



Australian Centre for the Study of Armed Conflict and Society Occasional Paper Series No: 4

Intelligence analysis: what is it good for?

The problem with probabilistic forecasting

Dirk Maclean



Occasional Papers Series

The ACSACS on-line publication series embodies UNSW Canberra's commitment to engaging in public discussion and, where possible, the development of public policy. This series was created in response to requests from Centre Fellows and Adjunct Lecturers for a vehicle to gain an audience for research and writing that relates to specialist topics that are not addressed in standard scholarly publications. Three broad categories of work are 'published' in the Series – Position Papers, Working Papers and Occasional Papers – each reflecting the length and purpose of the manuscript rather than its academic discipline.

Position Papers are 2,000-4000 words in length and seek to shape debate, direct discussion or outline a position on some aspect of policy. The emphasis is on highly topical work embodying the opinion and judgements of the contributor on matters of contemporary concern. Working Papers are 3,000-5,000 words and are intended to be 'work-in-progress'. Papers in this category are offered for comment from other scholars working in the area. These papers are 'first drafts' of more substantial pieces of writing and present interim conclusions. Occasional Papers exceed 5,000 words and constitute completed work. Papers in this category include high quality descriptive and analytical work that might be too specialised or too topical for a scholarly journal. There is no upper word limit for papers in this category.

These papers are available in identical HTML and print-ready PDF formats and include an author note, illustrative material and references for further reading. The series is promoted through UNSW Canberra and ACSACS Twitter accounts operated by the Centre Manager.

Submission Guidelines

Researchers interested in having their work appear in the series are encouraged to contact the Director of ACSACS in the first instance. Prospective contributors need to decide whether their submissions are to be assessed as Position, Working or Occasional Papers. Copies of the UNSW Press Author Pack containing guidelines on style and format are available from the Centre Manager.



About the author

Dirk Maclean was recruited into the RAAF in 2012 as an Intelligence Officer. He has a PhD in philosophy from La Trobe University and an MBA. He has lectured at Masters level in performance management and the characteristics of high performing organisations. He also acts as an industry consultant and trainer in the field of Incident and Crisis Management building on his experience as a career firefighter in Victoria. In November 2014 the Australian Defence Journal published his co-authored paper (with Charles Vandeppeer), 'The Changing Role of Intelligence Within Military Decision Making' which explored the intimate relationship between intelligence and operational decisions. His book, *Shoot, Don't Shoot: Minimising the Risk of Catastrophic Error through High Consequence Decision Making (HCD)* is currently in press with the Air Power Development Centre. Implementation of the HCD research is rolling out through a number of initiatives across Defence affecting training, doctrine, team functioning, and organisational culture.

Abstract

There are a number of challenges associated with trying to measure the value of intelligence analysis. One current solution that has gained popularity is to focus on predictive intelligence, and to use statistical techniques to test predictions against the actual course of events. I will demonstrate that this approach is not only fundamentally flawed in terms of method it is also dangerous because it gives priority to idle speculation about unknowable futures. I want to show that intelligence analysis is best measured by its ability to give decision-makers the broadest set of options and that its value resides in the outcomes of the actual choices that are made. My arguments and conclusions are based on two case studies: the Battle of Kursk during 1943 and the 2012 Malian coup. This approach will assist senior managers and strategic direction setters across intelligence agencies whose output includes predictive intelligence. It will be especially helpful to officials grappling with the problem of how to measure the quality and value of their intelligence analysis.

Introduction:

Country X and country Y are highly unlikely (0.1) to go to war with each other over the next 12 months... unless of course they do. The prediction was a highly unlikely (0.1) chance for any of these ten countries to suffer a military coup that year, and we were right, 9 of them didn't. Which of the ten was going to be the unlucky one, of that we had no idea, but our forecast was 100 percent accurate.

How can we measure the performance of intelligence analysts, and the value of their products? This is a question that has engaged the intelligence community in recent years, above all in the wake of the Iraq Weapons of Mass Destruction (WMD) affair. Among those efforts to improve the quality of intelligence analysis has been the drive to apply statistical methods in a form that allows for the rigorous measurement of forecasts. This in turn creates the possibility of analysing the qualities displayed by good analysts, defined by the accuracy of their forecasting.

Leading the way in this field has been the 'Good Judgment Project', financed by the Intelligence Advanced Research Projects Activity (IARPA), and directed by the prominent researcher Phillip Tetlock. Also useful for its insights has been the work of David Mandel and his colleagues in Canada. Both provide a useful jumping off point for a discussion on the merits of probabilistic forecasting as an analytical technique. The motives behind these efforts are understandable. If analysts are producing predictive intelligence then it seems useful to know how their forecasts compare to the unfolding reality. I will argue that this approach is fundamentally flawed on two counts. First, it is chasing a rainbow, strategic and political developments are unpredictable by their very nature. There is an absolute limit to the extent to which forecasting can be improved, seeking to reach beyond this is deluded and dangerous. Second, it defines the value of intelligence analysis in terms of forecasting accuracy, but this is a very poor metric, even if successful forecasts were possible. A better measure, I will argue, is the ability of intelligence to provide sound support for decision making in the present. Its value, in other words, lies in the quality of the decisions it inspires. This too can be measured.

The future is unknown, it carries both risks and opportunities, it contains possibilities for action, and situations where any effective intervention will prove to be impossible. One of the often-cited objectives of intelligence analysis is to 'reduce uncertainty', to penetrate the unpredictability of future events through the use of sound analytical techniques, including statistics, and the application of probabilistic reasoning.

Probabilistic forecasting has been pioneered in the field of weather forecasting. It is one means of dealing with inherent uncertainty. The great advantage of this method is that the accuracy of its forecasts can be measured. How this works can be seen in the following scenario. If we make the prediction based on our modeling that there is a 10 percent (0.1) likelihood of a storm tomorrow, and then produce 10 such predictions over a period of time, and one of these storms does eventuate, we can safely argue that our forecasting was accurate. Even if some uncertainty still remains, the model is never completely correct. We might determine that if 9 out of 10 storms are successfully predicted using this method, this is an acceptable result because the storm warnings issued prevented large-scale loss of life. This measurability has led to probabilistic formats being adopted by intelligence agencies over the past decade. The results appear impressive. According to Thomas Fingar, a senior figure in the American intelligence community, the Central Intelligence Agency (CIA) is currently scoring 85 percent in the accuracy of its forecasts. David Mandel's research into Canadian analysts arrives at a figure of 94 percent. It would seem that the future is not so unknown after all.

These numbers stand in sharp contrast to Philip Tetlock's earlier research that showed expert predictions to be no better than 'a chimp throwing darts'. In other words, it is a purely random approach. While it is true that major reforms into the way intelligence analysis is carried out have been introduced over the past decade largely in the wake of the Iraq WMD failure, is it really plausible that the performance of forecasters can have been transformed so dramatically as a result? Or is it more likely that these figures are simply 'too good to be true'? Fingar certainly expresses doubt over the matter. He remarks: 'These numbers worry me'. Mandel is also surprised at the results of his study. In my view, they are both right to be skeptical for reasons I will outline.

