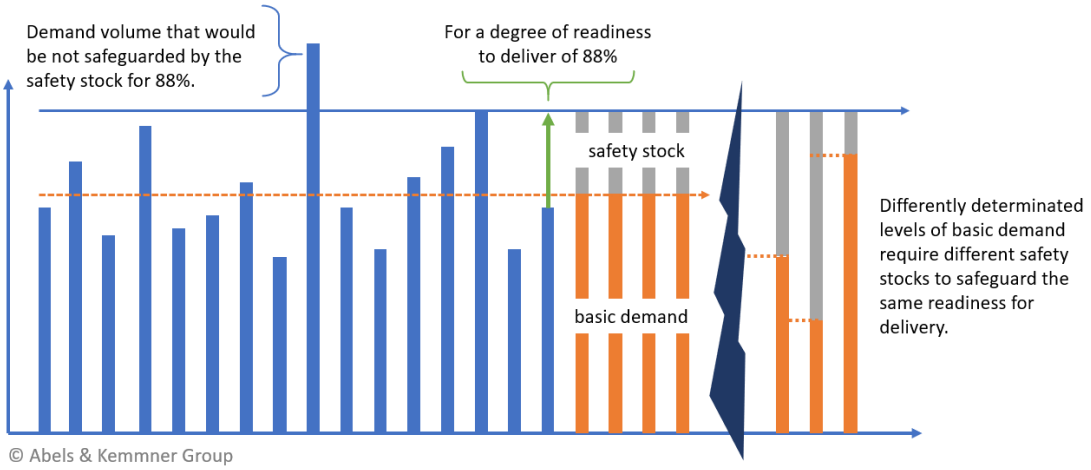

Figure 2: A qualitative analogue model of a forecast with the forecasting components of basic demand, degree of readiness to deliver and safety stock

It is important to understand that a forecast is not just made up of forecast figures, for example the expected monthly consumption over the next few months, but also has to include a statement about the uncertainty of this forecast. That is why we correctly talk about basic demands instead of forecast figures. The forecast's uncertainties then have to be protected by safety stocks. The safety stock's extent depends on the statistical certainty you want to be able to deliver. The higher the required readiness for delivery, the higher the required safety stock (cf. Fig. 2).

Best Practice Rule 5 is therefore:

A sales forecast is only complete if it contains details about the basic demand, safety stock and required readiness for delivery.

Basic principle 6: The demand for articles is almost never "normally distributed"

Perhaps you can still remember hearing about normal distribution or the Gaussian bell curve at school. First of all, normal distribution says that measured variables around an average value fluctuate symmetrically towards smaller and larger values. The fluctuation range is subject to certain universal laws, for example 68.27% of all values are around the mean with a fluctuation range of one "standard deviation".

You do not have to understand Gaussian distribution, but you do have to know that all classic forecasting methods, for example the sliding average or first or second order exponential smoothing for articles, rely on a normally distributed demand. Only then can these methods calculate statistically correct basic demands., Yet this only applies to approx. one in 20 articles in practice (cf. Fig. 3). As a result, approx. 95% of your demand forecasts are systematically wrong.

