

Sharpening Your Forecasting Skills

Foresight Is a Measurable Skill That You Can Cultivate

September 28, 2015

Authors

Michael J. Mauboussin

michael.mauboussin@credit-suisse.com

Dan Callahan, CFA

daniel.callahan@credit-suisse.com



“Beliefs are hypotheses to be tested, not treasures to be protected.”

Philip E. Tetlock and Dan Gardner¹

- Philip Tetlock’s study of hundreds of experts making thousands of predictions over two decades found that the average prediction was “little better than guessing.” That’s the bad news.
- Tetlock, along with his colleagues, participated in a forecasting tournament sponsored by the U.S. intelligence community. That work identified “superforecasters,” people who consistently make superior predictions. That’s the good news.
- The key to superforecasters is how they think. They are actively open-minded, intellectually humble, numerate, thoughtful updaters, and hard working.
- Superforecasters achieve better results when they are part of a team. But since there are pros and cons to working in teams, training is essential.
- Instruction in methods to reduce bias in forecasts improves outcomes. There must be a close link between training and implementation.
- The best leaders recognize that proper, even bold, action requires good thinking.

Introduction: The Bad News and the Good News

What if you had the opportunity to learn how to improve the quality of your forecasts, measured as the distance between forecasts and outcomes, by 60 percent? Interested? *Superforecasting: The Art and Science of Prediction* by Philip Tetlock and Dan Gardner is a book that shows how a small number of “superforecasters” achieved that level of skill. If you are in the forecasting business—which is likely if you’re reading this—you should take a moment to buy it now. You’ll find that it’s a rare book that is both grounded in science and highly practical.

Phil Tetlock is a professor of psychology and political science at the University of Pennsylvania who has spent decades studying the predictions of experts. Specifically, he enticed 284 experts to make more than 27,000 predictions on political, social, and economic outcomes over a 21-year span ended in 2004. The period included six presidential elections and three wars. These forecasters had crack credentials, including more than a dozen years of relevant work experience and lots of advanced degrees—nearly all had postgraduate training and half had PhDs.

Tetlock then did something very unusual. He kept track of their predictions. The results, summarized in his book *Expert Political Judgment*, were not encouraging.² The predictions of the average expert were “little better than guessing,” which is a polite way to say that “they were roughly as accurate as a dart-throwing chimpanzee.” When confronted with the evidence of their futility, the experts did what the rest of us do: they put up their psychological defense shields. They noted that they almost called it right, or that their prediction carried so much weight that it affected the outcome, or that they were correct about the prediction but simply off on timing. Overall, Tetlock’s results provide lethal ammunition for those who debunk the value of experts.

Below the headline of expert ineffectiveness were some more subtle findings. One was an inverse correlation between fame and accuracy. While famous experts had among the worst records of prediction, they demonstrated “skill at telling a compelling story.” To gain fame it helps to tell “tight, simple, clear stories that grab and hold audiences.” These pundits are often wrong but never in doubt.

Another result, which is related to the first, was that what mattered in the quality of predictions was less *what* the expert thought and more *how* he or she thought. Tetlock categorized his experts as foxes or hedgehogs based on a famous essay on thinking styles by the philosopher Isaiah Berlin. Foxes know a little about a lot of things, and hedgehogs know one big thing. Foxes did better than the dart-throwing chimp, and hedgehogs did worse.

It’s not hard to see the link between these findings. Most topics of interest in the economic, social, and political realms defy tight, simple, and clear stories. But imagine you are the producer of a television show that covers politics. Who do you want to put on the air, the equivocal guest who constantly says “on the other hand,” or the one who confidently tells a crisp and controversial story? It’s not a hard decision, which is why many hedgehogs are both famous and poor predictors.

While the conclusions of *Expert Political Judgment* were nuanced, they were on balance bad news for pundits. Despite how some read his results, Tetlock never believed in the extreme point of view that forecasts are useless. That foxes were better forecasters than the average of all experts provided a strong clue that foresight might be a real skill that could be identified and cultivated. Tetlock marked himself as an “optimistic skeptic.”

Expert Political Judgment is excellent scholarly research but is written in, well, scholarly prose. In *Superforecasting*, Tetlock collaborates with Dan Gardner, a journalist and author of a book about the failure of prediction. The result is great research that is easy to read.

Naturally, Tetlock is not the only one interested in learning how to make effective forecasts. The United States intelligence community was also keen to improve the quality of predictions, especially in the wake of the failure to anticipate the terrorist acts on September 11, 2001 and the overestimation of the probability of the existence of weapons of mass destruction in Iraq in 2003. An agency within the community, Intelligence Advanced Research Projects Activity (IARPA), was assembled to pursue high-risk research into how to improve American intelligence. IARPA decided to create a forecasting tournament to see if there might be a way to sharpen forecasts.

Tetlock and some colleagues launched the Good Judgment Project (GJP), one of five scientific teams that would compete to answer questions accurately. The teams could use whatever approaches they wanted to generate the best possible answers. Starting in September 2011, IARPA asked nearly 500 questions about various political and economic outcomes. The tournament garnered more than one million individual forecasts in the following four years. It is important to note that the time frames for the questions in the IARPA tournament, generally one month to one year, were shorter than the three to five years that were common in Tetlock's study of experts.

Now the good news: the GJP results beat the control group by 60 percent in year one. Results in year two were even better, trouncing the control group by almost 80 percent. In fact, the GJP did so well that IARPA dropped the other teams.

Of the 2,800 GJP volunteers in the first year of the tournament, the top 2 percent were called "superforecasters." To give you some sense of their acuity, the superforecasters performed about 30 percent better than the average for the intelligence community—people who had access to classified data—according to an editor at the *Washington Post*.³

Encouraged by the GJP's results, Tetlock came to a couple conclusions. The first is that foresight is a real and measurable skill. One test of skill is persistence. High persistence means that you do consistently well over time and are not a one-hit wonder. About 70 percent of superforecasters remain in those elite ranks from one year to the next, vastly more than what chance would dictate.

The second is that foresight "is the product of particular ways of thinking, of gathering information, of updating beliefs." Importantly, the essential ingredients of being a superforecaster can be learned and cultivated. The beauty of the GJP is that it was carried out with scientific rigor, which allowed the researchers to distill the elements of success. We explore these elements in this report.

Even though most people can improve their thinking and forecasting, there has always been resistance to change based on what Tetlock and Gardner call "illusions of knowledge." Intuition is one example. Intuition is a form of pattern recognition that works in settings with lots of "valid cues."⁴ But intuition is notoriously unreliable in unstable or nonlinear environments. An overreliance on intuition leads to poor decisions.

Another case is insufficient self-reflection. This is in part prompted by a module in our brain that seeks to rapidly close cause-and-effect loops. We show you an outcome and your mind quickly comes up with an explanation for it. As Tetlock and Gardner write, "we move too fast from confusion and uncertainty to a clear and confident conclusion without spending any time in between." This is related to the concept that Daniel Kahneman, an eminent psychologist, calls thinking fast.⁵

Keeping track of forecasts and the outcomes may not serve an expert's interests. If you are paid well to be a pundit, introducing a scoring system offers little upside and lots of downside. In extreme cases, experts are so sure that they are correct that they see no need to measure outcomes at all.

