CHAPTER

7

# The Hardest Question: What Is Going to Happen?

## Prediction and Warning in Analysis

We have been talking about the leader's need to ask the right question and avoid the "wrong" question. One type of question is in a class by itself for analysts: "What is going to happen?" The whole business of prediction is simply variations of that question. The question stands out as the single most difficult part of the analyst's work. Sometimes this is a "right question"; sometimes it's wrong. When it's "right," it taps into the analyst's expert, considered judgment about what is most likely to occur and how the customer can affect the outcome. When it's "wrong," the question is unanswerable but wastes time and encourages both the customer's and the analyst's delusions about our capabilities. When it is "right," it invites us to shine at least a dim light into the darkness that is tomorrow.

Writing *War and Peace* in the 1860s, Leo Tolstoy noted that even yesterday is dim:

> But all these hints at what happened, both from the French side and the Russian, are advanced only because they fit in with the event. Had that event not occurred these hints would have been forgotten, as we have forgotten the thousands and millions of hints and expectations to the contrary which were current then but now have been forgotten because the event falsified them. (Tolstoy 2010: chap. 1)

Tolstoy was tackling the tendency of historians to build a logical, persuasive narrative of cause and effect that completely misses the reality of chaos in great events. He describes how the Russians beat the French in 1812 despite, not because of, what the Russian generals were trying to do. If *history* is so difficult to get right, imagine how much harder is *forecasting*.

The whole business of prediction is a minefield for analysts. We cannot know the future. We can never know whether we have identified all the variables at play in a complex situation, much less weighed them correctly. We can never know when a wild card might be introduced, a game changer that arrives randomly. Also, prediction elicits the expert analyst's deep bias toward continuity; what has been happening for a long time always appears likely to continue, if you understand the drivers and the

drivers are still present. And at best, the analyst can only lay out what is *reasonable*, but we live in a world where the unreasonable, even the ridiculous, happens too frequently. Because of such hazards, our batting average appears low to outsiders. In the bright light of hindsight, our batting can even appear inept.

Because of such hazards, many intelligence leaders whom I respect say flatly, "We don't predict. We don't have a crystal ball." That is true technically—but only technically. We do a range of activities that look forward. We project, we forecast, we warn, we weigh likelihoods, we identify trends. We try hard to figure out what is driving events today in order to anticipate tomorrow. And because our customers think of all this as prediction, we must acknowledge that we are in the prediction business.

What would be nice is clarity as to what is and is not predictable. But no reliable formula is available. And tempting rules of thumb come crashing up against a long list of exceptions. Never predict the outcome of individual battles—but some battles are easy to call. Never predict when more than a dozen variables are at play—except when you can. Never predict an issue when your batting average on that issue is no better than the layperson's—unless your customer asks. So the clarity we need to provide will never amount to a menu we can hand the customer: "Sir, Madam, here is a list of the things we will be able to predict tonight, and might I mention that our election forecast is particularly fresh."

Neither is there a reliable formula for forecasting itself. It is easy to believe that data scientists who have written and tested an algorithm have found just such a formula. That work is enormously valuable, as I discuss in later chapters. But that work finds patterns in the data, and the algorithm presumes the pattern will continue. Saying that "the past pattern will continue" will certainly make any analyst right most of the time, but that rule of thumb leads to failure. Dr. Bruce Bueno de Mesquita (2009) prescribes a game theory approach, identifying the self-interest of each player in a situation and weighing the relative influence of each player, to almost calculate the future. His approach can be applied successfully to some very hard issues but, as he acknowledges, "is right for some problems but not all" (Kindle location 219). My own rule of thumb, which will serve you as well as any, is this: run from anyone who makes forecasting appear formulaic.

Rather, we must analyze each situation to *diagnose what is predictable* and *whether we have something useful to say about the future*. Some situations will require courage to make an unpopular call. Some situations will require us to have the courage to refuse to make a call. All situations will tap our ability to communicate clearly our predictions and their limits. And the test for our work will not be whether we manage to clearly describe an event before it happens. The test will be whether we genuinely help our customer cope with an unknowable future.

## Analyzing Predictability

I know of only a few situations when intelligence analysts are excused from a need to predict. The first year I arrived at CIA, my first boss told me, "Unless we get evidence of a plot, we don't predict assassinations." Three decades and many assassinations later, I never heard a policy customer fault us for this limitation. Prime ministers falling to votes of no confidence in parliamentary governments seem also to be accepted as bolts from the blue. The notion seems to be that if the prime minister, with all his inside knowledge and savvy, was surprised by the event, it would be unreasonable to expect more from us. Nor are we expected to predict the emergence of breakthrough technologies, although we often are expected to project their potential threat applications soon after they emerge.

Most high-stakes situations, however, must be examined more closely to judge predictability. Let's look at one. In September 2012, the Islamic world was ablaze with outrage at an anti-Islam video, *The Innocence of Muslims*, produced in America and posted on the Internet. Was it predictable that jihadists would opportunistically publicize such an offensive video and fan the flames? Absolutely. Was it predictable that more moderate Muslims would be stirred to outrage if they saw the video, even without prodding by the extremists? Yes. Was it predictable that this particular video *could* be a catalyst for violence? You bet.

But was it predictable that it *would* be the catalyst for violence? No. In sorting through possibilities, the difference between *could* and *would* is vast. There is plenty of other anti-Islamic bile on the Internet, posted by racists and other extremists, that gets little attention. Some of it would be considered not simply offensive but blasphemous by pious Muslims. But these other offensive postings did not touch off violence in September 2012. So far, we are unable to predict what goes from obscurity to viral in this context—a problem Internet denizens share in many other contexts. Nor are we able to pick a single catalyst out of a large array of equally possible catalysts for violence. Like the Forest Service does outside my hometown in the summer, we can judge when the fire danger is high, but we cannot predict which spark will touch off the forest fire.

How about predicting elections? One of my most important mentors in leading analysis used to say, "We don't predict elections." Her statement was clear and well-founded but wrong in several situations. I, and the rest of the world, confidently predicted every election Saddam Husayn and Hosni Mubarak ran in through the 1980s and '90s. She'd say, "Well of course, I didn't mean those charades; I meant real elections." But with intensive study and masses of data, Nate Silver (2012) made predictions of stunning accuracy about the US presidential and senatorial elections in 2008 and 2012. He describes his approach in *The Signal and the Noise,* a must-read for anyone in the prediction business. Illustrating just how maddeningly difficult

prediction is, Silver could not replicate his US success with the UK general elections in 2010 (Ball 2013). And though Silver did a creditable job of alerting voters to the fact that Donald Trump had a real chance of winning the 2016 presidential election, few observers actually give him credit for a good forecast.[1] Because he did not say a Trump victory was *likely*, shallow observers consider Silver's analysis wrong. This is a cautionary tale for forecasters. If the forecaster says there is a 1 in 4 chance of something happening and it happens one quarter of the time, the forecaster can justifiably claim accuracy, but the customer will look at the forecast in isolation and say, "You were wrong."