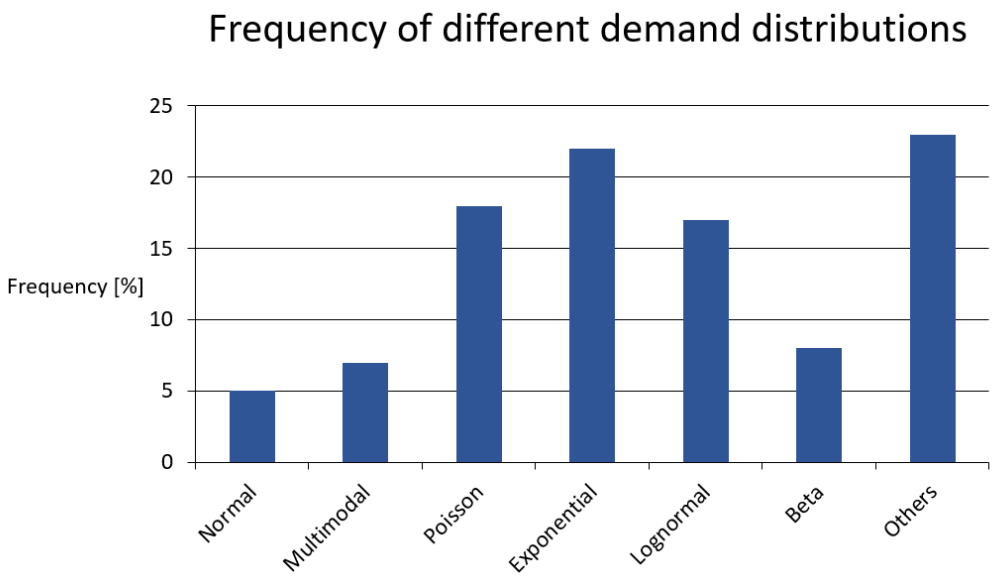

Figure 3: Most articles do not have "normally distributed" demand

Naturally, there are also forecasting formulas and safety stock methods for other types of demand distribution. However, they are not effective in practice: Demand distribution, i.e. the rule by which the demand for an article fluctuates, may change from month to month and no demand distribution can be identified at all for a good fifth of all articles.

There are three solutions to deal with this problem: The most common solution is to ignore the problem. This is the typical approach most ERP systems take.

Approx. 95% of your demand forecasts are systematically wrong due to a wrong statistical approach.

You can only obtain reliable forecasts if you check, with every forecast run for every article, which of the various basic demand and safety stock calculation methods would have produced the best results for you when forecasting a specific article in the past, to then apply this combination of basic demand and safety stock calculation methods to the forecasting of future sales volumes. This solution requires the simulation of basic demand and safety stock calculation and the resulting planning and scheduling depending on the actual demand over a specific period, ideally, for a previous year (cf. Fig. 4). Although this approach will not solve the problem of the forecast process being incorrect from a statistical perspective, it will at least help you to identify those basic demand and safety stock procedures that helped planning and scheduling to achieve the best results for this article over the past months

This approach is complex as a large number of simulations with different basic demand and safety stock calculation procedures have to be carried out for every forecast. It would bring an ERP system to its knees. Hence these calculations have to be carried out outside the host system in special sub-systems.

Another approach involves the application of so-called "distribution-free" methods. These are basic demand and safety stock calculation methods working without knowledge of the actual demand distribution, therefore working mathematically and statistically correctly again and leading to the required readiness for delivery. These distribution-free methods can only be found in very few sales forecast systems.