Tetlock and Gardner share the words of Galen, the physician to Roman emperors, who practiced in the second century. Of a particular cure he wrote, "All who drink of this treatment recover in a short time, except those whom it does not help, who all die. It is obvious, therefore, that it fails only in incurable cases." Fortunately, medical researchers have applied the scientific method more rigorously in recent centuries, but it's still easy to spot an overconfident, and unchecked, expert.

So what is the source of good forecasting? Tetlock and his colleagues found four drivers behind the success of the superforecasters:⁶

- **Find the right people.** You get a 10-15 percent boost from screening forecasters on fluid intelligence and active open-mindedness.
- **Manage interaction.** You get a 10-20 percent enhancement by allowing the forecasters to work collaboratively in teams or competitively in prediction markets.
- **Train effectively.** Cognitive debiasing exercises lift results by 10 percent.
- **Overweight elite forecasters or extremize estimates.** Results improve by 15-30 percent if you give more weight to better forecasters and make forecasts more extreme to compensate for the conservatism of forecasts.

The scientists measure these improvements using a Brier score (the appendix provides more detail on the calculation). A Brier score reflects the difference between a forecast and the outcome. Like golf scores, lower is better. There are a couple of ways to calculate Brier scores, but a common scale runs from zero to 2.0. Zero means that the forecast is spot on, 0.50 is a random forecast, and 2.0 means that the forecast is completely wrong.

By this scoring, a person who predicts a 55 percent probability of an outcome that happens receives a Brier score of 0.405. A subsequent forecast of a 65 percent probability of an event that occurs gets a Brier score of 0.245, nearly a 40 percent improvement.

There are a lot of details in finding the right people, effectively building and managing teams, and proper training. But Tetlock and Gardner offer a simple formula that is at the core of the whole process: "Forecast, measure, and revise: it is the surest path to seeing better."

How can we all become more like superforecasters? To answer that question, we discuss the profile of a superforecaster, the tools you will need to think like a superforecaster, the lessons for leadership, and the valid doubts that remain.

Find the Right People

Throughout the book, Tetlock and Gardner provide the pieces that allow you to construct the profile of a superforecaster. Because the GJP researchers took the time to run the forecasters through a battery of psychological tests, they were able to examine the personalities of the superforecasters. Further, the data

allow the researchers to avoid the error of first observing success and then attempting to find common attributes after the fact.⁷ The portrait of a modal superforecaster has four elements:

Philosophical Outlook. Superforecasters tend to be comfortable with a sense of doubt. Scientists sometimes sense that they know the truth. Good thinkers can feel the same way. “But they know they must set that feeling aside and replace it with finely measured degrees of doubt,” write Tetlock and Gardner, “—doubt that can be reduced (although never to zero) by better evidence from better studies.”

Recall that our minds are keen to assign causality. We want the case to be closed. But as Daniel Kahneman says, “It is wise to take admissions of uncertainty seriously, but declarations of high confidence mainly tell you that an individual has constructed a coherent story in his mind, not necessarily that the story is true.”⁸

Superforecasters are also humble, but not in the sense of feeling unworthy. Rather, their humility comes from the recognition that reality is profoundly complex. Indeed, it is possible to think highly of yourself and to be intellectually humble at the same time. Tetlock and Gardner note that, “Intellectual humility compels the careful reflection necessary for good judgment; confidence in one’s abilities inspires determined action.”

It is common, and often soothing, to attribute outcomes to fate. Superforecasters aren’t big believers in fate. On a one to nine “fate score,” where one is a total rejection of fate and nine is complete belief in it, the average adult American falls near the middle. The mean score for a student at the University of Pennsylvania is a little lower, the regular forecasters are below that, and the superforecasters are the lowest of these groups. Superforecasters don’t think that what happened had to happen.

Ability and Thinking Style. The first point is that superforecasters are not geniuses. The researchers tested the fluid and crystallized intelligence of all the GJP volunteers. Fluid intelligence is the ability to think logically and to solve novel problems. It doesn’t rely on accumulated knowledge. Crystallized intelligence is exactly what it sounds like: your collection of skills, facts, and wisdom, and your ability to use them when you need to.

Those who participated in the GJP were not a valid sample of the population—these are people who raised their hand to make lots of forecasts in return for a \$250 gift certificate from Amazon.com. The regular forecasters scored higher than about 70 percent of the population on intelligence tests. That translates roughly into an average intelligence quotient (IQ) of 108-110 where the average of the population is 100. The superforecasters scored higher than about 80 percent of the population, or an average IQ range of 112-114. There is a much bigger gap between the overall population and regular forecasters than there is between those forecasters and the superforecasters.

Keith Stanovich, professor emeritus of applied psychology and human development at the University of Toronto, distinguishes between IQ and what he calls “RQ,” or rationality quotient.⁹ The correlation coefficient between the two is a relatively low .20 to .35. Those with high RQ’s exhibit adaptive behavioral acts, efficient behavioral regulation, sensible goal prioritization, reflectivity, and the proper treatment of evidence. These qualities are very consistent with those of the superforecasters.

Jonathan Baron, a professor of psychology at the University of Pennsylvania and a colleague of Tetlock’s, coined the term “active open-mindedness.”¹⁰ Those who are actively open-minded seek views that are different than their own and consider them carefully. Tetlock and Gardner suggest that if they had to reduce superforecasting to a bumper sticker, it would read, “Beliefs are hypotheses to be tested, not treasures to be guarded.”

The Big Five is one of the most widely-accepted personality tests. Subjects are tested for five personality traits: openness to experience, conscientiousness, extraversion, agreeableness, and neuroticism. The tests revealed that superforecasters score high in openness to experience, which suggests a preference for cognitive variety and intellectual curiosity. Superforecasters are interested in the world and are willing explorers.

Superforecasters also spend time thinking about their own process and constantly seek to improve. When collaborating, the superforecasters often leave lots of comments in their online discussions, which allow them to recreate their thought processes and improve them when possible. Timely and accurate feedback is an essential element of improvement. Superforecasters embrace feedback.

The superforecasters rarely use sophisticated mathematical models to make their forecasts, but they are uniformly highly numerate. Comfort with numbers is a prerequisite for making good forecasts but fancy quantitative models are not.

Methods of Forecasting. Superforecasters, similar to the foxes in Tetlock's study of expert political judgment, tend to be pragmatic in their methods. Rather than looking at all aspects of the world through a single lens, good forecasting requires considering multiple points of view. Charles Munger, vice chairman of Berkshire Hathaway, captures this concept well with the mental models approach. Says Munger, "Well, the first rule is that you've got to have *multiple* models—because if you just have one or two that you're using, the nature of human psychology is such that you'll *torture* reality so that it fits your models, or at least you'll think it does."¹¹

Kahneman has popularized the notion of two systems of the mind. System 1 is fast, automatic, and difficult to train. System 2 is slow, deliberate, and purposeful. One of the important conclusions from this research is that we commonly rely on our fast system when we should recruit our slow system.¹² Tetlock and Gardner call the fast system the "tip-of-your-nose perspective," because it is unique to each of us. Superforecasters have a firm sense of when they need to engage System 2.

There is good evidence that the aggregation of diverse points of view, done correctly, improves the accuracy of forecasts. James Surowiecki provides ample illustrations of this idea in his bestselling book, *The Wisdom of Crowds*.¹³ You can gain from diversity by capturing the views of different individuals, for example, investors in a stock market, or by aggregating multiple views in your head.

