Webcast Transcript

Using Range Forecasting and Scenario Planning to Manage Risk

Tuesday, July 26, 2011

Joe Fleischer: Welcome to our one-hour Webcast, titled "Using Range Forecasting and Scenario Planning to Manage Risk," brought to you by CFO Publishing and by the Webcast's sponsor, IBM. I'm Joe Fleischer and I will be your moderator today.

I'd like to take a moment to tell you about our distinguished guest speakers. Joining us today are Steve Morlidge who is the founder of the consultancy Satori Partners, as well as the former controller with Unilever Foods, and a former chairman of Beyond Budgeting Round Table Europe, as well as Tony Levy who is with the product marketing area of IBM Business Analytics Software.

By way of background, Steve Morlidge has 30 years of practical experience in designing and running performance management systems at Unilever, including three years as the lead of a global change project. And having served as the chairman of the European Beyond Budgeting Round Table, Steve Morlidge now works as an academic and independent consultant for a range of major companies. He is also a, as you may be aware, a prominent public speaker drawing on his years of experience as the leading edge of performance management thought and practice. And later on during this Webcast we will also hear the perspective of Tony Levy of IBM.

But at this stage it is now my pleasure to introduce our very first speaker, once again, Steve Morlidge. Please give a warm welcome to him. Thank you.

Steve Morlidge: Thank you, Joe, and hello everybody and welcome to this short presentation. What I'm going to be talking to you about today is what I believe is one of the most important but least understood elements of forecasting, which is the forecasting of risk. And I call this presentation how to deal with always being wrong. Before I talk about this, I'd like to set the context a little by talking in more general terms about how to master forecasting, which is a topic I think many businesses struggle with. My general contention is that forecasting practice and dealing with risk in particular is in a mess and the reason is, I believe, is because people don't understand the principles behind forecasting. I also think that once you understand the principles, the practice of forecasting becomes easier or at least easier.

Before I go on, I should introduce myself. As Joe said, I've spent the best part of 30 years as a practicing manager with a major corporation and I first got into the subject of forecasting when I was a leader of global change project looking for an easy way to lead my business into the world of beyond budgeting, which is a bit scary for many businesses. To my mind helping to make business forecasts more reliable was a really good way to introduce people to the need for and the benefits of a more adaptive management approach. And one of the main problems with the

traditional approach to performance management that I was aiming to replace, which is traditional budgeting and variations on that theme, is little scopes we're admitting to risk.

One of the problems with the traditional approach is that there is an underlying assumption that the world is capable of being predicted and that the top and bottom of good management lies in securing compliance to that prediction, be that through multiple analysis of variances or the dispensing of awards and punishment. For me, the word risk is simply a word that describes our inability to predict the future. But as you will discover, I hope, not being able to predict the future precisely doesn't mean that you can't manage it, providing you've got the right approach and the right tools that is.

Risk and our failure to manage it properly has had a lot of press recently. As you can see from this slide, even the British queen has got in the act. Here she was commenting on the failure of the world's economists to predict the credit crunch. Of course you might say such things are always obvious in retrospect with the benefits of hindsight. But perhaps it might be worth thinking for a moment why that should be and what can you do to make things more obvious in the moment when things are actually just about to happen and when there is still time to do something about it. I'll come back to this question at the end and give you some tips on how to speed up your hindsight later on in the presentation. But first let's just focus on the challenges of forecasting a little closer to home in companies.

As you can see from this slide, improving -- which is from The Hackett Group, improving forecasting processes sit at the top of CFO's priorities, along with improvements to the budgeting processes. Moreover, this desire to improve forecasting is taking place at a time when it's becoming increasingly more difficult to forecast because post the recent banking crisis the world has become more volatile, demand, input prices, exchange rate, market prices are all much less predictable than they were even just a few years ago. But it's not just about volatility, the fact is that the pace of change generally is increasing.

For example, it took 50 years from the introduction of the telephone until it captured 50 million users, and the T.V. took 13 years, and the Web achieved that feat in just four years. As another example, which is very topical at the moment, many people believe that the company that's actually still got the highest share of the mobile phone market, Nokia, is already doomed. And if you think about it that is quite incredible that despite selling more phones than any other manufacturer it's already too late to address their strategic weakness. And why, because they didn't perceive the risk from smartphones.

But where could we look for solutions. In my experience there are two common misconceptions about practice and forecasting, which apply equally to forecasting risk. The first is that it's easy and it's all just common sense. If you believe this when it goes wrong it, therefore, must either because people are incompetent or they're dishonest. So the logic goes you need to give the job to somebody who knows what they're doing and wait until this lies and that solves the problem. That I believe is wrong. The other misconception is that forecasting is very complex and technical and that's probably what you would believe if you consulted the Web. For instance, if you typed in the word forecasting into Amazon, after about 360,000 books you get about 30,000

hits. And the vast majority of these contain detailed explanations of statistical extrapolation techniques, Monte Carlo simulations, and things of that sort.