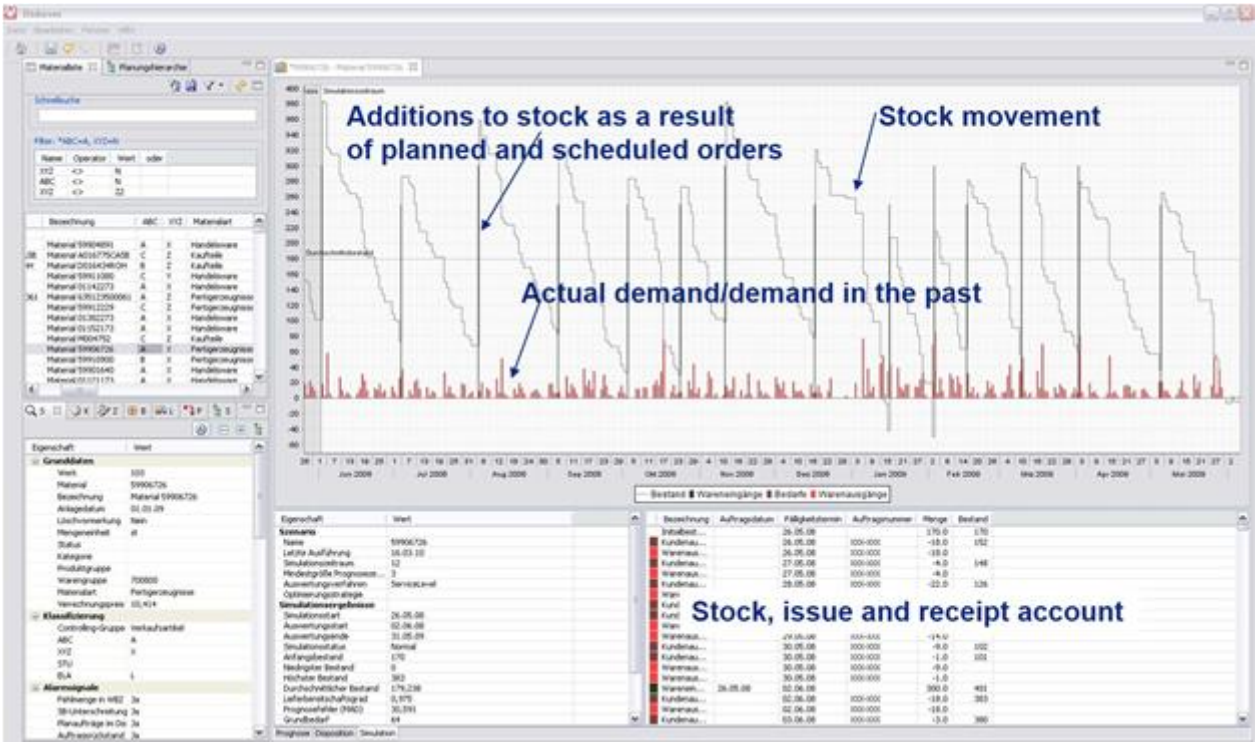

Figure 4: Example of the outcome of a forecast and planning/scheduling process simulation for a single article

Which brings us to Best Practice Rule 6:

To calculate reliable sales forecasts, you have to determine the most suitable basic demand and safety stock calculation methods for every article every month by simulation, to calculate the forecast for future demand. This process requires a large number of different basic demand and safety stock calculation methods. At the same time, you also need to use distribution-free methods. However, instead of working with distribution-free methods purely, you need to integrate the distribution-free methods into the simulation runs.

Basic principle 7: The quality of the sales

The quality of the sales forecast depends on the quality of the historic demand time series.

If you are producing a statistical sales forecast, then you try to base your statements about the future demand on a demand time series of the past. It is obvious that the more accurately the past time series maps the actual market demand in the past, the better the quality of the forecast. If you are only using warehouse issues as the past time series, it may be that there are gaps in the time series, when you were not in a position to deliver, and logically no warehouse issues could take place. If you continue to use this time series unchanged as the basis for your forecast, then you are assuming that the lack of warehouse issues was due to a lack of market demand. With this, you have already bent the truth.

You become a bit better at forecasting if you use the time series of agreed delivery dates and quantities; but these may also show delivery gaps that have nothing to do with lacking demand.

Using the time series of customers' requested quantities and dates will take you another step further, as it gets you very close to the customers' requirements and therefore market demand. If you were not able to deliver, informed your customers of this and they therefore didn't place an order, then the current market demand is not fully documented in the time series of customers' requested quantities and dates either. Thus, you should also capture “zero sales”, i.e. inquiries that did not proceed into a customer order, provided the order was not completed as a result of a lack of delivery capacity and not due to the price.

Best Practice Rule 7 therefore says:

For forecast at finished product level, if possible, use customers' requested dates and quantities as measure for the market demand history and record zero sales due to poor delivery situation. If you are forecasting at component or individual part level, derive the "corrected" demand from the past time series of finished products.

Even though you are now relying on the right demand history for your forecast, your work is still far from done. Do you have anything to do with promotions or project business? Then basic principle 8 must be observed.

Basic principle 8: Besides general demand

Besides general demand, past time series usually also include promotions and project business which are very difficult to forecast together.

There are generally three different kinds of demand. Firstly, in many industries there are "promotions" which are carried out by sales or have to be carried out as a result of customer requests. Meaning that either a certain product is offered for a short period only, or continuously sold products are temporarily ”on sale”. This typically happens in retail. In the first case, you can differentiate between products that are offered as a one-off and those that are offered repeatedly. In the first case, there is no past time series to rely upon, whereas in the second case a promotion-specific past time series from the last promotion(s) is available.

