You do not need a crystal ball to make better sales and operations plans

LOP.ai unveils

uncertainty in your Sales and Operations Planning







1. Introduction

We are kicking an open door by saying that today's businesses are subject to a high level of uncertainty. The difficult thing about uncertainty is that it is hard to grasp, often overlooked, and challenging to communicate, and convincing people about the possible impact of uncertainty can be problematic.

Uncertainties can be classified according to the famous statement made by Donald Rumsfeld, US Secretary of State for Defence, at a Defence Department briefing in 2002:



There are known knowns; there are things we know we know.

We also know there are known unknowns; that is to say we know there are some things we do not know.

But there are also unknown unknowns – the ones we don't know we don't know

Hence, we can classify uncertainties according to the quadrant below.

There is basically a big difference between occurrences we are aware we cannot predict, but which very often have a particular likelihood associated with them, and possibilities that are not even on our radar. On the extreme end of the last category we find "black swan events": events that no one anticipated and have no corresponding historical statistics, but have a fundamental impact on business or society. Examples are the 9/11 terrorist attacks, the rise of the internet, and maybe very recently the COVID-19 crisis (although the inventor of the term, Nassim Nicholas Taleb, does not recognize the latter as a true black swan event as it has long been predicted).



Meta- knowledge: do you realize that it can happen?	High	Known Knowns We are aware of it and we understand it	Known Unknowns We are aware of it but we do not understand its behaviours
	Low	Unknown Knowns We are not consciously aware, but we understand it	Unknown Unknowns Thing we are not even aware of not knowing
		High	Low

Knowledge level: do you understand the source of variability, and do you know about it upfront?

Figure 1: Classification of uncertainties according to knowledge level (do we know it can happen and what the impact can be?) and meta-knowledge level (are we aware of the risk element?)



In his book Operations Rules, David Simchi-Levi gives an overview of the spectrum of uncertainties that businesses are facing. On top of the list we find events that have a high level of surprise linked to them. At the bottom we find the elements of uncertainty that we are very much aware of. Typically, data has been gathered on these types of events over the course of history, so they can very often be expressed in a mathematical way through statistical distributions.



Figure 2: Classification of uncertainties in the known-unknown spectrum according to David Simchi-Levi

Despite the fact that these known unknowns can be described relatively well through their statistical properties, we see that this knowledge is barely used in business. The typical example is forecast accuracy: even though every supply chain manager knows that forecasts are by definition subject to error, there are still many companies that do not take into account the uncertainty that these forecasts entail when developing scenarios in their sales and operations planning (S&OP). As soon as forecasts are fed into the supply planning processes, they are being considered as certainties.

The unknown unknowns are of a different nature. These are either completely impossible to describe through statistical distributions, or they are subject to fat-tailed skewed distributions (distributions with very high volume, very low likelihood observations). It does not make sense to model these uncertainties mathematically, as their effects go so far in the likelihood/impact spectrum, that they belong much more in the business continuity decision process than in the operational planning process. If you want to learn more about these types of uncertainties, *The Black Swan* by Nassim Nicholas Taleb is an excellent thought-provoking book on the topic.

One of the reasons businesses do not take uncertainty into account explicitly, is that every well-understood element of uncertainty gives rise to a wide range of possible scenarios. And there are a lot of uncertainties like that.

If you do not have the concepts, technology or processes to deal with this, it gets very hard or even impossible to make these different possibilities transparent. Consequently, many companies deal with this uncertainty in a qualitative way, with decision-making often driven by gut feeling and politics.

The most obvious process for beginning to incorporate this knowledge is the S&OP process. This process lays the foundation for decisions that will ultimately have a significant effect on the performance of the company: decisions about capacity, investments, promotions, etc. Incorporating the notion of uncertainty, at least for those elements of uncertainty that can be described through mathematical formulas, would already be a major leap forward compared with the overly deterministic approach that is often followed today.

By translating the uncertainty into a set of metrics that highlight the risk level associated with a particular set of decisions, we can guide management in the right direction. We quantify the risk and make it very transparent. A major advantage is that these metrics have a very clear meaning. This ensures that their values can effectively be used as guidance towards better decision making.



2. S&OP as decision-making process

S&OP is a fundamentally simple process, yet it is one of the most performance-enhancing processes that exists in an organisation. Its basic setup is simple – but that does not mean that it is easy to implement properly. Many companies implement an S&OP process, but struggle to reap the full benefit of it. One of the reasons why implementation remains a challenge is that the element of uncertainty is rarely taken into consideration in a consistent way.