Now, mathematical techniques are useful to a degree, but if they were the silver bullet we should never be in the mess we're in given the amount of brainpower that's been applied to the subject over the last 50 years. I think the truth lies somewhere in between and I think the fact is that in real life we're actually very effective forecasters and we understand and deal with risk rather well and we've had to be in order to survive as a species. What we need to do is to learn how to apply our natural ability to new situations of the sort that we find in management.

As an example, think about how we teach our children to learn the craft of forecasting when we teach them how to cross the road. Estimating the speed of oncoming cars and extrapolating it is an important start, but that's not the end. We don't teach our children to cross the road by sending them to a class on differential calculus. Instead we teach them to make sound assumptions about what might happen based on an understanding of the environment and other road views, such as is the road wet, are drivers likely to be slowing down or speeding up, have they seen us, assumptions that are continuously tested and validated by observation. And we also encourage our children to consider the margin of error attached to their estimates, what might happen if those assumptions are wrong and what action they should take if they are wrong. And in this way we build their capacity to understand and manage risk.

Our problem is that when we forecast in our day-to-day life we tend to do it below the level of a consciousness so that we simply don't encourage -- acknowledge what it is we're doing. And so we fail to transfer our innate understanding of the process into the field of management when we're at work. I say that once we've exposed and understood these principles we can teach ourselves to design and operate effective and practical forecast processes, as well as discover what's wrong with the ones we've already got.

So our contention is that most of the problems with forecasting in business are a result of little thinking and the behaviors associated with this not a shortage of tools per se. and we've identified six principles to help you get your thinking straight. And once you've got your thinking straight, we believe that you can apply the expertise, most of which you already have or are readily available, in an appropriate way. The six principles also provide the framework for the book that we've written on the subject. What I want to do is to set the scene quickly by running through all six principles before we do a deep dive into risk, and we're going to start with number one, which is about purpose.

The purpose of forecasting, I believe, is to make decisions but what kind of decisions, that's the key question. To answer this we need to understand that there at least three sorts of forecasts which support three sorts of decisions. Once we understand the link between forecasting and decision making we can then start to specify our forecast process. Now, think for a moment about the future and our relationship with it. In the very short-term the future is likely to be like the past. Why might that be? Well, firstly there isn't much time for things to change, but there also isn't much time for us and anybody else in the same position as us to respond to anticipated changes to make things different, which we might want to do, for instance, if the forecast outcome was undesirable. So in the very short-term we can do more than attempt to predict the

future based on the recent past, i.e. trends, and react to it, for example, by producing enough product to meet demand. Mathematical techniques -- forecasting techniques are very useful here and, as we've seen, there are lots of them out there. In this space we handle risk by building stock and buffers like stock.

At the other extreme in the very long-term, the future is likely to be very different from the past, partly because of the elapse of time and partly because we and others might choose to make a difference. We can create and execute strategies. As a result, there is very little point in producing a single point forecast in the long-term. We could do no more than take a bet on which of the possible futures will win out, which isn't very productive. Instead we need to create multiple scenarios and use those to help shape the future based upon full -- an understanding of the full range of possibilities and the factors at play. Here, risk is effectively represented in the alternative scenarios. For example, in the publishing business two scenarios might be the iPad takes over the world and another paper books make a comeback.

In the medium term, however, we're somewhat but not completely restricted in what we can do. We don't have complete freedom because the future is to the degree constrained by the momentum of our existing business, our existing infrastructure, product, customers, and so on. A good way to think about what we're doing in this space is to compare it to navigating in a book, but we're constantly making decisions to steer the business in response to changing currency wins. And the risk that our forecast might be wrong is something we actively have to recognize and deal with. We call this kind of forecasting business forecasting, and unfortunately there are few tools available to managers to help them to do this well, which is where we first came in.

So we forecast -- we use forecast for at least three different purposes. In the short-term we predict the future in order to react better. In the long-term we aim to create the future we want using techniques like scenario planning to explore possibilities. In the medium term we shape the future by anticipating what might happen wherein to navigate the course.

Let's continue on the theme of navigating and steering and I think you'll find this will help us to be a little bit more specific about how to do this well. Before we start any journey it's natural to plan, and in business we often call this plan a budget. But in turbulent times we often find ourselves in a different place than where we intended to be, and at this point the original plan is of little use to us. In these circumstances our next step is to ask our navigator to produce a forecast to establish where we think we will end up if we carry on executing our current plans. Very often the result of this is that we'll end up -- we won't end up where we want to be, which is our target, unless we change our plans.

One important point that this example draws out is the reason why we forecast, which is to provide information to support good decision making. On the water we might move the teller or hoist a sail, in business we change our plans. Essentially we do one of six things, we start, stop, bring forward, delay or change an existing activity, or create a new one, an activity being anything from a price change, a promotion, a new product, a decision to hire, and so on. So the purpose of business forecasting is to make decisions to steer the business, and in steering the business we're constantly looking at the difference between where we want to go and the

direction where we're currently heading and making course corrections where there's a gap between the two.

