



Reviews & Essays

Reading Tarot on K Street

By Philip Tetlock

Ian Bremmer and Preston Keat, *The Fat Tail: The Power of Political Knowledge for Strategic Investing* (New York: Oxford University Press, 2009), 272 pp., \$27.95.

Bruce Bueno de Mesquita, *The Predictioneer's Game: Using the Logic of Brazen Self-Interest to See and Shape the Future* (New York: Random House, 2009), 272 pp., \$27.00.

George Friedman, *The Next 100 Years: A Forecast for the 21st Century* (New York: Doubleday, 2009), 272 pp., \$25.95.

We are awash in political-forecasting products: books, reports, updates, outlooks, data points, extrapolations, scenarios . . . and who doesn't want to have a glimpse into the future? Who doesn't want to know whether the Dow will close above ten thousand at year's end, whether Russia will bully more of its neighbors, whether North Korea will do something really "crazy," whether the Saudis can maintain their oil production . . . and who among us is not curious whether 2045 will be the year of the "singularity" when, according to some futurists, artificial intelligence surpasses human minds, technological advances skyrocket under the control of machine-human hybrids and those of us who can hang on long enough have our own shot at immortality? But if

wishes were horses, beggars would ride. There is a paucity of evidence—peer-reviewed scientific evidence—that forecasters know how to deliver the goods: reliably accurate political, economic and technological predictions. In fact, when I have staged competitions, many forecasters fail to outperform the proverbial dart-throwing chimpanzee—and most cannot outperform extrapolation algorithms that simply predict "more of the same."

Karl Marx hit the nail on the head when he quipped that, when the train of history hits a curve, the intellectuals fall off. But there is still great potential for mischief in the political-forecasting business. The demand for accurate predictions is insatiable. Reliable suppliers are few and far between. And this gap between demand and supply creates opportunities for unscrupulous suppliers to fill the void by gulling desperate customers into thinking they are getting something no one else knows how to provide.

On top of this, the laws of probability tell us that at least a few superstar forecasters will arise just by chance if enough pundits make enough predictions. According to the extreme skeptics, we should just wise up and give up waiting for our forecasting savior. Even such fabled figures as Warren Buffet, the sage of Omaha, and George Soros, the man who broke the Bank of England, may just be lucky—the human equivalents of coins that, if tossed long enough, eventually yield surprisingly long streaks of "heads" or "tails."

But extreme skepticism is too bitter a pill for most mortals to swallow, implying as it does that, when it comes to the big decisions—whether to reassure rather than deter

Iranian ayatollahs or North Korean dictators or bail out or let sink mega-financial institutions—we do as well in the long run by flipping coins as consulting experts. That is why, in part, I classify myself as a moderate skeptic. Hope springs eternal—and I cannot bring myself to write off the forecasting business completely. When new products pop up in the marketplace of ideas, I cannot help but wonder whether any of them might somehow fill the void.

I am not however a moderate skeptic for just desperate wish-fulfillment reasons. A good deal of research indicates that some ways of thinking (“cognitive styles”) do translate into somewhat more correct forecasts. When we score the accuracy of thousands of predictions from hundreds of experts across dozens of countries over twenty years, we find the best forecasters tend to be modest about their forecasting skills, eclectic in their ideological and theoretical tastes, and self-critical in their analytical styles.¹ Borrowing from philosopher Isaiah Berlin, I call them foxes—experts who know many things and are not finicky about where they get good ideas. Paraphrasing Deng Xiaoping, they do not care if the cat is white or black, only that it catches mice.

Contrast this with what I call hedgehogs—experts who know one big thing from which likely future trends can be more or less directly deduced. The big thing might be any of a variety of theories: Marxist faith in the class struggle as the driver of history or libertarian faith in the self-correcting power of free markets, or a realist faith in balance-of-power politics or an institutionalist faith in the capacity of the international community to make world politics less ruthlessly anarchic,

or an eco-doomster faith in the impending apocalypse or a techno-boomster faith in our ability to make cost-effective substitutes for pretty much anything we might run out of.

