

Improved forecasting

How to improve forecasting in the face of limited knowledge

By Matthew Leitch, 26 May 2011.

Think of forecasting in its widest sense. People make forecasts at work for a variety of reasons, including these:

- To make decisions, where the effect of each option under consideration is forecast and compared. This could be part of investment planning or 'risk management'.
- To look ahead for possible problems in future, again as part of business planning or as 'risk management'.
- To make a plan by predicting the future and planning to be ready for it.
- To keep people informed of what to expect.
- As a disguised exercise in (re)negotiating targets.

Forecasts can be useful or useless. They can also be easy or hard to create. Two ways to create forecasts stand out as being both hard work and, frequently, useless.

Budget holder forecasts: One common way to create forecasts is to ask lots of people to provide forecasts for the work they are responsible for then add the forecasts up. This process usually involves a certain amount of negotiating, because the forecasts become a form of promise or target, and further work is needed to search for and resolve obvious inconsistencies. Overall, this is timing consuming and people prefer not to do it frequently. It tends to be useless because the element of commitments leads to gaming and dishonest forecasting.

Sophisticated best estimates: Another labour intensive approach to forecasting is to build an elaborate model, representing activities in as much realistic detail as possible, and use it to project one possible future. This is typically what comes from systems dynamics models and some other types of simulation. Although fewer people are involved the individuals doing this kind of modelling are usually highly skilled and spend a long time doing it. Unfortunately, their predictions may lack credibility if their methods are not understood and their assumptions not accepted. More fundamentally, these forecasts only describe one possible future, and are almost always wrong.

If you are using either of these methods, or something similar, you should be able to produce more realistic, more informative, more useful forecasts, with less effort.

Ideas for improvement

There's a lot of scope for improving forecasting so that uncertainty (limited knowledge) is better handled. What changes will help in your case is something you need to find out by considering what is distinctive for you and trying ideas out.

On the next few pages some of the ideas that are leading contenders are briefly described. If any seem suitable for your situation then you can find out more using the suggested references at the end of this article.

The basic approach should usually be to create an automated, interactive forecasting system (e.g. an electronic spreadsheet) that allows you to see the effects of alternative courses of actions, and associated uncertainty.

Easier methods of forecasting

Changes to your forecasting methods that make forecasts easier to do will also make them better at dealing with uncertainty. This is because uncertainty leads to forecasts being out of date, so ideally you will revise them often. Nobody wants to revise forecasts often if they take a lot of work to produce or cannot easily take into account changing conditions. Here are some ideas for cutting the work involved.

Cut detail: Analysing numbers into narrower categories might lead to more accurate and reliable forecasts but often it does not and it usually involves more work. Summarise to a higher level and check if accuracy is any worse. If not stop going into detail.

Extrapolate statistically: If you know what happened in previous periods of time you can often use this to extrapolate statistically. Relatively stable activities can be predicted at the click of a button with no danger of manipulation by budget holders anxious to manage expectations. Consequently, these mindless, effortless statistical forecasts can be a lot more accurate than more complicated predictions by people who are supposed to really understand the business. This method can be tested before use if you have enough past data because you can test to see how well a statistical rule would have performed in the past.

Analyse only changes: Statistical extrapolation only tells you what to expect if you carry on acting as you have in the past. Obviously, if you think you can do better and decide to act differently (e.g. hire someone extra, double sales effort, change prices in a way you haven't in the past) then the forecast should be adjusted to reflect that. This is easy and a spreadsheet can be set up to tabulate changes and their estimated impact on the forecast. (It can also remove unrepresentative events from the past history using the same method.) The spreadsheet can then combine these adjustments with the statistical extrapolation to give the final forecast. Meetings to discuss the forecasts and business plans can focus on just those changes, knowing that the massive detail of business-as-usual is already reflected in the statistical extrapolation.

