AN AUDIENCE WITH...

Philip Tetlock

Drug discovery is notoriously unpredictable, and as a result extremely expensive. Even small improvements in the ability to forecast how drug candidates will fare as they move through clinical trials could save the industry billions of dollars. Recent research by Philip Tetlock, a Professor of Psychology at the University of Pennsylvania and author of Superforecasting, shows that the art of predicting the outcomes of incredibly complex one-off events can be taught and fostered. He spoke with **Asher Mullard** about the lessons his work holds for the pharmaceutical sector.

© Eric Mencher

Q *Your work focused on how contestants, in an open competition, fared over time at forecasting the outcomes of complex geopolitical events of the sort that US intelligence agencies grapple with. What did you find?*

When we started this work it was not immediately obvious that it is possible to improve the probability estimates of these types of events. They seem to be one-of-a-kind events. Will Greece leave the Eurozone? What will Putin do next in Ukraine? Will Arctic sea ice expand or contract? What will happen with Ebola? They are very heterogeneous questions. And yet we had a group of people who got progressively better at assigning probability estimates to these questions. The project is pushing the boundaries of what we thought was possible in probability estimation.

We found several other things as well. We found that we could train people to be better forecasters. And we found that there were ways of organizing top performers into teams that were more than the sum of their individual parts. We also found that it doesn't really matter how smart you are, but that you do have to believe that subjective probability estimation is a skill that is cultivatable and is worth cultivating. If you don't believe this, you are not going to try and you are not going to get better at it.

Q *Do you think these findings are applicable to drug discovery?*

The nature of the tournament was interesting because the questions were so heterogeneous. If you ask me in what domains have we demonstrated that objective probability estimation is cultivatable and worth cultivating, I'd say either everything or nothing.

Q *A key part of forecasting is the ability to break large questions down into smaller, more tractable, questions. Scientists are already adept at this kind of approach. What else can they do to become better forecasters?*

When I think about the scientific method, I think at least two types of cognitive method are engaged. There is causational reasoning and there is statistical reasoning. And I think those two communities probably have difficulty communicating in drug discovery. They certainly have difficulty meshing in other areas of enquiry.

Researchers on the biochemistry side of things, for example, use causal reasoning. And I think it is probably associated with over-confidence: you can build up a very good causal case for a lot of things — for example, that a particular molecule should be able to stop a neurodegenerative process. Whereas if you step back and take the statistical view, and look at the base rates of success when a scientist has convinced themselves that they have a molecule that will help cure a particular disease, well those base rates are very low and discouraging. You'd never try anything if you looked at things just purely from the statistical base rates. So you need to be able to alternate back and forth, balancing the inside view, the mastery of the biochemical specifics of that drug–organism interaction, and the somewhat demoralizing statistical prospects of success.

There is also value in keeping score. You'll never discover how close you are to the optimal forecasting frontier in your domain if you don't start exploring.

Q *Your work suggests that corporate culture can impact forecasting ability. What can companies do to foster better forecasting?*

I'm reluctant to give off-the-shelf advice, because I think different organizations are making different mistakes. But there is a tendency towards over-confidence in human judgement and organizational judgement. So, most organizations are probably in the over-confident end of the continuum. But there may be some that are under-confident and are too quick to write off potentially promising projects. Similarly, there are probably some companies that show excessive rigidity in their forecasting, never updating their forecasts, and there are others that show excessive volatility.

We really don't have a clear idea of what types of errors we are more or less prone to make unless we keep systematic score. I would suggest that there is long-term value in running internal forecasting tournaments to keep track of who is better at forecasting, and to measure how well calibrated they are. This way you can give greater weight to the judgement of people who prove to be better calibrated.

Also, a corporate culture that punishes people for being on the wrong side of maybe is a corporate culture that is guaranteeing that its analysts are not going to learn to make better probability estimates.

Q *You found that teams tend to be better forecasters than individuals. How should we structure teams to make the most of this?*

We find that our amalgamation algorithms work better when we have diverse inputs into teams. We call this viewpoint diversity. I don't know drug discovery well enough to say what the relevant viewpoint diversity space would look like, but I could imagine you'd want people who take a statistical perspective, a biochemical perspective, a human physiological perspective and so on, all on the same forecasting team.

Is it possible to have too much diversity on a team? Sure. It is possible to have too much of almost anything. But, on balance, teams usually don't have enough viewpoint diversity in them and they don't usually manage the viewpoint diversity very well.