What experts think—where they fall along the Left-Right spectrum—is a weak predictor of accuracy. But *how experts think* is a surprisingly consistent predictor. Relative to foxes who are less encumbered by loyalties to an all-encompassing worldview, hedgehogs offer bolder forecasts and, although they hit occasional grand slams, they strike out a lot and wind up with decidedly poorer batting averages.

Readers now know my biases: a deep belief in the need for independent, objective scoring rules for gauging expert accuracy and a deep skepticism of big-idea hedgehogs. Nothing in the books under review—merely the latest tomes to sit on the forecasting shelves—has altered the first belief but, as we shall see, one of the books has shaken the second.

The authors are all entrepreneurial futurists, but each offers a strikingly distinctive approach to prediction. I organize these approaches under three headings: the superpundit model in which readers take it, more or less on faith, that the forecaster has a pipeline into the future not available to ordinary mortals (a category into which I place George Friedman’s *The Next 100 Years*); the technocratic-pluralism model in which the authors never get around to making falsifiable predictions of their own but do offer readers a pretty

¹ I conducted exercises and gathered evidence on this for my 2005 book, *Expert Political Judgment*.



comprehensive survey of forecasting mistakes and an inventory of tools for avoiding them (a category into which I place Ian Bremmer and Preston Keat's *The Fat Tail*); and the scientific-reductionist model in which the author embraces a particular theory from the social sciences and shows how, if you apply that theory thoughtfully to real-world contexts, you can derive surprisingly accurate forecasts (a category into which I place Bruce Bueno de Mesquita's *The Predictioneer's Game*).

Reading these three books, it is easy to feel like a frustrated shopper wandering aimlessly down the forecasting aisle in the supermarket of ideas. The products on offer are packaged well—but we have no objective benchmarks, no trusted *Consumer Reports*, against which to gauge performance. We have no idea whether we would be better-off paying one of these consultancies gobs of money for their proprietary forecasts or simply downloading the latest odds from a high-profile prediction market that culls individual bets on world events such as Tradesport. Indeed, would we do as well relying on the dart-throwing chimps or mindless extrapolation rules, like “Predict the most recent rate of change”? So, caveat emptor.

First, the superpundit model. George Friedman, the author of *The Next 100 Years*, is also the founder and CEO of STRATFOR, “the world’s leading private intelligence and forecasting company.” Of the three books, his was the least persuasive. But Friedman does his readers at least one good turn. He repeatedly warns them of how easy it is for the conventional wisdom, anchored as it is in the status quo, to underestimate cumulative change over multidecade spans of time. How many in 1945 would have predicted that Japan and Germany would in thirty-five years be the second- and third-most powerful economies in the world, or that China in 1970—still in the convulsive grip of Mao’s Great Proletarian Cultural Revolution—would be the fourth-biggest economy in 2007 and moving up the rankings fast? The answer is “extremely few.” Yet these outcomes now have the whiff of inevitability about them. Appreciating the power of hindsight bias makes us think harder about how we think: if we can “explain” the past so easily and predict the future so poorly, perhaps we should reconsider how we go about doing both.

Friedman unfortunately does not develop this critical theme. He is in the forecast-

ing business and he leaves the impression that he is one of the select few today who can see deeply—and in amazing detail—into the twenty-first century. He relies heavily on plate-tectonics metaphors—earthquakes and Huntingtonian fault lines—to organize the recent past and project into the future. He



sees China being torn into regional fiefdoms by roughly 2020, a Russian military collapse shortly thereafter, a surge in American growth rates around 2040, and a world war in the mid-twenty-first century in which the United States uses its technological superiority with space-based weaponry and hypersonic aircraft to squelch rising regional powers such as Turkey and Japan.

