


Predicting the future of society

Matthew J. Salganik

 Check for updates

Predicting the future is something that humans have tried to do – in various ways – for a very long time. A paper by Grossmann et al. tests the ability of social scientists to predict societal change and finds that they are not particularly good at it.

For thousands of years, experts have made predictions about what is going to happen to an individual or a society, but these predictions have not usually been subject to rigorous evaluation¹. This changed in 2005, with the publication of a landmark book evaluating expert prediction of geopolitical events². One of the most well-known – and troubling – findings of this book was that supposed experts were not much better at predicting than were dart-throwing chimpanzees. However, subsequent work has shown that there might be some hope for expert social prediction. For example, researchers have found that intelligence analysts show some skill in making certain kinds of social forecasts^{3,4}. Against this background of previous work, the paper by Grossmann et al.⁵ adds to our understanding of expert forecasting in social systems and raises some important new questions.

Grossmann et al. ran two forecasting tournaments in which social scientists were asked to forecast social indicators such as political polarization. In the first tournament, researchers were provided with 39 months of historical data and asked to predict the next 12 monthly values (May 2020 to April 2021). Six months later, teams were invited to participate in a second tournament covering only six months. In both tournaments, teams were asked to forecast 12 different social indicators in a single country (USA), but from a variety of domains (such as life satisfaction, sentiment on social media and gender or racial bias). A wide range of social scientists entered the tournaments: 86 teams entered the first tournament, and 120 teams entered the second. The teams were compared to two types of benchmarks: (1) simple statistical models and (2) a nonexpert crowd.

When the forecasts were compared to the true outcomes, Grossmann et al. discovered that the forecasts of the social scientists were not particularly impressive. They conclude that “social scientists’ forecasts were on average no more accurate than simple statistical models”⁵, a finding that is consistent with previous research². Grossmann et al.⁵ also found that “for most domains, social scientists’ predictions were either similar to or worse than the [nonexpert] crowd’s prediction”. However, these findings are also consistent with a slightly different comparison that makes the social scientists look a bit better: in five domains, social scientists beat the nonexpert crowd; in one domain, the nonexpert crowd beat the social scientists; and in six domains, it was not possible to declare a clear winner. But no matter how you summarize the results, it is hard to conclude that the social scientists are especially good at making these forecasts.

Although these horserace-style comparisons might be what many people will remember from this paper, some of the other findings



may turn out to be more important scientifically. The design of the forecasting tournaments means that predictions were made by many teams about many social indicators at many forecast horizons. Grossmann et al. take advantage of this structure to explore three key dimensions of forecast accuracy. They find that the more-accurate teams tended to be interdisciplinary, use simpler models, base predictions on prior data and have prior experience with forecasting tournaments. Focusing on prediction targets, they do not find a clear substantive pattern about which types of social indicators are more difficult to predict (although they did find that indicators that showed more statistical variability during the training period were more difficult to predict). Finally, the authors find that longer forecast horizons led to more accurate predictions, a surprising finding that is at odds with other fields such as meteorology, where weather forecast accuracy decreases as the forecast horizon increases⁶.

Similar to many exciting papers, this one creates more questions than it answers: there are two that I would like to highlight. First, in doing the study the authors had to make several design decisions, and it is important to understand how these results could change if the authors had made different – but still sensible – decisions. For example, teams were making forecasts about the USA between May 2020 and April 2021. This was an incredibly turbulent period that included the COVID-19 pandemic; the Black Lives Matter protests sparked by the murder of George Floyd; a presidential election during which the losing candidate refused to concede; and the 6 January riot at the US Capitol. How much would the results be different if people were predicting in a more stable time? On the one hand, a turbulent time could give social scientists an advantage, because this is a setting in which theory might be most valuable. On the other hand, one could wonder whether existing social theories would even apply in especially turbulent times. Ultimately, as Grossmann et al. note, more work is needed to understand how forecasting in turbulent times compares to forecasting in more stable times.

A second important question is what, if anything, these results tell us about other kinds of phenomena that social scientists might want to predict. In addition to predicting aggregate social trends, a social

scientist might want to predict outcomes for individual people (that is, rather than predicting the birth rate, they might want to predict which specific people will give birth). Or social scientists might want to predict collective outcomes, such as whether a country will fall into civil war. I hope that future research can explore the similarities and differences between these various types of social predictions.

In addition to raising questions, there are two important misinterpretations of these results that we should try to avoid: one from social scientists and one from policy makers. Social scientists might conclude that their poor performance is proof that forecasting is a pointless task. In other words, if we are not good at it, then it cannot be important. This kind of thinking would be a mistake. Rather, I hope we – as a community – take the opposite approach. To me, the results of Grossmann et al. suggest that we should increase our efforts to rigorously measure and understand our ability – and inability – to predict the future.

A second possible misinterpretation could come from policy makers, who might conclude from these findings that social scientists do not know anything. But that is not quite right. For many tasks of critical policy importance, I expect that social scientists can make important contributions – or at least do better than a simple statistical model and a nonexpert crowd. For example, some social scientists almost certainly have expertise in designing and evaluating policy interventions intending to change social indicators, rather than merely forecast them. Even in the realm of forecasting, social scientists can probably make important contributions to forecasting the probabilities of long-term, existential risks. Of course, these speculations would have to be tested empirically.

In addition to its specific findings and open questions, the paper of Grossmann et al. also gives us an important reminder of the power

of large-scale collaboration in the social sciences^{7–10}. Had Grossmann and colleagues undertaken independent studies, the results would have invariably ended up incompatible because of different design choices. It was only by working together that these researchers were able to produce such important results. Thus, these forecast tournaments remind us that there are things researchers that can accomplish collectively that none of us can accomplish individually. If we want to make progress on the many interesting questions raised by Grossmann et al., we are going to need to work together.

Matthew J. Salganik  

Department of Sociology, Princeton University, Princeton, NJ, USA.

 e-mail: mjs3@princeton.edu

Published online: 9 February 2023

References

1. van Creveld, M. *Seeing Into The Future: A Short History Of Prediction* (Reaktion Books, 2020).
2. Tetlock, P. E. *Expert Political Judgment: How Good Is It? How Can We Know?* (Princeton Univ. Press, 2006).
3. Mandel, D. R. & Barnes, A. *Proc. Natl Acad. Sci. USA* **111**, 10984–10989 (2014).
4. Tetlock, P. & Mellers, B. *Proc. Natl Acad. Sci. USA* **111**, 11574–11575 (2014).
5. Grossmann, I. et al. *Nat. Hum. Behav.* <https://doi.org/10.1038/s41562-022-01517-1> (2023).
6. Bauer, P., Thorpe, A. & Brunet, G. *Nature* **525**, 47–55 (2015).
7. Salganik, M. J. et al. *Proc. Natl. Acad. Sci. USA* **117**, 8398–8403 (2020).
8. Milkman, K. L. et al. *Nature* **600**, 478–483 (2021).
9. Open Science Collaboration. *Science* **349**, aac4716 (2015).
10. Makridakis, S., Fry, C., Petropoulos, F. & Spiliotis, E. *Inf. J. Data Sci.* **1**, 96–113 (2022).

Competing interests

The author declares no competing interests.