Occasionally, we can predict the outcome of wars or battles. When conventional military forces line up against one another, comfortable judgments can be made about the relative numbers of tanks, experience levels, demonstrated capabilities, and the like. When there is a clear imbalance, a strategic call can be made that might be useful to a policy customer. In the 1982 Lebanon war, for example, the imbalance between the Israeli and Syrian air forces was so clear that it would be no challenge to predict decisive Israeli air dominance. In the event, the Syrians lost more than eighty aircraft, compared to two for the Israelis, according to published accounts (see, e.g., Lambeth 1984: 11). Intelligence analysts shouldn't be expected to predict the score in that air contest, but they should be expected to predict a lopsided Israeli win.

We can also help predict the occurrence of specific events *if we have significant evidence*. Before Operation Desert Storm, for example, my analysts were asked to write a paper predicting what Saddam would do in the first two days of the war. I was livid at the time, arguing that we had already written much about Iraqi capabilities but that to predict precise actions was simply a way to be wrong. I underestimated my analysts. Based on what they could see and their deep knowledge of Saddam's war-fighting habits, they confidently judged, for example, that Iraq would launch Scud missiles against Israel and release oil into the Persian Gulf. Operation Desert Storm kicked off January 17, 1991. Iraq fired seven Scuds at Israel that day and released a massive oil slick into the Gulf the next week (National Guard Bureau 2000).

But in both of these last two examples, there were clear limits to our ability to predict. In the 1982 Lebanon war, for example, the ground balance was more even than the mismatch in air forces. Local terrain in ground battles would be critical, but there would be no way to determine precisely where those ground battles would be fought or who would hold the high ground. Certainly tanks could be counted, and certainly Israeli air superiority over the battlefield would matter. But a significant difference between the Israeli and Syrian tolerance for losses would trump such simple arithmetic. In Desert Storm, we could predict that Saddam would send

[1] Silver's (2016) *FiveThirtyEight* website showed Trump with a 50.1 percent chance of beating Clinton on July 30, 2016, and a 28.6 percent chance on election day.

Scuds into Israel, but we could not even guess where they would land. Whether one landed in a field or on a hospital was nearly random, given the Scud's technology, but that outcome could dictate the Israeli response.

Even in conventional military conflicts, variables pile up quickly, making some predictions a fool's game. Often troop morale or a commander's creativity—the most human of human factors—will be critical. Now we are not only talking about variables but unmeasurable variables of real, yet indeterminate, impact. Even if you know how forces are arrayed on a battlefield, battles can turn on variables that are simply unknowable in advance. One of my favorite old video games is the complex, detailed, and realistic *Sid Meier's Gettysburg*. Individual units would have their morale and effectiveness drained if the battle swept them into a disadvantageous position. I was endlessly fascinated replaying the scenario of Little Round Top, with a different outcome each time—I knew an enormous amount about each side in the contest, but I still could not reliably predict the outcome.

Move to unconventional warfare and predictions become even more challenging. A fascinating account of the 2006 Lebanon war is presented by Joshua Cooper Ramo (2009) in his *The Age of the Unthinkable*. He shows how "fewer than 500 Hizb'allah fighters had frustrated a 30,000-man Israeli attack, including one of the most extensive air campaigns in Middle East history" (187–90). My own confident but unhelpful prediction at the start of that war—unpublished, since I was not responsible for that area at the time—was, "This is going to be messy." Even years after that war, we can still argue about who won; messy indeed.

Policy makers are desperate for help at least framing their expectations before such contests. We can, indeed, help them somewhat, *depending on their receptivity to informed speculation*. One approach they have found instructive, in my experience, is to game scenarios. We might tell a policy customer, for example, "We have gamed this conflict three times, Blue won twice, and casualties ranged from 10,000 to 50,000." Such a perspective helps them think about possible costs, and often uncovers surprising twists, without claiming a prescience we clearly don't have. And games, of course, can be designed not just for military scenarios, but for diplomatic maneuvers, trade negotiations, energy markets, and zombie apocalypses.

This introduces a key concept in the prediction business: *usefulness*. When weighing possibilities about the future, most attention naturally goes to forecasts and whether a forecast turned out to be right. But a more important metric is whether analysis about the future was *useful* to the customer. Did the analysis "narrow the range of uncertainty" for the decision makers?[2] In the war game example I just used, the game results usually aren't

[2] Analytic tradecraft pioneer Jack Davis used this phrase to define the very role of intelligence analysis. He said, "The role of intelligence analysis is to narrow the range of uncertainty for decisions that must be made."

intended to stand as a prediction about the future. The implied forecasts in a game—when Red Commander did X, it prompted Blue Commander to do Y—are far less important than what the actual Blue Commander learns about key dynamics of the situation, the interplay between variables, or the unexpected results of some of his favorite options.

Policy makers frequently ask a question like, "How close is Country X to having a nuclear weapon?" One of my veteran nuclear analysts told me, "Every time I do one of these predictive timelines, the only thing I know for sure is that it's wrong." In the absence of compelling evidence, it would be unreasonable to predict, "Country X will have a nuclear weapon by 2025." That would, indeed, be a prediction and, after 2025, it might be easy to judge whether it was right or wrong. But to be *useful,* we would use a formulation something like, "Based on what we know of their capabilities, the program start date, and the physical realities of enriching uranium, we believe Country X could have everything it needs to assemble a nuclear weapon as early as 2025." With all the proper qualifiers and caveats—all those "weasel words" we must use fearlessly—such assessments help give policy makers a sense of whether they have a year or a decade for their counterproliferation efforts to work. That is useful.

If you look closely at the judgment in the last paragraph, you also see why those senior intelligence officers sometimes are correct when they insist, "We don't predict the future." Look again at the statement, "We believe Country X *could* have *everything it needs* to assemble a nuclear weapon *as early as* 2025." It is a judgment about what is within Country X's reach, not a prediction they will reach it. This far in advance, a judgment about whether Country X *will* reach that goal might amount to informed speculation. If your customer is actively trying to stop Country X's march toward a nuclear weapon, your look forward would not only have to analyze Country X but also consider your customer's counterproliferation effectiveness—including the effectiveness of things your customer has not even decided yet. Speculation would be the best you could offer in this scenario.