S&OP is highly impactful as it is primarily designed as a decision-taking process involving the most senior levels in the organisation (president, CEO, managing director, COO, CFO, etc.). Senior management often prefers clear, concise and easy-to-digest information to base decisions on. This makes it very challenging to include the uncertainty element in the decision process, because it is difficult to present it in a coherent way.

S&OP is a decision-making process that balances demand and supply at the aggregate level. This happens typically through a Rough Cut Capacity Planning (RCCP) model that evaluates, at a product group level, whether demand can be fulfilled with the available capacity at machine group level, taking into account the inventory position and evolution. This happens in most cases in a monthly cycle, with some possible mid-month adjustment whenever major changes occur. This RCCP is typically made with deterministic assumptions and we already know that this condition is very flawed for a number of reasons:



Demand will not happen as forecast. Despite all the efforts that have been made over the last decades to improve forecasts, there are still substantial deviations between actual sales and forecasts in a lot of businesses. We also know that forecast accuracy decreases the further we look into the future. Bearing in mind that S&OP is a process that typically looks 12 to 18 months into the future, the level of uncertainty can be very substantial towards the further end of the horizon.



Machines will not produce at a constant rate. Production rates are influenced by machine downtimes, and variations in the quality of incoming materials, operator skills, and product mixes. Although intuitively one expects production rates to be relatively stable, we have frequently observed substantial levels of variation.



Suppliers do not deliver as promised. Lead times of components and raw materials can also deviate from what was promised. Missing materials may lead to unused capacity and unfulfilled downstream requirements: an hour of capacity lost can never be recuperated. Today, many of these uncertainties are typically tackled through safety stocks. But very often these safety stocks are only based on demand uncertainty or on simple rules-of-thumb. The complex interaction between all the elements of uncertainty can be pretty difficult, if not impossible, to incorporate into a safety stock formula. In any case, safety stock is only one of the levers to mitigate uncertainty.



- The purpose of S&OP is to explore all the options, such as:
- Overtime
- Extra shifts
- Extra machines
- Outsourcing part of the production
- Alternative suppliers
- Stock building





3. Risk-based RCCP

World-class companies have begun to understand that they need some uncertainty-related metrics to help them to define the best possible course of action. In demanding markets, like a consumer goods market or a pharmaceutical market, it is imperative that you can follow demand, whatever is happening in your upstream or downstream supply chain. If you are not able to fulfil that demand, you are likely to incur penalties, or even worse, lose the market.

As we have seen, uncertainty in RCCP comes from different angles and creates very complex interactions. If you have an RCCP that runs for over 18 months, with 10 product families and 5 machine groups, that leads to 180 demand elements and 90 capacity figures that are subject to uncertainty. Adding some critical component lead times to that makes this a puzzle that cannot be solved through purely analytical methods, let alone by gut feeling.

LOP.ai addresses this challenge using simulations. The system generates different versions of the reality. By evaluating enough of these versions, we can arrive at a statistic for the expected outcome of a set of decisions. This is visualized in the curve shown in Figure 3, where all versions of a scenario have been ranked according to their unfulfilled demand: best performance / least service losses at the bottom, worst performance at the top. The y-axis lists what percentage of the replications in our simulation have an unfulfilled demand below what can be found on the x-axis. The curve splits the surface of the graph into two parts. As both the horizontal and vertical axis go from 0 to 1, the surface of the entire graph is equal to 1. The area under the curve (also called the AUC) is what primarily interests us, and we typically prefer it to be as high as possible. The area above curve (the blue part, equal to 1-AUC) indicates the service losses that we expect to be facing.



Figure 3: Chart where replications of the simulation are ranked on the y-axis according to their service losses on the x-axis

If this curve runs flat at 100%, whatever happens, you will always be able to fulfil demand. This may be a trigger to downsize your production machine, as it probably indicates that you are performing at overcapacity. In that case the surface under the curve corresponds with the entire surface of the graph, or we say that the AUC is 100%. Although the curve itself reveals more information about the uncertainty profile than the metric, it is a very useful number to summarize the expected performance of a scenario. In the context of RCCP risk assessment, this AUC metric is renamed "probability to execute" or "PTE". But that term doesn't fully describe its meaning: it is the expected value of the percentage of demand that will be fulfilled on time.

Deviations from that 100% come in two dimensions, expressed through the two axes of the graph: the intersection with the y-axis tells you what the **likelihood** is that you will not be facing any shortage over the considered horizon. The position on the x-axis then gives you the corresponding **impact of the shortage**. In the example shown in Figure 3 – where the PTE is a very poor 92.6% – you can see that there is an 80% chance that the shortage will be less than 10%, and a 90% chance that the shortage will be less than 20%. This example is obviously a very high-risk scenario, one that few companies would be willing to accept.