Is the future predictable?

The first and most fundamental objection to this approach flows from the question: is the future predictable? More specifically, are geopolitical events of such a nature that it is possible to predict them? If the answer is negative, there is simply no escaping the reality that we are better employing chimps and buying a dartboard instead of investing millions of dollars seeking reliable forecasts.

This is a deep ontological and philosophical question. It has been debated for centuries, perhaps even since the Enlightenment in modern Europe. The argument against predictability rests on the assertion that historical events are unique and unrepeatable in part because human beings have free will. They can intervene in any situation and change the outcome. Those who reject this argument must respond with one of two possible counters. The first is to demonstrate a causal mechanism that drives human affairs and shapes what will happen. This was the claim of Marxism. It understood history as driven by the development of the productive forces and the growing contradiction between these forces and a more conservative social order. This 'historical materialism' has long been discredited after being taken to task by the philosopher of science Karl Popper, among others. Nevertheless, some theoretical explanation of what shapes history is needed if we are not to conclude that events are random and unpredictable by their very nature.

This constitutes an importance difference between political predictions and, for example, predicting the weather. In the case of weather forecasting, what gives this activity a scientific basis is not the presence of sophisticated computer models but a set of theories as to why the weather acts in some ways and not others. These inform both forecasts and the inputs to the models themselves.

The second response to this is to rely on statistics and to generate statistical probabilities that extrapolate into the future based on current conditions. The promise of ‘Big Data’ and super computing is that it will succeed in making such extrapolations more and more accurate. This is a dangerous illusion. It rests on a number of assumptions that have no basis in fact.

	FOR THE FUTURE TO BE PREDICTABLE...
1	PAST PATTERNS MUST REPEAT
2	INITIAL CONDITIONS MUST BE REPRODUCED
3	EVENTS MUST FOLLOW A NORMAL DISTRIBUTION
4	VARIABLES MUST BE CALCULABLE
	TABLE 1

The first of these assumptions is that patterns discernable in the past can be applied to predict the future. Again, we need a causal mechanism for this proposition to have any force. We can make a case that the life expectancy of an African dictatorship has tended to be, say, two years, drawing on past events. But unless we have a plausible explanation as to why this has been the case, and show the same conditions to be present now, we have no grounds on which to argue that a current regime is likely to be replaced over the next 24 months. Secondly, if historical events are unique, then it is ruled out that exactly the conditions that led to the overthrow of previous African dictators in the past will apply in the future. This is easy to demonstrate with a quick review of what has changed irreversibly across the continent in recent years. Thirdly, and this is Nicholas Taleb’s argument, is that statistical methods assume a normal distribution curve in order to arrive at their probabilities. However, many historical developments do not follow this pattern, but instead form ‘fat tails’, or ‘Black Swans’ - ‘low probability high impact’ events. Finally, there is the problem of complexity. This problem was clarified by the mathematician Henri Poincare as early as the nineteenth century. Even if certain trends can be identified and shown to have predictable properties, such as the urbanisation of the world’s population, to convert these into forecasts specifying individual events is mathematically impossible. There are simply too many variables to include in the calculation.

Better forecasters, or just 'less bad'?

In his interview with Russell Roberts, Phillip Tetlock seems to concede this point. The future is unforeseeable by its very nature. Nevertheless, his 'Good Judgment' project is guided by the attempt to circumvent this obstacle by focusing not on the unpredictability of events but on those forecasters who appear to navigate this successfully. Not all forecasters are equal. Some are remarkably effective. Others are hopeless. Tetlock's aim is to detect what distinguishes the good from the bad in the hope that by identifying those key qualities that separate the top 2 percent, the 'super-forecasters' from the pack. The same results can be reproduced by the wider analytical community.

This is an attractive proposition. It is a way of 'cheating the system'. Nevertheless, its essential premise is false. If future events are by their nature unpredictable, the only possible explanation for the success of some forecasters is statistical. If we throw a large enough number of guesses, informed or not, at a series of random events, probability dictates that some of these will sweep the board. This does not mean, however, that those forecasters have some magical gift the rest of us lack. It simply means they are lucky.

There is a response to this argument. It rests on the discovery that the likelihood of whether a forecaster will be in the 'lucky' group or the 'unlucky' can be predicted. There do appear to be some skills that allow us to separate the two. Tetlock's earlier research contrasted the results obtained by 'hedgehogs' against 'foxes', where he found that the 'foxes', who have a more open-minded, multi-disciplinary approach, consistently outperformed the 'hedgehogs'. The 'Good Judgment' project appears to back this conclusion up. Regrettably, this does not overcome the problem of unpredictability. While it is impossible to improve on dart-throwing chimps when predicting the unpredictable, it is possible to do worse. Foxes turn out to be 'less bad' than 'hedgehogs', at least for those events that follow a normal distribution pattern ... which is most of them. In other words, a probabilistic approach, drawing on solid analytical techniques, can be 'less bad' than another that draws on unsound techniques. This true most of the time. This observation does not help us. We are still better off with our chimps. Random events will be better forecast by random guesses than any rational process. There is simply no way around this aspect of our reality. The system cannot be cheated.

We can get a sense of this reality if we conceive of a 'Bad Judgment' project. Such a project is one that takes not the best but the worst 2 percent of our forecasting sample. If there is any causal relation between the accuracy of their forecasts and the techniques they use, all we need to do is reverse their predictions, place a minus where they put a plus, and we will instantly match anything our top 2 percent can achieve. Training analysts to be terrible will work just fine.

The problems do not end here because even the performance of our chimps is not as good as it first appears. Why? Because the unpredictability of events defies a purely statistical approach. It does not follow a normal distribution curve. This is the challenge posed by 'fat tails' or 'Black Swans'. If we adopt a probabilistic approach, even if we succeed in calculating the probabilities correctly through either good analytical techniques or pure chance, we will stumble every time a Black Swan occurs, ie., an event that defies the odds. As Nicholas Taleb points out, this is more often than we think.

What does a probabilistic forecast mean?

The attempt to cheat the system comes at a price. For while it is true that the accuracy of probabilistic forecasts can be measured, this is not sufficient to answer a more fundamental objection against this method. This is the argument that such forecasts are meaningless.

This becomes apparent once we shift our focus back from the forecasters to the forecasts. Let's say we arrive at a typical kind of prediction that reads: 'It is probable (0.67) that sections of the armed forces in country X will launch a military coup within the next 12 months'. We might also make another forecast with the reverse probability: 'It is unlikely (0.33) that country Y will break off diplomatic relations with country Z in the next year'. The question that surrounds such predictions is this: what do they mean? What, for example, would confirm or deny their accuracy? The answer is nothing because these are probabilistic statements and whether the events take place, or do not take place, either eventuality is covered by them. In Karl Popper's words, this makes them unverifiable and therefore unscientific, they are 'pseudo-science'.

