

Editor's Note: In celebration of Math Awareness Month on “The Future of Prediction,” we provide essays by Lisa R. Goldberg and Olle Häggström, and greetings from some of our fellow mathematical societies around the world.

Lisa R. Goldberg

Prediction Is a Young Science

Predictions go back at least as far as the Oracle of Delphi and the astrologers of the Chinese Han Dynasty. Arguably, they are as old as the human race. However, our recently acquired capability to collect, store, and analyze data has emphasized statistics and elevated prediction to a science that pervades virtually everything we do. Despite recent advances, prediction is a young science. We are just beginning to explore its limits.

Scientific Prediction Can Be Humbling and Confusing

Time to eat some crow.
—Nate Silver

On July 8, 2014, Germany won a stunning 7-1 victory over Brazil in the World Cup semifinals. This was a problem for superstar statistician Nate Silver, whose Poisson distribution-based Soccer Power Index (SPI) model had forecast a win for Brazil and had set the odds of Germany beating Brazil by six goals or more at 1 in 4,000.

Silver has been using statistical models to predict outcomes of sporting events since the early 2000s, but he is best known for predicting elections. In 2008, Silver and his colleagues at FiveThirtyEight correctly called every state in the US presidential election with the exception of Indiana and the Second Congressional District in Nebraska, which awards its own electoral vote. In 2012, FiveThirtyEight correctly called the US presidential election in all fifty states (including the toss-up in Florida) and the District of Columbia. Silver's book on statistics, *The signal and the noise*, became an international bestseller. “Triumph of the Nerds” and other articles celebrating data science appeared, exceptionally, in the mainstream media.

Lisa R. Goldberg is founding partner of Berkeley Associates LLC, Director of Research at Aperio Group, and Co-director of the Consortium for Data Analytics in Risk at the University of California at Berkeley. Her email address is lrg@berkeley.edu.

This article has benefitted from conversations with Bob Anderson, Jeff Bohn, Alan Cummings, Nick Gunther, Ola Mahmoud, Barry Mazur, Caroline Ribet, Ken Ribet, Stephanie Ribet, Alex Shkolnik, Philip Stark, and Roger Stein.

For permission to reprint this article, please contact:
reprint-permission@ams.org.

DOI: <http://dx.doi.org/10.1090/noti1377>

In response to questions about SPI's disastrous Brazil-Germany prediction, Silver explained:

Statistical models can fail at the extreme tails of a distribution. There often isn't enough data to distinguish a 1-in-400 from a 1-in-4,000 from a 1-in-40,000 probability.

Statistical models rely on estimated distributions, outcomes, and likelihoods to make predictions. The tails of a distribution are populated by rare or never-before-seen outcomes. It is virtually impossible to assess the credibility of the estimated likelihood of a rare event.

Silver was not the only one to misforecast the outcome of the 2014 World Cup. Many of the big investment banks had used statistical methods to predict a win for Brazil. The *Financial Times* demoted data science from hero to goat with headlines like “Brazil's World Cup eludes banks' best minds”. On the other hand, Microsoft was delighted to point out that its prediction engine, Cortana, had correctly predicted the outcome in 15 of 16 World Cup matches, including the Brazil-Germany semifinal.

Extreme Events May Seem Obvious in Hindsight

On June 22, 2007, the investment bank Bear Stearns was quietly working to rescue two of its hedge funds: the Bear Stearns High-Grade Structured Credit Fund and the Bear Stearns High-Grade Structured Credit Enhanced Leveraged Fund. These funds were trading securities backed by sub-prime mortgages, which involve loans to homeowners by banks who readily acknowledge the possibility that some of the homeowners might default. However, relatively few considered the possibility that lots of homeowners might default at the same time.

The rescue failed since the sub-prime market had begun to collapse in response to a rapid rise in mortgage payment delinquency. By mid-July, Bear Stearns disclosed that the two funds had lost nearly all their value. Smooth and upward-trending equity markets began to rock. This was the preamble to the global financial crisis that dominated the ensuing two years and which has been blamed for the destruction of trillions of dollars, the loss of millions of jobs, and the suicides of more than 10,000 individuals.

