

# A conversation with Tyler Cowen on 8 April 2014

---

## Participants

- Tyler Cowen—Professor of economics at George Mason University
- Nick Beckstead—Research Fellow, Future of Humanity Institute at Oxford University; Board of Trustees, Centre for Effective Altruism

## Summary

**Purpose of the conversation:** I contacted Tyler to learn about his perspectives on existential risk and other long-run issues for humanity, the long-run consequences of economic growth, and the effective altruism movement. I believed Tyler was likely to have interesting views on these issues that the effective altruism community was generally unaware of. Tyler has written about some policy issues for people who care deeply about distant future generations, we have many friends in common, and he supports GiveWell.

We had a very wide-ranging conversation, but two major themes were caring about the distant future and the effective altruism movement. Tyler agrees with Nick about the importance of long-run outcomes for humanity and the importance of path-dependent outcomes in that framework. However, Tyler expressed significant skepticism about the value of philosophical work on these problems and the value of speculating about the impacts of future technology, at least in comparison with more historically grounded, tractable issues such as innovation, international conflict, geopolitics, and pandemics. He also expressed much more pessimism about humanity's prospects for surviving and thriving for a very long time (e.g. millions of years or more). Tyler supports effective altruism, but doesn't see it as extremely important in the grand scheme of things; others, especially innovators, will be much more important in his view. We also discussed a few philosophical issues related to historicism, literature, and rational choice ethics.

*These notes reflect the main points made by Tyler in the course of the conversation.*

## Tyler on caring about the distant future

### Tyler's views on the long-run future vs. Nick's views

Nick had a look at Tyler's paper "Caring About the Distant Future," which argues that, from a long-run perspective, the main public policy objective should be maximizing the sustainable rate of economic growth. Tyler had a look at Nick's dissertation prior to the conversation, but not a close one. He agreed that "optimize path-dependent aspects of the distant future" was a helpful handle for thinking about

how we should hope to change the distant future. This is consistent with how Tyler would now articulate his view, and thought some of the differences in how we framed our arguments reflects the fact that Tyler writes for economists whereas Nick does not.

## **Long-run consequences of increasing the rate of economic growth**

We discussed this issue, but the recording was damaged and I haven't attempted to reconstruct it.

## **Pessimism about the future**

Tyler is optimistic about growth in the coming decades, but he doesn't think we'll become uploads or survive for a million years. Some considerations in favor of his views were:

1. The Fermi paradox is some evidence that humans will not colonize the stars.
2. Almost all species go extinct.
3. Natural disasters—even a supervolcano—could destroy humanity.
4. Normally, it's easier to destroy than to build. And, in the future, it will probably become increasingly possible for smaller groups to cause severe global damage (along the lines suggested by Martin Rees).

The most optimistic view that Tyler would entertain—though he doubts it—is that humans would survive at subsistence level for a very long time; that's what we've had for most of human history.

## **Existing efforts to reduce existential risk**

People doing philosophical work to try to reduce existential risk are largely wasting their time. Tyler doesn't think it's a serious effort, though it may be good publicity for something that will pay off later. A serious effort looks more like the parts of the US government that trained people to infiltrate the post-collapse Soviet Union and then locate and neutralize nuclear weapons. There was also a serious effort by the people who set up hotlines between leaders to be used to quickly communicate about nuclear attacks (e.g., to help quickly convince a leader in country A that a fishy object on their radar isn't an incoming nuclear attack). This has been fixed in other countries (e.g. US and China), but it hasn't been fixed in other cases (e.g. Israel and Iran). There is more that we could do in this area. In contrast, the philosophical side of this seems like ineffective posturing.

Tyler wouldn't necessarily recommend that these people switch to other areas of focus because people motivation and personal interests are major constraints on getting anywhere. For Tyler, his own interest in these issues is a form of consumption, though one he values highly.

## **Artificial intelligence: could trying to shape the future of AI be an effective way to change the long-run future?**

Tyler expects that artificial intelligence will have substantial effects on wages and GDP. But he doesn't think it will fundamentally transform the human condition. Tyler wouldn't be shocked if we had 5-10 more years of Moore's law and improvements started to slow down. People treat Moore's law as if it were a law of nature, but it isn't.