It does not matter that the accuracy of all our 0.67 and 0.33 predictions can be lumped together, measured against their outcomes and our forecasters given a round of applause if, in fact, they hit the required strike rate. For this prediction, for any of them taken individually, there is nothing that can separate true forecasts from false. They are neither 'correct' nor 'incorrect'. They have no meaning whatsoever. This becomes clearer if we consider this question: is one event more likely than the other - the 0.67 more than the 0.33? The answer is: no. If political developments of this kind are historically unique, subject to complexity, and defy statistical probabilities, then the reality is that they are both equally likely and unlikely. This is the essential point to understand. One may look more likely than the other from our present vantage point. But it is not. It only seems to be so. In reality, the improbable happens as much as the probable does. This is the problem with probabilistic forecasting. It gives the appearance of saying something of substance while actually saying nothing at all.

Weather forecasters are acutely aware of this problem, and act with great care not to overstep the boundaries of what their field is able to deliver. Weather patterns do display a level of predictability, a 24 hour forecast of rain within a region has a large chance of being accurate. A tornado, on the other hand, cannot be predicted with any prospect of success more than one hour in advance. Weather forecasters accept these limits and work within them while continually seeking to improve their models and their understanding of how the weather works. Ironically, probabilistic forecasting has its greatest value when predicting rare severe weather events such as storms or cyclones, where the benefits of a 72-hour warning far outweigh the costs of false alarms. Furthermore, weather events are clearly definable and resemble one another. Classification systems such as the Beaufort scale for cyclones have real meaning and utility. This is by no means so for geopolitical developments, whose category and meaning cannot be pre-determined. An impending cyclone will always trigger a warning, rather than a debate as to whether its arrival should be welcomed or feared. As we shall see below, this is just not the case when it comes to human affairs. Weather forecasting has one additional advantage. It is a highly repetitive activity with rapid feedback provided day-in-and-day-out. Therefore, models can be constantly refined to improve their predictive ability which, in turn, is easy to measure. The results have been impressive. The accuracy of weather forecasts has improved dramatically over the past two decades. Nevertheless, there are strict limits beyond which the weather will always remain unpredictable.

With intelligence analysis, it is not clear that the same kind of discipline is present. This is because the use of probabilistic forecasting is in truth an attempt to cheat, to over reach what is actually possible. This is similar to many so-called 'predictive sciences', such as the kind of political forecasting which was busy building models to predict the 2016 United States Presidential election, based on the patterns of voter behaviour in previous contests. The futility of the effort seems to be outweighed by the money gullible users are willing to pay for the output of such models in much the same way as the careers of financial market analysts always seem to survive the consistently dismal results of their forecasting. On the other hand, even these can be rated 'foxes' when compared to economists, who are in a whole class of their own.

Weather forecasters avoid the pitfalls of these 'pseudo-sciences'. A large part of the science of meteorology seeks to uncover major patterns such as the El Nino effect, and then to put knowledge of the mechanisms behind such phenomena to practical use by integrating them into their models, to the extent this is viable. As it happens, the same possibility does exist for intelligence, as long term trends such as demographic shifts, urban growth, even climate change, can be usefully incorporated into estimates and projections, so long as proper caution is used. The argument here is simple: where there is predictability, predictions are possible. Where there is no predictability, they are not. Forecasting specific events such as the outcome of an election 12 months out is nothing more than having a punt. It is guesswork. The rigour of the analysis has no bearing on the prospect of success.

Acclaimed 'super-forecaster' Nate Silver recognizes this truism. His great achievement, accurately predicting all 50 states in the 2012 United States Presidential race, was made on the morning of the election. His presuppositions precluded a Black Swan event such as the assassination of a candidate. He attributes much of his success to luck, in particular, for states such as Florida where the two candidates were polling neck and neck. The secret of his success was not so different from taking a look at the dartboard as his wonder forecast was largely derived from a 'poll of polls'. This conclusion is supported by the failure of every forecaster, including Silver, to forecast the 2016 election result. This is not to say that political science, including the study of voter behaviour, cannot be of use. But this is not because it can predict the outcome of an election. It can't. Its real value is to campaign managers and their deliberations on what issues to run on, initiatives to announce, where to spend their money, and how to respond to their opponents' moves. Its merit lies not in forecasting the result but in shaping it. This is a crucial difference and one I will return to later.

Cheating the system? Playing it safe

If this objection stands, then how to explain the high success rate noted by Fingar, Mandel, and Tetlock, which seems to leave our chimps far behind, back with the hedgehogs and economists. Here it turns out that it is possible to cheat the system, not in the sense of predicting the unpredictable, but in framing forecasts in such a manner as to guarantee a high score for accuracy. This was Fingar’s concern over the 85 percent result for the CIA. The problem is, however, that this ‘success’ is an illusion. Fingar himself expressed his worry with these words: ‘I would prefer a success rate of 35 percent, not because I want us to get things wrong, but because I fear we are avoiding the hard questions’. Fingar is correct then the results are rigged. We can get an idea of how this works by looking at the two sample forecasts cited by David Mandel in his research. They are,

1	“It is very unlikely [0.1] that either of these countries will make a strategic decision to launch an offensive war in the coming six months”
2	“The intense distrust that exists between Country X and Group Y is almost certain [0.9] to prevent the current relationship of convenience from evolving into a stronger alliance.”
TABLE 2	

Several features of these two forecasts stand out. The first is that they both represent a very high degree of certainty, 0.9 or 90 percent. It is not clear exactly how many of the 1,514 forecasts examined by Mandel were of this nature. It seems the number was high. This alone would be enough to skew the results. If we make a series of 0.9 predictions, then we would expect to achieve a success rate of 9/10. In fact, this would be the ideal outcome. The actual score of 94 percent would reflect a slight under-confidence bias among our analysts, a conclusion that Mandel notes. In a revealing comment, Mandel praises this willingness to make strong predictions. He concedes: ‘There were few uninformative forecasts near maximum uncertainty (probabilities of 0.4–0.6). If analysts were only concerned with playing it safe, we would likely have seen a bulge rather than a trough in that region’.

That is one interpretation. There is, on the other hand, an entirely different one we can reach reflecting Fingar’s concern. An alternative interpretation is that analysts are concentrating their forecasts on those areas where a 0.9 probability is possible in order to improve their accuracy scores. If we can pick and choose our forecasts in this way, it is quite easy to achieve a very high score. This is apparent when we take any number of situations in the world today. It is common to say we live under times of great turbulence and uncertainty. In reality, this is not true everywhere. It is not difficult to find areas of calm and stability. Here are two quick examples written in January 2016. They are based not on research but general knowledge.