The other important insight this example provides us with is perhaps the single biggest contributor to successful forecasting, in my view. It is the recognition that there's a fundamental and necessary difference between a target, which is what we would like to happen, and a forecast, which is what we think will happen given current assumptions about the world and our plans. In my experience what many business call their forecast is in fact a mixture of the target and the forecast, in reality just an updated budget. And as a result it fails to properly guide performance or to provide a reliable estimate in future outcomes.

One consequence of separating forecasting from target setting is there will more often then not be a gap, and to close this gap we make decisions, i.e. we make decisions -- we make changes to what we intend to do, in other words, to our plans. And the greater the rate of change in the environment the quicker the gaps will open and the faster we'll have to adjust our plans. This contrasts with traditional budgeting. So the forecast is the source of information upon which we base decisions. From this understanding we can now draw the definition for a good forecast, and based on this definition we can specify effective forecasting practices.

In our view, and in this order, there are five key attributes of a good forecast. It needs to be timely, actionable, reliable, aligned and cost effective. Unfortunately, the first letter of each of these doesn't make a good mnemonic so to help you remember we suggest you reverse the letters so we have carat, which is a measure of something precious. Each of these letters describe one of our other five principles of effective forecasting. Let's just briefly explore each in turn.

One of the most important decisions you make in designing a forecast process is to decide how far ahead you need to look in order to make sure you're able to make timely decisions. The decision about how far to look into the future depends on the length of time between taking a decision and the impact of that decision coming to fruition, i.e. your decision making lead times. So, for example, if you were a captain of a supertanker you would need to look further ahead than if you were a driver of a speed boat. Second, the forecast needs to be actionable. This means that it should provide us with enough data on an incremental impact of our planned activities for us to make effective decisions about whether and how to change our portfolio of plans. Information that's irrelevant to decision making merely needs to be reliable, detail isn't required. In practice this means we have to make informed decisions about what type of forecast model we are to use, based on an understanding of their relative strengths and weaknesses.

Third, forecasts need to be reliable, which means they need to be unbiased and free from unacceptable levels in variation. The next principle is that of alignment. Since a forecast is not and cannot be a prediction, it's dangerous to align a single point forecast. At the other end of the scale, having a plethora of competing views to the future will lead to dysfunctional chaotic decision making. So we say we need a single forecast in which to base our decisions, but at the same time recognize there will always be error attached to the forecast, in other words risk. And we need to develop a good appreciation of the nature of this risk so our decisions are robust in all our likely scenarios and we have contingency plans, which are simply alternative sets of decisions, in place to deal with the less likely ones. We'll develop this theme in the next slide.

Finally, the forecast needs to be cost effective, which we believe follows naturally from the observance of the other principles.

Now, for risk. There is only one thing you can ever be sure of when you're forecasting, and that is you're going to be wrong. Any single point forecast is only one version of many possible futures, at best a good estimate and an average outcome. And as such, it suffers from what Sam Savage, who is an expert on risk management and measurement, calls the flaw of averages. So, as in this case, a river may be on three -- three foot deep on average, but you can still drown in it. In business, as in many walks of life, it's often not the average outcome that's important but the range of possible outcomes, the degree of probability attached to them, and the potential consequences if they were to materialize.

Before we do a deep dive into the detail specifically about how to measure risk, let's just take a step back and get a few basic principles sorted out. As we've discussed, any forecast represent only one of the large range of possible outcomes. But it's not practical or sensible for a business to create work with a multiple set of forecasts. Would you like to sail on a boat where every member of the crew had their own forecast? Of course not, but why? The reason is that in these circumstances an outcome across this in businesses might recorded eight separate forecasts in one business. The reason is that every member of the crew or every function would be making different or misaligned decisions based on their own view of the future, and as a result we'd have no confidence in the ship as a whole is under any kind of control. As I say, you might think it's a silly example, but in many businesses this is the norm.

For example, the sales function comes up with a forecast that marketing ignores because they believe that the impact to their new marketing campaign is being underestimated. Finance produce their own forecast because they don't believe either sales or marketing and they want to err on the safe side, and the guys in the supply chain produce their own forecast because they think everybody else is playing games. Sound familiar? Instead we recommend that you fix on one central forecast, ideally the one that's proved to be most reliable, and make an [estimate] to the scale of risk attached to it, the degree of variation that any natural process exhibits. As you can see from this example, the degree of risk attached to our central forecast is such that it wouldn't lead us to change our decision. But if the level of risk was twice as big, as shown by the red lines there, we probably would not try to steer a course between the two rocks because we couldn't guarantee our safety.

But this doesn't tell us the whole picture. The fact is that scholars recognize there are two possible types of risks attached to any kind of forecast. The first is the type we show here. The likelihood or more accurate rate the inevitability that any forecast will be wrong because of random variation or a noise. Now, this kind of risk can be dealt with relatively easily providing we have some historic data. For example, insurance companies can make very accurate estimates of the risk of a particular type of individual falling ill or having a car accident because they have lots of data about what's happened in the past and the patterns of behavior are relatively stable at the time. Ironically, therefore, risk when taken as a whole at least is predictable and as a result measures can be taken to protect oneself against it.