Some analogies and categories can help you assess the predictability of the situation your team is examining. For example, you can quickly ask, Is this situation *mechanical?* By that I mean, are major elements of the situation bound by strong if-then relationships? If the price of oil drops by $20 per barrel, then Borostan's annual revenue declines by $4.2 billion. If a small aircraft flying 150 knots disperses 200 gallons of nerve agent X over Washington, DC, and the wind conditions are Y, then this many people would be in the lethal zone. A limiting factor in making such predictions would be whether key variables are known or can be collected. For some customers' purposes, some variables can be reasonably assumed (perhaps the average wind speed and direction over Washington in June is sufficient).

Is the situation like *short-term weather forecasting*? Because we understand weather dynamics very well, and we have excellent evidence about today's weather, some forecasts are easy (the cold front will move in tomorrow morning). Some weather dynamics are chaotic, so I can forecast scattered showers in your area but not how much your particular lawn will receive. Judging likelihoods in some conventional military conflicts has similar elements. The dynamics of some force-on-force contests are well enough understood to confidently predict that one force will prevail. But chaos plays its part in the contest, preventing forecast of casualty rates, especially when battlefields have not even been chosen.

Is the situation like *long-term weather forecasting*? It helps to have a sense of what divides the short term from the long term. For weather forecasting, five days might be the limit of the forecaster's confidence. A forecast looking two weeks out, on the other hand, might be considered long term. That far out, the forecaster won't even be confident whether the temperatures will be in the 40s, as opposed to the 50s, much less which day it will rain. Still, he will know that he is looking at March, so there is little chance temperatures will reach the 80s, and maybe the customer wants to know that. Depending on the customer's needs and how esoteric the topic, framing broad parameters of what is and is not likely in long-term forecasts might be useful.

Is the situation one in which a *pattern prevails*? Patterns are wonderful for forecasters, as long as they hold up. They can be hidden in masses of data. This is a large component of how Nate Silver is able to discern how to convert masses of disparate polling data into a forecast for US presidential and senatorial elections. Or, the pattern might be a correlation of events that experts have noticed in previous situations; perhaps, for example, Waritania's deployment of reconnaissance aircraft to the frontier typically occurs two days before it tests its long-range missiles. This isn't quite mechanical; the aircraft deployment didn't cause the missile test. But anyone closely watching repeated situations will learn to recognize that when X happens, they can expect Y to follow.

Of course, all patterns end eventually. Some day (day?) the sun will not rise. But part of being expert is to (1) learn to recognize patterns and (2) analyze each pattern to determine whether its underpinnings are intact. The first of those comes naturally to experts and can seem effortless. The expert "reads" the situation and, sometimes unconsciously, recognizes something in it and says, "Ah, this again." The second requires more intellectual diligence, and overlooking it is a prime factor at work in many intelligence failures. As leader, you'll need to enforce the diligence. You may be the one requiring your expert to freshly assess whether the key drivers of the pattern remain present in today's situation.

Is the situation one *driven by human characteristics*? This is a particularly challenging arena for forecasting. *Social* sciences are at work and allow

much less certainty than any of the other sciences. You often have the chaos of multiple actors interacting independently. But even with human decision making, we often are able to make useful projections. When we are talking about humans at the macro level, characteristics (behavioral patterns) can be highly relevant and predictive. The expected behavior of a culture, sect, or a demographic slice of society might be relatively predictable if you know the variables that matter.

When we are talking about humans as individuals—a leader, for example—some characteristics run deep and have predictive value. As noted, my analysts predicted that Saddam would fire Scud missiles at Israel and dump oil into the Persian Gulf at the start of Desert Storm. These calls, in part, were based on their understanding of Saddam as a warfighter, spotting characteristics he demonstrated over the course of Iraq's war with Iran. But human complexity (complexity that includes whim and caprice) is such that even a well-understood leader will deliver dramatic surprises sometimes. Saddam's hidden dismantlement of his WMD arsenal certainly did not fit his characteristics as we understood them. Tenacious defiance and deception did, which is why many of his own generals believed to the last that Saddam had retained at least some of those weapons.

Holding your situation up to these models and analogies sometimes will disappoint you. You will consider weather forecasting, patterns, and the rest and conclude, "Well, it's not like any of this." For other situations, you will find there are bits of several paradigms that seem to apply. You'll also notice the paradigms overlap. You find patterns, for example, in weather and in human characteristics. Yes, the world is complicated and the business of forecasting does not always divide neatly into a handful of common paradigms. But as you consider and reject, say, the weather analogy, you and your analyst will come up with other analogies that seem more applicable to the situation at hand. Some will be quite current. Does the progress of ISIS in Syria, for example, shed light on what might be predictable about its progress in Afghanistan? Looking backward at the Syria case, can we find forecasting indicators we missed the first time but can use in the Afghan situation?

The point is to get you started at analyzing the predictability of the situations you are watching and then identifying the elements of those situations that are predictable. I will task my analysts to assess when the cold front will move in and whether there will be scattered showers. I will not ask them to forecast how much rain will hit a particular lawn. This analysis is key to focusing the attention of your analysts on productive forecasting, asking them the "right" questions, and not wasting their time on the unpredictable.

This analysis will also help you guide your analysts in the business of forecasting. Keep in mind that you are not doing this analysis of predictability by yourself. You need to do it with your analysts—they usually know

the situation better than you do. Work with them to determine whether the situation is mechanistic, fits known patterns, is like weather forecasting, or is influenced by known characteristics. Work with them to determine whether other forecasting parallels apply. In the process, they will develop a sense of what is and is not predictable in the situation at hand. They also will develop clarity on what they need to know to make their predictions. For mechanistic issues, what you need to know is clear, and whether it is obtainable is often clear. And once you have what you need, you can make your projections with high confidence. For the mushier world of human decision making, the range of relevant variables is much broader and often less collectable, so high confidence in forecasts rarely is warranted. All of this might seem a bit "meta" to some of your subordinates. In a way, you are working with them to analyze analyzability. But doing this together with them on real issues will help them become skilled in the more general discipline of forecasting.

Your analysis must not stop with analyzing the situation and identifying its predictable elements. You and your team must also analyze your customers' needs relative to the situation. What is useful to them? What they want to know is whether to water their lawn today. They *want* a simple yes or no, and you know you cannot give it to them. But it will be *useful* to tell them they will see scattered showers tomorrow. It will be even more useful to tell them there is a 70 percent chance their lawn will receive rain. Part of analyzing usefulness is to clearly identify which forecasts are *not* useful even though they are correct. Every leader of analysis sees a remarkable number of these. I still hear from analysts whom I supervised decades ago that they look back fondly on my frequent annotation "N.S." in the margin of their drafts. It was short for "no shit," and would be my way of telling the analyst that they were wasting the reader's time when they said something like, "Prime Minister X will take into account the political repercussions before approving this budget measure."