<p>UKRAINE CEASEFIRE HOLDS 0.9</p>	<p>The position in Donbass is widely described as a ‘frozen conflict’, and rightly so. There are powerful constraints on each of the main actors that suggest the current ceasefire will hold, at least for the next few months. There is no certainty to this, the war could erupt violently at any time, but a forecast that gave the status quo continuing a strong probability would be a reasonable bet.</p>
<p>LEBANON COUNTRY IMPLODES 0.1</p>	<p>All the sectarian tensions that led to the Civil War in the 1970s and 1980s are still present and unresolved, the deadlock over the Presidency, the Iran-Saudi rivalry, and the Syrian refugee problem are exacerbating these. Nevertheless, all of the key players have a strong aversion to a renewal of the ‘bad old days’ before Taif, a confident forecast that there will be no implosion in Lebanon during 2016 is also a pretty safe call.</p>
<p>TABLE 3</p>	

Forecasts of this nature can be banged out by the dozen if all that counts is accuracy. A 90 percent strike rate is no great miracle if predictions are restricted to this kind of relatively stable situation. In reality, we hardly need to employ and train expert analysts to derive these kinds of forecasts. Practically anyone can do it and do it quite well. This conclusion is supported by the success of crowd sourcing initiatives. Furthermore, it is hard to see how such forecasts achieve their stated aim of ‘reducing uncertainty’, as they avoid all the really uncertain situations, the ones that would rate 50/50 or 60/40. There is an argument that these are precisely the ones where intelligence expertise is in greatest demand but our forecasting mindset encourages us to stick instead to the ‘almost certain’ or ‘very unlikely’.

When we examine Mandel’s two samples with this in mind, then we find that they are in fact of the same type. Both in essence, say the same thing: nothing significant is going to change. Both are negative statements. They predict the future to be an extrapolation of the present. It will be ‘business as usual’ in other words. This is true even though one is written as a positive prediction, 0.9 likely, and the other as a negative, only 0.1 probable, and therefore appear on Mandel’s statistical chart in two separate categories. In substance both forecasts are fundamentally the same, they are relatively safe bets that life will continue as we know it. If we reflect on this a moment, it leads us to ask, what conclusions follow from predictions of this nature, what decisions do they entail for policy makers? The answer is none. Their clear implication is ‘forget about it, no action required’. The same applies to our Ukraine and Lebanon forecasts.

What if we are wrong?

This points to an obvious problem, first highlighted by Taleb. What if we are wrong? His argument is that the consequences of error and success are not even. If we are correct, and nothing of significance happens, then that is nice for our analysts, perhaps they will get a bonus or an award. Beyond that, who really cares? But if we are wrong, if we are caught out by improbable and unanticipated developments, the two countries go to war, Ukraine erupts, Lebanon implodes, the consequences are far and away more serious. We may well achieve our 94 percent strike rate, but the fall out from our 6 percent of errors is disastrous.

This is a major effect of probabilistic forecasting. It leaves us wide open to Black Swans, low-probability high-impact events. This is a severe failing. What we need, in another words, is not simply a metric that tells us how accurate our forecasts are but one that measures how valuable they were, what consequences flowed from both our successes and failures. What is needed is a qualitative as much as a quantitative measure.

We can see a further limitation to our accuracy score in another example. A forecast was made at the start of 2015 by a private sector forecaster over Russia's involvement in the Middle East. One year later, they evaluated their predictions and awarded themselves an 'A'. This was on the basis of a number of statements relating to al Assad, his opponents, other regional actors, and the United States. Most of these statements turned out to be reasonable descriptions of how the year played out. The problem, however, is not what the forecast said but what was overlooked. The intervention of the Russian Air Force proved to be a game changer for the entire course of the war in Syria.

In order to evaluate our forecasting, therefore, we need to include another category, that is, the forecasts that were not made at all. Our success rate may be 94 percent for the ones we did make but it is this larger estimate that tells us something of much greater value as to how well our intelligence analysis has performed. We also have the famous 'Lucas critique'. It describes the distorting effect of our own actions on prediction accuracy. This particularly applies to the more mid-range forecasts. We may predict that a favourable outcome is likely (0.6) but not overwhelmingly so. This encourages a redoubling of political efforts to ensure we get the result we want; we succeed, and the prediction is confirmed. It is not clear how great an effect on the statistics this has, as it can work in the opposite direction. Nevertheless, it does undermine a more naïve measurement that only looks at accuracy. It means for example, that we may look positively on predictions of unfavourable outcomes, which then didn't eventuate, because active steps were taken to prevent them coming about. The forecast was inaccurate but served as a warning. Job well done.

Less prediction, more analysis

Placing the emphasis on forecast accuracy, because this can be measured, has further negative consequences. If the value of intelligence lies in ‘reducing uncertainty’, and this is understood as making probabilistic predictions about the future, then forecasting emerges centre stage as a key activity and forecasts as a key product. This comes at the expense of analysis, which looks not only at what might happen, but why. This is an important shift. The reality is that decision makers do not only want advance notice of important developments in the world, if such were possible, they also need to understand their significance, what they mean. This is not always obvious, and is a critical task for intelligence analysis. This especially applies for developments that are improbable, events that strike out of the blue, sudden shifts in a situation, precisely those in other words, that a probabilistic forecast tends to dismiss as ‘highly unlikely’ (0.1) and therefore safe to ignore.

If we take our two countries that are ‘very unlikely’ (0.1) to go to war, we can certainly decide to evaluate this forecast in terms of whether they do in fact go to war in the future. But this is not all we want to know. We are also interested in why this prospect seems so unlikely at present, what are the factors that are holding either nation back from taking the fateful decision to launch an all out attack? We want to know this, because in actual fact, this gives us much better information about the future than does the attempt to predict the unpredictable.

If we play with this scenario, then the point becomes clear. Let’s say our analysis, in support of this forecast, reads as follows:

	<p>‘Both countries are heavily engaged in developing a nuclear capability, neither has achieved a decisive lead over the other at this stage, each is between 3-5 years away from an initial test of a nuclear device. The progress of their competing nuclear weapons programs dominates the strategic thinking of the leadership in both nations’.</p>
	<p>TABLE 4</p>

This analysis is not a prediction. It is actually a statement about the present. There is a predictive component but that is about a highly complex technical process involving a number of stages, is relatively objective in nature, and therefore predictable to some extent. The benefit of such an analysis such is that it now allows us to say something important about the future:

	<p>“If either country were to pull significantly ahead of the other in their nuclear weapons programs, the risk of a preventative offensive war launched by the lagging nation would be extreme’.</p>
	<p>TABLE 5</p>

This is not a forecast either. It is a conditional statement about the future. The value of this statement consists of the fact that it tells us what to look for, it advises us of the key warnings and indicators that will give us notice that the risk of war is escalating rapidly, and will soon be extreme if things continue along this path. This is useful information, it is actionable intelligence, it can inform collection priorities, set in motion diplomatic initiatives, and guide operational planning for the military.