If continuously sold products are being used in promotions, the demand history consists of demand from the promotional phases and demand from the non-promotional phases. Generally, a company or sales, should be able to plan promotions in good time so the promotion's demand can be planned and scheduled on time; although there is admittedly the problem of defining the right quantities for the promotion. Looking back at promotions in the past might be helpful here, as mentioned before.

To create a forecast for the ongoing non-promotional demand, the promotions have to be taken out of the demand history if they stand out from the background noise with regards to their size. You can do this manually for a few articles but not for a large number of articles. Automatic mechanisms are called for here to separate the promotion's demands from the background noise.

Projects must be looked at separately from promotions. Projects exist where larger quantities of one or several articles are sold to individual customers and there are explicit negotiations with the customers about quantities, prices or delivery conditions. Typically, you are battling with competitors for projects. Projects are often decided at short notice after long forward planning and products often have to be delivered at short notice as well. Projects also have to be automatically filtered out of the article's demand history using suitable filtering mechanisms if they stand out substantially against the background noise. If this is the case, then sales is required to submit estimates, including all uncertainties, with regards to the forecast of projects.

Said promotions and projects are noticeable due to the "peaks", i.e. short-term demand peaks, and can in principle be identified and eliminated using so-called outliers clearing processes. The challenge lies firstly in a good process and secondly in setting the filter correctly. This can only be done reliably by simulation the consequences of different levels of selectivity for the forecasting and material planning behavior.

Best Practice Rule 8 is therefore:

Project business and promotions that stand out substantially from the background noise must be separated and forecast separately. The right selectivity to automatically filter out projects and promotions from the background noise must be identified regularly for each article using simulations.

Basic principle 9: The demand for the same article

The demand for the same article develops differently in different sales channels / sales regions.

It's a pretty banal fact that an article sold to several key accounts or in different regions will not be in demand equally in all sales channels. In California, for example, demand is very even but at a low level, while in Georgia sometimes larger quantities and sometimes smaller quantities are purchased very irregularly. If both states are supplied from one production site, you might think forecasting the total demand for this article would be adequate to keep the central warehouse stocked sufficiently. This might lead to a forecast for this article being requested from the California sales staff and the Georgia sales staff. By generation separate forecasts for the sales in both states, it is no longer necessary for California to submit forecasts for this article, as a technical forecast can offer sufficiently accurate and objective figures. If the national sales divisions have their own regional warehouses, separating the forecasts into different sales channels offers Logistics the advantage of it

being clear which replenishment quantities have to be delivered to the individual national sales divisions' local warehouses and which safety stocks have to be kept there.

Best Practice Rule 9 therefore demands:

Forecasts should be divided into sales channels / regions.

Do you supply your customers via several distribution levels, i.e. from a central warehouse to different regional or national sales divisions' warehouses? Then another principle has to be observed:

Basic principle 10: The real market demand

The real market demand in a supply chain is the end consumers' demand

The question of who your customer is arises whenever you supply your finished products to end consumers through several distribution levels. Is it the national subsidiary or the end consumer in a specific country? Evenly distributed end consumer demand may turn into irregular or intermittent demand at the central warehouse through a chain of replenishment decisions, as each warehousing level collects the demand and then replenishes intermittently. From a transport costs point of view, it can be, but is not necessarily, cheaper to proceed in this way. For sales forecasting, resp. replenishment it definitely means higher safety stocks and a distortion of the actual market demand. If you regard the national subsidiary as your customer, you have to forecast its demand, based on its possibly stochastic demand in the past to guarantee the requested readiness for delivery.

However, by seeing the national subsidiary and not the end consumer as your customer, you are deviating from optimizing your entire supply chain and practicing segment optimization, resulting in poorer forecasting and higher stocks in the supply chain.

In different countries, your company might have to cooperate with vendors like a hardware store or a sales agent since you do not have a subsidiary located there. So you will have to treat this intermediary like your end consumer, even though they are technically your distributors. In this case, segment optimization is your only choice.