Constructive simplicity: This is an idea from Chris Chapman and Stephen Ward that helps you control the amount of effort that goes into developing a forecasting model. The idea is to start very simply and add refinements iteratively, but there are some crucial details that make this effective. Start by making initial forecasts using a *very* simple but fully quantified model that also captures your initial uncertainty (using probabilities in some way). You can then analyse the contribution of each uncertain variable to the forecast and see which are the most important. Those are the areas of your model that deserve to be developed a bit further in the next version, with more thought and perhaps complexity, and more data gathering and validation.

More realistic forecasting

The usual human bias in forecasting is to think that we are better at it than in fact we are. We also imagine we have more control of events than we really have, and when we are surprised by outcomes we are able to explain them away so that we still think we are good at forecasting. To make more realistic forecasts we need to open our minds.

Distributions: The single most important change is to stop trying to make a best estimate prediction and instead make a forecast that identifies all possible outcomes and says how likely each one is. This immediately makes uncertainty explicit and provides more information. It allows you to avoid two common errors in calculations arising from working with only one possible future (the Flaw of Averages and Failure to Value Options). There are many ways to represent distributions of possible outcomes. It's important to understand that the probabilities here are not real world frequencies but expressions of belief, just as you would use for betting on sports events, for example.

Propagating uncertainty: Most quantitative forecasts involve a number of estimated input numbers, some calculations, and a resulting output forecast. Propagating uncertainty is the technical name for working out what the probability distribution of your output forecasts should logically be, given your uncertainty about the correct values for the inputs (also expressed as probability distributions). Mathematically this would be complicated if weren't for a technique called Monte Carlo simulation which is very easy provided you have a computer capable of doing the forecast calculations thousands of times in a few seconds – which any low budget laptop can today. You don't even need special software, but a low cost tool such as XLSim does make it easier.

Statistical spreads: An alternative to propagating uncertainty is to use past experience as a guide to the distribution to use for forecasts. If you know how far out your past best estimate forecasts have been you can use the distribution of those past errors as a guide to how accurate a new forecast is likely to be.

Programmed decisions: One reason that best estimate forecasts fail is that they cannot show the value of decisions that can be taken in future once more is known (sometimes called real options). In a typical spreadsheet forecast with columns representing time periods it is quite easy to include formulae using the if() function that test how things have been going so far and make decisions on actions such as changing investment levels. These can make quite a big difference to outcomes.

Consider the relatively certain context: Although a lot of things about the future are very hard to predict there are some features of the future that are fairly certain. For example, demographic changes and fundamental physical limits can be useful for bounding forecasts. For example, although the financial side of pensions and care for the elderly is hard to predict, demographic trends make it clear that in future there will be more elderly people to be looked after and fewer people able to do the looking after.

Forced consideration of ranges: When thinking about what might happen in future, either with or without numbers, it is important to think a bit more widely than seems necessary. This is because we have a stubborn tendency to think too narrowly and cannot feel when we are doing it. Go through important drivers of results and imagine what might happen if these were unusual favourable, or unfavourable. How would you respond? What would be the knock on implications? This technique is usually part of scenario planning methods.

Forced consideration with memory joggers: Another way to force yourself to think more widely is to work from a list of things to think about. Make the list a structured one that ensures you rigorously consider all possibilities. In particular, when we imagine a plan of action we tend to think of what could happen by starting from the plan and working outwards. Make sure you also try working inwards from major events that could happen in the environment.

Rich predictions: Instead of just trying to predict, say, overall costs, or overall time required for something, make and consider more implications of your plan and its evolving environment. If you sold a lot of product in a week then that would mean a certain amount of revenue, cost of sales, and profit, but that's just the headlines. What would it mean for storage space? For stocks of packing materials? For overtime? For accidents? For reputation? Is there a limit to how much you can handle in a week? Examining these predictions might point out flaws in your predictions (e.g. where you haven't realised that there *is* a limit to how many sales can be handled in one week). Detailed aspects of the forecast also help more people with their planning decisions, such as when to order more packing materials and de-clutter the stock room. (This is in conflict with the desire to keep forecasting high level, so focus on forecasting details that can inform identified decisions.) Forecasting more implications (e.g. safety, reputation) gives a fuller picture for outcomes that are hard to translate into money.