The differences between an analysis of this kind and a forecast are worth highlighting.

INTELLIGENCE ANALYSIS OVER PREDICTION	
1	the analysis is of the present, which is largely knowable, rather than of the future, which is not
2	the focus is on the actions of key players, not possible future events
3	the analysis employs the familiar categories of intent and capabilities
4	the analysis not only describes the current position, it explains it
5	the analysis provides some key warnings and indicators that the position may be changing, it tells us what to be looking for
6	the analysis covers a Black Swan event, it helps to prepare for such a contingency
7	the analysis is actionable intelligence, it implies certain steps and measures be taken by military and civilian leaders
8	the analysis informs future collection priorities

TABLE 6

Each of these eight points are capable of generating useful metrics as to the value of this intelligence assessment. This means a picture can emerge that is multi-dimensional. Some aspects will prove more effective than others. Areas for improvement will be highlighted. From a management perspective, this offers a far richer set of tools for evaluating the performance of analysts and of an agency. Analysis is also needed when the significance of an event is ambiguous. How it comes about is as important as whether it comes about, because depending on the circumstances it can be a risk, or an opportunity. It can also be both. This is another advantage for the approach that understands historical events as singular, unique, unpredictable, and ambiguous, over one that seeks to subsume them in universal categories that often contain ideological biases and hidden value judgments.

If war breaks out between our two sample nations, is this a good or a bad thing? We may believe that war is evil in its very nature, but this is a moral stance not an intelligence assessment. From an analytical perspective, the most appropriate initial response is: 'it depends, we need to know more'. Even if the war goes nuclear, the answer remain: 'it depends, after all we might be talking about the United States and Japan in the early 1940s, from a Chinese standpoint, in which case the dropping of two atomic bombs is something to be welcomed.' This is the essential problem, prediction without analysis, even if the forecasting can be relied upon, is useless, we need to know what the event means, in order to determine what our response should be. War may break out because one of the nations has joined our alliance, and betrayed our mortal enemy, or it may break out because one of the nations has betrayed us, left our alliance, and joined forces with our chief rival. Break out the champagne or run for the hills, this is the choice we want our intelligence community to inform us on.

An example highlighting this point and drawing together many of the threads in the foregoing discussion is the March 2012 military coup in Mali.

The 2012 military coup in Mali

The coup is of interest because it was 'missed' by the experts. This is the claim made by Jay Ulfelder. He asserts that he was able to predict the event using a statistical model. Ulfelder's criticism is widely held although it may be unfair. Of more interest is why the analysts might have been caught out. Mali had enjoyed political stability for 20 years. There was no history of its military launching coup attempts during that period nor any outstanding personalities who made likely coup leaders. Mali, on the other hand, did display many of the risk factors associated with previous military coups across the African continent. Ulfelder built a statistical model using information from 40 coups from the previous 60 years. His model placed both Mali and Guinea-Bissau in the top ten of high-risk nations across Africa. These were the two countries that did suffer military coups in 2012. This is the source of Ulfelder's claim to have beaten the experts.

On closer examination, however, a different picture emerges. First of all, if the risk analysts were using probabilistic reasoning, as most likely they were, then their assessment that the risk of a coup in Mali was 'very unlikely' (0.1) was not unreasonable, given the country's track record over the previous two decades. If their forecasts had averaged a 0.2 probability for the top ten African nations, with Guinea-Bissau included alongside Mali, then they could counter with the argument that in fact their forecasting was 100 percent accurate, including their prediction for Mali. They did not get it wrong, they got it right. This would remain so even if they had also rated the risk for Guinea-Bissau at 0.1. Not only that, they might have outperformed Ulfelder's statistical model by a large margin. While his model did put Mali and Guinea-Bissau in the top ten, its actual probability estimate was way off, averaging around 0.08, or about 1 in 12. In other words, his prediction was that neither coup would happen. Ulfelder and the 'experts' can argue all they like. My point is that the argument is futile. Both methods are useless and a waste of time.

In reality, the difference between the two is not as great as it seems. Ulfeder's approach appears more rigorous because he has used a statistical model based on historical data, and run the numbers. But when we examine the factors he used as inputs into his model, we find that they are highly subjective, his selection of just four criteria that are no more than his personal interpretation of what was critical in earlier African coups. The rigour is an illusion. When we consider that a major cause of the coup in 2012 was the unprecedented collapse of Libya the previous year in the wake of the Arab Spring, a collapse that was set against the background of a global jihadist movement with an active presence in North Africa, the limitations of trying to base a prediction for Mali on African coups running back decades is apparent. The coup was a direct product of historical circumstances unique to the present century, not the repeat of a longstanding African tradition.

Even worse, if decision makers had followed the logic of Ulfeder's model, they would have been seriously misled. Their decisions would have been poorer, not better and not only because they would have discounted the likelihood of a coup in any African country. His rankings placed Niger and Guinea-Bissau far and away the frontrunners with a risk rating no less than three times higher than for Mali. Available resources and attention would have understandably been directed to these two standouts, leaving Mali even more likely to be 'missed'. Although it did enter the top ten, its rating was only marginally different from Senegal, the least likely country in Africa to make the rankings and some 30 places lower on the ladder. On the other hand, this was a Black Swan event. A probabilistic prediction based on the last 20 years of political stability would not have helped us either. This is the key point. It does not matter that a coup was 'very unlikely', the improbable happens. When the improbably does happen, it is hardly a consolation to know that it was a surprise for the 'experts'.

Ulfeder and the experts were both wrong, but on their own terms they were also both right. This is why we argue probabilistic forecasting for geopolitical events is meaningless, it is both right and wrong at the same time, it predicted the coup successfully and it didn't predict it at all. By concentrating our attention on forecasting in this way, we miss what decision makers really needed to know at the start of 2012. This was not whether a coup would happen but what it would mean if it did. The focus is on its significance. Ironically, there is a relationship between this significance and its degree of probability. Highly unlikely events are generally very significant when they do take place. They mark a break in the situation and the opening of a new period. This is another reason why probabilistic forecasts throw us off track. Military coups in both Niger and Guinea-Bissau have been commonplace over the past few years. This explains why they topped the rankings. But for the same reason they are also not necessarily very significant. In reality they are just more of the same. A coup in Mali, on the other hand, would bring to an end two decades of stability and political neutrality on the part of the military. It would represent a major turning point in the nation's history.

Equally important, the 'experts' were correct in believing that the military leadership had a strong aversion against taking political power. Their forecast may have been wrong, but their understanding was quite good. This meant that if the military did step in and overthrow the elected civilian president, they would only do so under the most extreme circumstances. As it happened, the coup leaders themselves had no intention of launching their coup just days before they did, they were as much caught out by the turn of events as anyone else.

