

A conversation with Elizabeth Iorns on 02/26/13

Participants

- Elizabeth Iorns — Science Exchange, Founder and CEO
- Alexander Berger — GiveWell, Senior Research Analyst

Note: This set of notes was compiled by GiveWell and gives an overview of the major points made by Elizabeth Iorns.

Summary

Elizabeth Iorns founded Science Exchange, an online marketplace for scientific experiments. Science Exchange played a major role in launching the Reproducibility Initiative, which is a project that aims to incentivize researchers in biomedical research to do careful experiments with findings that can be reproduced. GiveWell spoke with Dr. Iorns to learn about the Reproducibility Initiative as part of our investigation of the cause of meta-research.

Academic incentives in biomedical research and reproducibility

One of the biggest problems in biomedical research is that researchers are not rewarded for generating findings that are reproducible.

Low incentives for doing reproducible research

Researchers are rewarded primarily for publishing papers in prestigious journals such as *Nature*, *Science* and *Cell*. These journals select for papers that report on surprising and unusual findings. Papers which report on unsound research which is apparently exciting are more likely to be published than papers which report on less exciting research that is sound.

There is little post-publication check on the soundness of papers' findings, because journals, especially prestigious ones, generally don't publish replications, and there is little funding for performing replications.

Low reproducibility as a problem

Because many published results are not reproducible, it is difficult for scientists to use the published literature as a basis for deciding what experiments to perform.

Pharmaceutical companies rely on the biomedical research academic literature for inspiration for drug candidates. These companies are unhappy with the low reproducibility rate, because it is wasteful for the industry to expend resources

pursuing experiments that do not reproduce academic papers' findings.

As things stand, the pharmaceutical industry does replications, however, these are generally unpublished. Because a given lab doesn't know whether other labs have found that a study fails to replicate, labs duplicate a lot of effort. Making the results of replications public could save people time spent replicating studies that are not reproducible.

The scale of the phenomenon

Pharmaceutical companies such as Bayer and Amgen have studied the frequency with which studies are reproducible by trying to reproduce them, and they have found that about 70% of published papers in the areas that they considered don't reproduce.

The ALS Therapy Development Institute attempted to replicate 221 studies in the field of research on amyotrophic lateral sclerosis and found that none of them were reproducible.

Some people claim that papers that are unsound are retracted, so that there's already a mechanism that addresses the fact that some papers don't reproduce. However, the retraction rate is only 0.1%, so the vast majority of papers that do not reproduce are not retracted.

The Reproducibility Initiative

The premise

The Reproducibility Initiative aims to reward scientists who produce research that is reproducible. The Reproducibility Initiative does this by independently replicating studies, methods and reagents that are submitted.

The hope is that signing up for this service and having one's findings replicated will serve as a badge of credibility. If this is so, researchers will be motivated to sign up for it when they believe that their findings are solid, and will be motivated to ensure that their findings are solid so that they can submit them for replication with confidence that their findings will be reproduced.

The implementation

The Reproducibility Initiative asked the authors of every paper indexed by PubMed in 2012 if they'd be willing to have their study be validated by an independent lab. The authors of 1,892 studies answered that they would like to have their studies' results independently validated. The Reproducibility Initiative is hoping to facilitate the replication of some of these studies: initially perhaps 100 studies as a proof of

concept, but perhaps all studies eventually.

Science Exchange is a network of about 1,000 research facilities, which operate outside of the usual academic system. Labs in Science Exchange will perform the replications. Because the labs are outside of the usual academic system, the lab providers won't have personal stake in the outcomes of the experiments, so the labs that are performing the replications will be relatively objective.

The Reproducibility Initiative will share the data from the replications via figshare, which is a data-sharing platform. The papers resulting from the replications will be published in an open access journal (PLOS ONE).

The cost-effectiveness of the replications

The cost of doing the replications that the Reproducibility Initiative is coordinating will be relatively small (10%-15% of the cost of the original experiments) because:

- Much of the cost that goes into the publication of a paper is spent on initial exploratory research. The cost of the final experiments that a paper reports on is a small fraction of the cost of producing the paper. Only the final experiments need to be repeated in a replication (and even then only key results are replicated).
- The labs that will do them have workers who are specialists in the techniques that were used for the experiments.

Although they vary by orders of magnitude, a typical experiment to be replicated might cost around \$100,000, so replication might be in the \$10,000 to \$15,000 range per study.

Research into characteristics of reproducible studies

There's currently not enough data on the reproducibility of biomedical research for it to be possible to determine which subfields have the worst reproducibility problems.

The Reproducibility Initiative's hope is that studying which papers are reproducible will ultimately facilitate the construction of an algorithm for judging the likelihood that a paper's findings are reproducible, using variables such as the country that the authors live in, the institution that the authors work at, the lab that the authors are from, the diseases that the paper studies and the methods that were used.

Funding for the Reproducibility Initiative

The Reproducibility Initiative has not yet secured funding to commission

replications. Potential funders were initially skeptical that the project would find a significant number of researchers who want their papers to be replicated. Now that the project has found many scientists who want their papers to be replicated, it is again searching for funding.

Assorted issues surrounding data collection and data sharing

Data deposition as a signal of quality

It's often the case that a journal has an official requirement that people who publish papers deposit the raw data in a repository, but in practice this often doesn't happen. There was a recent study that found a strong correlation between data being deposited and the data being reproducible.

A trend toward data sharing

There is momentum behind the cause of increasing open access to publications and to the data that the publications report on. The US government recently issued a statement that it will require that grantees make not only their publications open access but also make their data open access. It seems likely that the open access problem will be solved in the near future.

Representativeness of data

It's generally the case that the data that researchers report is not representative of the data that they collected. Only about 1% of experiments that are ever done are published. It would be hard to enforce a policy of researchers sharing all of the data that they've collected. However, there is some progress in the direction of researchers recording more of their data in machine-readable format, with a shift toward electronic lab books, which would help address this problem.

Robotics and Automation

If robots generated biomedical research data, then it would be possible to track the data produced in the course of experiments, because it would be collected in machine-readable format. Robots could also track relevant variables that are not currently tracked, such as light conditions, humidity and temperature. Having robots do experiments would also reduce variability in the data coming from differences between experiments across labs: if the robots were built identically and the code that they executed was identical, then the experiments by the robots would be much more similar than the experiments done by different human experimenters.

As such, a shift toward robot-performed experiments could substantially improve the quality of biomedical research. It seems possible that such a shift will occur over

the next decade.

People for GiveWell to talk to

- C. Glenn Begley — the former vice-president of Hematology and Oncology Research at Amgen. He coauthored a study titled *Drug development: Raise standards for preclinical cancer research* which attempted to replicate 53 “land-mark publications” in preclinical cancer research and found that only 11% of them replicated.
- Richard Klausner — The former director of the National Cancer Institute. Dr. Klausner has expressed concern about the incentive structure in academic research.
- Bruce Booth — A partner in the life sciences group of Atlas Ventures. Dr. Booth has studied the amount of money that has been wasted as a result of papers’ low reproducibility rate.

All GiveWell conversations are available at <http://www.givewell.org/conversations>