## When History Pivots

What about predicting history's tipping points? Nassim Nicholas Taleb (2007, 2012, 2013) argues persuasively that we should just quit kidding ourselves. In *The Black Swan* and elsewhere, he makes the case that predicting such events is beyond human capabilities. He describes as "black swans" those high-impact events that fit no pattern and have no compelling precedent. (If all you have seen are white swans, you cannot predict the existence of black swans.) He puts the 9/11 attacks, the rise of Hitler, the precipitous collapse of the Soviet bloc, and the rise of Islamic fundamentalism in this category (2007: xviii). Don't let yourselves be fooled by the fact that evidence existed before these events; before the event, that evidence

would not have stood out from the evidence for a myriad of realistic alternatives (Tolstoy's "thousands and millions of hints and expectations to the contrary").

I believe Taleb's insight is sound in several aspects. There are real limits to our ability to predict, and major historical discontinuities exceed the limits of both analyst (to forecast) and audience (to listen). Some big events occur only after virtually random—certainly not *inevitable*—predicates align. Some human decisions, and the dynamic mix of their decisions, cannot be forecast far in advance.

People who expect analysts to provide early forecasts of the pivot points of history sometimes point to a so-called visionary who said something deeply prescient. Usually, those "seers" imagined that something would happen eventually, and—*eventually*—it did. Sometimes, they guess right with stunning accuracy. Each of us has done this once or twice. Sometimes we manage to put our finger on the one overriding truth. But guessing or making a statement of ideological faith is not what we are talking about here. What we are talking about is what is reasonable to expect in the day-to-day business of analysis.

But with acute awareness of our limitations in this arena, let's look at what we can do to be *useful*. Analysts are not helpless in the face of history's tipping points. We can, first, imagine that big change is possible, and second, we can look for whether we have something useful to say. Take, for example, the performance of CIA's intelligence analysts before the collapse of the Soviet Union. They applied their imaginations responsibly, had something useful to say, and said it well before the event.

They did not—*could* not—forecast Soviet collapse ten years before its occurrence. (Others who wish to take credit for identifying the Soviet system of governance as morally bankrupt and economically unsustainable can do so, but who among them in, say, 1985 made a convincing case that it wouldn't at least limp into the twenty-first century?) Gorbachev's profound influence on the collapse of the USSR, and his unique mix of vision and blindness, brilliance and ineptitude, could not be predicted even with his ascendance to power. But CIA's analysts weighed in and, according to their White House customers, mattered.

In his insider's account of the end of the Cold War, *From the Shadows*, Robert Gates (1996) details what CIA said and what the policy customer did about it. Gates, who was then Deputy National Security Advisor, says, "We knew early in the [George H. W.] Bush administration that change was coming fast in the Soviet empire, so fast that we worried about an explosion or widespread instability. Thanks to analysis and warnings from CIA, we at the White House began in the summer of 1989 to think about and prepare for a Soviet collapse" (525). (Remember, the dissolution of the USSR came in December 1991.) In July 1989, crediting "a stream of reporting and assessments I had seen from CIA," Gates recommended to the

president that "we should very quietly begin some contingency planning as to possible U.S. responses, actions and policies in the event of leadership or internal policy changes or widespread ethnic violence and repression" (526). The president approved, and a study group of very senior policy makers was discreetly set up that summer. "This group commissioned a number of studies by CIA and used them in reviewing and planning U.S. options" (526). This is a nice example of analytic projections catalyzing policy action—and the dialogue between policy and intelligence officials helping both sides do their work.[3]

In an important lesson for leaders of analysis, some of this influential analysis was published *before CIA analysts reached consensus on likelihoods*. Gates (1996) mentions being struck by a 1989 paper assessing that "Conditions [in the Soviet Union] are likely to lead in the foreseeable future to continuing crises and instability . . . and perhaps even the localized emergence of parallel centers of power. . . . [Instability would] prevent a return to the arsenal state economy that generated the fundamental military threat to the West in the period since World War II" (514). (Talk about a pivot point in history.) Gates notes, "A number of other analysts in the Soviet office disagreed with the paper, saying it was much too pessimistic. And so it carried a caution to readers that it was 'a speculative paper drafted by a senior analyst'" (514). Gates says, "What was important was that the paper was issued. It made a difference" (515).[4] Many leaders of analysis would advise waiting until a more powerful consensus can be delivered to decision makers, especially on such a high-stakes issue. That hesitance would have been a shame in this event.

Of course, the more distant the future, the more speculative our analysis must be. A good example of speculation about what, to intelligence analysts, is the distant future can be found in the National Intelligence Council's (NIC) 2012 paper, *Global Trends 2030: Alternative Worlds*. It highlights what it calls "megatrends . . . which are virtually certain, exist today, but during the next 15–20 years . . . will gain much greater momentum." It forecasts that China's economy will probably grow larger than that of the United States "a few years before 2030" and that

[3] Another insightful accounting of what CIA published about the failing Soviet Union is by former CIA DDI Douglas MacEachin (2007).

[4] President Bush (1998) might also have been referring to this paper when he wrote about this period: "I found the CIA experts particularly helpful, if pessimistic. One analysis paper concluded that Gorbachev's economic reforms were doomed to failure, and that his political changes were beginning to cause problems he might not be able to control. . . . Based on those conclusions, some people in the NSC began to speculate that Gorbachev might be headed for a crisis which could force him to crack down in the Soviet Union to maintain order, or might even force him out of power" (154).

"regional players such as Colombia, Indonesia, Nigeria, South Africa, and Turkey will become especially important to the global economy" (iv). These judgments sound supremely confident, but they are balanced by the identification of what the authors label "game changers," which are far less predictable but would have enormous impact. For example, the paper notes, "[M]any countries will be zig-zagging their way through the complicated democratization process during the next 15–20 years. Countries moving from autocracy to democracy have a proven track record of instability" (vii). The paper adds, in this context, "China . . . is slated to pass the threshold of US $15,000 per capita purchasing power parity (PPP) in the next five years, which is often a trigger for democratization" (vii). The net effect leaves the reader with the impression that *here are the forces that* will *shape the next two decades, here's what you are* likely *to see, and here's what you* might *see, and in our increasingly globalized world, brace yourself for major events that* cannot *be foreseen.*

To further convey that intelligence analysts cannot "see" the future but do have something useful to offer, the NIC authors adopted the unorthodox use of fictional scenarios. They were saying to the reader, in effect, here is some deep thought about the future rather than an intelligence-based assessment of the future. They said, "We have fictionalized the scenario narratives to encourage all of us to think more creatively about the future. We have intentionally built in discontinuities, which will have a huge impact in inflecting otherwise straight linear projections of known trends." They created four scenarios:

- **Stalled Engines**—a scenario in which the risk of interstate conflict rise owing to a new "great game" in Asia . . . illustrating the most plausible "worst case."
- **Fusion** is . . . what we see as the most plausible "best case." This is a world in which the specter of a spreading conflict in South Asia triggers efforts by the US, Europe, and China to intervene and impose a ceasefire. . . .
- **Gini Out-of-the-bottle** . . . is a world of extremes. Within many countries, inequalities dominate—leading to increasing political and social tensions. Between the countries, there are clear-cut winners and losers.
- **Nonstate World.** In this world, nonstate actors—nongovernmental organizations (NGOs), multinational businesses, academic institutions, and wealthy individuals—as well as subnational units (megacities, for example) flourish and take the lead in confronting global challenges. (NIC 2012: xii–xiv)

I see four tests for such a product. First, does the reader understand clearly that this is informed, rigorous thought about *possible* futures and not a forecast of *the* future? Second, is it written with the customer clearly in mind? By that I mean, does it understand the customer's average expertise, prejudices, and concerns? And third, is it worth the customer's time? As a leader of analysis, it is your job to make sure that the product passes the first two tests. Your answer on the third test, however, will be your own best prediction: *Will* this product prove to be worth the time of a very busy policy maker? To me, a product like the NIC's attempt at "Global Trends 2030" would be appreciated by *some* policy makers *between crises*. Most of the policy makers I know would put such a study on their ever-rising stack of papers to be read later. But the few who spend a weekend with it would have been well served.

The fourth test is not aimed at the customer; it is aimed at the analysts: Does this investment of their attentions make them better analysts? For the NIC project, for example, even if few policy customers spend the time to read the paper, the authors' focus on what the year 2030 might be like deepens their own understanding of the strategic forces at work in the world today. That deeper understanding will enrich their own work on many topics, providing a payoff to the customer in other products down the road. In simplest terms, the analyst grows when forced to wrestle with the customer's recurrent question, "What is going to happen?"

## Humility, an Open Mind, and Practice Required

But are we making a fundamental mistake here? A debate has been going on for many years among intelligence professionals. Some argue that we are kidding ourselves and our customers about the usefulness of our forecasts. With so many limits to our ability to look forward, they say, we should simply refuse to look beyond the immediate future. They say this with particular conviction when considering something like the NIC's attempt to consider the year 2030. When none of the over-the-horizon judgments depend on classified information, why use intelligence resources for such a project? Others take an even more stark view, saying we might even be doing harm when projecting so far ahead. Some of our speculation might be flat wrong but still could trigger a policy decision. For us to weigh in on such issues—worse, for us to *initiate* such analysis—is hubris, according to this argument.

To me, the most important word in that argument is *hubris*. Whether forecasting the immediate or distant future, where analysis does disservice is when it forecasts without humility. Overconfidence tainting forecasts is

by no means unique to intelligence analysis. Both Nate Silver's (2012) *The Signal and the Noise* and Daniel Kahneman's (2011) *Thinking Fast and Slow* are full of anecdotes about bad predictions confidently made.

Too often, on the journey of crafting their judgments, analysts move from a recalcitrant cry—"How can I possibly know what the future holds?"—to supreme confidence—"I have pierced the darkness!" As Heuer (1999), Kahneman (2011), and many others have shown, analysts are human, prone to delusion about our own capabilities, seduced by our own logic. The human brain is wired to look forward, to anticipate, to develop working models to make sense of what is ahead. When the model explains evidence neatly, it takes a firm grip on the brain. And when the model holds up for an extended period, our faith in it grows. The longer a model holds up, the more time we have to form a compelling narrative supporting it and to describe the future the model predicts. There is no more persuasive forecast than one that comes with a confident, compelling narrative (see Taleb 2007, especially chap. 6). The history of intelligence failures is a history of confident forecasts persuasively argued.

It is your job to teach and enforce analytic humility throughout your products. It is your job—concerning the topics for which you are responsible—to analyze carefully what is and is not predictable in the particular situation you are examining. When your analysts are looking forward, it is your job to make sure they never lose sight of their uncertainty, and that they always convey that uncertainty to the customer. And it will be your job to dial back the confidence of your bolder analysts, changing their "almost certainly" to "probably," and their "will" to "might," when they overreach their evidence.

Nothing teaches humility better than examining our track record on prediction. Part of this is teaching your institution's successes and failures. CIA does this fairly well, especially in mining the lessons of our intelligence *failures*. And I am proud that we routinely do this unflinchingly, well before the finger-wagging outsiders weigh in. For the analysts, the story of every failure carries a huge dose of "That could have happened to me!" Significantly, the story of every intelligence *success* comes with a subtext of how close it came to not happening.

But even more powerful than teaching about *our* history of forecasting is insisting that each analyst keep track of *his own* history. Even though forecasting is the single most difficult thing an analyst does, as individual analysts we all presume we are better at it than we are. Having each analyst periodically go back and score the accuracy of his own predictions will improve his next prediction—at least, it should reinforce the analytic humility I am calling for.

An important recent academic study lends support to my call for analytic humility, while supporting the notion that forecasting is not feckless

(Mellers et al. 2015). A team, largely from the University of Pennsylvania, conducted a two-year forecasting tournament under the sponsorship of Intelligence Advanced Research Projects Activity (IARPA). The tournament focused on geopolitical forecasting, which in my experience is the most challenging arena for prediction, requiring the analyst to untangle the motivations and intentions of interacting collections of humans. Moreover, they used real-world topics, "ranging from whether North Korea would test a nuclear device between January 9, 2012, and April 1, 2012, to whether Moody's would downgrade the sovereign debt rating of Greece between October 3, 2011, and November 30, 2011." Participants were free to choose whether or not to make a forecast, so they were not forced into blind guessing. By the time the tournament was complete, they had a significant body of data, including "150,000 forecasts of 743 participants on 199 events."

The team found that participants did significantly better in their forecasts than random guessing. "We developed a profile of the best forecasters; they were better at inductive reasoning, pattern detection, cognitive flexibility, and open-mindedness. They had greater understanding of geopolitics, training in probabilistic reasoning, and opportunities to succeed in cognitively enriched team environments. Last but not least, they viewed forecasting as a skill that required deliberate practice, sustained effort, and constant monitoring of current affairs" (Mellers et al. 2015).

The team's point about open-mindedness is more than just a call for objectivity. They note, "Actively open-minded thinkers . . . have greater tolerance for ambiguity and weaker need for closure," not feeling compelled to rush to conclusion (Mellers et al. 2015). A good forecaster constantly asks what she might be missing, weighs how she might be wrong. Analytic humility.