This is why what we need in these kinds of situations are not predictions but analysis. By early 2012, a coup was still ‘very unlikely’. Yet, it was a real possibility. Mali was entering a ‘perfect storm’ of ecological, economic, political, and military upheaval that threatened the total collapse of the nation, its territorial integrity, institutions, and economic base. The strains were clearly visible. It was against this background that the significance of a coup would emerge, it would mark the explosion into the open of an already present, but up to now still simmering, crisis overtaking the country.

It would mean that the crisis was now coming to a head. For this reason, a risk analysis, as opposed to a forecast, would have reversed the rankings of Mali and Niger. A coup might well have been three times more likely for Niger but the significance of such an event would have been at least nine times higher for Mali. As a guide to where attention and resources should be devoted, analysis beats forecasting hands down, there is simply no contest.

Was the coup a good or a bad thing?

Forecasting is also deeply flawed because it assumes categories in advance and places a value on them whereas, in reality, this value is often precisely what is in question. The Mali coup is a good example of this too.

If we ask the question: ‘will a military coup take place in Mali during 2012?’, there is a powerful tendency to build into the question a value judgment – that this would be a bad thing. If we categorise this possible coup as a ‘risk’ or ‘threat’, then the value statement is more explicit. The same applies to the category ‘political instability’ which we would naturally, and not incorrectly, associate with the military coup should it happen. In general, from a Western perspective, it is pretty safe to make the following equations:

MILITARY COUP	=	BAD
ELECTED CIVILIAN RULE	=	GOOD
POLITICAL INSTABILITY	=	BAD
POLITICAL STABILITY	=	GOOD
		TABLE 7

These are not unreasonable associations. The problem is, however, that they are general value statements not an assessment of a concrete event. In real life, things don't always fit our preconceived categories. Situations arise when our normal value judgments need to be stood on their head. Mali presented just this kind of dilemma in 2012. We can see this inversion if we reproduce our categories with this statement inserted above them:

IF OUR NATION, A DEMOCRACY, IS ON A SURE
PATH TO DESTRUCTION ...

MILITARY COUP	=	?
ELECTED CIVILIAN RULE	=	?
POLITICAL INSTABILITY	=	?
POLITICAL STABILITY	=	?
		TABLE 8

Before we can answer these questions and evaluate whether the 2012 coup was going to be a good or a bad thing, we first need to understand what is going on, and what the coup might represent in a broader setting. This is the key point: the categories are not pre-ordained. We cannot say in advance what will be good and what will be bad. It all depends on the circumstances. This is another reason why our forecast of a coup, even if it turns out to be accurate, is useless. This is because it is meaningless, it does not tell us anything.

If an analysis of the situation is able to clarify what is at stake should a coup take place, we have a much more helpful outcome. But even this outcome does not answer the most pressing questions. This is because no answer is possible in advance. Only time will tell whether it is true or false, good or bad, it all depends on what happens next. This is why the correct answers to our questions are neither a forecast nor an analysis, they are:

IF A COUP TAKES PLACE IN MALI THIS YEAR (2012),
FOLLOWED BY INSTABILITY...

MILITARY COUP	=	? WE'LL SEE
POLITICAL INSTABILITY	=	? WE'LL SEE
		TABLE 9

This form of answer would appear to be even more useless than a forecast. In fact, the opposite is the case for two reasons. First, it is telling it like it is. It is an accurate description of the position immediately following a military takeover. A coup would indicate both that a period of 'maximum uncertainty' had now opened up, and would exacerbate the position by introducing political instability into an already critical situation. The question mark in our response is a realistic assessment. In predictive terms we would class it as 50/50 or 0.5. Second, it tells us something extremely important, that if a coup does happen, this is the time to take action.

Supporting decision-making, or decision-avoiding?

If we remember that Mali in 2012 was a nation sinking ever deeper into crisis, in the face of a number of seemingly intractable problems, insurgency in the North, economic ruin in the South, which the elected civilian government was proving incapable of resolving. Under these circumstances, the effect of a military coup would be to blow the whole situation wide open, at the same time massively increasing the prospect of state collapse, and the possibility of a change in direction that might save the day. It represented not only a risk but an opportunity as all crises do. This is the meaning of 'maximum uncertainty'. It indicates a moment when anything is possible and outcomes are entirely unpredictable. But it also means that if we have a preferred option, if we want to shape events, this is precisely the point in time when we have the best chance for doing so. Is this useful for decision makers to know? Of course, it is. It is exactly what they need to know.

This why defining the goal of intelligence analysis as 'reducing uncertainty' is highly problematic. It reflects a profoundly passive view of the world. It assumes that uncertainty is a bad thing, a risk, rather than an opportunity. In reality it is both. The issue is not simply about forecasting, it runs deeper, it is about passivity, about using intelligence forecasts to justify doing nothing. We can see this in the tendency to produce 0.9 probability predictions that tell us things are going to continue basically as they are, we can take no notice of them and still sleep safely in our beds at night. This is a conception of intelligence that serves not decision making, but decision avoiding. It is also a complete delusion. It equates putting a number on the degree of uncertainty posed by a situation with reducing this uncertainty. It does nothing of the kind; it simply puts a number on it. To claim that forecasting, accurate or not, 'reduces uncertainty' is to adopt a fallacy and risk drifting into fantasy.

Learning from battlefield intelligence

We can find a very different understanding of the role played by intelligence analysis if we look at its role on the battlefield (battlespace). This is captured in the 2013 Joint Intelligence publication JP 2-0. In this document, intelligence is also defined as 'reducing uncertainty', but it soon becomes clear that something very specific is meant by this phrase, and it has nothing to do with probabilistic speculations about the future. In fact, probabilistic forecasts do not make any appearance in JP 2-0. Predictive, or estimative intelligence, do appear. What a commander wants to know is Intel's assessment of how an adversary will respond to a course of action we are considering. If the intelligence function is able to do so with confidence, then this does remove some uncertainty from the commander's decision making. There is an important difference, however, between this and the content of forecasts in strategic intelligence. Battlefield intelligence is about the enemy, the forecasts we have been discussing are about events.

Predicting what an adversary will do next is not easy nor is it a science. Unlike the future, however, which is unpredictable by its very nature, the difficulty with an enemy is a practical one, not one of principle. We may have limited sources of information, an adversary may be wildly inconsistent in their actions, they might do something extremely reckless or irrational that catches us out. Nevertheless, coming to understand an opponent and predict their moves is a realistic objective and a degree of success is entirely achievable. It is also possible to have an awareness of how likely we are to have a correct understanding at any one point in time, and also how predictable our target's behaviour is likely to be. In any case, a commander has no choice. They must factor in some calculation as to what an adversary will do, whether this is based on blind guesswork or a wealth of good knowledge.