Should this "intermediary", however, be for example a national sales company, it belongs to your corporation. You should then consider forecasting on the basis of end consumer demand (point of sale demand) to comprehensively optimize your supply chain. The historic demand time series of the point of sale demand of different regional and national warehouses can be consolidated at different distribution levels. You can produce much more than forecasts at regional and national warehouse level with this data: You can also consolidate the demand to central warehouse level and forecast the basic demand and safety stocks up the value stream. But please be aware that you are probably removing part of the "intermediary's" control over planning and scheduling with this.

Best Practice Rule 10:

To optimize the entire supply chain, you have to produce the required forecasts for different distribution levels, based on end consumer demand and not on the basis of your intermediaries' demand (requests). This applies even more so if the intermediary belongs to your own corporation.

Basic principle 11: Many new articles have predecessors and similar articles

It is part of the routine in practically every company to regularly launch new products on the market. And each time you are, once again, faced with the problem of forecasting the expected demand for this article.

If a successor article is merely a technical update of its predecessor, transferring the predecessor article's demand history to the successor article in a predecessor/successor relationship is usually sufficient.

As soon as a successor article has technical compromises or design changes compared to its predecessor, quantifying the future demand for the new article gets more complicated . Despite this, adopting the demand trend of a suitable predecessor article or representative article is yet a helpful approach.

A successor article does not always simply replace a predecessor article. Often a predecessor article is discontinued while the successor article is already being run in. In these cases, a pure predecessor/successor relationship is not enough to produce a sales forecast for both articles. In addition to the predecessor/successor relationship, an overlap period must be defined and the course of its substitution curve taken into consideration.

If a new article expands the current range of products, it will probably attract part of the already distributed articles' market demand. This is called "cannibalizing the demand for existing articles". An efficient sales forecast takes this into account as well as the fact that the cannibalization effect builds up over a certain period of time.

Predecessor/successor relationships, like substitution relationships, "disappear in time", as a sufficiently long demand history builds up for the "new" article after some time.

Best Practice Component 11:

Document predecessor/successor and substitution relationships to forecast the demand trend of new articles and the demand cannibalization of similar continuing articles.

Basic principle 12: Some forecast situations evade classic approaches

We are all familiar with the curious articles whose sporadic market demand affects us by chance. A typical example of this group are spare parts. It seems to be a universal rule: Customers always ask for a (spare) part that has not been required for months or years just after it has been scrapped.

Almost even more exotic are the end-of-life forecasts in the electronics industry, for example, to ensure enough supplies of important costly or irreplaceable components your suppliers have discontinued while your articles are still in production or warranty period for quite some time.

The same happens in fashion industry, which is regularly required to determine the demand volume for new collection pieces.

Sometimes it helps to develop a specific forecasting method in these cases, as we did for example for the end-of-life forecasting of components for Medion. With the help of a special "forecast exchange", hsx.com, a company in Los Angeles, provides forecasts for the box office returns of new movies.

If you don't manage to get a grip on forecasting, either sales are lost as a result of insufficient readiness for delivery or excess stock and scrapping generate high costs. The only thing that helps reduce these costs is to intelligently align the corporate business model with the limits of the logistics business model.

Thus the operative problem of a poor sales forecast turns into a strategic challenge for the company.

Best Practice Rule 12:

Special situations may require special forecasting methods. If this should fail or the consequences be too expensive, the only choice is to intelligently align the corporate business model with the limits of the logistics business model.

The deeper you delve into the details of a sales forecast, the more best practice rules can be established. Not every best practice rule can be applied in every case; especially not if you are trying to solve the problem with in-house means.

Please do not be discouraged by the seeming complexity of these rules. If you are using an effective and well configured sales forecast system, it will obey most of the afore listed rules automatically. Which will take you much closer to a best practice solution and less complex sales forecast processes.

Finally, it should be noted: Anyone shooting with a home-made gun and bent barrel should not be surprised when the gun misses its target. One can be surprised though when the hobby gunsmiths conclude from this that professionally produced guns basically miss too.