Causal links: Underlying several of the above ideas is the need to look at the causal connections between things that might happen. For example, if a project falls behind schedule what does that mean for the effort needed to complete the project? Often, if the project starts off harder than expected it continues to be harder than expected. Also, once things start to go wrong the pressure often leads to mistakes and knock on problems from the mistakes. Some people get less motivated and cooperative once things start to go wrong. Failure to think about causal links is one reason our forecasts tend to be too narrow. It's not difficult to draw some arrows on a flipchart and start to build a better sense of how events may be connected.

Progression: Forecasting is difficult and initial attempts tend to be poorly defined and structured as well as based largely on reasonable but untested guesses. View forecasting as an ongoing challenge and, over the weeks, months, and years, keep on trying to learn more about how things really work and improve the structure and definition of your forecasting models. Detailed discoveries, such as quantifying the strength of a causal link, can lead to important strategic insights. As Constructive Simplicity shows (see above) each forecast can provide information about what uncertainties are most important to results and the realism of forecasts. There will be actions you can take to achieve useful results for stakeholders *and* find out more about how the world works that will improve your forecasts and strategies. This is true even if your models are not fully quantified.

Comparing actuals with forecasts: Over time, forecasting ability can of course be improved by learning from actual results. However, since a probabilistic forecast is being used you will not be able to estimate a prediction error from just one result. You need lots so that you can see how your predicted distributions compare with reality.

Cross validation: This is a statistical model fitting check where you divide past data into two parts. One part is used to build and fit the model. The other part is then used to test how well the model predicts.

De-biasing experts: Forecasting sometimes involves asking 'experts' to provide input estimates. Several ways to do this that try to reduce our usual biases have been devised. One is the Standard/SRI Assessment protocol, which involves several

stages designed to check for conflicts of interest, educate about possible cognitive biases, and encourage use of data.

Interactive elicitation software: Software is available that makes it easier for people to translate their views into specific numbers. For example, the package @RISK, used for Monte Carlo simulation, has drop down menus of distribution shapes and as you tweak parameters of distributions it immediately displays numbers that help you judge if the distribution is what you want. A similar effect can be achieved with home made spreadsheets.

Private elicitation: When asking for expert opinions to use in a forecast, do not allow people to know what others have said before giving their own view. This prevents a bias called anchoring, where people tend to stay near to the first number mentioned (even if they know it was selected at random).

Plous's method: Professor Scott Plous has shown that, if you want estimates for the range of an uncertain quantity, then it is a good idea to ask a small group of people to give high and low estimates individually, then take the highest high from anyone in the group, and the lowest low and use those as the range.

Range-first effort estimation method: For estimating effort involved to do project work one idea from NASA is to get a best estimate then ask for the probability that the actual effort will be twice that or more, and the probability that the actual effort will be half that or less. This method usually gives less biased results than just asking for a range with given probability.

Prediction markets: A cutting edge technique is to encourage lots of employees to participate in 'markets' that reward them for making good predictions. Different designs of market now exist to deal with situations where lots of people make just a few predictions, and where a few people make lots of predictions (which is possibly the most common need.)

More informative forecasts

Realistic forecasts should be presented well so that all that useful work can be understood and used. Tasteful, data-dense information graphics are always a good idea. Here are some more specific suggestions.

Tornado graphs: These are usually used in conjunction with Monte Carlo simulation. They show how important each uncertain input estimate is to the uncertainty around the output forecast. There is more than one way to calculate this connection for display as a Tornado graph.

Sets of cumulative probability graphs: The most flexible style of graph for showing probability distributions is the cumulative probability graph, but they don't get used enough and take a little while to get used to. Alternatives in a decision can be compared using these graphs, and they can also be used to show how important each input uncertainty is to the output, as with Tornado graphs.