These forecasts are entertaining. Friedman can spin a riveting good yarn far into the future. I particularly enjoyed his tale, circa 2080, about Mexican secessionists running successfully and simultaneously for election to the American and Mexican Congresses. And, judging by the book blurbs and internet commentary, Friedman has attracted a following: “*Barron’s* consistently has found STRATFOR’s insights informative and largely on the money—as has the company’s large client base, which ranges from corporations to media outlets and government agencies.” Indeed, the *New York Times Magazine* is apparently tempted, when it is “around George Friedman, to treat him like a Magic 8-Ball.”

Glib public-relations chatter aside, one does not need to be an extreme skeptic to refrain from joining the dust-jacket chorus. To his credit, Friedman himself admits that, in 2007, he completely missed the economic implosions of 2008. If he can be that far-off on something that big and that close, why should we assign anything appreciably above a zero probability to his longer-range predictions? Here, Friedman’s defenders must rally around his geophysical metaphors—and argue that although we cannot predict a major earthquake in California in any given year between now and 2050, we can predict with substantial confidence that a major earthquake will occur in one of those years—and we might even be able to justify the claim that the likelihood of a major political earthquake increases the longer we go without one. Readers must judge the merits of these metaphors for themselves. But based on my experience scoring the accuracy of many expert predictions, I would wager against most,

Many forecasters fail to outperform the proverbial dart-throwing chimpanzee—and most cannot outperform extrapolation algorithms that simply predict “more of the same.”

if not all, of Friedman’s ten- to fifty-year projections coming true.

It is worth stressing that my skepticism is not aimed at Friedman personally. It is grounded in my broader transideological skepticism of the superpundit model of forecasting. It matters not if the superpundit is a traditional realist (e.g., Henry Kissinger) or a moderate conservative (e.g., David Brooks) or a neoliberal (e.g., Tom Friedman) or a left-leaning economist (e.g., Paul Krugman). I have seen too many broad-brush, loquacious experts get it spectacularly wrong to place much confidence in any of them.

The best thing I can say for the superpundit model is likely to annoy virtually all of that ilk: they look a lot better when we blend them into a superpundit composite. Aggregation helps. As financial journalist James Surowiecki stressed in his insightful book *The Wisdom of Crowds*, if you average the predictions of many pundits, that average will typically outperform the individual predictions of the pundits from whom the averages were derived. This might sound magical, but averaging works when two fairly easily satisfied conditions are met: (1) the experts are mostly wrong, but they are wrong in different ways that tend to cancel out when you average; (2) the experts are right about some things, but they are right in partly overlapping ways that are amplified by averaging. Averaging improves the signal-to-noise ratio in a very noisy world. If you doubt this, try this demonstration. Ask several dozen of your coworkers to estimate the value of a large jar of coins. When my classes do this exercise, the average guess is closer to the truth than 80 or 90 percent of the individual guesses. From this perspective, if you want to improve your

odds, you are better-off betting not on George Friedman but rather on a basket of averaged-out predictions from a broad ideological portfolio of George Friedman-style pundits. Diversification helps.

Turning to Ian Bremmer and Preston Keat’s *The Fat Tail*, we find authors of an altogether-different intellectual temperament. These are the technocratic pluralists. They eschew bold, guru-ish forecasts and, depending on one’s own disposition, one can take that reticence as a sign either of hard-earned modesty or of rank cowardice. I lean toward the more charitable interpretation. Bremmer and Keat encourage their readers to look closely at how they think about the future, deploy a wide range of tools to assist them in this task and warn readers against common pitfalls. They are largely in the business of prying open otherwise-closed minds—Bremmer is president and Keat is director of research at the Eurasia Group, another “leading global political risk research and consulting firm.” And they use provocative thought experiments, scenarios and historical analogies to alert readers to the need for humility—and the potential for abrupt discontinuities.

