

A conversation with Robin Hanson on 9 October 2013

Participants

- Robin Hanson—Associate Professor, Department of Economics, George Mason University; Research Associate, Future of Humanity Institute at Oxford University
- Nick Beckstead—Research Fellow, Future of Humanity Institute at Oxford University; Board of Trustees, Centre for Effective Altruism

Summary

Purpose of the call: I organized this call to learn about government policies that Robin Hanson believes may reduce global catastrophic risk (GCR), for the purpose of identifying areas for deeper investigation later. I was especially looking for proposals about how we could better use prediction markets and related institutions and policies, and proposals about how to increase the probability of survival if there is a global catastrophe.

Why this person: Robin Hanson has published work on prediction markets and global catastrophic risk. I also had an existing relationship with him.

Topics discussed: We discussed refuges (to help people survive global catastrophes), disaster-contingent pricing for key infrastructure (to increase incentives to keep key infrastructure running in disaster scenarios), refuge markets (to help predict catastrophes), government prizes for people who prevent or help during major disasters (to provide better incentives to prevent and help during disasters), and research testing prediction markets (to generally improve predictions).

Policies that increase the chance of eventual recovery from a global catastrophe

Refuges

By *refuges* Robin means places where:

1. people know they can go to if there is a global catastrophe
2. people can remain safely until danger passes
3. there are a lot of resources that will help them eventually survive and eventually recover

Many technologies cannot be produced and maintained unless society is big enough, even if the technical knowledge necessary to understand how to build and maintain the technology is available. Because of this, Robin believes that a global catastrophe which reduced global population to very low

levels would set substantially set back technological progress until the size of the human population became very large again.

Where are we now?

Most refuges are military or government, and it's hard to get information about them. Many people with private refuges don't want to share information about their refuges because they wouldn't want other people to be able to find their refuges if there were a catastrophe that made the refuges necessary. However, private refuge websites do make claims about how many units they've sold, specify designs, and state recommended stocking of their refuges. Robin blogged about private refuges at <http://www.overcomingbias.com/2010/09/disaster-shelters.html>.

Nuclear submarines are constantly staffed, stocked, and ready to spend a few months in isolation. They could also restock in naval shipyards if necessary.

Could we be more prepared?

Robin doesn't know which areas we could be more prepared in, but suggested looking at the highest levels of coverage along specific dimensions and looking into extending coverage. For example, we might want to make sure that there are some places that have five years of food; could survive for five years on a clean, independent air supply; had a very large stockpile of materials that could be used by a group of survivors; or had people trained to operate the refuge and assist with recovery following a catastrophe.

Possibly, naval shipyards could have larger stockpiles of key resources which nuclear submarines could use to restock in the event of a global catastrophe.

People to talk to

I asked Robin who had written the most detailed discussion about refuges. He said there that private companies that sell refuges would probably be the best people to talk to. He suggested there may be military documents on refuges, but suspects the most helpful ones would be classified.

The refuge company Robin blogged about was Radius Engineering. They were covered on p. 112 of the October 2010 edition of *Wired*.

Disaster-contingent pricing for key infrastructure

What is it?

Key aspects of our infrastructure (such as telephones and the electricity) are more fragile than they need to be. Prices for some of these goods and services do not go up very much in disaster settings, often because regulations would not allow them to increase their prices. For example, a large enough geomagnetic storm might destroy transformers in electrical grid. These transformers would take months to replace, and the costs from the interruption would be extreme. For relatively small amounts of money, spare transformers or ground-induced current blocking devices could be purchased to mitigate this risk. However, electricity providers must accept prices from regulators, and, given those prices, currently don't see it as in their financial interest to make these preparations.

One strategy to harden the infrastructure would be with specific regulatory requirements. For example, a simple requirement might be that, prices for essential services go up by a factor of 10 if a disaster is officially declared by an appropriate government body, unless this requirement is explicitly negotiated out of contracts. If an electricity provider thought they could charge much higher prices for providing key services in a disaster scenario, they might make more appropriate preparations for such risks. For instance, perhaps electricity providers would see it as worthwhile to prepare better of geomagnetic storms. And other providers of key infrastructure might make analogous preparations for many risks that regulators wouldn't be able to think of.

How would it help with global catastrophic risk?