Rich predictions: It may be that your model already predicts some useful variables, but you haven't been showing them to forecast users because they are just part of getting to the headlines. Consider showing more.

Interactive display: If there's a lot of information to show then use software to let people choose what they want to see from a forecast, such as different time periods and different variables.

Set programmed alerts: Set software rules to check for particular conditions that you need to respond to, such as a likelihood of over-booking, running out of cash, or missing deadlines.

Better use of forecasts

If you are used to forecasts being made from time to time but can't remember the last time any action flowed from them other than a demand for a more acceptable forecast then there is perhaps scope for improvement in the way forecasts are used.

Of course, the way you expect to use forecasts should guide the way forecasts are made. Think about possible courses of action early on, and keep coming back to them. Make sure the forecasting system can show the results of alternative courses of action.

Forecast more: Since forecasts don't have to identify one future outcome with confidence, but instead should be probabilistic, why not try forecasting more things? Try forecasting things you've never tried to forecast before because of lack of data or understanding. For example, if there are more people you can try to sell to than you have time to sell to, and if you can identify a set of people with twice the usual likelihood of buying, then on average you can double your sales productivity even if most sales attempts are still a failure.

Forecast useful things: Forecast things that can inform useful decisions, like exactly who will be sold out and who will have too little work unless you take some action to rebalance work. There is a limit to this because more detailed forecasts tend to be less reliable than aggregated forecasts, so you need to bear this in mind and not make large, frequent adjustments on the basis of weak indications.

Forecasts as early warnings: Each forecast is an assessment of what might happen if a particular course of action is taken. Periodically forecasting on the basis of the current plan (even if the plan is just to 'carry on as usual') can reveal future situations that might need some adaptations. For example, if your forecast suggests a 10% chance of running out of cash, is that enough to prompt you to start looking at where you could get a loan, or perhaps to invest less aggressively, or find a way to react more quickly to early evidence of cash problems developing?

Forecasts as planning tools: The final forecast tends to be less useful than the forecasts that led to the plan that is the basis of the final forecast. Developing plans and developing the ability to forecast their possible results should go hand in hand.

Interactive models: Sometimes it is possible to build a computerised model that is so easy to use that alternatives can be tried out in a few second during meetings as people think of alternatives and wonder what difference they would make. Realistically, this works up to a point and then people start suggesting things that need more backroom work before a forecast can be produced.

Forecasts for mental preparation: Some forecasting systems are multi-time period simulations that you can interact with. They show you what they predict (which might be an outcome selected at random from more than one possibility) and you can react with an action, see what might happen, and so on. This can be a useful experience for mentally preparing people to respond to real events in future. Clearly there are problems with this if the simulation consistently behaves in a way that is unrealistic. It might be rewarding behaviours that in the real world are a poor choice. On the other hand, a system that's realistic enough can be very useful. Another approach is to use a competitive game, where all participants are learners and simulate behaviour for each other (e.g. in an imaginary market place or battle).

Finally – do not duplicate effort in risk management

In the world of 'risk management' a forecast using distributions instead of predicting just a single future is, logically, the same as a 'risk analysis' or 'risk model'. This is

obvious with financial risk models of the value of a portfolio of investments, where the model provides a probability distribution for the value of the portfolio at a point in time. It's also obvious where quantitative models of projects are used to produce probability distributions for things like the overall cost and final completion date of large projects. It is less obvious where a risk register is used with little or no additional workings because here the quality of thinking tends to be poor.

Rather than duplicate effort between 'risk management' and other forecasting work, do the job once, properly, perhaps taking the best ideas from each effort and combining them.



About the author

Matthew Leitch is an author, educator, consultant, and researcher specialising in the way we deal with uncertainty at work.