The very title of their book gives a sense of their intellectual predilections. We expect high-impact events to happen very rarely—i.e., as the thin tails on a normal distribution curve. But in fact they happen surprisingly often—hence the idea of the fat tails. One never knows when a revolution of some sort is just around the corner: whether it be Bolshevism or Nazism, nuclear physics or the internet, Islamic terrorism or viral pandemics . . . So, be prepared to be shocked by pos-

sibilities lurking in the intellectual shadows—Rumsfeld’s unknown unknowns. Bremmer and Keat’s main thrust is on keeping an open mind, paying attention to contrarians and unexpected events, and allowing for a broad spectrum of opinion and analysis.

One should indeed beware of mathematical models that assume well-defined normal distributions of risks and benefits flowing from policy options. Borrowing a page from risk theorist Nassim Nicholas Taleb’s marvelous *The Black Swan*, Bremmer and Keat emphasize that you should reconsider your probability estimates of financial phenomena when one-in-a-quadrillion-years events start popping up a few times every decade. If anything, and even barring the promise of their title, the authors do not draw enough from Taleb because the problem is not just fat tails in statistical distributions of political outcomes. The problem is that the distributions can suddenly take on shapes that a few years earlier were quite unthinkable. Economists call these distribution-altering events—the Bolsheviks, Nazis, atomic bombs, al-Qaeda, stray asteroids—exogenous shocks. Medieval mapmakers used more evocative language for expressing the depth of their ignorance: they would affix the label “hic sunt dracones” to regions unknown.

Like George Friedman, Bremmer and Keat urge readers not to forget how often elite conventional opinion has gotten it wrong in the past. They remind us of the humbling of the Nobel-staffed hedge fund Long-Term Capital Management in the wake of the Russian debt default of August 1998. A lot of smart money assumed that a sovereign entity as large as Russia was just too big to fail—so financial instruments were not properly priced

to take into account the sorts of case-specific risk factors that savvy foxes tend to be good at picking up: the alcoholic haze through which Boris Yeltsin governed the country, the widespread corruption in the government, the number of well-placed Russians who stood to make lots of money from devaluation and, on the Western side, the defective regulation of financial markets that permitted risky borrowing to finance investments. Bremmer and Keat also tease the *Wall Street Journal* for praising Yugoslavia as the most promising place to invest in Eastern Europe—in May 1990, just as the country was sliding into a bloody civil war that festers even now.

The authors point out too that just because something—like geopolitical risk—is hard to quantify does not give you license to ignore it. Rough approximations are better than tacitly treating the risk as zero. Ask the energy companies that rushed into Yeltsin’s Russia in the 1990s to make large fixed-capital investments and were then compelled by Putin in the next decade to “renegotiate.”

This means we need to value contrarian thinkers, even though they can be a royal pain and slow down the adoption of policies preferred by insiders. And so the authors suggest we might consider one much-hyped, but still-useful, method of prying open otherwise-closed minds: scenario planning. This requires policy advocates to envision worlds in which things don’t work out as planned—a world in which we are not greeted as liberators in Iraq; or a world in which deregulation leads not to greater market efficiencies but rather to excessive borrowing that severely destabilizes financial markets.

History rarely overtly repeats itself but it

often rhymes—and there is an advantage to those who can pick up the more subtle similarities. Saddam Hussein bore some resemblance to Hitler, but he also bore some to Ceauçescu, Mussolini, Kim Il Sung and Stalin, all of whom played a far more risk-averse geopolitical game. The case for the 2003 invasion of Iraq loses much of its rhetorical force when we use historical analogies in a more nuanced fashion.

The authors are even nimble enough to see that although there are many settings in which foxes like themselves have an advantage—they are slower to jump to premature conclusions and are quicker to change their minds in response to new evidence—hedgehogs are still sometimes remarkably prescient. As far back as 1933, Churchill classified Hitler as a grave threat to Great Britain—a category into which, incidentally, he also tossed Gandhi. Similarly, the most bearish and bull-

ish financial pundits occasionally have their days in the sun.