Tetlock and Gardner use the metaphor of a dragonfly's eye. Each eye has up to 30,000 individual lenses aimed in slightly different directions that provide the dragonfly's brain with massive input. The result is extraordinary visual acuity, allowing the dragonfly to nab small, fast-moving insects.

The authors find that superforecasters, similar to the dragonfly, are able to consider and synthesize multiple points of view. They also emphasize that "aggregation doesn't come to us naturally." We are generally content with our own beliefs and see no reason to entertain alternative thoughts. Sometimes the answer is to survey the views of others and ask them to criticize your view. Other times you can simply think about the same topic at different times and create a crowd within your head. No matter how you get there, taking in the views of others is valuable.

Good forecasters think in probabilities, but there's nothing easy about that. Tetlock and Gardner propose that we come out of the factory with three settings on our dial of probability: it's going to happen; it's not going to happen; and maybe. They suggest this worked fine for our ancestors. "Is that a lion? YES = run! MAYBE = stay alert! NO = relax."

Forecasters commonly use 50 percent to represent “maybe,” and sure enough the forecasters who used 50 percent the most frequently were less accurate than the average.

Superforecasters provide more finely detailed forecasts than the other forecasters. Instead of 60, 70, 80 percent they are more likely to use 70, 75, 80 percent, or even 70, 71, 72 percent. This precision was not for show: the more granular forecasts were more accurate than the less granular ones.

Philosophers suggest a distinction between cases where you don't know the outcome but the possibilities are knowable—the roll of a die for instance—and cases where you don't know and the alternatives are unknowable.¹⁴ Superforecasters recognize that the cloudier the outlook for the answer, the more it benefits them to stay in the zone of “maybe.”

We will come back to probability when we discuss tools, but it is again worth underscoring the value of feedback. One of the surest ways to improve calibration, the alignment of subjective and objective probabilities, is through timely and accurate feedback. Part of the reason that weather forecasters are better than financial forecasters is that they quickly see whether their predictions are accurate.¹⁵

We live in a dynamic world, so new information that should change our assessments of probabilities arrives all the time. Superforecasters have more accurate initial predictions on average than the regular forecasters do, but they also update their views more often. Such updating requires an open mind, but also comes with the risk of under- or overreacting.

Tetlock and Gardner offer three reasons that forecasters underreact to new information. To start, sometimes we are so busy that novel information merely slips our attention. We also may take our eye off of the original question and dwell on a simpler or slightly different one. So the new information may not appear relevant to the question in our minds even though it is relevant to the question at hand. Finally, and probably most likely, is belief perseverance. This is typically accompanied by confirmation bias—actively seeking information that supports our view and dismissing views counter to it.¹⁶

But it's also possible to overreact to new information. One reason is that we take into account irrelevant information. People may base their initial estimate on solid reasoning, but subsequently place weight on additional information that has no bearing on the issue at hand. A second reason for overreaction is a lack of commitment. Say you buy a mutual fund after having done research convincing yourself that the portfolio manager is skillful and that her strategy will do well over time. If you sell at the first bump in performance, you are showing a lack of commitment.

Use of Bayes's Theorem is a formal way to update probabilities, but it turns out that the superforecasters generally don't use it. This is true even for those steeped in the math of the theorem. There is no easy way to correctly update views, but we know that superforecasters spend a lot of time thinking about how to do it well.

We think intuitively most of the time. Our intuitive system tends to use rules of thumb, or heuristics. For example, the availability heuristic suggests that we deem events that are easy to remember, because of vividness or recency, to be more numerous than events of equal frequency that are not as simple to recall. Other heuristics include representativeness and anchoring.¹⁷

Heuristics are wonderful because they save us a great deal of time. But they also come with biases that can undermine the quality of our judgments. For example, people are unjustifiably more fearful of flying after having heard of an aviation accident.

Superforecasters have an above-average awareness of these biases and try to manage them. Both the thinking styles and forecasting methods of superforecasters help address bias. The melding of humans and computers is another path to improved decisions. Done correctly, the output of man plus machine can exceed man or machine in certain domains.

David Ferrucci is an artificial intelligence expert who was in charge of creating Watson, the computer that beat champions Ken Jennings and Brad Rutter in the game show *Jeopardy!* He sees a role for human judgment even with the rise of the robots. But he also believes that computers can help overcome human bias. “So what I want is that human expert paired with a computer,” said Ferrucci, “to overcome the human cognitive limitations and biases.”

Work Ethic. Carol Dweck, a professor of psychology at Stanford University, is best known for her work on mindset, or a way of thinking. She suggests individuals can be placed on a continuum, with a “fixed mindset” at one extreme and a “growth mindset” at the other, based on their implicit belief about the source of ability. People with a fixed mindset believe that ability is innate and therefore can’t be budged. People who say, “I’m just bad at math,” have a fixed mindset. Those with a growth mindset believe that ability is the result of hard work and effort and therefore can improve over time.

Superforecasters fall on the growth mindset side of the continuum. They believe that there is always room for improvement and seek ways to get better. In discussing the results of one of her experiments, Dweck noted, “Only people with a growth mindset paid close attention to information that could stretch their knowledge. Only for them was learning a priority.”¹⁸

Another quality that the superforecasters have is grit, a term popularized by Angela Duckworth, another one of Tetlock’s colleagues at the University of Pennsylvania.¹⁹ Grit is perseverance in the service of long-term goals. It entails the ability to overcome failure and obstacles along the way of achieving an objective.

Combine a growth mindset and grit and you have an outstanding formula for personal development and improvement. Tetlock and Gardner call the combination “perpetual beta.” A product in beta is nearly complete but has room for improvement. Perpetual beta suggests a desire for ongoing improvement. Exhibit 1 summarizes the composite portrait of a superforecaster.

Exhibit 1: Composite Portrait of the Modal Superforecaster

- **Philosophic Outlook**
 - **Cautious:** Nothing is certain
 - **Humble:** Reality is infinitely complex
 - **Nondeterministic:** What happens is not meant to be and does not have to happen
- **Abilities and Thinking Styles**
 - **Actively open-minded:** Beliefs are hypotheses to be tested, not treasures to be protected
 - **Intelligent and knowledgeable, with a “need for cognition”:** Intellectually curious, enjoy puzzles and mental challenges
 - **Reflective:** Introspective and self-critical
 - **Numerate:** Comfortable with numbers
- **Methods of Forecasting**
 - **Pragmatic:** Not wedded to any idea or agenda
 - **Analytical:** Capable of stepping back from the tip-of-your-nose perspective and considering other views
 - **Dragonfly-eyed:** Value diverse views and synthesize them into your own
 - **Probabilistic:** Judge using many grades of maybe
 - **Thoughtful updaters:** When facts change, they change their minds
 - **Good intuitive psychologists:** Aware of the value of checking thinking for cognitive and emotional biases
- **Work Ethic**
 - **A growth mindset:** Believe it’s possible to get better
 - **Grit:** Determined to keep at it however long it takes

Source: Philip E. Tetlock and Dan Gardner, *Superforecasting: The Art and Science of Prediction* (New York: Crown Publishers, 2015), 191-192. Used by permission.