As a user you obviously want a more fragmented view of this risk profile. A first indication is the risk evolution over time: is the risk higher in the early part of the horizon (when there is often still little one can do on a S&OP level) or is it more apparent towards the end? The heatmap on the right, that highlights the risk level in our example for the next 18 months from left to right, indicates that the risk is particularly high from month 3 to month 6.

Stockout Prediction For each time bucket



This will help you to identify the right actions for making this an acceptable scenario. This brings us to the most challenging management decision: how much risk is your company willing to accept? Is it a different story if the risk is mainly located at the end of the horizon? How will you express the risk? Is it the PTE-value that will drive your risk strategy, or will the intersection with the y-axis be the dominant KPI? These are all fundamental questions that have no universal answer. Just as risk is a concept that is very difficult to grasp, the best strategy to deal with it is hard to define. But that does not make it any less essential to the success of the supply chain. And with this approach, you at least have all the information at hand and can make an informed decision. Appropriate reflection on how risks impact your business is a mandatory step towards assuring that you deal with the uncovered risk profile properly.



4. The technology: LOP.ai

For global companies that typically have an extended supply chain network, lots of different products, or a complex technical architecture, a risk-based RCCP approach is impossible to execute without the proper technology. LOP.ai is a platform that supports companies to move towards end-to-end (E2E) supply chain orchestrations based on a supply chain digital twin. Risk-based RCCP is one of the solutions that is embedded in the platform. Below, we will outline the components that enable it.



Digital Supply Chain Twin (DSCT)

One of the fundamentals for a useful risk-based RCCP approach is a solid understanding of the reality of the landscape. This means understanding how real demand behaves versus predicted demand, how operations behaves versus planning, the variances that occur during execution, and the internal or external factors that can influence certain events.

Having a DSCT is therefore essential in a risk-based RCCP approach. As defined by Gartner: "A DSCT is a digital representation of the physical supply chain that can be used to create plans and make decisions. It replaces the traditional supply chain model and sits at the heart of any planning solution with a model (twin) that is built from granular data and is near real-time in its reflection of the real world."

The last sentence in the definition is crucial. Firstly, it is built from granular data, meaning transactional and event driven data. Therefore, LOP.ai is able to capture a lot of actual information, generate insights and run predictions. Secondly, it's a reflection of the real world, meaning that this digital twin is not based on static master data, assumptions, old-school design exercises, etc. It is representing reality as it is happening.

In a risk-based RCCP the digital supply chain plays a crucial role in two ways. One, the digital twin contains a huge set of data which helps to capture and understand the variances that typically occur in your network in an automated and data-driven way. And two, the digital twin enables simulations that can demonstrate how resilient the proposed supply chain will be to these variations and uncertainties.

Cloud and Processing Power

When moving from a static approach towards a fully dynamic or risk-based RCCP approach, the technical requirements for the platform change drastically. By relying on a hyperscale cloud platform and the latest technologies, LOP.ai can take advantage of the scalability.

In the world of simulation this scalability is a real must. LOP.ai runs thousands of scenario versions in the background, each of them assuming a different reality, which ultimately result in a scenario whose risk profile is presented to the user. From a calculation effort, this is obviously significantly different from comparing only a dozen scenarios, as happens in a traditional S&OP setting.

It goes without saying that an LOP.ai user doesn't wait 1000 times longer for all that information. Thanks to our cutting-edge infrastructure, the simulation is fully complete in only a couple of minutes!



Multiple Interacting Digital Skills

Having speed, power and scalability is of course a pre-requisite from a technology perspective, but that is not enough to create a service of this calibre. To function at this level requires a number of specific big data skills.

The virtual twin is generated automatically using data mining and crawling skills. They create links between all types of transactional data and events to build your data-driven model from the bottom up. Big data handling skills are required to feed and connect the (internal and external) supply chain events to the model.

Algorithmic processing and AI/ML are also important components in risk-based RCCP. They can convert the unknown uncertainty into known variability, by making parameter predictions and identifying which factors are correlated to the target variable. Finally, optimization and simulation techniques enable true mimicry of the different versions of reality, and allow the system to set up a resilient supply chain that buffers against uncertainties, while simultaneously proposing actions to improve the probability of execution.

User Experience

Lastly, we have to tackle the topic of user experience and adoption.

As you've probably understood by now, the concept of risk-based RCCP is built on very strong data, technology and mathematics. The biggest pitfall of the risk-based RCCP approach is that the system is so complex that the end-user or manager cannot follow the process.