Despite plentiful data, the elaborate and expensive risk management systems in place at many financial institutions in 2007 failed to alert investors to the impending crisis. There were exceptions. In *The Big Short*, Michael Lewis tells the tales of a few investors who noticed that

housing prices were inflated, acted on their observations, and profited. For the vast majority of us, however, the crisis signals were obscured by the noise of everyday life.

I recall a San Francisco limousine driver telling me early in 2007 about the houses he was flipping in Georgia. It seemed odd, but I did not ask myself what it implied. I could have known. That nagging I-could-have-known feeling that comes in the aftermath of a disaster is an example of hindsight bias, and it depends on our selective memory of what turned out to matter. In his bestselling book *Thinking, Fast and Slow*, Daniel Kahneman explains the behavioral roots of hindsight bias. These include the availability heuristic, our tendency to rely on information that easily comes to mind; and the representativeness heuristic, our tendency to profile.

A Correct Prediction Can Also Be a Bad Prediction

The Monty Hall Problem goes like this:

There is a sports car behind one of three closed doors, but the other two doors hide goats. You get to select a door and keep whatever is behind it. After you make an initial choice, Monty Hall opens one of the two unselected doors, revealing that it does not hide the sports car. He offers you the opportunity to switch to the other unopened door or stay with your original choice. Should you stick or switch?

If you stick with the door you originally selected, you may be lucky and win the car. But the probability that the car is behind the door you originally selected is $1/3$, while the probability that the car is behind the unselected door Monty Hall failed to open is $2/3$. The prediction that the car is behind the door you originally selected is bad, even if it turns out to be correct.


A skilled predictor may never make a bad prediction, but she will inevitably make some predictions that turn out to be incorrect. This conundrum is concisely summarized in Nassim Nicholas Taleb's *Table of Confusion*, which pairs terms that tend to be mixed up: luck and skill, belief and knowledge, theory and reality, signal and noise, randomness and determinism.

Scientific Prediction Learns from Mistakes

In 1948, major polls in the US famously predicted that New York governor Thomas Dewey would beat Harry Truman in the presidential election.

The predictions were based on quota sampling, which targets subsets of the population that are prespecified rather than representative. The predictions were also stale. As George Gallup, co-chairman of the Gallup organization explained, "We stopped polling a few weeks too soon." Modern election polls use random sampling rather than quota sampling, and they run up to the last minute.

Today, scientific predictions tell us which movies, books, and restaurants we will enjoy, which diseases will afflict us in the future, whether a cancerous growth on one of our feet will prove fatal, how drugs will interact, who




Incorrect banner headline on the front page of the *Chicago Daily Tribune* on November 3, 1948.

will be our soulmates, and when we might be wiped out by rising sea levels and ocean acidification. It can be difficult to distinguish good predictions from bad ones, since a correct prediction need not be good. A prediction might be good if it is based on ample data and an algorithm that learns, if it fits in-sample and is validated out-of-sample, and if it is interpretable and appealing to common sense. It may be impossible to test predictions of the time and location of the next big earthquake or the efficacy and safety of a cure for a rare disease or the time and cause of the next financial crisis.

But some scientific predictions have demonstrated track records. Advanced warnings of hurricanes and tornadoes have allowed communities at risk to evacuate before storms hit. Airline travel is remarkably safe, in part, due to accurate weather forecasts. When Google Navigation tells you that, despite some traffic, you will reach your destination by 3:12 pm, there are reasons to believe it.

About the Author

After starting her career in academia in topology and complex dynamical systems, Lisa moved to mathematical finance in 1993 and hasn't looked back. She and her husband, AMS President Elect Ken Ribet, have two daughters. An avid swimmer, Lisa averages two kilometers per day and has swum a lifetime distance of 25,000 kilometers.



Lisa R. Goldberg.

Photo courtesy of Ken Ribet

Olle Häggström

Our Desperate Need to Predict the Future

In 1814, Pierre Simon de Laplace envisioned a demon who could pinpoint the exact present positions and movements of all particles and then calculate all of their future trajectories, thereby accurately predicting everything that will ever happen. He understood, however, that this task would be forever beyond mere human capabilities. Due to what we now know of as sensitive dependence on initial conditions, this brute force approach to predicting the future does not work (other than in certain limited situations with limited time horizon). Instead, we'd like to find regularities in the past and assume "by induction"

What happens when we succeed in creating a machine that surpasses human capabilities?

that they will continue into the future. Unfortunately, there seems to be no way to defend induction—the extension of past regularities into the future—without reverting to circularity by pointing out that induction has served us well in the past and thus can be expected to do so in the future.