But there is another kind of forecast error. Frank Knight, the economist who first made this distinction, called the second type of error uncertainty. Uncertainty is the consequence not of random variation or random mean, but a consequence of patterns of behavior changing dramatically. So in our example as we sail along tides or patterns of currents might cause us to veer off in a completely deferent direction. Unlike what Knight called risk, uncertainty can't be predicted and, therefore, presents a much bigger management challenge, which is on the upside as well as on the downside. A particular challenge is the fact that even if we had a very good idea that an uncertain event, a discontinuity, is very likely to happen we can never be sure when it will happen and what the impact will be.

For example, take a look at this graph which plots the sterling and euro exchange rate, which is a very important variable for many U.K. businesses. The euro was introduced on the 1st of January 2002 and for most of the next five years it traded at around 1.45 to the pound. In 2007, the pound suddenly weakened against the euro and over the next two years there was something like three of the major discontinuities in the exchange rate, none of which were predicted or predictable, at least in detail. So at the current time sterling trades at something like about 1.15 to the euro and has been for the last two years, but it's almost certain this won't continue. For example, U.K. interest rates are at their lowest level since the Bank of England was set up in 1694, and many economies in the euro zone are in turmoil. Sovereign debt is a very high probability. But when will sterling appreciate the euro and how big will the adjust be? We simply do not, we cannot know. So unlike risk we can't manage uncertainty arithmetically. Instead what we have to do is establish what might happen, keep the situation under close watch and respond quickly as soon as something becomes clear. And as soon as it becomes clear there could be a major shift in behavior, a discontinuity that is occurring.

Here is a slide that provides an illustration of what we mean. On the left you can see a forecast for the path of Hurricane Ernesto, which was made on the 26th of August in 2006, as well as the most likely forecast you can see an estimate was made of the risk attached to the forecast shown in white, and based on this the island of Jamaica was put on alert, as you can see by the blue outline around it. On the right you can see what actually happened. Instead of hitting Jamaica the hurricane, shown in yellow, suddenly veered off to the north, hit Cuba and then Florida. Of course meteorologists knew that this could happen but they also knew that despite their very sophisticated models based on solid scientific theories many hundreds of years old, they couldn't predict that it would happen. All they could do was keep a close watch on the path of the hurricane and react quickly as soon as it became clear that it's changed course putting into action the hurricane's contingency plans already in place in Cuba and Florida.

The big challenge for meteorologists, business people and anybody else dealing with uncertainty, is to determine when a discontinuity has occurred. To be precise, as well as failing to spot a discontinuity, what's scientists call a Type 2 error or false negative, we need to be careful to avoid calling a discontinuity when all we're really seeing is random variation. This is called a Type 1 error or a false negative. Both kind of errors have a cost attached to them, as anybody who has suffered misdiagnosis with cancer can tell you.

Now we've clarified some of the key concepts surrounding risk, let's address the much more practical issue of how we measure it because without some addressing of the potential impact we

can't sensible manage it. Broadly speaking there are three ways in which we can measure uncertainty. The first way is to analyze past performance. I should have said, by the way, three ways in which you can measure risk or uncertainty. The first way is to analyze past performance. Now this works very well when we're looking to measure risk over the short-term because we'll have a lot of data but also a much more confidence that what has just happened is a good indicator of what is just about to happen.

Over the longer term none of these conditions apply, we have less data and less confidence that what data we do have is relevant. Nevertheless, it's a good place to start. At least it gives us a size -- a sense of the likely size of risk, which is particularly important since human beings we have a systematic tendency to underestimate risk. A good tip, I think, is to start by measuring the size of the average forecast error ignoring the sign and use the simple rule of thumb, which is that 70% of the risk lies within 1.25 of this number, i.e. the average forecast error, and 90% of the risk within the multiple of 2. The multiple for 95% is 2.5. This is all based on this being a rough approximation to a standard deviation. By definition, this approach assumes that the future is like the past, which as we've said is often not the case.

Another approach to quantification we need to use is based on creating multiple alternative futures. This approach is based on a recognition that, as with our example of crossing the road here, any forecast is based on assumption and that by varying these assumptions we can create multiple alternative futures. These alternative futures are often called scenarios and they can be very powerful, particularly in exploring possible discontinuities. In the banking world such an approach is often called stress test and is used very extensively right now where the balance sheets of financial institutions are subject to stress test. The weakness of this approach is that there is a limit to the number of scenarios you can create and analyze, and by definition you're reliant on the judgment of the individual or the individual's creating them. You also find it very difficult to ascribe probabilities to the scenario so it's possible, for example, that you could create a scenario that's very implausible but fail to create one that's very likely.