Tyler is very skeptical about uploads (aka “whole brain emulation”), especially claims about uploads in the next few decades. He’d maybe give them a 1 in 1000 chance of happening. In his view, uploads are just an idea that some people came up with, most ideas don’t work, and most institutions are dysfunctional. Those truths seem more important for thinking about the distant future than any complicated arguments for the feasibility and importance of uploads.

Tyler doesn’t think that trying to change the development of AI is likely to be an effective way of trying to change long-run outcomes for humanity. If he thought it was more likely that AI would fundamentally transform the human condition, then he might change his opinion about this, though it’s unclear to him how anyone could control AI in such a scenario.

### **What types of issues would be more important than the issues that people in existential risk think about?**

International conflict, weapons of mass destruction, geoengineering, and pandemics may be more important.

Within international conflict, we could go for the obvious trouble spots—e.g. India-Pakistan, Israel-Iran, North vs. South Korea, sanctions policy on Russia, biological warfare, etc. People can make progress on these issues. E.g. many trouble spots have gotten better, such as conflict in Ireland has improved or Africa has become more peaceful. Tyler believes this kind of progress reduces existential risk.

Within pandemics, we could try to improve disease surveillance.

### **What about politics?**

Only a small percentage of people are good at politics, and in that respect it’s similar to trying to innovate. But being good at politics is negatively correlated with certain type of philosophical thoughtfulness. This is less true outside of the US, e.g. in Canada.

## **Effective altruism**

*Question: What are your thoughts on the effective altruism movement in general—how familiar are you with it? If you are, are there things you wish we were doing that we aren’t doing or things you wish we were doing differently?*

Tyler likes it, supports GiveWell on his blog, and donates to GiveDirectly. But it’s small potatoes in comparison with, say, innovation.

Tyler’s intuition is that improving marketing is a key issue for effective altruism, rather than fine-tuning where people should be giving.

## Other philosophical issues

### The problem of “fanaticism”

In Tyler’s view, the constructs underlying the problem of fanaticism and the Repugnant Conclusion and at best work locally, and not globally. In sufficiently exotic settings, these constructs fall outside the scope of meaningful discourse. (These are issues discussed in Nick’s dissertation, and that’s why they were discussed.)

### Historicist philosophy

Tyler thinks about the future and philosophical issues from a historicist perspective. When considering the future of humanity, this makes him focus on war, conquest, plagues, and the environment, rather than future technology.

He acquired this perspective by reading a lot of history and spending a lot of time around people in poor countries, including in rural areas. Spending time with people in poor countries shaped Tyler’s views a lot. It made him see rational choice ethics as more contingent. People in rural areas care most about things like fights with local villages over watermelon patches. And that’s how we are, but we’re living in a fog about it.

### Rational choice ethics and the “Straussian truths of the great books”

The truths of literature and what you might call “the Straussian truths of the great books”—what you get from Homer or Plato—are at least as important as rational choice ethics. But the people who do rational choice ethics don’t think that. If the two perspectives aren’t integrated, it leads to absurdities—problems like fanaticism, the Repugnant Conclusion, and so on. Right now though, rational choice ethics is the best we have—the problems of, e.g., Kantian ethics seem much, much worse.

If rational choice ethics were integrated with the “Straussian truths of the great books,” would it lead to different decisions? Maybe not—maybe it would lead to the same decisions with a different attitude. We might come to see rational choice ethics as an imperfect construct, a flawed bubble of meaning that we created for ourselves, and shouldn’t expect to keep working in unusual circumstances.

## Discussions topics sent to Tyler prior to our conversation

1. The ideas in your “Caring About the Distant Future” and the extent to which they are sensitive to assumptions about long-run economic growth (e.g. whether it is more exponential or logistic), and your opinions about the importance of mitigating global catastrophic risks as a way of helping distant future generations.
2. Your thoughts on how contingent/inevitable you think sustained economic growth is in the distant future—e.g., what could/couldn’t realistically permanently curtail long-term economic growth?

3. Your thoughts on which issues are very important from a utilitarian perspective that counts all future generations equally—are there specific policy issues or career areas that you think deserve much more attention from this perspective?
4. Your thoughts on the effective altruism movement in general—how familiar are you with it? If you are, are there things you wish we were doing that we aren't doing or things you wish we were doing differently?