Nevertheless, we rely on induction to make predictions. We simply do not know any other way. Good applied mathematicians and good scientists know that the extent to which these predictions are reliable depends on how closely the state of the system we are trying to predict will remain within the envelope of what has already been observed. The further we push atmospheric CO₂ levels above those of the last several million years, the less reliably we can predict the future climate. Moore's law—the exponential curve that fits several decades of computer hardware development so well—eventually predicts the physically impossible.

Today we face unprecedented challenges of the following kind. We are beginning to develop new technologies that have the potential to take us very far outside the envelope of the observed and the familiar—technologies that can bring us enormous benefits, but also catastrophes on all scales, including the extinction of humanity. In my recent book *Here Be Dragons*, I discuss several such technologies and argue for the need to act with foresight. Examples abound in biotechnology as well as in

nanotechnology. Here, for concreteness, let me focus on artificial intelligence (AI). What happens when we succeed in creating a machine that surpasses human capabilities in terms of general intelligence? Will it enter the kind of rapid spiral of self-reinforcement known as an *intelligence explosion* or *the Singularity*? Will we be able to remain in control when we are no longer the smartest creatures on the planet? Much has been written about this, for instance by Ray Kurzweil, emphasizing the wonderful prospects brought by the Singularity, and by Nick Bostrom, emphasizing the risks. Kurzweil likes to emphasize that we have always lived with double-edged technologies and coped well: "Fire kept us warm and cooked our foods, but also burned down houses," Kurzweil told CNN in June 2015. The implicit suggestion here is that things will be fine with AI as well, but as physicist Max Tegmark recently pointed out in a panel discussion with Kurzweil, there may be a crucial disanalogy: we learned to use fire using trial-and-error, while if we do not get the first AI superintelligence right (i.e., beneficial to humanity), we may not get a second chance.

These scenarios are so far outside what we are used to thinking about as proper areas of extrapolation and prediction that any attempt in that direction may seem reckless. Yet, I claim, the amount of value at stake compels us to try, as the alternative tends to involve implicit and unfounded assumptions that nothing drastic will happen and is tantamount to running blindfolded into a minefield. Induction on its own is entirely inadequate for making useful predictions here, but in combination with deduction, useful conclusions perhaps can be drawn, analogously to climate modelling, where empirically well-understood physical laws and regularities are combined to predict the future (along with uncertainty estimates). Predicting the consequences of an AI breakthrough is far more difficult than predicting the climate in 2100, but this should not stop us from doing the best we can and accompanying the predictions with the right amount of epistemic humility. Without a conscious effort to predict, we will rely on ill-founded and erroneous implicit predictions of the status quo.

One more caveat. When I talk of the need for predictions here, I really mean the need for *conditional* predictions, predictions conditional on various actions on our part. A prediction like "a robot Apocalypse by 2100 has probability 20 percent" is of little use in itself; what we really need as a guide to policy and action is to understand, e.g., how promoting or restricting various AI research directions affects that probability. This is analogous to how IPCC projections are conditional on different greenhouse gas emission scenarios, thus guiding us in how choices of emission policies affect future climate. As modelling in economics and other social sciences becomes more sophisticated, it becomes increasingly tempting to include social aspects as endogenous to climate models. Just as the models encompass physical feedback mechanisms (such as warming leading to reduced snow covering, leading to more warming), it is also possible to include social

Olle Häggström is professor of mathematical statistics at Chalmers University of Technology, a member of the Royal Swedish Academy of Sciences, and the author of four books. His email address is olleh@chalmers.se.

For permission to reprint this article, please contact: reprint-permission@ams.org.

feedback mechanisms (such as warming leading to air conditioning, leading to more warming). While such studies can be very interesting, there is also the risk that building human reactions into the system of differential equations can lead to a deterministic view in which our free will seems to disappear. To be able to work effectively towards improving our prospects for a bright future, we need to retain a mindset according to which different actions and different policies are both possible and consequential.