The third approach comes from -- overcomes some of the shortcomings of these first two approaches. This involves combining using Monte Carlo simulation tools, the probability distributions attached to the key drivers of future performance in order to create a summarized distribution for the variables that we're looking to forecast. A strength of this approach is that the results presented in the form of a set of probabilities to guide action. Also, the inputs to the model driver distributions can be derived either from history where that's more appropriate or from judgment where history isn't available or is inappropriate for some reason. A down side of this approach, however, it is technically quite complex and so carries a higher risk of error. In addition, it's apparent sophistication is a little superficial in that it's vulnerable to failure in human judgment and weaknesses in the historical record as the bankers who relied upon this type of model and belatedly discovered during the recent financial crisis.

A very common problem we encounter when people are attempting to quantify risk however they do it is a failure to understand some of the issues associated with probability math, which I just want to share with you know. Take this example as a set of stepladders. Before you step on them you will probably give them a shake and based upon what you feel you're going to make an estimate of risk, the risk of them collapsing. Let's say we make an estimate of 10%. What's the

risk of two sets of similar stepladders collapsing? Well, the answer is 1% since we multiplied the two sets of probabilities, 10% by 10%. So, therefore, it's a rare event for both sets of stepladders to independently fail at the same time.

But what would happen if there's a plank of wood spanning them? The answer is very different. Now the risk of them both failing is 20% because in these circumstances we add the probability, 20 times more likely. Why should this be? The reason is because specifically if the wood was securely fastened to both ladders either stepladder collapsing would bring the other down so the fate of the two stepladders is no longer independent. Technically speaking we say that the state of the two ladders is positively correlated, which is another way of saying if one thing happens then there's an increased chance of another thing happening. In business people very often tend to consider risks in isolation and, therefore, drastically underestimate the likelihood of extreme events. Again, witness what happened in the recent financial crisis where problems in one part of the system increased the risk of problems in another because we have a cascade affect.

In practice, of course, the issues involved are a lot more complicated than this. Drivers may be weakly correlated, indeed they're very often negatively correlated. So another -- an example of this might be where a retailer can pass on a price increase that they have received to their consumer. And so because price and cost move in opposite directions there might be no negative impact on the bottom line at all. In addition, probability can only be combined accurately using Monte Carlo simulation techniques. Despite this complexity, however, I think it's really important that anybody involved in analyzing risks understands the issues involved with the risk of probability and makes an attempt to incorporate them in the analysis, even if only judgmentally. For those of you wanting to explore this issue more deeply, I'd recommend Sam Savage's book, which is very easy to read and is very accessible.

So what might we expect to see from a well executed analysis of risk? Perhaps, as you can see on the left, we might start off by flexing our assumptions around key drivers, such as volume of cost based on our understanding of the market, to get a feel of the degree of normal variation we might expect. This kind of exercise is typically called a sensitivity analysis. Perhaps we might test this against what we have observed in the past just to check that is what we might expect say a 90% level of confidence. On the right you can see the cumulative impact of this variation as shown by the dotted lines set against a forecast profit profile shown by the blue bars. We might now want to stress test our forecast by introducing a couple of scenarios. In the example shown on the right we can see that scenario A, which assumes a speedy withdraw of a product following a failed product launch, exposes this to a significant bigger downside, even after taking mitigating action. Perhaps because we want to know quite how sensitive -- because we know how sensitive our future plans are for this product launch we might want to understand the risk associated with it better by performing a Monte Carlo type analysis on that project. On the other side of the equation, even if we succeed in landing the major account we've been chasing for the last year as assumed in scenario B, the upside isn't great.

So now we've measured risk, what do we do with this information? Basically there are two approaches to risk management. The first approach is a passive management of risk. This involves creating something that will buffer you against the risk. Say you might build a cash reserve for a rainy day or build redundancy into your business model so that a failure in one part

of the system will be covered by another, which is the reason why as human beings we've got to kidneys and why aircrafts have multiple fail safe systems. You could also build contractual buffers, so you might take out an insurance contract or hedge your risk. All these approaches work well but they're not without difficulty. For a start they all cost money and they're also inflexible because you have to identify the right risk against which to buffer. Also, they only protect the downsides. It's often forgotten that risk and uncertainty can be upside and logically we should be as anxious to exploit the upsides as we are to mitigate the downsides.

The second approach, the active management of risk, avoids some of these downsides. Essentially active risk management involves a speedy reaction to events as they occur rather than passive protection irrespective of their occurrence. Contingency planning, which we've talked about already, is an example of this. In our previous example the mitigating actions we took in response to failed product launch referred to might be an example. It's also possible to stretch your decision so they have create options to change courses, events and fold. So perhaps the downside risk associated with our failed product launch might have been mitigated if we had add source production to start up with rather than building an expensive new plant. There are many, many examples of this kind of thing but we don't have time to explore them further.

What I want to do before I finish is to make good on my earlier promise to explain why things are always obvious after the event. In the process I'll demonstrate to you, hopefully, what I believe is the most important benefit associated with incorporating risk into your forecast process, a benefit that has absolutely nothing to do with quantification or planning in any way. Take a look at this next picture. We've all been in this position in a crowded place perhaps in a strange city. What do you see? Just a sea of strange faces. But then, as so often happens, somebody comes up to us, pokes us in the ribs and say, "Hey, what are you doing here? Why didn't you see me?" It's your best friend and of course you had seen him but you hadn't recognized him. But why hadn't you recognized him? The answer to this question is that you didn't recognize him because you hadn't expected to see him, and because you hadn't expected to see him you didn't see him. This tells us something very profound and important about the way you see the world, which I think is highly relevant to forecasting a business.