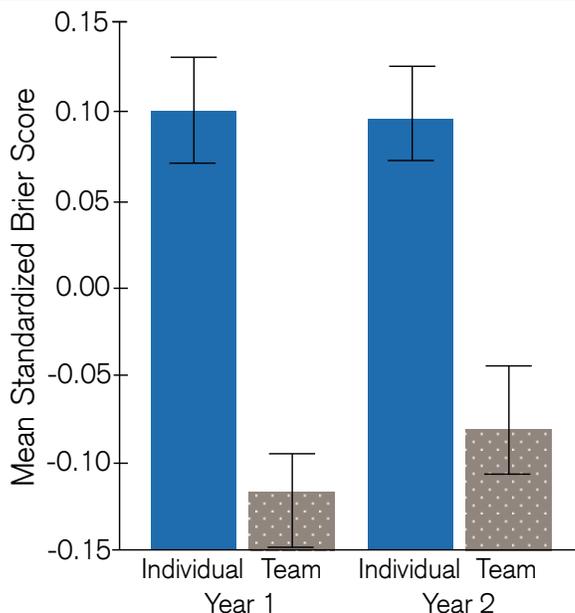
These elements of a superforecaster may be valuable for adding structure to hiring processes and performance evaluation. Further, leaders of organizations in the forecasting business should give careful consideration to creating an environment conducive to good judgment.

Manage Interaction

The objective of the IARPA tournament was to make accurate forecasts, and the GJP had already proven that it could do that. The next question was whether working in teams would improve accuracy. To find out, in year one the GJP researchers randomly assigned some forecasters to work in teams and provided them with tips on how to work together effectively. Others were to work alone. This was before the scientists had identified the superforecasters.

The results were clear: teams were on average 23 percent more accurate than individuals. Exhibit 2 shows the statistical results for the first two years. Recall that a lower Brier score is better than a higher one.

Exhibit 2: Forecasting Teams Outperformed Individuals



Source: Based on Barbara Mellers, Lyle Ungar, Jonathan Baron, Jaime Ramos, Burcu Gurcay, Katrina Fincher, Sydney E. Scott, Don Moore, Pavel Atanasov, Samuel A. Swift, Terry Murray, Eric Stone, and Philip E. Tetlock, "Psychological Strategies for Winning a Geopolitical Forecasting Tournament," *Psychological Science*, Vol. 25, No. 5, May 2014, 1106-1115.

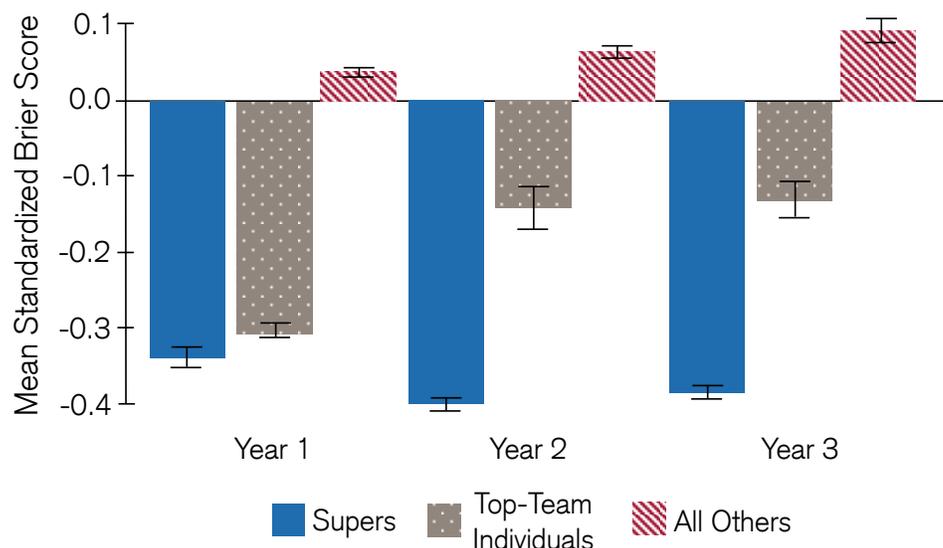
Note: Error bars represent plus and minus two standard errors.

Most organizations use teams to get their work done. In some sectors, including investment management, the trend shows an increasing use of teams.²⁰ But the research also reveals that there are pros and cons to working in teams. The pros are that individuals can share information, and aggregation tends to lead to more accurate forecasts. The cons are that team members may be tempted to loaf, and the team may fall into groupthink, failing to capture the value of cognitive diversity.

In year two, the researchers created teams of superforecasters. Each team had a dozen members, but about five or six individuals did most of the work. Once again, the teams were trained on how to work together, and because the members did not meet face-to-face the project coordinators created forums for them to communicate.

The outcome was remarkable. Those who were good enough to achieve superforecaster status in year 1 were 50 percent more accurate, on average, in year 2 as part of a team. Year three was more of the same. (See Exhibit 3.) Indeed, these superteams were 15 to 30 percent better than prediction markets, a high standard to exceed.

Exhibit 3: Superforecasters Are Even Better in Teams



Source: Based on Barbara Mellers, Eric Stone, Terry Murray, Angela Minster, Nick Rohrbaugh, Michael Bishop, Eva Chen, Joshua Baker, Yuan Hou, Michael Horowitz, Lyle Ungar, and Philip Tetlock, "Identifying and Cultivating Superforecasters as a Method of Improving Probabilistic Predictions," *Perspectives on Psychological Science*, Vol. 10, No. 3, May 2015, 267-281.

Note: Error bars represent plus and minus one standard error.

This is not to say that all went smoothly. Initially, the superforecasters were cautious, restraining comments that would be perceived to be too critical of the assessments of others. And not all of the superforecasters had great social skills. In most cases, however, the teammates figured out ways to engage one another in constructive confrontation that allowed the groups to improve performance. The superforecasters were invited to get together in person after years two and three, and many of them reported that the added human dimension and sense of commitment were helpful.

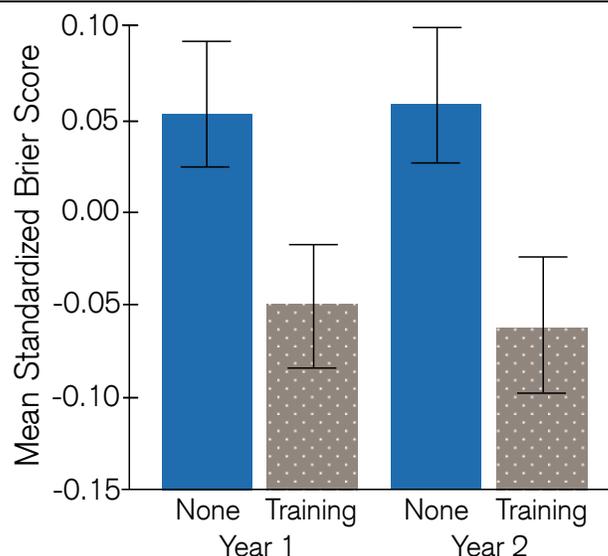
Notwithstanding the success of the superteams, the GJP researchers have reservations about how well the formula would apply in a corporate setting. Identifying individuals as "super," getting employees from different parts of the firm to work together, and interpersonal dynamics might all create friction.