The team also strongly endorses the need to track one's predictions, and the tournament facilitated this. Participants made an average of 121 predictions and could see how they were doing over the two years. "These conditions enabled a process of learning-by-doing and help to explain why some forecasters achieved far-better-than-chance accuracy" (Mellers et al. 2015). Both practice at predicting and tracking your success make you better at this tricky business.

One of the authors of the study, Philip Tetlock, tells its story and lays out its lessons in the superb book, *Superforecasting*. In studying the most successful of the forecasters in the tournament, those he refers to as "superforecasters" shared some common attributes. Tetlock shows that "superforecasting demands critical thinking that is open-minded, careful, curious, and—above all—self-critical. It also demands focus. The kind of thinking that produces superior judgment does not come effortlessly" (Tetlock and Gardner 2015: 20).

One of many things Tetlock does well in *Superforecasting* is to clarify the distinction between being expert and being a skilled forecaster.[5] Too many critics give the impression that being expert is actually a handicap in forecasting.[6] Far from considering expertise to be a handicap, Tetlock disowns those who used his earlier research for such a position. He says, "The message became 'all expert forecasts are useless,' which is nonsense. . . . My research became a backstop reference for nihilists who see the future as inherently unpredictable and know-nothing populists who insist on preceding 'expert' with 'so-called'" (Tetlock and Gardner 2015: 4). The issue, as Tetlock demonstrates, is not that expertise is a handicap in forecasting. It helps, especially in near-term forecasting. But expertise in a substantive topic and expertise in the craft of forecasting are two different things. Substantive experts can improve their skill in forecasting with training, practice, and attention to the attributes of Tetlock's "superforecasters."

There are some good reasons why substantive expertise does not automatically come with forecasting strength. One might think that any analyst who has become deeply expert would build up skill in the analytic discipline of forecasting. But I worked with many expert analysts whose knowledge was almost exclusively tapped to diagnose or identify "what is going on," rather than to project "what is going to happen." They could go a year or more without ever being asked to project beyond the immediate. Even analysts who work on inherently turbulent regions like the Middle East might forecast less than you would imagine. The very turbulence means that they are often busy analyzing tactical developments and might make only one or two strategic predictions in a year. That experience often is enough to teach them humility but not enough to build forecasting skill.

## Prediction Is Always a Gamble

For all the examples I have given so far, you'll not find simple rules for when to predict, when to speculate, and when to simply decline. This is one of the areas in which the leader of analysis needs to ply her own best analytic skills. The leader, together with the analysts involved, must examine the situation in all its complexity. She must examine the needs of the customer and whether they have something useful to say. Then the leader of analysis must be willing to gamble.

[5] Tetlock had made quite a splash in his 2005 book, *Expert Political Judgment*, noting that many so-called experts did no better than "dart throwing chimps" at forecasting. In describing this more recent tournament to *Harvard Business Review*'s Walter Frick (2015), Tetlock said, "The best forecasters are hovering between the chimp and God."

[6] For example, even such usually savvy observers as Chip and Dan Heath (2013), citing Tetlock's earlier research, say "trust experts about base rates [meaning descriptions about the current situation] but not predictions" (Kindle location 2226).

In his book *The Signal and the Noise,* Nate Silver (2012) provides a very useful section on gambling, especially describing the popular poker game Texas Hold'em. The game requires the player to make predictions—judging the prospects for his own hand and judging the skill of his opponents—and to refine predictions as more data (cards and bets) come in. And the game requires the predictor—the gambler—to put his money where his mouth is. Like predictions in intelligence analysis, the game involves both skill and luck. Some professionals are good enough at the game to make six-figure incomes—Silver was one of them for a while. He says, "Skilled poker players are probably better than 99.9 percent of the population at making reasonably good probabilistic judgments under uncertainty" (Kindle location 5200). As a leader of analysis, you are expected to be the "skilled poker player."

Given everything you know about the prediction game you are in, you will be the one deciding when to fold, check, call, or raise. Use your best analytic skills to judge which are the safe bets, good bets, and silly bets. With your analysts, assess what you know in the situation and what you can reasonably infer:

- Judge the odds in this "hand"—that is, judge the trends, precedent, available data, obtainable data, and the degree of order or chaos in this particular situation.
- With your analysts, form a clear-eyed assessment of whether *what you know* has any diagnostic value in determining what lies ahead.
- Don't be seduced by the stakes on the table; they don't influence the odds of the next card being a winner. That is, just because the situation is really, really important—just because your customer is screaming for knowledge about the future—doesn't affect whether you can make a useful prediction.
- Analyze the "players at the table": Are your analysts prone to overconfidence? What biases have they displayed? What might your customer find useful even if it's not the clear prediction he wants?
- Don't let your last bet drive your next bet. New cards are being dealt—this new data might completely change the picture.
- Like the best gamblers, be a devoted student of the uncertainty game (probabilistic reasoning).
- And, like the best gamblers, study yourself—ignoring matters of luck, what does your track record tell you about your own boldness, biases, and wisdom?

Some leaders of analysis will push back at this notion of gambling. They will say, "If you communicate clearly your uncertainty, no gamble is involved." By this they mean, if your forecast says, "The Israeli prime minister *probably will succeed* in his effort to form a unity coalition," you have left room for uncertainty. "Probably will succeed" logically implies that he might not succeed. The weight of your forecast is with success, but if the prime minister fails, your statement was not wrong.

I think this reasoning is at the same time *important* and *misjudges our customers*. It is important because we always must carefully convey the limits of our certainty when communicating our judgments. There is a significant difference between "probably will" and "will," and we must never apologize for obsessing about such nuance. And our most experienced customers learn to appreciate the care we take with our qualifiers. But let's not kid ourselves. Even our most experienced customers will, after the fact, think we were *wrong* if the Israeli prime minister failed to form that unity coalition. The "probably" and its logical implication will only occur to the customer if he goes back to read our careful forecast, which he rarely has the time to do. Instead, he will judge us by his recollection of our forecast. And his recollection will be that we led him to expect a unity coalition.

Finally, our acceptance of the reality that we are gambling helps us connect better to our policymaking customer. The policy maker is always, inescapably, gambling with the decisions he makes. And he frequently, at least in part, is basing that gamble on what we have told him. It is both arrogant and off-putting for us to say to him, in effect, "We know the stakes are high here, and we know you are going to bet on our advice, but even if you are wrong, you'll find *we* were not wrong."

## Prediction Is the *Leader's* Responsibility

Perhaps because of the dangers and difficulty, many analysts shy away from prediction. At the other end of the spectrum are analysts too quick to make too grand a prediction with too little evidence. Both types of analyst—and those in between—have roles to play in prediction. But the responsibility for prediction falls to the leader of analysis.