Psychologists tells us that the signals that reach our brains that reach our eyes only make up about 30% of what we think we see. The other 70% is filled in by our brain based upon what we expect to see. Now this makes a lot of sense from an efficiency point of view, it cuts down a lot of processing in the brain but it does expose us to the risk of misperception as magicians and conmen have known for millennia. It all explains -- it also explains why things are always obvious after the event. Before the event we might not have recognized something was possible, therefore, we literally do not see the signs. After the event, of course, it's obvious.

So do you see your best friend now? My guess is that his face jumps out at you, even though you've only known him for about a minute. What this tells us is that one of the most importantly reasons why we should explicitly incorporate risk into our forecast processes is that it helps to prime our perception. And by making us conscious of what might happen, it helps us spot what is happening quickly and to avoid the false negatives and false positives and take swift action to mitigate the risk or exploit the opportunity.

Well, that's all the time we have for this subject right now. I hope you found this talk useful and insightful. Now I'll hand you back to Tony and thank you and good luck.

Tony Levy: Thank you, Steve. You've stated the challenges and opportunities for improved forecasting well. My name is Tony Levy and I represent IBM's business analytic software solution and in this portion of the session we will focus on the leverage of modern technology to enable and sustain some of the principles that Steve outlined.

As we reflect on these principles, let's first consider and revisit the question what is holding finance back. One barrier is that today's core forecasting processes were designed as part of an overall resource allocation process, the budgeting process for a world long ago when volatility and uncertainty were comparatively lower. Sixty years ago when you could sell everything you manufactured why forecast at all. Indeed push-based batch-oriented resource allocation processes characterize today's performance management systems. These work well in those days when variability was low, but in today's economy where unexpected events are now part of a normal course of business we need more dynamic decision processes.

In addition, these processes were designed for a world where the pace of business was slower. I can remember working in the computer electronics industry in the early '90s we made 10 critical decisions per quarter and we wanted 100% of them to be correct. Today in that exact same business you make 10 critical decisions a week and you win if only 60% are right. The speed of business is 10 times faster today in that industry than 15 years ago and this is growing exponentially in all industries and all geographies of the world. Instead in this new normal where variability and speed of business remain high, we need pull-based demand-driven resource allocation processes that are continuous and responsive to the marketplace. The principles Steve outlined are intended to transform outdated forecasting approaches and help us manage in the face of greater variability and pace of business.

And these principles need to be enabled and sustained with modern technology. Today we use spreadsheets, email and manual techniques to enable our forecasting process. This is a barrier. While spreadsheets is an individual productivity tool, it was not designed for dynamic enterprise collaboration. There is no workflow to help automate the rollup of forecasts or help with the review and approval process. There is no metadata management, no central data management and limited security, no audit capabilities. We refer to this as the problem of the spreadsheet in the middle. The opportunity is to replace the spreadsheet in the middle with a dedicated planning and forecasting system that provides robust multidimensional modeling, workflow, processes management, integrated reporting and what-if analysis capabilities. One important impact is that by doing this you will be able to free up the capacity of your financial analysts to spend more time on value added analysis to support the business. For instance, as Steve outlined, to assess normal variation around key metrics as a measure of risk and design compelling scenarios to account for uncertainty. It should be no surprise then that the forecasting processes are under extreme pressure.

For the majority of companies today the process looks something like this. It starts with a creation of next period forecast. Spreadsheet-based models and templates need to be updated to reflect any changes. This straightforward copy and paste of metadata has not changed, but if

product or sales organization hierarchies and dimensions have changed then these model updates can be error prone and time consuming. With next period models finalized, we update last period actuals from ledgers and operational data sources, such as HR and manufacturing. And here IT staff is needed to help connect to the latest data sources and pull the data into flat files and then uploaded into our spreadsheet. This can take time and is error prone.

Our internal controls organization is concerned with data governance due to uncontrolled access to information flowing through a steady stream of disconnected spreadsheets. The spreadsheet templates are distributed through shared network drives and by email. Your financial analysts then send out emails and phone calls to remind the responsibility center managers to enter their forecast on time. There is limited visibility into the true status of the forecast since the spreadsheets are sitting on someone's email box -- inbox or hard drive as work in progress.

The forecast then are returned to finance staggered in various states of review. You struggle to aggregate them into corporate P&L balance sheet and cash flow projections in two to three weeks for the average organization or longer for the majority of us. You then scramble to validate the data and create relative analysis and management reports. And if you're lucky you barely have time for a value added analysis to help the business owners gain deeper insights into the business, for instance, into profitability of products, services, customers and channel. This is in fact the analysis that business owners are hoping you can help gather for them.