The main lesson is that interaction among a diverse group, especially those with the profile of a superforecaster, can be very effective if managed properly. Companies that seek diversity should also be prepared to train employees in how to manage diversity.²¹

Train Effectively

One remarkable finding from the GJP is that relatively little training, about an hour, can improve results by 10 percent. (See Exhibit 4.) For individuals, the training focused on sharpening probabilistic reasoning and removing cognitive biases. One example is how we tend to use the "inside" versus the "outside" view. With the inside view, we tend to rely on our own information and perception. The outside view considers a problem as an instance of a larger reference class and appeals to the base rate of past occurrence. While both are important in good judgment, psychologists have shown that we commonly rely too much on the inside view. Making individuals aware of the outside view can help reduce this bias.

Exhibit 4: Forecasters Benefit from Debiasing Training



Source: Based on Barbara Mellers, Lyle Ungar, Jonathan Baron, Jaime Ramos, Burcu Gurcay, Katrina Fincher, Sydney E. Scott, Don Moore, Pavel Atanasov, Samuel A. Swift, Terry Murray, Eric Stone, and Philip E. Tetlock, "Psychological Strategies for Winning a Geopolitical Forecasting Tournament," *Psychological Science*, Vol. 25, No. 5, May 2014, 1106-1115.

Note: Error bars represent plus and minus two standard errors.

For groups, the training addressed how to work together. The goal is to strike a balance between conflict and harmony. Too much conflict and the group dynamics break down. No one wants to interact. Too much harmony leads to a false consensus, or even groupthink. The right balance is an atmosphere of constructive criticism.

Overweight Elite Forecasters or Extremize Estimates

The simplest way to capture the wisdom of crowds is to aggregate the estimates of a large number of people with diverse views. For example, you can let lots of people examine a jar filled with jelly beans and ask each of them to provide an estimate of the total. The collective prediction will beat the average individual predictions and will generally be very close to the actual number of jelly beans.

The statisticians working on the GJP figured out two ways to improve the quality of forecasts.²² The first is to place more weight on what the superforecasters say. The intuition behind this is straightforward; give those who predict more accurately a greater voice.

The second is the application of an algorithm to "extremize" answers. Imagine five individuals with different points of view who all estimate the probability of an outcome to be 75 percent. The extremizing algorithm would push that probability closer to 100 percent. Likewise, a group estimate with a low probability would be adjusted closer to zero percent. How far the algorithm pushes the aggregate forecast toward 0 or 100 is a function of the diversity and sophistication of the pool of forecasters.

The idea is to capture unshared information.²³ Here's one way to think about it. Each of our diverse forecasters who come up with a probability of 75 percent uses some information that's common to all of them and some that's unique to them. What would happen if each of them knew all of the information? It would strengthen their confidence, moving their collective estimate toward 100 percent. So the extremizing algorithm captures what would happen if diverse individuals could share all of their information with the group.

A Superforecaster's Tools

We have covered what makes for a superforecaster and the elements of an accurate prediction. But we have glossed over some of the essential tools necessary to come up with good forecasts. We now cover those tools, including defining a good question, keeping score, and approaches to forecasts.

Good Questions. The first step in coming to useful answers is to ask good questions. A good question has a clear outcome within a specified period of time and addresses an issue that is relevant to the world. Let's look at each of these components.

Because keeping score is essential to developing the skill of foresight, questions must have clear answers. Consider the simple example of weather. Weather forecasters predict temperature, precipitation, and winds. They can measure each of these and compare the prediction to the outcome. Likewise, questions in the GJP and those from the intelligence community should have clear answers.

Questions must also have a set time frame. One key difference between Tetlock's prior research on expert judgment, which showed poor results, and the latest results from the forecasting tournament, which paint a more optimistic picture, is the time frame. In the prior research, the questions went out three to five years whereas in the IARPA tournament they average about a few months to a year. Predicting outcomes in the political, economic, and social realms over three to five years is very hard for anyone. Predictions in shorter time frames are more feasible and yet still relevant for business and policy decisions.

Open-ended predictions are of limited value. Tetlock and Gardner discuss a letter sent in November 2010 to Ben Bernanke, then chairman of the Federal Reserve. Signed by nearly two dozen economists, investors, and commentators, the letter suggested that quantitative easing should be "reconsidered and discontinued" and noted the risk of "currency debasement and inflation." In the fall of 2014, reporters at *Bloomberg* contacted some of the signees and found that "all of those who commented stood by the letter's contents." With no time frame or objective way to score the prediction, the letter provides no way to keep score.

Finally, good questions are also relevant. Some questions may be a subset of a group of questions that allow you to assess a larger, but more difficult, issue. Tetlock and Gardner suggest a useful question should pass the smack-the-forehead test: "when you read the question after time has passed, you smack your forehead and say, 'If only I had thought of that before!'"

Here's an example of a question from the forecasting tournament. Note that it has a clear outcome, specified time period, and relevance.²⁴

"Will Italy's Silvio Berlusconi resign, lose reelection/confidence vote, or otherwise vacate office before 1 January 2012?"

Tetlock and Gardner make an additional point on questions. They suggest that the skills in asking good questions may be different than the skills in answering them. The superquestioners might not be the same as the superforecasters. Thoughtful organizations should seek to cultivate skills in both capacities.

Keeping Score. We have already mentioned the Brier score, the main method that psychologists use to measure results. If you have a prediction with a specific probability and an outcome, you are in a position to grade the predictions. Since Brier scores reflect the distance from a forecast to the outcome, lower figures are better. Superforecasters were those in the tournament who earned the lowest Brier scores.

Keeping score is crucial because it provides feedback and therefore an opportunity to learn. Forecasters want to improve in two ways. One way is called “calibration,” which means that your forecasts line up with the outcomes. For example, you are well calibrated if you say that certain events will occur with a 40 percent probability and they actually happen 40 percent of the time.

Another way to improve is what the authors call “resolution.” Resolution means that when you are sure something is not going to happen, it doesn't happen, or when you're sure it will happen, it does. It's a measure of conviction. Good calibration and resolution are correlated, but they are distinct. You want to sharpen your skills in both ways.

Tetlock is fond of the phrase “vague verbiage slows learning cycles.”²⁶ In our day-to-day conversations we use lots of phrases—maybe, possible, probable, might happen—that are ambiguous. In one famous illustration, in 1951 the U.S. intelligence community produced a report saying that the Soviet Union's invasion of Yugoslavia was a “serious possibility.” Sherman Kent, a professor of history at Yale who was then serving at the Central Intelligence Agency, asked some members of his team what “serious possibility” meant. Even though they all agreed to use the term in their report, one analyst said it meant an 80 percent probability and another said 20 percent.

The lesson from this story is to use specific probabilities and time horizons whenever possible and to keep track of those forecasts. Numerical probabilities dismiss the risk of misinterpretation or misunderstanding that Kent faced and provide forecasters with the feedback they need to improve.

Approaches to Forecasts. In studying how superforecasters approach their task, Tetlock and his colleagues noticed a few methods that may be useful beyond the tidy confines of the forecasting tournament.

The first is the idea of question triage. In times of medical emergency such as war, doctors and nurses sort casualties based on their injuries. This process is called triage. Patients who require immediate care receive priority over those who are expected to live and those who have no chance of survival. Similarly, you can sort questions. Some are too easy and others too hard. The priority should go to questions that are between those extremes where “effort pays off the most.”

Superforecasters also use specific techniques to answer some bigger, but answerable, questions. One is associated with Enrico Fermi, a professor at the University of Chicago who worked on the Manhattan Project and won the Nobel Prize in physics. It is effectively a back-of-the-envelope calculation. Fermi used the approach to estimate the strength of an atomic bomb test, but you can use the technique for any question for which there is an answer.

