A conversation with Allan Steinhardt on June 13, 2013

Participants

- Allan Steinhardt – former Program Manager and Chief Scientist at the Defense Advanced Research Projects Agency (DARPA), now an Executive Advisor at Booz Allen Hamilton
- Holden Karnofsky – Co-Founder and Co-Executive Director, GiveWell
- Dario Amodei – Postdoctoral Research Fellow, Stanford School of Medicine
- Cari Tuna – Co-Founder, Good Ventures

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

Summary

Allan Steinhardt is a former DARPA Program Manager and Chief Scientist. GiveWell spoke with Dr. Steinhardt about DARPA's approach to funding research, how it chooses projects and what has made it successful at sponsoring innovative work.

Characteristics of DARPA

The absence of tenure helps DARPA foster creativity and productivity. People spend limited time there, so it's not populated by people whose entire careers have been in the same place. The expectation is that people make something happen and then move on – this limited time horizon combined with the freedom that the agency allows means that researchers terminate unproductive avenues and have the flexibility to try new approaches.

Additionally, DARPA doesn't use a peer review system, instead having people who are technically knowledgeable and experts in their field make funding decisions. The program managers who are responsible for assigning funding are experts who have done similar research themselves and therefore are able to make informed judgments. This is different than the typical model, where an administrator relies on the competitors of the person looking for money to review the proposal. There is a lot of evidence that the method of peer review harms innovation because of the way researchers take turns evaluating each other's work.

DARPA is a relatively lean organization. It has roughly 100 employees, far fewer than other government organizations with similar budgets. Program managers have significant power because they have large budgets, which has a multiplicative effect: being able to wield large sums of money towards a single cause can influence the funding decisions of universities and private industry. Setting up prizes is one way in which this happens.

DARPA is more proactive than academic research, going out and seeking to do things rather than making publications and waiting for a response. Such an approach requires having the ability to make good decisions, so DARPA relies on the technical savvy of its employees. The lack of tenure forces people to bring out their best decisions in their limited time.

The interplay between industry and academia is important for DARPA, and both sectors are crucial for much of the research that gets done. For example, the creation of the Internet involved work in networking theory by academics and knowledge of phone infrastructure from industrialists.

Investment Decision Process

At the highest level, Congress controls DARPA's budget. The director of DARPA and office managers, along with the Department of Defense, propose broad research goals and Congress has final say over these broad goals. This is done at a very high level, and the end result of a project often does not match the initial direction. The goals set will be a mix of things, and often very basic concepts, not necessarily related to defense or national security. For example, the Internet started as an aim to create a secure communication network in case of a nuclear war. Goals like that are set without prescribing a specific way to achieve them, so farther down the organizations there's more and more autonomy.

Choosing projects

DARPA intentionally looks for projects that aren't being undertaken elsewhere, making a point not to do things that are already being done. Furthermore, they don't do any projects that absolutely have to succeed, always acknowledging that failure is a possibility. Thinking that failure is not an option attracts much more conservative approaches and stifles innovation.

To choose projects, the people responsible for assigning funding talk to a variety of entities working in a field – boards of different research organizations and military branches, etc. – and see what issues they are dealing with, and choose projects based on these. DARPA has also funded short-term think tanks, bringing together a variety of relevant experts to comment on the problem they wish to investigate. Doing this gives a good perspective of the current situation and helps direct funding decisions.

The organization is flat, with only three levels – program manager, office director, and director. To get funding, program managers pitch their proposals to the director, the office directors, and a handful of scientists, allowing debate and discussion in a way that is not possible with a peer review process.

Prizes vs. grants

Once a program is approved, the program manager is allocated funding and decides how to spend it. Different methods are suited to different situations. Being aware of the incentives of the people who can be expected to complete the task is important in choosing whether to assign the money as a prize or as a grant.

The barriers to entry are an important consideration when choosing between a prize and a grant. Prizes are good at filling in the gaps in incentives to bring in players from outside the mainstream – for example, the X-Prize was intended to garner entries from people outside of the aerospace industry. A grant, on the other hand, limits the participants, as only certain people who are willing and able to go through all the work of contracting with the government and applying for the grant will put in the effort. This method can work well if you already have an idea of everyone who might be able to accomplish your task, as you can set up the grant to incentivize those specific people. If you have no idea who might do it, it is difficult to know what motivations people will act on and so a prize is a good idea.

Getting the right level of specificity with requirements is also important. Under-specification can lead to the result not being what you're looking for, and over-specification can stifle creativity. Prizes must be extremely specific, since the rules of the contest can't be changed once launched.

Potential ideas for philanthropic funding

There are areas where philanthropists have the potential to be highly impactful in ways the government cannot, because philanthropists are not similarly constrained by public policy and legal issues. The Gates Foundation, for example, has been able to think about the health issues in the third world in an unconstrained way, which has given them a lot of flexibility and freedom. Some other areas this idea may apply to are privacy and international collaboration.

One example is the risk of a global disease pandemic. It would be difficult for the government to do relevant research on people's traveling patterns without raising red flags, but philanthropic organizations aren't as restricted by public fears of government overreach.

Another idea is to try to create a system that encodes a lot about what researchers are working on: what algorithms they're using, what ideas they're testing, etc. People might fear a loss of privacy but the knowledge of what others are working on and thinking about could be very helpful. Privacy issues could be an obstacle to DARPA's working on this, but a philanthropic organization may be better suited.

Other people for GiveWell to talk to:

- Arati Prabhakar **–** director of DARPA
- John Fredrickson – currently working on a government project to improve electronic health care records
- Dr. Steinhardt can recommend someone to talk to at the Potomac Institute, which advises congress on matters of science and technology and investment