About Working In Uncertainty

The idea is that, by making incremental improvements to the way we work that help us deal with uncertainty, we can achieve more with less stress. This applies to individuals and organizations. The starting point is usually to learn more about techniques that work well and try them. For more information please visit www.WorkingInUncertainty.co.uk.

Further reading

Here are some suggestions for further reading.

Improvement idea	What to search for
Cut detail	Hope, J.D., and Fraser, J.R.T. (2003) <i>Beyond Budgeting: How Managers Can Break Free from the Annual Performance Trap</i> , Boston: Harvard Business School Press, ISBN 1-57851-866-0
Extrapolate statistically and Analyse only changes	Leitch, M (2007) <i>How the United Kingdom Accreditation Service improved its financial forecasting</i> available here http://www.workinginuncertainty.co.uk/cases_UKAS.shtml
Constructive simplicity	Chapman, Chris and Ward, Stephen (2002) <i>Managing project risk and uncertainty: A constructively simple approach to decision making</i> , Chichester, UK, John Wiley & Sons.
Distributions, Propagating uncertainty, and Statistical spreads	Leitch, M. (2010) <i>A pocket guide to risk mathematics</i> , John Wiley & Sons, Chichester. Leitch, M (2010) <i>Visualizing an uncertain sales pipeline at Z/Yen</i> , available here http://www.workinginuncertainty.co.uk/cases_ZYen.shtml

	Savage, S.L. (2009) <i>The flaw of averages</i> , John Wiley & Sons.
Programmed decisions	-
Consider the relatively certain context	-
Forced consideration of ranges	Morgan, M.G. and Henrion, M. (1992) <i>Uncertainty: A Guide to Dealing with Uncertainty in Quantitative Risk and Policy Analysis (second edition)</i> , Cambridge University Press. Schoemaker, P.J.H. (2003) <i>Profiting from Uncertainty: Strategies for Succeeding No Matter What the Future Brings</i> , Free Press.
Forced consideration with memory joggers	Ward, S.C. (2004) <i>Risk management: organization and context</i> , Witherby Seamanship International Limited. (see chapter 4)
Rich predictions	-
Causal links	Leitch, M. (2008) <i>Intelligent Internal Control and Risk Management: Designing High-performance Risk Control Systems</i> , Gower.
Progression	Leitch, M. (2008) <i>Progressive risk control integrated with strategy and performance management</i> , available here http://www.internalcontrolsdesign.co.uk/progressive/index.shtml
Comparing actuals with forecasts and Cross validation	Leitch, M. (2010) <i>A pocket guide to risk mathematics</i> , John Wiley & Sons, Chichester.
De-biasing experts and interactive elicitation software	Morgan, M.G. and Henrion, M. (1992) <i>Uncertainty: A Guide to Dealing with Uncertainty in Quantitative Risk and Policy Analysis (second edition)</i> , Cambridge University Press.
Private elicitation and Plous's method	Plous, S. (1995). "A comparison of strategies for reducing interval overconfidence in group judgments". <i>Journal of Applied Psychology</i> , 80, 443-454. Available here: http://www.socialpsychology.org/pdf/jap1995.pdf?logged=true
Range-first effort estimation method	Jørgensen, M. (2004) <i>Realism in Assessment of Effort Estimation Uncertainty: It Matters How You Ask</i> , available here http://simula.no/research/se/publications/SE.4.Joergensen.2004.e/simula_pdf_file
Prediction markets	-
Tornado graphs	-
Sets of cumulative probability graphs	Chapman, C.B. and Ward, S.C. (2003) <i>Project Risk Management: Processes, Techniques and Insights (second edition)</i> , John Wiley & Sons.
Rich predictions	-
Interactive display	-
Set programmed alerts	-

Forecast more	-
Forecast useful things	-
Forecasts as early warnings	-
Forecasts as planning tools	-
Interactive models	-
Forecasts for mental preparation	Schoemaker, P.J.H. (2003) <i>Profiting from Uncertainty: Strategies for Succeeding No Matter What the Future Brings</i> , Free Press.