There are several things the leader needs to do to get it right—or as wise as we can make it when it is guaranteed frequently to fail.

- Make sure analysts at the two ends of the "boldness spectrum" are talking to each other. Their instincts create a useful balance. Their dialogue can help hone everyone's critical thinking.
- Make sure analysts of different disciplines are talking to each other. The most strategic predictions usually cross analytic

boundaries. Take perhaps our most important prediction: forecasting war. Too many military analysts will say a leader's decision to go to war is ultimately a political decision, and therefore the responsibility for predicting that decision falls to political analysts. Too many political analysts say they can never predict a leader's decision to go to war without the military analysts' thorough understanding of force capabilities and balance. To some extent, both schools are correct, but they will rarely bridge these positions on their own initiative. The leader of analysis must understand the gifts and hesitance of both types of analyst and draw them together.

- Make sure we are not just analyzing the situation, but also analyzing *predictability* in the situation. With your analysts, identify key variables, weigh the diagnosticity of the evidence at hand, and try hard to identify whether you have something useful to say.
- Make sure your analysts understand what we are asking from them. We are asking analysts to offer something useful about the future. We are asking for the best expertise and critical thinking to be applied to a forecast. We are asking our workforce to be professional gamblers. We are not asking for infallibility.

## What's the Worst that Could Happen? Leading Warning

I called prediction the most difficult part of the analyst's work.

One type of prediction, warning, may be the most *important* thing done in intelligence. The identification of a potential threat before it arrives is the reason the United States launched intelligence organizations in the first place. Protecting Americans from deadly threats—that is why they pay us. This expectation does not just apply to intelligence analysts. Any CEO who has invested in an analytic cadre at least hopes they will warn of developments that threaten the company.

Albeit vital, warning is a tricky business. Sometimes, warning is a prediction, with all the elements that make prediction the most challenging thing analysts do. Sometimes it is not a prediction per se but the identification of a vulnerability inherent in the situation. Beyond the challenge of enlightening a customer, which all analysis strives to do, warning tries to stir the customer to urgent action. And it is a tricky business because it is easy for the diligent and imaginative analyst to get caught in a death spiral of warning about endless threats—threats *conceivable* but *not real*.

Warning is best when analysis is least. By that I mean the most effective warnings come with clear and compelling *evidence*. The evidence often speaks for itself, and it is sometimes the analyst's job simply to get out of its way. Quite appropriately, the intelligence collectors get the credit for this type of warning.

But when the evidence is vague or contradictory, the analyst's challenge is significant. Even when the analyst has put the pieces together brilliantly, and is confident about her assessment, she still needs to convince the policy maker. Understandably, the more vague and contradictory the evidence, the more reluctant the customer will be. In response to a warning, the customer's choices often are to do something expensive—and visible to all critics—or to gamble that you are wrong. You affect that calculus by *earning the customer's trust over time*.

Vague and contradictory evidence comes with another burden to analysts. Knowing that we are seized by the urgency of the threat we have identified, it is powerfully tempting to communicate that threat starkly and unequivocally. Sometimes you just want to grab the policy maker by the collar and shake him. But our responsibility to accurately convey our uncertainty, to acknowledge what we don't know, pertains to all analysis, and no exception can be made for the warning arena. Our responsibility is to enlighten but not manipulate the policy maker in this area, as in all our work. Our credibility depends on this, and nowhere is our credibility more important than when we are warning.

Let me illustrate with an example from the 1990s, when we received evidence of a coming action that would jolt US policies in the Kurdish provinces of Iraq. The evidence ran strongly counter to behavior we had seen for years. I can't provide more detail here, but I can tell you I felt sorry for National Security Advisor Anthony Lake, who was going to have to decide whether to trust us.

My analysts ran to me one morning with some intelligence about an ugly imminent threat in the Kurdish provinces. I got the word up to CIA's seventh-floor leadership, and George Tenet said, "Let's get in the car." We headed to Tony Lake's office in the West Wing, to brief him and Mark Parris, the president's special assistant for Near East and South Asia. I had worked with Lake during my stint at the National Security Council, briefed him and Mark several times since, and have enormous respect for both. We laid out the evidence we had, took great pains to explain why we believed it, and predicted what we thought was likely to happen. We also laid out alternative possibilities, taking time to explain why they were less reasonable. Lake accepted the prediction and, with Parris, quickly formulated a plan of action. As we were walking out of Lake's office, Parris leaned over to me and said, "Ok, we're going ahead, but this is all on you guys." I took this to mean that they were running with our call, gambling that we were

right. Our warning was vindicated the next day. Lake's trust in us was also vindicated, and he would likely consider that the next time—but he would never escape the reality that each time is a gamble. I think it helped him in this event that we presented our information and assessment rationally, reasonably, and anticipating his concerns.

It also helped, I believe, that our warning to Lake was not a weekly occurrence. We sounded *this* alarm, but he didn't think of us as alarmists. This would have been a much more difficult gamble for him if he was new to the job and did not have a sense of our abilities. I often wonder whether the Bush administration would have been more receptive to George Tenet's summer 2001 warnings of a major attack by Al Qa'ida had we been another year into the administration (Tenet 2007: chap. 8). They didn't yet have Lake's understanding that when Tenet's hair is on fire, you stop what you are doing and listen carefully.

One of the challenges of warning is that vindication rarely comes as quickly as it did in these examples. Sometimes, if your warning is heeded, vindication never comes at all. There were times, when I was in CTC in 2002 and 2003, when we issued warnings that you heard about because the national threat level publicly was raised. We would brief our evidence and assessment to the policy makers, and they would decide, among other things, whether to adjust the threat level, say from Yellow to Orange. Our warning was based on intelligence that pointed to an imminent attack, but the attack never came. Was our evidence incorrect? Was our analysis bogus? Or was the attack deterred? I drafted some of the public threat announcements myself and, even in retrospect, consider them sound. But to this day I do not know for sure. Eventually, the system of nationwide threat warnings was replaced with warnings more specific to geography or sector. But that wasn't because we were taught to do our jobs differently. Rather, I saw the refined approach as the policy makers' willingness to gamble a bit more that we probably didn't need people in Des Moines to react to a threat we thought might be aimed at Los Angeles.