Battlefield intelligence deals with this kind of uncertainty, not with probabilistic forecasts, but by assessing the level of confidence that can be had in any particular judgment. This is an important difference. 'Almost certain' (0.9) forecasts have a similar affect. This is part of their appeal. But as discussed earlier, when dealing with the unpredictable this probability is an illusion. A 0.9 probable is not more likely to happen than a 0.1 probable. It only seems more probable from where we are standing now. The level of confidence, on the other hand, is a realistic judgment based on defined criteria and is communicated to the commander so it can be factored into their decision-making. If the confidence rating is low, because information sources are poor, then remedying this position is a question of practicality, assigning resources to improve the intelligence picture. The illusion of 'reducing uncertainty' by placing a number on it does not play any part in the process. Instead, the problem is acknowledged and real world solutions are sought.

In JP 2-0, the confidence scale is limited to three categories – low, medium and high. There is no reason why more sophisticated or granulated metrics cannot be constructed for strategic intelligence. The High Consequence Decision-Making program (HCD), which is concerned with minimising the risk of catastrophic error (such as a friendly fire incident) employs a tool for measuring the level of difficulty posed by any particular 'information management' and 'situation assessment' task. It roughly corresponds to the usual 'collection' and 'analysis' functions within intelligence.

level of difficulty posed by a situation for high consequence decision making

lack of information leading to poor situational awareness	1	2	3	4	5
---	----------	----------	----------	----------	----------

information is not available, gaps exist, no sources					
information volume is overwhelming					
information is out of date, overtaken by events					
information is unreliable, sources are unproven					
information integrity is in doubt, suspicion of deception					
information is conflicting					
information is ambiguous					

INFORMATION MANAGEMENT - OVERALL DIFFICULTY					
--	--	--	--	--	--

difficulty in assessing that information	1	2	3	4	5
--	----------	----------	----------	----------	----------

situation is dynamic, complex, non-linear					
situation is out of control, minimal ability to influence events					
situation is unique, without precedent					
situation is unassessed up to now, and unexpected					
situation is unfamiliar, lack of personal experience					
situation appears chaotic, patterns are hard to find					
the significance of information pieces is hard to evaluate					
multiple assessments are possible and credible					
time for assessment is limited and excludes structured techniques					
assistance with the assessment is unavailable					

SITUATION ASSESSMENT - OVERALL DIFFICULTY					
--	--	--	--	--	--

Tools of this kind can be used to arrive at a 'confidence level'. In an agency, rating the assessment task can be given to one analyst as another actually produces the analysis. This can assist managers in determining where best to allocate available resources to both collection and analytical activities. This is one way to cope with uncertainty but it does not reduce it. For uncertainty to be reduced, something else is required. Here again, the battlefield model is helpful because it focuses on an adversary and poses the problem as a practical challenge. Reducing the uncertainty surrounding an opponent's action is not simply an intelligence task; it is an operational objective. This task is done not through knowledge, but by taking the initiative and achieving control. If an enemy's options can be taken away so they effectively have one viable course of action, and even better if it is a bad option, uncertainty ceases to be a problem.

Reducing uncertainty by taking control

This is a long-standing principle of warfare. Once again it is an active rather than a passive approach to the challenge. If we are going to make predictions, the ideal is for them to be self-fulfilling because we are the ones in control. We know the future because we will determine what it looks like. There are countless historical examples of commanders who have achieved this mindset. A good one to cite is the 1943 battle of Kursk, because it was an intelligence coup that proved decisive for the outcome. Long before the battle opened, the Soviet High Command received a detailed report of the German army’s intent, where and when they would launch their summer offensive. This allowed the Soviets to plan and prepare their defences. In this instance it was for the defence of the Kursk salient. They did so by constructing strong points along the front lines and a series of defensive layers running back to the rear. They did not, however, reinforce the front evenly. Instead they deliberately left some areas weak, knowing that the Germans would be able to break through at these points. This was exactly what the Soviets wanted. It would channel their opponents into well-prepared killing zones. This left the Germans with a limited set of options. These were:

	GERMAN OPTIONS AT KURSK
1	Attack right across the front, probing for weak points, and following through on any penetrations achieved HIGH ATTRITION RATE EXPECTED
2	Refuse to be channeled by the Soviets and concentrate on the strong points EVEN HIGHER ATTRITION RATE EXPECTED
3	Concentrate only on the weak points, where these are known, to break through to the rear APPEARS THE BEST OPTION, BUT COULD BE A TRAP
4	Call off the operation and leave the front as it is SURRENDERS INITIATIVE, VULNERABLE TO COUNTER OFFENSIVE
	TABLE 10

Uncertainty remained as to which option the Germans would chose. This did not matter because all of the options were bad. The Soviets had used their intelligence reports to achieve control over the battlespace and shape the contest to their advantage.

Advance notice of the Kursk operation was not obtained through probabilistic forecasts. No one sat down and calculated the likelihood that this was the point the Germans would select for their main thrust in 1943. Instead it came from reliable information on the intentions of the German High Command, including details of their planning. It turns out that the secretary who typed the minutes of the meeting that decided on the 1943 summer offensive was a Soviet spy. There was a risk of deception. Minimising this possibility was a practical intelligence task and could be given the highest priority over the months that followed.

In the absence of Human Intelligence (HUMINT) such as this, the Soviets would have had little choice but to make a 'low confidence' assessment of German intentions. If they had wished, they could have converted this into a probabilistic format. It is hard to see how this would have helped as the numbers generated would have been around the 0.5 mark. If we do not know what the enemy is going to do, they would have mused, then we just don't know and we will have to prepare accordingly. In the case of Kursk, the Soviets did know, and put that knowledge to good use. Intelligence played a critical role in battlefield success. This was its merit. It allowed decision-makers to take control of the situation and determine the outcome to their advantage. Intelligence as decision support, its value measured by the quality of the decisions it enabled.

Opportunities like Kursk do not come along every day. It is unusual to find oneself in a position where it is possible to decisively influence a situation. If such an opportunity does arise, it must be recognised and exploited before the moment is lost. This is where the parallel with the Mali coup can be made. It shows how an intelligence analysis along the lines described above could have done precisely that at the beginning of 2012, preparing decision-makers months in advance for the possibility. If a coup did take place, this would be the best time to intervene and counter the growing prospect of the country becoming a failed state.

In Mali, this proved to be not such an improbable idea. The incoming military junta was unprepared for its new role, hesitant and uncertain as to what to do. It proved relatively straightforward for representatives from the Economic Community of West African states (ECOWAS) to intervene and negotiate an agreement with the coup leaders undergirded by the imposition of sanctions. Things did not return to their earlier position, however, the previous government was too discredited by this stage. As part of the settlement Mali's President resigned and a transitional government of national unity set up. In the meantime, the balance of forces in the North had shifted decisively, it was clear a military victory over the insurgents was off the cards, and a serious dialogue between the two sides opened for the first time. One year later, when the French were invited by the government to intervene militarily in order to assist Mali ward off the jihadi threat, Tuareg rebels served as a key ally in this fight.

