Psychologists' betting market hints at most reliable research findings

Markets beat opinion polls in predicting which studies will be replicated successfully.

Erika Check Hayden

09 November 2015



Andrew Shurtleff/The New York Times/Redux / eyevine

When psychologists are asked which findings in their field will pass a replication attempt, their predictions prove little better than flipping a coin. But if they play a betting market on the same question, the aggregate results give a fairly good prediction of which studies will hold up if repeated, a prediction experiment suggests¹.

The study was run in parallel with a large, crowd-sourced replication effort, which reported² in August that fewer than one-half of 100 findings in the field could be replicated by other teams. Last year, while that effort was underway but before its results were revealed, some forty-odd researchers involved in it filled out a survey indicating — for more than 40 studies — how likely they thought that study's main finding was to be successfully replicated.

The same researchers also played a betting game. They were given US\$100 to buy and trade contracts in the replication efforts that were going on. The contracts each paid \$1 if the key finding was eventually replicated, and nothing if it was not. (Researchers were not allowed to bet on studies that they were involved in reproducing).

A study published on 9 November in the *Proceedings of the National Academy of Sciences*¹ reveals the results. The survey predictions were not significantly more accurate than a coin flip. But a model developed from the betting market accurately predicted the results of 71% of the replication efforts, significantly outperforming the simple survey. The market was a little optimistic, however: the average contract price was \$0.55, suggesting that players thought 55% of the studies would be replicated.

That suggests that research communities are fairly good judges of which research will stand up to replication, say other scientists. "The results show that a collection of knowledgeable peers do have a good sense of what will replicate and what won't," says behavioural economist Colin Camerer at the California Institute of Technology in Pasadena. "This information is in the judgements of peers but has never been collected and quantified until now."

Researchers at the Center for Open Science in Charlottesville, Virginia, coordinated a team effort to replicate the findings of 100 psychology studies — and ran a betting game to see if those results could be predicted. (From left: Johanna Cohoon, Mallory Kidwell, Courtney Soderberg and Brian Nosek.)

Place your bets

The market approach might have proved more accurate than the survey because money was on the line, or because the market enables people to learn from others' biases and to adjust their bets, notes Simine Vazire, a psychology researcher at the University of California, Davis. Still, the market bets were inaccurate enough of the time that "it would be unwise to use prediction markets to draw conclusions about the replicability of individual studies", she says. "But I think they could provide useful information about moderators of replicability — what types of studies do prediction markets anticipate to be more likely to replicate?"

Using the details of the studies and their later replication experiments, as well as results from the betting market, the researchers also developed a model (using Bayesian statistics) to estimate how likely a hypothesis is to be true after another study affirms it.

The median probability of a hypothesis proving accurate after publication of an initial finding was 56%. In other words, after one publication affirms a hypothesis, "the probability of the hypothesis being true does not rocket up to near certainty", says psychologist Brian Nosek of the University of Virginia in Charlottesville, who is one of the co-authors of the study. After a successful replication, however, the probability of a hypothesis being true approaches 100%, the study found.

The authors suggest that markets could be used to decide which studies should be prioritized for replication. "If these results are reproducible, then we can use markets to help estimate the credibility of current findings, and prioritize which findings are important to investigate further because of their uncertainty," Nosek says.

Nature | doi:10.1038/nature.2015.18754

References

- 1. Dreber, A. et al. Proc. Natl Acad. Sci. USA http://dx.doi.org/10.1073/pnas.1516179112 (2015).
- 2. Open Science Collaboration. Science 349, 6251 (2015).