In these counterterrorism examples, after a time at an elevated threat level, the policy makers would ask whether the threat had passed. This was a completely reasonable question. It was our warning that had triggered the alert, after all. But, as frequently is the case with many types of threat, we had received *evidence of a threat* but, even after weeks, we had no further evidence—certainly *no evidence the threat had gone away*. Ideally, the alert (and the myriad actions triggered by the alert) averted the threat, but how could we know? How could we be confident that the terrorists were not simply waiting for the security forces to stand down? Intelligence owns the alarm bell but rarely owns the all-clear whistle. And the same policy maker who asked whether we are safe now is entitled to wonder whether the threat was real at all. If you install high security locks on your doors and

windows because of a rash of burglaries in the neighborhood and you are burglary-free after that, was it because the locks did their job, or did you waste your money? Even without an answer, you leave the locks in place because you already paid for them. In the alert business, however, leaving extra security in place comes at a steep cost, including the tax on public patience.

Both we and our customers would do well to understand that an intelligence community tuned to warn of threats will not be strong at recognizing the expiry of threats. Many analysts were slow to acknowledge the collapse of the Soviet threat. We didn't know that Iraq had secretly destroyed its WMD arsenal. And you can bet we will struggle to declare the ISIS threat dead even if years pass with no attack. Knowing this is an inherent weakness, leaders of analysis should compensate however they can. They might, for example, periodically assign devil's advocate analysts to make the case that a threat has passed, and see how that case looks to fresh eyes.

You should never lose sight of how irritating this can all be for your customer. It is easier for us to warn than it is for him to act, and it is easier for us to stay "on alert" indefinitely because that is our job. Some customers suspect that we often warn simply to cover our asses. "If some unlikely bad thing happens," they complain, "you'll be able to say it wasn't your fault." If you were the one who had to decide whether to cancel shore leave and move the fleet, you'd be cranky, too.

As frustrated as they get being forced to gamble one way or another when they have heard our warnings, they get even more frustrated when we warn repeatedly of something they feel helpless to address. They feel nagged. "I heard you, dammit . . . what do you expect me to do about it?" This is not simply a *cri de coeur*; this is a legitimate question from your customer. We owe it to him to provide our clearest thinking about the minimum level of action that might avert the threat. In the directorate of analysis, we call this "opportunity analysis." It should be a required part of warning.

Finally, a word about threshold. I have referred to the policy makers' gambles once they receive our warning. Again, as the leader of analysis, you must make a gamble as well. You have to decide what threshold to set for warning in the first place. Any good analyst can imagine ten threats that are conceivable-but-remote possibilities. If you hand all ten warnings to your customer, they will tune you out. But the customer doesn't want to be surprised either, so do you pick the most likely two out of the ten? Three? There is no formula for getting this right. You will always be gambling about when to warn and when to just chill.

But I can give you one loud piece of advice and one rule of thumb. The advice is talk to your customer about your warning threshold. If he wants you on a hair trigger for warning, ok. If he wants you to dial it back,

fair enough. If he wants you to lean forward on threats to his latest policy initiative but lean backward on another area, you can fine-tune that, too. And check back with him periodically, to make certain you are in touch with his current appetite. But whatever guidance he provides is just that: guidance. His guidance doesn't relieve you of your responsibility to give him the warning he *needs*. You have his back regardless of his appetite. To meet that responsibility, I had success with this rule of thumb: I don't scare easily, but when *I* am scared, I'll do my best to make sure my customer is, too. I won't—I cannot—hype the threat. But I will do my best to communicate the reasons for my concern.

## What Is Reasonable in an Unreasonable World?

Let me end this section with a caution for our customers and ourselves. As hard as we work to peer into the future, as diligent as we can be with both prediction and warning, there is an inherent limit to our success. Analysis can only deliver an assessment of what is reasonable, but the world is not always reasonable. There will always be a degree of randomness in human affairs. There will always be irrational impulses at least tempting—sometimes driving—leaders.

Can we allow for this randomness in our forecasts? Yes and no. No, we cannot write a compelling piece of analysis that requires one or several actors to do something illogical unless there is evidence for it. Yes, we can be humble in making our predictions. We can remind ourselves and our readers that, in the play of events, dynamic interaction often matters more than intent and capability. We can remember that, despite a vast difference in abilities, tonight the worst team in the NBA might beat the best.

When we do look forward with humility and communicate our uncertainty with clarity, I believe we can bring useful insight to the customer about what tomorrow might hold. When weighing the *near* future, when the customer must make decisions, those decisions are likely to be wiser if informed by the best thinking of expert analysts. The farther away the future we are considering, the less expert our analysts are and the more speculative our projection must be. For a distant future, perhaps it would be useful to remind the customer that we are not *forecasting* but *considering* the future, and we are inviting the customer to share our thoughts. On a case-by-case basis, we will have to make hard-nosed decisions about whether to invest analytic resources on issues when informed speculation is the best we can offer. And the customer will make a similar resource decision about whether to invest his time reading it. If he does, the customer will be smarter about the future than he was before.

## KEY THEMES

Forecasting is the analysts' most difficult function. It is also usually the most important. Because our customers can only affect the future, that is where the analysts' interest must be. The future is uncertain, but we work to narrow the range of uncertainty for decision makers.

Run from anyone who makes forecasting sound formulaic. Where forecasting can be easily modeled (Should blackjack players hit on 17?) is not where forecasters struggle. In more complex situations, leaders of analysis must figure out not which formula to use but what we have to say that would be *useful* to our customers.

With your analysts, analyze *predictability* in the situation you are watching. How confident we can be looking forward will depend on the situation:

- Is the situation *mechanical* (meaning there are clear if-then relationships, with strong evidence, and deductive reasoning applies)?
- Is the situation like *short-term weather forecasting* (with well-understood dynamics and abundant evidence)?
- Is the situation like *long-term weather forecasting* (clear historical trends but plenty of exceptions)?
- Is the situation one where *long-term patterns* are evident? Patterns are great and hold up most of the time. But they all end someday.
- Is the situation *dominated by human characteristics*? Humans are complex—individual humans even more so. You may be obliged to forecast what an individual actor is likely to do in a high-stakes situation even before he has decided for himself. Confidence is rarely warranted here, so lower yours and work hard to determine what your analysts can say that is useful.
- Is the situation like none of these? That tells you something.

Humility is required in forecasting.

- Never fall in love with your forecast. Actively look for contrary evidence.
- Restrain your overconfident analysts.
- Work hard to choose the appropriate term of uncertainty ("probably," "may," "a slim chance") and to preserve those terms through the editing process.
- Forecasting is always a gamble (and you won't win them all). The leader of analysis bears the responsibility in this dangerous game.

Warning—flagging threats or vulnerabilities—is arguably the most important type of forecast. ("What do I pay you guys for if you don't warn me of imminent dangers?!"—any CEO who has an analytic unit.) The most effective warning starts with a relationship with your customer:

- It is based on credibility built over time.
- It is sensitive to the customer's appetite for alerts and tolerance for false alarms.

To be effective, warning must be

- Credible
- Actionable
- Heard