A symptom of this process is that analyst time is inefficiently applied to data collection and validation rather than value added analysis, as illustrated by the results from an APQC survey taken in January of 2009. On average only 23% of our analyst time is spent on added value analysis during critical planning and forecasting periods. Now if you were to lay this along a one-week continuum for illustrative purposes only, these results imply approximately two days spent on collecting and validating data, one and a quarter days on spreadsheet maintenance, one day on developing reports, and half a day on performing analysis. The point here is to highlight the obvious. When we hire smart MBAs into our FP&A teams didn't we promise them in the job description and interview process that they would be conducting exciting value added work to help drive the business. So tell me, exactly where in this weekly timeline does that fall. It's no surprise that one of CFO's major internal concerns is retaining talent.

Manual processes correlate with higher cost of finance as a percent of revenue. From IBM benchmark data in collaboration with the APQC, we see that the median cost of finance for organizations is 1.8% for those companies that predominantly use spreadsheet manual processes for their planning environment versus 1.46% for integrated environments using dedicated purpose built software to automate the planning process. It should be no surprise that companies using spreadsheet manual processes require more staff effort to operate their planning processes than companies with integrated dedicated software. Furthermore, the impact of spreadsheet manual processes is not limited to cost it impacts timeliness and reliability. Again from the same APQC survey mentioned earlier, we see that companies that predominantly use spreadsheets are less likely to provide real-time access to information, less likely to realize planning accuracy within 5% of original assumption, and less capable of completing budgets within two iterations or less.

So is there a better way, an alternative to manual spreadsheet processes. The answer starts by moving beyond spreadsheets. Replace a spreadsheet in the middle of a dedicated planning and forecasting system that provides robust multidimensional modeling capability, automated enterprise workflow in process management, integrated reporting and analysis capabilities, and use spreadsheets primarily as a user interface into a robust planning and forecasting system. This will help you free up valuable capacity for your financial analysts to spend more time on assessing normal variation and key metrics and executing more thoughtful scenario plans to support the business, which is the reason they joined your FP&A group in the first place.

Dedicated planning and forecasting reporting and analysis software can help you automate data extractions from ledger systems and operational data source as the use of data in duration tools that your financial analysts can use. Support from IT is limited to access provisioning of the data sources. Indeed your financial analyst can model the data and automate the extractions on their own. Workflow, automated alerts and process reporting can help you automate the distribution, collection and aggregation of the forecast from all the right participants in your enterprise. Security and access controls helps automate and coordinate submittals and approvals. You can run process reports to provide real-time visibility into the status of forecasts, either they are not started, incomplete, in progress, or successfully submitted. Your financial analyst can centrally manage the models and control metadata changes quickly and easily. Driver-based, multidimensional models can be created, modified and maintained easily by your analysts without the need to wait days for IT to run complex scripts. The result is that time is spent on collecting and validating data is cut by two-thirds.

Reports can be run quickly and distributed securely to all the right people in any format they require. Users have different skills, different needs, different access points and may require everything from a simple report in PDF to a dynamic chart rendered on their smartphone. We need to deliver across all the different user requirements. Our analysts and business users now have more time to conduct scenario analysis based on more reliable and more relevant data. The net result is that financial analysts can now spend 65% of their time in discussions with the business owners and decision makers focusing on important decisions rather than spending 65% of their time collecting and validating data.

We're proud of the way our software solutions enable complete planning, reporting and analysis capabilities to help you automate and transform your entire planning, budgeting and forecasting process. IBM Cognos 10 solutions for finance are able to source data from all of your ledger and transaction applications. They're not tied to any one transaction system so you can continue to leverage your corporate standard data sources, as well as Legacy data sources. The outcome is more effective management meetings and decision making.

Hopefully, this short presentation gives you an idea of how modern technology can help you create more timely, actionable, reliable, aligned and cost effective forecasts, as Steve mentioned earlier. So don't wait, raise the business case for change now, improve your process guided by the six principles that Steve outlined earlier and update your technology. Automate and structure your planning and forecasting processes so that you can reduce cycle time by up to 70%. Improve staff's leverage so that financial analysts spend 65% of their time on creating value added analysis versus collecting and validating data. And then sustain driver-based rolling

forecasts that involve all the right participants as frequently as needed while at the same time conducting range forecasting and scenario planning to manage risk and uncertainty. There's resulting improvements in accountability and forecast reliability, as well as the effectiveness of your resource allocation process, will help you drive profit and revenue growth and better return on investment capital.

In summary, the ability to forecast timely, reliable results should be one of the top concerns for your CFO in this new normal. With renewed planning and forecasting process finance leaders can help their business anticipate and shape business outcomes more effectively. IBM is uniquely capable of helping you automate and structure and sustain best practices in planning and forecasting. Thank you for your time. With that, I will turn the microphone back to you, Joe.

Joe Fleischer: Well, thank you very much, Tony, and thank you very much, Steve. We have a number of great questions from our audience and we will now give you the opportunity to, not only pose your questions, but we will take the opportunity to address them. Here's a question for Steve. The attendee asks who within an organization should ultimately have responsibility for calculating determining risk?