In *Superforecasting*, Tetlock and Gardner use a popular example of a Fermi problem: How many piano tuners are there in Chicago? The tactic is to break the big question into a series of smaller, and easier, questions. You might estimate the population of Chicago, assess the number of households, judge how many pianos there are per household, consider the number of schools and places of worship, determine how often pianos need to be tuned, and contemplate how long it takes to tune a piano and how many hours a piano tuner might work. It turns out there are about 250-300 piano tuners in Chicago.²⁷

Another technique is called question clustering. The idea is that the big question we want to answer—“Will North Korea initiate a war?”—is too big, and the small questions—“Will North Korea launch a multistage rocket by some date?”—are too small. Tetlock and Gardner suggest that a cluster of small questions, each of which you can update, provide insight into answering the bigger question. They use the metaphor of the

painting technique, pointillism, which consists of adding dots to the canvas. No dot by itself means much, but together they paint an evocative picture.

To gauge the likelihood of North Korean aggression, they propose that you might ask questions about missile launches, nuclear tests, cyber-attacks, and artillery shelling. The patterns of probabilities provide a path to tackling some of the big questions. Here again, it's good to remember that those who come up with the questions and those who answer need not be the same.

Leadership

Tetlock and Gardner suggest that if you ask people to list the qualities of a strong leader, you will hear adjectives such as confident, decisive, and visionary. But as we have seen, those descriptions don't fit the superforecasters very well. Leaders are supposed to be in perpetual action mode, while superforecasters seem to be in perpetual learning mode.

The authors refer to the wisdom of Helmuth von Moltke, the chief of staff of the Prussian Army for three decades and considered by many to be the inventor of the modern method of directing troops in the field. "In war, everything is uncertain," said Moltke, recognizing the limits of knowledge. And acknowledging the necessity of flexibility, he observed, "No plan of operations extends with certainty beyond the first encounter with the enemy's main strength."

Moltke's point, which was eventually embraced by the German army, is that leaders must think and change their course of action as necessary. Orders could be questioned, even criticized, if there was a better way. The message was to be actively open-minded. A German military manual notes, "The art of leadership consists of the timely recognition of circumstances and of the moment when a new decision is required."

David Petraeus, a retired four-star U.S. Army general, embraced many of Moltke's themes in a contemporary context. He supported sending his officers to top universities to pursue graduate studies in order to hear points of view that were different from their own. This creates vital mental flexibility and sounds a lot like what superforecasters do.

But Petraeus also recognized that leading requires thinking and doing. He said that a leader "needs to figure out what's the right move and then execute it boldly." Humility should make a leader think carefully about what he or she is doing, and confidence should give an individual the strength to act.

Ultimately, we want our leaders to be confident, decisive, and visionary. But the lesson of *Superforecasting* is that thinking is an important precursor to action. Further, circumstances change and it is vital to update views and change strategic course appropriately.

Doubts

Phil Tetlock is highly respected in his field and, like the superforecasters he studies, is in constant search of improvement. Two of his friends and colleagues, Daniel Kahneman and Nassim Taleb, offer some challenges to the skill Tetlock has identified in forecasting.

Kahneman's challenge relates to "scope insensitivity." The idea is that the questions we face often evoke a feeling of emotion. Scientists can measure that affect with money.

For example, researchers might ask you how much money you would be willing to pay to save 2,000 migratory birds from dying from hazards, including oil spills, wetlands destruction, or residue from herbicides and pesticides. Subjects say about \$80. You feel bad about the situation and are willing to pay some amount of money to help.

Here's the issue: When researchers asked subjects how much they would be willing to pay to save 20,000 birds, the answer was \$78. And for 200,000 birds the response was \$88.²⁸ The dollar sum reflects how people feel about the scenario, not how much they would be willing to pay to save each bird. There is no cost-benefit analysis. Subjects think using System 1 and as a result are insensitive to the scope.

Kahneman thought that a similar type of scope insensitivity might affect the GJP forecasters. Specifically, he thought scope insensitivity would have to do with the time frame. For example, if you are asked how likely it is that a particular dictatorial regime will fall you will focus on the probability and fail to properly consider the time frame.

Running the numbers, the GJP scientists found that indeed many of the forecasters were scope insensitive. When asked the probability that the regime of Syrian President Bashar al-Assad would fall, regular forecasters assigned a 40 percent probability that it would happen within 3 months and a 41 percent chance within 6 months. They assigned essentially the same probability to a period twice as long as the first.

But the superforecasters did much better. Their probabilities were 15 percent for 3 months and 24 percent for 6 months. That's not perfect sensitivity, but it is a lot closer than the regular forecasters. Tetlock and Gardner interpret this result as the ability of superforecasters to engage System 2 thinking more readily than the majority of forecasters. Indeed, they argue that superforecasters have internalized this way of thinking to the point that it has become automatic.

Nassim Taleb has popularized the concept of a black swan event, which comes as a surprise, is consequential, and is explained after the fact. Tetlock and Gardner say that "Taleb insists that black swans, and black swans alone, determine the course of history." If so, IARPA's forecasting tournament is of little value.

Tetlock and Gardner offer some thoughts to counter this view. The first has to do with the definition of a black swan. If we know what a distribution of outcomes looks like, even those with extreme events, then we are dealing with "gray swans," not black swans. Scientists have done a lot of work classifying the distributions of certain phenomena, including earthquakes, terrorist acts, and power-grid failures. While predicting a specific outcome is very difficult, scientists have a general understanding of how these systems behave. Taleb calls these "modelable extreme events" and they capture much of what is of interest to us.²⁹

Black swans and gray swans are highly improbable outcomes. As a result, there are simply not enough data from the forecasting tournament to conclude that superforecasters are either good or bad at spotting them. In

fact, even if the tournament ran for decades it would be hard to know. So those seeking accurate forecasts of black or gray swan events will have to look elsewhere.

But the fact is that the forecasting tournament does address a lot of issues that are relevant for executives, investors, and policymakers. While black swans have a large impact on the world, so do the accumulation of lots of small events. Thinking about the future is not black swan or nothing; there is plenty of room to consider outcomes that have smaller impact.

One point on which Tetlock, Kahneman, and Taleb agree is that forecasting many years into the future is very difficult. There are severe limits on predictability beyond a certain period of time. The essential issue is whether there is value in sharp forecasts for horizons less than a year. Most professionals would answer with a resounding yes.

Summary

Humans have a deep-seated desire to anticipate the future. The demand for forecasts is met by a supply of seers, pundits, and experts expounding on what will happen next. But we generally don't measure the quality of predictions, and when it has been done the results are unimpressive.