It is hard to quarrel with foxes like Bremmer and Keat who see the virtues of a hedgehog approach to reality. But unfortunately, for all their open-mindedness and sensitivity to state-of-the-art tools for managing political-risk analysis, for all their sly, foxlike bobbing and weaving among the forecasting brambles, the authors remain timid foxes. Bremmer and Keat show as little interest as Friedman in keeping rigorous score of their predictions or cultivating the fine art of well-calibrated judgment.

There is only one way for deep-pocketed corporate clients to determine whether granting Friedman status as a “Magic 8-Ball” or following Bremmer and Keat’s executive-MBA lists of best practices has a positive payoff—and that is by conducting controlled experiments in forecasting accuracy among their executives.



I have seen too many broad-brush, loquacious experts get it spectacularly wrong to place much confidence in any of them.

Almost nobody seems interested in doing that—except perhaps the author of the third book.

Of the three books, I should be wariest of Bruce Bueno de Mesquita's *The Predictioneer's Game*. Bueno de Mesquita is an unapologetic hedgehog—but he is a hedgehog who must be taken seriously. He is a world-class political scientist who appreciates the pivotal role of transparent scientific methodology in advancing human knowledge. He makes strong claims about how far he can see into the future—with the aid of the right expert inputs and game-theoretic algorithms. He claims to find strikingly simple patterns underlying the superficial chaos of world politics. And he maintains that he offers a demonstrable competitive edge to any comers willing to toss aside reliance on expert intuition and join him in calculating the likely trajectory of events. Here we enter the realm of the scientific-reductionist model.

Forecast consumers should take notice when confronted with a high-credibility sales pitch of this sort—and judging by how busy his consulting firm is, many have.

The basic logic of Bueno de Mesquita's prediction machinery is straightforward. Imagine you are a consultant to the Department of Defense who has been hired, yet again, to explore scenarios for inducing North Korea to behave better on the nuclear front. To start, you need to define “better” by identifying the options available to the regime: refuse to negotiate about nuclear weapons; negotiate but then cheat on any agreement as soon as it becomes advantageous (Kim's preferred approach); slowly cut back the nuclear program

in exchange for various American economic and security concessions; abandon the program, conditional on various concessions; or eliminate the program unconditionally (America's most, and North Korea's least, preferred outcome). Next, you need to map out the options on the American side and then make inferences about Pyongyang's likely responses if, say, the United States publicly targeted nuclear missiles at North Korea or if it guaranteed North Korea's security within its borders.

Bueno de Mesquita declares that, once we have mapped the option space, we simply need to follow his four-step formula for making accurate predictions. First, get the best-possible experts to identify every individual or group with a “meaningful” interest in trying to influence the decision. Second, get the experts to estimate as accurately as possible which options each of the identified players is advocating in private—that is, what they want. Third, get experts to estimate how big an issue this is for each of the players—how motivated they are to prevail. Fourth, get experts to estimate the relative political clout or influence of each player in this issue domain.

As a game-theory hedgehog, Bueno de Mesquita believes that is *all* you need to know: no need to consider the specific history or culture of the state or the personality traits of the individuals. And, although Bueno de Mesquita does not describe all of the details by which he combines the expert inputs to derive his forecasts, it is clearly a variant of expected-utility theory: for each policy option and context, multiply each player's preferences by each player's motivation-to-prevail by each

player's political clout, and then predict that the regime will choose the option with the most "points."

It might seem here that the experts are doing all of the heavy lifting—so why do we need the model? But Bueno de Mesquita makes the eye-catching claim that "according to a declassified CIA study," his forecasting model hits the "bull's-eye" about twice as often as the government experts who provided Bueno de Mesquita with the input data for his formulas. And he offers a comparative-advantage explanation for this result: the experts are wonderful sources for case-specific inputs, but they are not experts on how people make choices—which is the value added from the model.