Steve Morlidge: I think for me the answer is that everybody in the organization I think needs to be involved. Again, if I can go back to my example of a ship, a sailing ship, I think we'd all agree that it's the more people that are on the lookout for things that might to wrong the better. And indeed risk is not purely an issue for the CFO in an organization where he's worried about having to deal with investors, it's a practical issue for local business units.

So I think risk involves everybody and I think it's particularly important that even the consumers of risk information have a good understanding of what's involved so that they have a sense of how robust the risk analysis is, how much they can be relied upon, and so on and so forth rather than just having blind trust. For example, in the recent banking crisis one quote came from a head of a merchant bank, I can't remember which one, who renown's the fact that what had happened was a sort of perfect storm because there were seven sigma outcomes, which is usually sort of 1 in 10 million year type event, which really showed that they didn't actually understand a critical assumption upon which their risk model was based, which was only based on five years worth of data.

So I think the simple answer is it's something that ought to be a concern for everybody and certainly there should be a good sense -- everybody involved in decision making should have a good sense of what assumptions were made and an idea about what assumptions would break the business. That's a general statement. At a practical level, obviously, there will need to be kind of well specified processes in place in order to assess risk, to measure risk in a formal sense. And I think the only thing I'd say there it's important that that takes place at all levels in business and not just at the top. Risk is not simply an issue for the CFO. I think really that's all I've got to say on that subject, Joe.

Joe Fleischer: Well, thank you very much, Steve. Tony, we have a few questions for you. You talked about the limitations an organization imposes on itself when it only employs spreadsheets,

but there's a question from an attendee. The tools you described, can they accommodate spreadsheets as at least an input as a source of information for planning?

Tony Levy: Sure, good question. As an input mechanism all of our business analytic software solutions permit the use of spreadsheets as a browser and an input interface into a controlled and auditable workflow. So in short, yes, spreadsheets are an important input and output mechanism for planning and business analytics generally and we support them with our solutions.

Joe Fleischer: Thank you, Tony. And a few other questions, so while we're on the subject go accommodating various sources, an attendee asks can these same tools work with third party systems, such as say, in this example, ledger systems from SAP.

Tony Levy: Great question and a common question. European financial ledgers are a vital source of information and getting reliable access to this and all appropriate data, whether in a single instance of SAP or varied instances of SAP or from a wealth of other relevant sources and warehouses is absolutely essential. Specifically on SAP we use access methods that are prescribed and certified by SAP that are fast, scalable and appropriate to reach these sources and other sources relevant for performance management.

Joe Fleischer: Thank you, Tony, and just a few more questions. Following up on the discussion of planning, can the tools we've discussed allow an organization, an attendee asks, to link multiyear strategic plans with annual, say, operating plans and forecasts. On other words, can the tool accommodate the linkage between longer term strategic plans and annual operating plans and forecasts, for example?

Tony Levy: Absolutely. Other software vendors might sell you the idea that you require different technologies for each of the different elements of your enterprise planning domain process, a separate strategic finance application, a separate workforce or capital planning application each carrying a license and an implementation maintenance price tag. We believe that these are part of the same coordinated process and thus require robust technology that can both logically connect these enterprise planning domains, but also accelerate the flow of information between them. For instance, updates to operational annual plan should not require a lift and shift to a strategic financial model. So we strive to enable one environment, robust, dimensionally rich, shared data rules and lower cost with less effort.

Joe Fleischer: Thank you very much, Tony. And, finally, an attendee asks a definitional question. FTEs, does that refer to full time equivalent within a finance team, for example, in the context of your presentation?

Tony Levy: That's correct, full time equivalent.

Joe Fleischer: Thank you very much, Tony. I think we maybe have time -- we have a lot of good questions, but I think we have time for one more question for Steve. What -- an attendee asks when we talk about range forecasting, the title of the Webcast, what are your recommendations on addressing risk assessments that may be skewed to begin with.

Steve Morlidge: Well, I think the first thing is to recognize that it is -- skewed risk is a likelihood rather than assuming that all risk attached with the forecast is symmetrically distributed. In fact, it's more likely to be the norm that are skewed. So I think the first thing is to recognize it. The second thing is to make sure that you assimilate that skewed forecast properly in the sense that you combine it with other estimates of risk of other drivers in an appropriate way. So using either some of the techniques I talked about earlier, like Monte Carlo simulation, if you want to go down that route. And for one I believe that in 5, 10 years time that will be much more prevalent than it currently now is or doing a judgment, which is probably much more likely in the short-term. So the first thing is recognize that it's likely. Second thing is make sure that, as I described earlier when I talked about the arithmetic of probabilities, you combine that risk with other risks, whether they be skewed or un-skewed, in an appropriate way.

Joe Fleischer: On behalf of our distinguished guest speakers, once again, Steve Morlidge and Tony Levy, we very much appreciate your joining us for our Webcast, "Using Range Forecasting and Scenario Planning to Manage Risk," brought to you by CFO Publishing and by today's sponsor, IBM. We thank you for your time and we hope you enjoy the rest of your day.