Superforecasting shows that prediction may not be so futile after all. The Good Judgment Project, part of a forecasting tournament sponsored by the U.S. intelligence community, revealed that some forecasters are not only good but consistently good. Using the best ideas from psychology and careful measurement, the GJP team has been able to provide essential lessons for anyone in the prediction business. Here are some of the main conclusions:

- **Forecasting skill exists.** The researchers found that a small percentage of the forecasting population were much more accurate than average and consistently so. They were able to find the superforecasters because they welcomed a large sample of forecasters, asked questions with time frames of a year or less, and kept track of the responses.
- **Way of thinking is vital.** Closer analysis of the superforecasters shows that they are bright, but not extraordinarily so. What distinguishes them from regular forecasters is the way they think. Superforecasters are actively open-minded, nondeterministic, intellectually humble, numerate, thoughtful updaters, and hard working.
- **Teams.** When superforecasters interact with one another, their predictions improve. But there are pros and cons to working in teams. The pros include more information and the ability to harness the power of aggregation. The cons are social loafing and the risk of groupthink. Teams that do the best have members who have been given instruction and training on how to work together.
- **Training.** Training in methods to de-bias forecasts and to collaborate effectively improves outcomes. Proper training has content that is valuable and processes that can be implemented immediately. Training without implementation is a waste.
- **Leadership.** At first blush, the qualities associated with leadership appear antithetical to those of the superforecasters. The way to reconcile the two is to acknowledge that proper, even bold, action requires good thinking. And the best leaders recognize that even the best laid plans need to be constantly revised based on the conditions.

Appendix: Keeping Score with Brier

Psychologists commonly use the Brier score as a method for gauging the accuracy of probabilistic forecasts. Glenn Brier, a meteorologist, developed the score in the 1950s.³⁰ In its simplest form, the Brier score measures the square of the forecast error, or $(\text{forecast} - \text{outcome})^2$. For binary events, the value of the outcome is 1 if the event occurs and 0 if it does not. As in golf, a lower score is better.

You can express a Brier score either on a scale of 0 to 1, or 0 to 2, depending on the calculation. We follow Brier’s original approach and place our results on a scale of 0 to 2. When calculating the Brier score this way, you consider the squared forecast error for both the event and the non-event.

Exhibit 5 shows a meteorologist’s probabilistic forecasts for whether it will rain over the next four days. For example, on Day 2, she forecasts an 80 percent probability that it will rain. Likewise, we can say she forecasts a 20 percent probability that it will *not* rain. Because it did rain, we place a 1 in the outcome column below “Rain” and a 0 in the “No Rain” column. Her Brier score for that day was 0.08. For multiple forecasts, the overall Brier score is the mean of the scores for each forecast. The meteorologist’s overall Brier score comes to 0.25.

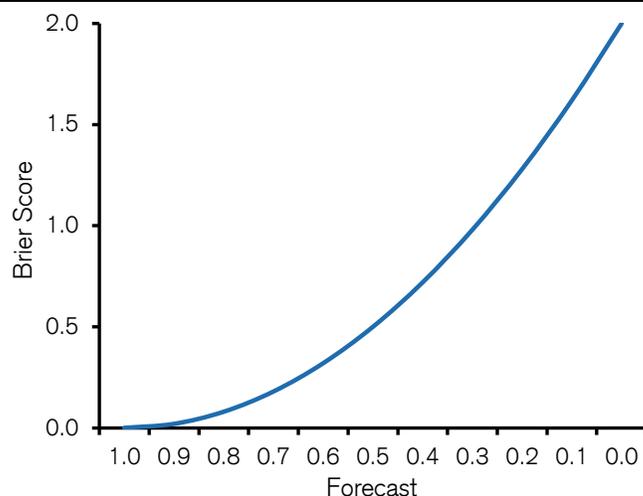
Exhibit 5: Calculation of Brier Score

| Day | Rain | | No Rain | | Brier Score | |
|-------------|----------|---------|----------|---------|---------------------------|-------------|
| | Forecast | Outcome | Forecast | Outcome | Calculation | Result |
| 1 | 30% | 0 | 70% | 1 | $= (0.3-0)^2 + (0.7-1)^2$ | 0.18 |
| 2 | 80% | 1 | 20% | 0 | $= (0.8-1)^2 + (0.2-0)^2$ | 0.08 |
| 3 | 60% | 0 | 40% | 1 | $= (0.6-0)^2 + (0.4-1)^2$ | 0.72 |
| 4 | 100% | 1 | 0% | 0 | $= (1.0-1)^2 + (0.0-0)^2$ | 0.00 |
| Mean | | | | | | 0.25 |

Source: Credit Suisse.

The scale from 0 to 2 has a nice feature. Random guesses have a Brier score of exactly 0.50. Exhibit 6 shows the Brier scores for an event that occurs (“Rain”) for subjective probabilities from 0 to 100 percent. Superforecasters have Brier scores of around 0.20 – 0.25 and in some exceptional cases can achieve scores in the teens.

Exhibit 6: Brier Scores of Event That Occurs for Various Subjective Probabilities



Source: Credit Suisse.

Endnotes

- ¹ Philip E. Tetlock and Dan Gardner, *Superforecasting: The Art and Science of Prediction* (New York: Crown Publishers, 2015), 191.
- ² Philip E. Tetlock, *Expert Political Judgment: How Good Is It? How Can We Know?* (Princeton, NJ: Princeton University Press, 2005).
- ³ David Ignatius, "More Chatter Than Needed," *Washington Post*, November 1, 2013.
- ⁴ Daniel Kahneman and Gary Klein, "Conditions for Intuitive Expertise: A Failure to Disagree," *American Psychologist*, Vol. 64, No. 6, September 2009, 515-526.
- ⁵ Daniel Kahneman, *Thinking, Fast and Slow* (New York: Farrar, Straus and Giroux, 2011).
- ⁶ Philip E. Tetlock, "The Good Judgment Project: Can We Improve Probability Judgments of Possible Futures?" *Presentation at the Credit Suisse Thought Leader Forum*, June 11, 2014.
- ⁷ To learn more of this mistake, see Phil Rosenzweig, *The Halo Effect . . . and the Eight Other Business Delusions That Deceive Managers* (New York: Free Press, 2007).
- ⁸ Kahneman, 212.
- ⁹ Keith E. Stanovich, *What Intelligence Tests Miss: The Psychology of Rational Thought* (New Haven, CT: Yale University Press, 2009). See also Michael J. Mauboussin and Dan Callahan, "IQ versus RO: Differentiating Smarts from Decision-Making Skills," *Credit Suisse Global Financial Strategies*, May 12, 2015.
- ¹⁰ Jonathan Baron, "Beliefs about Thinking," in James F. Voss, David N. Perkins, and Judith W. Segal, eds., *Informal Reasoning And Education* (Hillsdale, NJ: Lawrence Erlbaum Associates, Inc., 1991), 169-186.
- ¹¹ Charles Munger, "A Lesson on Elementary Worldly Wisdom as It Relates to Investment Management & Business," *Outstanding Investor Digest*, Vol. 10, No. 1 & 2, May 5, 1995, 49-63.
- ¹² Michael J. Mauboussin, *Think Twice: Harnessing the Power of Counterintuition* (Boston, MA: Harvard Business Press, 2009).
- ¹³ James Surowiecki, *The Wisdom of Crowds: Why the Many Are Smarter Than the Few and How Collective Wisdom Shapes Business, Economies, Societies, and Nations* (New York: Doubleday, 2004).
- ¹⁴ This is consistent with the difference between risk and uncertainty proposed by the economist Frank Knight.
- ¹⁵ Tadeusz Tyszka and Piotr Zielonka, "Expert Judgments: Financial Analysts Versus Weather Forecasters," *Journal of Psychology and Financial Markets*, Vol. 3, No. 3, 2002, 152-160.
- ¹⁶ Raymond S. Nickerson, "Confirmation Bias: A Ubiquitous Phenomenon in Many Guises," *Review of General Psychology*, Vol. 2, No. 2, June 1998, 175-220.
- ¹⁷ Max Bazerman, *Judgment in Managerial Decision Making, 4th Ed.* (New York: John Wiley & Sons, 1998).
- ¹⁸ Carol S. Dweck, *Mindset: The New Psychology of Success* (New York: Random House, 2006), 18.
- ¹⁹ Angela Duckworth, *Grit: The Power of Passion and Perseverance* (New York: Scribner, 2016).
- ²⁰ Michael J. Mauboussin and Dan Callahan, "Building an Effective Team: How to Manage a Team to Make Good Decisions," *Credit Suisse Global Financial Strategies*, January 8, 2014.
- ²¹ Elizabeth Mannix and Margaret A. Neale, "What Differences Make a Difference? The Promise and Reality of Diverse Teams in Organizations," *Psychological Science in the Public Interest*, Vol. 6, No. 2, October 2005, 31-55.
- ²² Jonathan Baron, Barbara A. Mellers, Philip E. Tetlock, Eric Stone, and Lyle H. Ungar, "Two Reasons to Make Aggregated Probability Forecasts More Extreme," *Decision Analysis*, Vol. 11, No. 2, June 2014, 133-145.
- ²³ Garold Stasser and William Titus, "Pooling of Unshared Information in Group Decision Making: Biased Information Sampling During Discussion," *Journal of Personality and Social Psychology*, Vol. 48, No. 6, June 1985, 1467-1478. Also, Jennifer R. Winquist and James R. Larson, Jr., "Information Pooling: When It Impacts Group Decision Making," *Journal of Personality and Social Psychology*, Vol. 74, No. 2, February 1998, 371-377.
- ²⁴ Barbara Mellers, Eric Stone, Terry Murray, Angela Minster, Nick Rohrbaugh, Michael Bishop, Eva Chen, Joshua Baker, Yuan Hou, Michael Horowitz, Lyle Ungar, and Philip Tetlock, "Identifying and Cultivating