Of course, efficient change management can play an important role in mitigating this. We will discuss this further in the next section. However, the platform itself also has a big role to play in making the shift as smooth as possible. During the simulation process for example, LOP.ai will not bother the end-user with the 1000s of different versions of the plan. Instead it will only recommend those versions that are most relevant to the user and inform them about key metrics such as PTE, potential stock-outs and the most significant risk areas. From there, a user can drill down to the next level of detail and discover what occurs where and when. Also, although these are all automatically generated scenarios, the user still has the possibility to create manual what-if scenarios to stress test the impact of certain decisions and predict the risks that go with them.

5. The importance of managing change

The fact that risk is quantified and made more tangible also implies that the organization needs to think very differently about the way they set up S&OP and RCCP. If you are coming from a tradition of using gut feelings or rules-of-thumb to determine how much excess capacity is required, the shift to this new approach doesn't happen overnight. In the beginning, there will always be a tendency to verify whether the old way of working and associated metrics gave a result that is more or less acceptable. This is most likely not the case. Experiment by setting up scenarios the way you would traditionally have done, and see where your approach generates too high (or too low) a risk level. Building a deeper understanding of why your old thinking did not give the optimal results, is a first step towards a new mindset.

But it's not only the champions behind the new way of thinking who need to be convinced. Typically, the risk strategy, even if it has never been made explicit, has permeated the entire organization and acts as a driving force in the decision-making process. It will require intensive knowledge-transfer and dissemination throughout the organization to generate the necessary shift in thinking about risk. It will pay off eventually, even though this too may be hard to prove.

The challenge with risk is that, if the strategy works, you do not really observe it. When you planned for a scenario with no disruptions, and there are no disruptions, there are few arguments to say that you did not follow an optimal strategy. But it is very possible that you would otherwise have set your buffers too high, and paid the associated costs. On the other hand, if you have a disruption, it does not mean that the decision you did take was not the best one at that moment. It requires a leap of faith to trust the figures and the calculations behind them.



Your RCCP exercise should be centred on the PTE-value; this is the pivotal point in the search for a set of capacity decisions to be taken during the S&OP process. This does not mean that you should accept the outcome of the search for an acceptable scenario straight off. It is possible that this scenario does not fit your utilisation or cost objectives. If this is the case, instead of moving to a more risky scenario, dig into the sources of uncertainty and reflect on whether there are options to reduce these uncertainties:





Can you improve the quality of the forecast? Can you improve the demand planning process, remove biases, incorporate other information, involve other people, structure the process differently?



Can you make production more reliable? Is there an opportunity to launch Six Sigma projects to reduce variability and hence increase reliability?



Can you make your suppliers more reliable? Or can an inventory increase for some critical materials mitigate the risk you are facing?

Although there is hard evidence that a quantified risk approach leads to better decision-making, it is not necessarily a simple move. With the right guidance, training and communication, however, this new approach can be successfully implemented.



6. Conclusions

Uncertainty is a fact of business. You can either deny it or you can embrace it. If you embrace it, it is better to do it properly. Rather than relying on gut feeling, make sure that uncertainty is quantified properly. Once you know how uncertainty can impact your business, you have all information at hand to make the appropriate decisions to deal with the risk.

Not having these insights can lead to two things:

- You underestimate the uncertainty. This can lead you into deep trouble later on and is definitely the scenario you want to avoid. Unfortunately, cost concerns often lead to this situation. Although management senses that trouble is at hand, they do not want to take the appropriate decisions today, and would rather avoid the associated costs. Giving too much priority to cost-saving can ultimately threaten your profitability.
- 2. You play on the safe side. In this case you will avoid running into service problems, but at a cost which is probably excessive compared to the risk you are dealing with. This too, will ultimately impact your profitability.

A platform like LOP.ai can help make risk tangible for you, so that you can mitigate it in the most cost-effective way. It offers you the possible courses of action, and makes you aware of the service and cost impact associated with them. It will make your decision-making more streamlined, more efficient and more reliable. The fact that the scenarios and their evaluation are fact based, will lead to easier, more efficient decision-making and fewer discussions based on intuition and sentiment.



Authors

References/ bibliography



Koen Cobbaert, Director Supply Chain Planning, PwC Belgium

Valérie Vandenbroucke, VP, LOP.ai



https://medium.com/datadriveninvestor/kno wn-knowns-unknown-knowns-and-unknown -unknowns-b35013fb350d

Nassim Nicholas Taleb: *The Black Swan* (2007), Random House

David Simchi-Levi: *Operations Rules* (2010), The MIT Press