I find Bueno de Mesquita's proposed intellectual division of labor persuasive, but I am not persuaded that Bueno de Mesquita is right about "all" we need to know for optimal forecasting or that he is working with the right model of human choice. And a remarkably simple game illustrates my reservations—the guess-the-number game. To predict the outcome of even this extremely stylized game, we need to go beyond Bueno de Mesquita's list. We need to know how rational the other side is—and how rational the other side assumes us to be. And once we admit these messy complexities of human nature back into the equation, we lose the radical-hedgehog simplicity of Bueno de Mesquita's model—and we are back in fox-friendly territory.

Consider what happened when, many years ago, the *Financial Times* ran a guess-the-number competition for its readers (promising the winner an all-expenses-paid trip on the Concorde). The task was deceptively simple: predict a number between zero and one hundred

such that your guess is as close as possible to two-thirds of the average guess of all other players. Some readers guessed $33\frac{1}{3}$ —and were classified as strategically naive. They assumed that others would pick numbers randomly between zero and one hundred, which averages out to 50, and two-thirds of 50 is $33\frac{1}{3}$. These forecasters made the beginner's mistake of failing to factor in the incentives at work. Other readers guessed zero. They were too clever by half. They assumed that everyone knew as much game theory as they did—and they quickly reasoned through the deductive sequence: everyone knows that everyone knows that the first-order answer is $33\frac{1}{3}$, but if everyone converges on that answer then the correct answer is $22\frac{1}{6}$, and if everyone converges on that answer then the correct answer is two-thirds of $22\frac{1}{6}$, but if everyone converges on that answer . . . The theoretically correct answer—the Nash equilibrium—is zero. But the real answer for readers of the *Financial Times*, a pretty savvy group, was around 18, roughly halfway between the strategically naive answer of $33\frac{1}{3}$ and the too-clever answer of zero.

Bueno de Mesquita might reply that such complexities matter in contrived lab settings but the real-world proof of the forecasting pudding is in the eating—and he claims a delicious 90 percent hit rate based on the aforementioned CIA report. Good social scientists are, however, Missourians: they insist on tasting the pudding themselves. The debate must unfold in peer-reviewed outlets—and there must be open, level-playing-field competition across approaches. Impressive though the numbers cited in Bueno de Mesquita's book are, they are also—without

getting into nitty-gritty technical details—devoid of significance. A 90 percent hit rate is, for example, no great achievement for meteorologists predicting that it will not rain in Phoenix. And it is no big deal even to achieve a 100 percent hit rate of predicting X—no matter what X may be—if doing so comes at the cost of an equally high false-alarm rate. Anyone can predict every war from now until eternity by simply predicting war all the time.

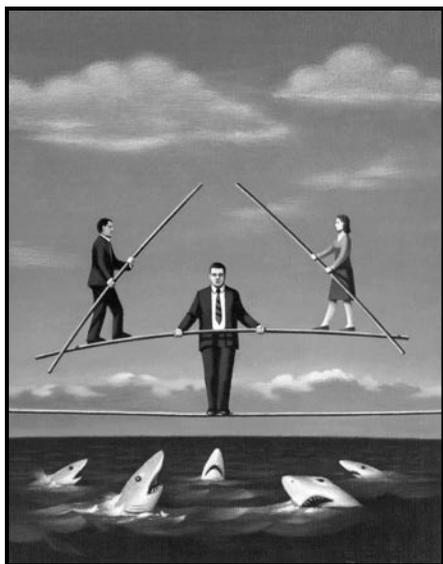
I do not mean here to trivialize Bueno de Mesquita's predictive track record: he is rare among social scientists in keeping score and his performance is impressive—even to one as jaded as I. That said, I have reservations. It is unclear how Bueno de Mesquita would counter the argument that outperforming the individual experts is no grand feat. As already noted, it is not hard to beat individual experts in the forecasting game. From this perspective, Bueno de Mesquita's model may

be accomplishing no more than what averaging routinely does—and even dart-throwing chimps can occasionally pull off.