Superforecasters as a Method of Improving Probabilistic Predictions,” *Perspectives on Psychological Science*, Vol. 10, No. 3, May 2015, 267–281.

²⁵ Caleb Melby, Laura Marcinek, and Dani Burger, “Fed Critics Say ’10 Letter Warning Inflation Still Right,” *Bloomberg*, October 2, 2014.

²⁶ Tetlock (2014).

²⁷ See https://en.wikipedia.org/wiki/Fermi_problem.

²⁸ William H. Desvousges, F. Reed Johnson, Richard W. Dunford, Kevin J. Boyle, Sara P. Hudson, and K. Nicole Wilson, *Measuring Nonuse Damages Using Contingent Valuation: An Experimental Evaluation of Accuracy, 2nd Ed.* (Research Triangle Park, NC: RTI Press Publication, 2010).

²⁹ Nassim Nicholas Taleb, *The Black Swan: The Impact of the Highly Improbable* (New York: Random House, 2007), 272.

³⁰ Glenn W. Brier, “Verification of Forecasts Expressed in Terms of Probability,” *Monthly Weather Review*, Vol. 78, No. 1, January 1950, 1-3.

Important information

This document was produced by and the opinions expressed are those of Credit Suisse as of the date of writing and are subject to change. It has been prepared solely for information purposes and for the use of the recipient. It does not constitute an offer or an invitation by or on behalf of Credit Suisse to any person to buy or sell any security. Nothing in this material constitutes investment, legal, accounting or tax advice, or a representation that any investment or strategy is suitable or appropriate to your individual circumstances, or otherwise constitutes a personal recommendation to you. The price and value of investments mentioned and any income that might accrue may fluctuate and may fall or rise. Any reference to past performance is not a guide to the future.

The information and analysis contained in this publication have been compiled or arrived at from sources believed to be reliable but Credit Suisse does not make any representation as to their accuracy or completeness and does not accept liability for any loss arising from the use hereof. A Credit Suisse Group company may have acted upon the information and analysis contained in this publication before being made available to clients of Credit Suisse. Investments in emerging markets are speculative and considerably more volatile than investments in established markets. Some of the main risks are political risks, economic risks, credit risks, currency risks and market risks. Investments in foreign currencies are subject to exchange rate fluctuations. Before entering into any transaction, you should consider the suitability of the transaction to your particular circumstances and independently review (with your professional advisers as necessary) the specific financial risks as well as legal, regulatory, credit, tax and accounting consequences. This document is issued and distributed in the United States by Credit Suisse Securities (USA) LLC, a U.S. registered broker-dealer; in Canada by Credit Suisse Securities (Canada), Inc.; and in Brazil by Banco de Investimentos Credit Suisse (Brasil) S.A.

This document is distributed in Switzerland by Credit Suisse AG, a Swiss bank. Credit Suisse is authorized and regulated by the Swiss Financial Market Supervisory Authority (FINMA). This document is issued and distributed in Europe (except Switzerland) by Credit Suisse (UK) Limited and Credit Suisse Securities (Europe) Limited, London. Credit Suisse Securities (Europe) Limited, London and Credit Suisse (UK) Limited, authorised by the Prudential Regulation Authority (PRA) and regulated by the Financial Conduct Authority (FCA) and PRA, are associated but independent legal and regulated entities within Credit Suisse. The protections made available by the UK's Financial Services Authority for private customers do not apply to investments or services provided by a person outside the UK, nor will the Financial Services Compensation Scheme be available if the issuer of the investment fails to meet its obligations. This document is distributed in Guernsey by Credit Suisse (Guernsey) Limited, an independent legal entity registered in Guernsey under 15197, with its registered address at Helvetia Court, Les Echelons, South Esplanade, St Peter Port, Guernsey. Credit Suisse (Guernsey) Limited is wholly owned by Credit Suisse and is regulated by the Guernsey Financial Services Commission. Copies of the latest audited accounts are available on request. This document is distributed in Jersey by Credit Suisse (Guernsey) Limited, Jersey Branch, which is regulated by the Jersey Financial Services Commission. The business address of Credit Suisse (Guernsey) Limited, Jersey Branch, in Jersey is: TradeWind House, 22 Esplanade, St Helier, Jersey JE2 3QA. This document has been issued in Asia-Pacific by whichever of the following is the appropriately authorised entity of the relevant jurisdiction: in Hong Kong by Credit Suisse (Hong Kong) Limited, a corporation licensed with the Hong Kong Securities and Futures Commission or Credit Suisse Hong Kong branch, an Authorized Institution regulated by the Hong Kong Monetary Authority and a Registered Institution regulated by the Securities and Futures Ordinance (Chapter 571 of the Laws of Hong Kong); in Japan by Credit Suisse Securities (Japan) Limited; elsewhere in Asia/Pacific by whichever of the following is the appropriately authorized entity in the relevant jurisdiction: Credit Suisse Equities (Australia) Limited, Credit Suisse Securities (Thailand) Limited, Credit Suisse Securities (Malaysia) Sdn Bhd, Credit Suisse AG, Singapore Branch, and elsewhere in the world by the relevant authorized affiliate of the above.

This document may not be reproduced either in whole, or in part, without the written permission of the authors and CREDIT SUISSE.

© 2015 CREDIT SUISSE GROUP AG and/or its affiliates. All rights reserved