The 2012 coup was a Black Swan, a low probability high impact event. In JP 2-0, preparing for these kinds of situations is incorporated into the operational planning process. Potential scenarios are evaluated, not only for their likelihood, which might be minimal, but their consequences, which might be enormous, and steps outlined to mitigate these if they are negative, or take advantage of them if they are positive. The idea is not to be able to safely ignore them, but to be fully prepared in case they eventuate. Building scenarios and contingency planning are the two critical activities that in practice 'reduce uncertainty' for a commander. Within them, intelligence analysis plays a decisive role by developing as full as possible an understanding of adversaries and other key actors. The same applies in strategic intelligence, where control over major geopolitical events is harder to achieve. Rather than deriving a false sense of certainty by assigning a probability score to a forecast, one that has no meaning in reality, the task is to identify and prepare for those scenarios where action is possible, where the outcome can be influenced.

In essence, this is no different from the approach of political 'scientists' when they apply their insights in order to advise campaigners. The value of their input is not hard to determine. It is the extent to which they assist the campaign to win. It lies in the quality of the decisions they inspire and how these decisions contribute to a victory. The 2016 Presidential campaign bears out this contention. Unlike the great majority of pundits, this author was confidently stating as early as September that Trump was heading for a strong victory, if not a landslide. The significance of this assessment did not consist of its predictive value, even though as a prediction it did turn out to be correct. More important, however, was the understanding that lay behind this statement, the fact that it captured the actual dynamic of the election contest, what was really happening on the ground. In other words, its value lay in the quality of its analysis, in explaining not only that Trump was in front, but why. This contrasted with the self-delusion that gripped the Clinton side right up to election night, when their world suddenly came crashing down. The analysis, which was shared by the Trump campaign team, directly informed their strategy, including the allocation of its most precious resource, the candidate's time, hence Trump's surprise appearances in Wisconsin and Michigan during the final days of the contest.

In the same way, the value of intelligence analysis consists not in its ability to predict the future, but to identify and support the decisions that will shape it. It does so above all by understanding the present, by 'knowing what's going on', and by focusing on adversaries, actors, rather than events. It seeks out not minimum, but maximum uncertainty, for the opportunities this contains, and where it does concern itself with predictions, it does so because these will be self-fulfilling, they are the decisive moments where control can be exerted, and desirable outcomes obtained. Its value ultimately lies in precisely those outcomes.

Bibliography

Ken Opalo, <<https://kenopalo.com/tag/jay-ulfelder/>> accessed 30 January 2017.

Ian Bremmer and Keat Preston, *The fat tail: the power of political knowledge for strategic investing*, New York, Oxford University Press, 2009.

Stephen Dubner, <<http://freakonomics.com/podcast/how-to-be-less-terrible-at-predicting-the-future-a-new-freakonomics-radio-podcast/>> accessed 30 January 2017.

Thomas Fingar, *Reducing uncertainty: intelligence analysis and national security*, Stanford, CA, Stanford Security Studies, 2011.

Philip Tetlock, <<http://www.goodjudgment.com/>> accessed 30 January 2017.

IARPA, <<https://www.iarpa.gov/index.php/newsroom/iarpa-in-the-news/2015/439-the-good-judgment-project>> accessed 30 January 2017.

Joint Publication 2-0 Joint Intelligence http://www.dtic.mil/doctrine/new_pubs/jointpub_intelligence.htm accessed 30 January 2017.

Dirk Maclean, *Shoot, don't shoot ; minimising the risk of catastrophic error through High Consequence Decision Making*, RAAF Air Power Development Centre, forthcoming.

David Mandel and Alan Barnes, 'Accuracy of forecasts in strategic intelligence', *Proceedings of the National Academy of Sciences*, 2014.

Henri Poincaré, *Science and method*, New York, Dover Publications, 1952.

Karl Popper, *The Logic of Scientific Discovery*, New York, Basic, 1959.

Thomas Fingar, <http://www.youtube.com/watch?v=CkQvRKRRcLY> accessed 30 January 2017.

Nicholas Taleb, *The black swan: the impact of the highly improbable*, New York, Random House, 2007.

Philip Tetlock, *Expert political judgment: how good is it? How can we know?*

Princeton, NJ, Princeton University Press, 2005.

Philip Tetlock and Dan Gardner, *Superforecasting: the art and science of prediction*, New York, Crown Publishers, 2015.

Jay Ulfelder, <https://dartthrowingchimp.wordpress.com/about/>> accessed 30 January 2017.

¹ See Philip Tetlock, <<http://www.goodjudgment.com/>> accessed 30 January 2017

² David Mandel, D. R. and A. Barnes. "Accuracy Of Forecasts In Strategic Intelligence". Proceedings of the National Academy of Sciences 111.30 (2014): 10984-10989. Web.

³ Thomas Fingar, Reducing uncertainty: intelligence analysis and national security. Stanford, CA, Stanford Security Studies, 2011,

⁴ Thomas Fingar, Reducing uncertainty: intelligence analysis and national security. Stanford, CA, Stanford Security Studies, 2011,

⁵ David Mandel and Alan Barnes. Accuracy of forecasts in strategic intelligence, Proceedings of the National Academy of Sciences, 2014

⁶ Philip Tetlock, Expert political judgment: how good is it? How can we know? Princeton, NJ, Princeton U Press, 2005,

⁷ Nicholas Taleb, The black swan: the impact of the highly improbable, New York, Random House, 2007,

⁸ Henri Poincaré, Science and method, New York, Dover Publications, 1952,

⁹ See Stephen Dubner, <<http://freakonomics.com/podcast/how-to-be-less-terrible-at-predicting-the-future-a-new-freakonomics-radio-podcast/>> accessed 30 January 2017

¹⁰ Karl Popper, The Logic of Scientific Discovery, New York, Basic, 1959,

¹¹ Thomas Fingar. <http://www.youtube.com/watch?v=CkQvRKRRcLY> accessed 30 January 2017

¹² David Mandel and Alan Barnes. Accuracy of forecasts in strategic intelligence, Proceedings of the National Academy of Sciences, 2014

¹³ Jay Ulfelder, <<https://dartthrowingchimp.wordpress.com/about/>> accessed 30 January 2017

¹⁴ Joint Publication 2-0 Joint Intelligence http://www.dtic.mil/doctrine/new_pubs/jointpub_intelligence.htm accessed 30 January 2017

¹⁵ Dirk Maclean, Shoot, don't shoot ; minimising the risk of catastrophic error through High Consequence Decision Making, RAAF Air Power Development Centre, forthcoming,

Contact Us

If you would like further information, please contact:

ACSACS

Telephone: +61 2 6268 8845

Fax: +61 2 6268 8879

Web: www.acsacs.unsw.adfa.edu.au

UNSW Canberra

PO Box 7916, Canberra BC ACT 2610

Cricos Provider Code: 00098G • CMU 170507



UNSW
CANBERRA



ADFA
AUSTRALIAN DEFENCE FORCE ACADEMY



GROUP
OF EIGHT
AUSTRALIA