About the Author

Olle Häggström, author of four books, with the first advance copy of *Here Be Dragons: Science, Technology and the Future of Humanity*, Oxford University Press, 2016.

Greetings from Some of Our Fellow Math Societies

C. Ciliberto

L'Unione Matematica Italiana

Each year in April the Joint Policy Board for Mathematics of the US sponsors the Mathematics Awareness Month to stress the importance of mathematics in our society and culture through a series of initiatives that highlight mathematical developments and applications. The 2016 Mathematics Awareness Month will celebrate its 30th anniversary since this initiative started off in 1986.

This year the theme of the Mathematics Awareness Month will be "The Future of Prediction". Human beings have been trying to foretell the future since ever. In the far past, this was the job of charlatans, astrologists, prophets and oracles, who appealed to superstitions and pseudoscientific tools (the interpretation of the flight of birds, the form of tea leaves or coffee grains at the bottom of a cup, the sound of a thunder, tarots, etc.) to convince people that future could be predicted. In more recent times this has been taken over by scientists, who have been trying to make predictions based on the current model of interpretation of the nature. Natural sciences (like physics, chemistry, biology, etc.) and socio-economical sciences (like economy, sociology, etc.) have been playing of course a crucial role in this respect.

But how can scientists anticipate the evolution of the climate in the next twenty years if even a three-day forecast is a serious challenge? How can we predict the development of epidemics? How can we accurately foresee our financial future? To make accurate predictions about the behavior of highly complex systems, we need a reference model to appeal to, and to collect, interpret and manipulate an enormous amount of data: of course, more a model is mathematically refined, more it is robust and amenable to provide the right framework for reliable predictions. In the last century or so, thanks to the development of more and more sophisticated tools, mathematics widely extended its domain of applications from physics (which was the classical one) to all other subjects mentioned above, including human sciences. Never in the history of our civilization, is mathematics as




Photo by Giulia Ciliberto.

**C. Ciliberto,
President of
L'Unione
Matematica
Italiana.**

C. Ciliberto is president of the L'Unione Matematica Italiana. His email address is ciliberto@mat.uniroma2.it.

*For permission to reprint this article, please contact:
reprint-permission@ams.org.*

immanent as [it is] today in all our activities. In this sense, never as nowadays, mathematics is involved in the art of predicting the future. Therefore the subject of the 2016 Mathematics Awareness Month is extremely timely.

The “Unione Matematica Italiana”, which is the reference association of the Italian mathematicians, looks at this fascinating theme with the greatest interests and wishes to this initiative the best success.

Jose Maria P. Balmaceda

Mathematical Society of the Philippines



Jose Maria P. Balmaceda,
President of
Mathematical
Society of the
Philippines.

Warmest greetings to the AMS from Mathematical Society of the Philippines! We join the AMS, other reciprocating societies, and the worldwide mathematical community in celebrating Mathematics Awareness Month in the US and laud the choice of this year's focus on the 'Future of Prediction'. Being in a region highly vulnerable to natural disasters, environmental hazards and economic uncertainty, our members are increasingly being relied upon to develop models and techniques for forecasting and risk management, and provide insightful

analysis of large amounts of data. We look forward to the *Notices'* excellent articles and commentaries on these topics.

Jose Maria P. Balmaceda is president of Mathematical Society of the Philippines. His email address is jose.balmaceda@gmail.com.

For permission to reprint this article, please contact: reprint-permission@ams.org.

Francis E. Su

Mathematical Association of America

The power of mathematics is, to a greater extent than ever before, being harnessed all around us in the information revolution, with mathematics and statistics as the power tools of a new economy. Phones now predict when you will get home. Search engines predict the spread of disease. Numerical models predict the weather. Mathematics makes new inventions possible. I hope you'll use the theme of this year's Math Awareness Month as a golden opportunity to amplify, for your students and your friends, the role of mathematics at the center of life in the twenty-first century.



Courtesy of Francis E. Su.

Francis E. Su,
President of the
Mathematical
Association of
America.

Francis E. Su is president of Mathematical Association of America and Benediktsson-Karwa Professor of Mathematics at Harvey Mudd College. His email address is su@math.hmc.edu.

For permission to reprint this article, please contact: reprint-permission@ams.org.