If there is a risk that a global catastrophe would destroy key infrastructure in a large portion of towns or cities, it will be very important that some towns or cities maintain key infrastructure. Disaster pricing would provide a stronger incentive for providers of key infrastructure to ensure that if such a global catastrophe occurs, key infrastructure continues to function. Robin believes this would improve recovery efforts.

Where are we now?

Currently, possible responses to many disasters are centrally planned by government agencies. Robin believes his market-based approach would provide additional flexibility.

People to talk to

Many economists have argued that stopping "price-gouging" in disaster contexts is bad. Robin doesn't know of anyone who has advocated this form of disaster-contingent pricing. Robin doesn't know whether he would know about these people if they existed. Some people in the disaster response community might have written about this.

Refuge markets

What is it?

Robin Hanson has advocated using prediction markets as an input into decision-making. Standard prediction markets are much less helpful for predicting global catastrophic risk than other events because dead people can't pay out on losing bets.

Robin advocates using *refuge markets* where purchasing a ticket allows one to use a refuge if a certain event happens. The prices of these tickets for different events would imply specific ratios of probabilities for different events because a refuge ticket for an event with probability x should cost $1/10$ th as much as a refuge ticket for an event with probability $10x$.

Robin explained this idea in more detail, together with a few refinements, here:

<http://www.overcomingbias.com/2008/07/refuge-markets.html>.

How would it help with global catastrophic risk?

Refuge markets would help assess the probabilities of different disasters, which would assist with planning. Policy-conditional refuge markets could help determine which policies would do the most to decrease catastrophic risk.

People to talk to

Robin is confident that this has no other prominent advocates.

Government prizes for people who prevent or help with major disasters

After disasters or possible disasters, a committee could identify people who did a great job of preventing a disaster or helping in a disaster scenario, and these people (or their families) get large rewards. For example, people like Stanislav Petrov, who helped prevent an erroneous retaliatory nuclear attack against the United States, could receive such rewards. This would help reduce global catastrophic risk by providing stronger private incentives to do so.

People to talk to

There is an economics literature on prizes in general. Presumably some humanitarian or other organizations give out (possibly non-cash) awards.

Research to test prediction markets

I asked Robin about the value of research testing prediction markets. To my surprise, he was not enthusiastic about funding such research. In his view, there is strong evidence that prediction markets would improve predictions, but policymakers don't want to use them.

He described an example of a missile test agency in the federal government. A key thing to forecast in these cases is to forecast whether tests will actually happen. Prediction markets were tested against conventional forecasting methods and outperformed the conventional methods. Yet, prediction markets are not used. Robin's explanation of why this happens is that people are not interested in accurate forecasts.

Robin would be enthusiastic about field trials of prediction markets. He is not enthusiastic about more math papers, lab studies, or software engineering.

For \$1M, Robin would like to set up a prediction market that sold bets on what would happen to different companies' stock prices if they fired their CEOs. The attraction is that people would have appropriate incentives to use this information if it were available, and it would be a prominent public test of whether prediction markets work in important practical settings.

People to talk to

1. Iowa Electronic Markets
2. Consensus Point (a prediction market company where Robin is Chief Scientist)

Questions discussed in our meeting

1. What could a government do to more effectively use prediction markets, decision markets, or other IARPA prediction tournament-type stuff?
 - a. For each proposal, we'd like to discuss:
 - i. What is it?
 - ii. Where are we now?
 - iii. Where would you like us to be?
 - iv. How could we get closer to where we should be? What roadblocks are there?
 - v. What could I read to learn more about this?
 - vi. Who else could I talk to?
2. What could a government do to increase the probability that, were there a global catastrophe like a nuclear war or a massive collapse of the human population, civilization would eventually recover (e.g. refuges, seed banks, wood-to-cellulose conversation)?
 - a. For each proposal, we'd like to discuss:
 - i. What is it?
 - ii. Where are we now?
 - iii. Where would you like us to be?
 - iv. How could we get closer to where we should be? What roadblocks are there?
 - v. What could I read to learn more about this?
 - vi. Who else could I talk to?
3. What other policies or institutions would you like to see in place, especially policies relevant to managing global catastrophic risks?
 - a. For each proposal, we'd like to discuss:
 - i. What is it?
 - ii. Where are we now?
 - iii. Where would you like us to be?
 - iv. How could we get closer to where we should be? What roadblocks are there?
 - v. What could I read to learn more about this?
 - vi. Who else could I talk to?