How then can we produce the most accurate forecasts? The answer is not obvious: right now all we can say confidently is that no one can be 100 percent confident about which approach would win if we were to run a series of level-playing-field forecasting tournaments stretching out to, say, 2020.

But if the market seems largely indifferent to our plight, who might rescue us? There is one potential savior on the horizon: a big institutional purchaser of forecasting services that has the financial clout and technical-support staff ready to run forecasting tournaments that would shed light on the relative performance of competing approaches—a big player that also has powerful incentives to discover superior analytical strategies, for even small improvements in its prediction accuracy can translate into billions of dollars and millions of lives saved. And that player is the Office of the Director of National Intelligence.

Unfortunately, although intelligence agencies have been heavy buyers of forecasting services, they have not used their massive purchasing power to their full advantage. They have allowed the diverse interest groups in the intelligence community to choose freely from private-sector forecasting products. On the one hand, this is commendably open-minded. On the other hand, there has been no integrative effort to assess the relative value added of each product. Indeed, intelligence agencies seem as allergic as private-sector forecasters to being held accountable to public accuracy metrics.



Even such fabled figures as Warren Buffet and George Soros may just be the human equivalents of coins that, if tossed long enough, eventually yield surprisingly long streaks of “heads” or “tails.”

To appreciate why, it is helpful to think like a game theorist—to take the forecasters’ perspectives and consider the incentives confronting them. Both private- and public-sector prognosticators must master the same tightrope-walking act. They know they need to sound as though they are offering bold, fresh insights into the future not readily available off the street. And they know they cannot afford to be linked to flat-out mistakes. Accordingly, they have to appear to be going out on a limb without actually going out on one. That is why (with the interesting exception of Bueno de Mesquita), they so uniformly appear to dislike affixing “artificially precise” subjective probability estimates to possible outcomes—the only reliable method we have of systematically tracking accuracy across pundits, methods, time and contexts. It is much safer to retreat into the vague language of possibilities and plausibilities—things that might or could happen if various difficult-to-determine preconditions were satisfied. The trick is to attach so many qualifiers to your vague predictions that you will be well positioned to explain pretty much whatever happens. China will fissure into regional fiefdoms, but only if the Chinese leadership fails to manage certain trade-offs deftly, and only if global economic growth stalls for a protracted period, and only if . . . And if you venture specific policy recommendations—such as invading Iraq or deregulating financial markets—make sure to leave yourself the fallback position: “Well, of course, my recommendation was fundamentally sound, but how was I to know that the idiots in charge would implement things so badly. If only they had . . .”

Having mastered this subtle balancing act, why should these private- and public-sector pundits open their reputations and livelihoods to the unpredictable risks of competing against each other in level-playing-field forecasting exercises? Why stray from the cozy vagueness-zone equilibrium, which seems to be working well enough for the providers of forecasting services, even though the societal outcome is decidedly suboptimal?

Sadly, the game-theory cynics may get the last laugh: we may be stuck in a suboptimal equilibrium—stuck because we are incapable of rising above our immediate, narrowly defined self-interest. But there is a wild-card escape possibility. It requires leadership. Players high up in the political system—who really do want the best-possible forecasts—could decide that it is worth investing a nontrivial share of their intelligence agencies’ budgets into a series of long-term forecasting tournaments designed to distinguish the more from the less promising forecasting approaches across policy problems. Strictly speaking, this would not be a rational choice for these leaders to make—because the budgetary reallocations will evoke howls of protest, and those responsible will be out of office long before the forecasting tournaments start yielding practically useful results. But building a robust capacity for learning—and learning how to learn—is not a bad legacy. □

Philip Tetlock is the Mitchell Endowed Professor at the University of California, Berkeley, and the author of *Expert Political Judgment: How Good Is It? How Can We Know?* (Princeton University Press, 2005).