

A conversation with Dr. Kevin Esvelt, March 31, 2015

Participants

- Kevin Esvelt, PhD – Technology Development Fellow, Wyss Institute for Biologically Inspired Engineering, Harvard University
- Nick Beckstead, PhD – Research Analyst, Open Philanthropy Project

Note: These notes were compiled by the Open Philanthropy Project and give an overview of the major points made by Dr. Kevin Esvelt.

Summary

The Open Philanthropy Project spoke with Dr. Esvelt of the Wyss Institute as part of its investigation into the potential ecological risks of synthetic biology. Conversation topics included the potential applications and risks of gene drives, recoded organisms, and other bioengineering applications.

Potential risks from synthetic biology

In terms of the severity of threat, Dr. Esvelt is most concerned about:

1. **Risks from pathogens** – This includes pathogens directed at crop systems and humans. Intentionally and naturally evolved pathogens pose roughly equally large risks. An outbreak of synthesized smallpox would pose a serious threat to civilization.
2. **Artificial photosynthesis** – If humans are able to build something that will be more efficient and effective than nature and has the capacity to move across species, it will pose serious ecological risks.
3. **Gene drives** – Gene drives don't pose a catastrophic risk because they are generally reversible and it would be difficult to target major crop systems or humans directly.
4. **Mirror organisms and recoded organisms** – It will be easy use molecular containment techniques to control mirror organisms and other recoded organisms.

Synthetic biology poses more of a threat to crop ecosystems than to wild ecosystems. It is more difficult to permanently alter wild ecosystems. Gain-of-function research on pathogens poses more of a threat than most synthetic biology applications.

Gene drives

Gene drives are unlikely to pose a catastrophic risk to global ecosystems because they are detectable and reversible. Gene drives also require multiple generations to have an effect. The fewer animals that are released, the longer a gene drive will take to eradicate a population. However, a smaller release also makes the gene drive more difficult to detect. Small organisms that reproduce quickly are not necessarily

more vulnerable to gene drives because they also recover more quickly and tend to have large population numbers.

It is unlikely that gene drives could ever be an effective weapon and it's not clear how they could be used to target a single nation. National security will probably have to pay attention to gene drives in the next decades, but they are unlikely to pose a major threat. Dr. David Gurwitz, a geneticist and biochemist at Tel Aviv University, has argued that gene drives should be kept confidential because of the risk that someone might engineer a mosquito that could deliver a toxin. While this project is technically feasible, it is on par with the most ambitious bioengineering projects worldwide. It would be very difficult to engineer a mosquito that a) produced a toxin, b) did not poison itself, c) delivered the toxin to humans effectively, d) spread this capability using a gene drive. Even if someone succeeded with this, it would be fairly easy to reverse. With gene drives, defense trumps offense.

If researchers detect an unwanted gene drive, it can be easily sequenced. With that information, it is possible to build an immunizing reversal drive that would override and block an unwanted gene drive. The reversal drive would spread through the unaffected members of the population and immunize them against the unwanted gene drive.

In a worst-case scenario, someone might use a gene drive to target an unexamined population. If no one else detected the gene drive until the altered organisms had reproduced enough to make up 50% of the population and ecological effects were already visible, it might be difficult to reverse.

It is exceedingly unlikely that anyone would accidentally build or release a gene drive that could crash a population. If this were to happen, the main concern might be the negative public reaction. Safe guards can be used to limit risks of negative effects. For example, scientists can:

- Release organisms that have the desired genetic alterations but without the actual gene drive to see if there is an effect in the wild.
- Have a reversal drive ready.

There is a possibility that an unwanted side effect (e.g., a gene drive affecting a keystone species) would occur only after a certain threshold is reached. However, existing biological control efforts, such as introducing new predators to an area, are arguably at least as risky as altering a species with a gene drive. Anything done to an existing species with a gene drive is probably less risky than introducing a new species.

Gene drives in crops

In the developed world, crops are fairly immune to gene drives. In order to spread, gene drives require wild mating. 80 – 90% of major crops are grown from engineered seed, which is the product of controlled breeding or precision genetic engineering. Large companies, such as Monsanto and Syngenta, keep the genetics of

their major crops under tight control and run seed farms that produce most of the seeds for the following year.

In order for a gene drive to spread in this environment, it would have to infiltrate a seed farm and remain undetected for multiple generations. The large companies that run seed farms devote a significant portion of their research and development budget to surveying genetics and complete full genome sequencing often. A gene drive would be detected before seeds for the next generation of crops were sold.

Crops in the developing world are more vulnerable to gene drives because farmers are more likely to reuse seeds. However, the developing world produces proportionally less of the world's food supply.

Gene drive regulation

The Wyss Institute has been pressuring the National Institutes of Health (NIH) to develop a set of regulations and rules for gene drives. The Gates Foundation has also funded a National Academy of Sciences (NAS) Study and Evaluation to generate recommendations for gene drive usage. The NIH is likely change regulations when the NAS report is completed.

A research group recently developed a *Drosophila* (fruit fly) gene drive using only conventional barrier confinement rather than employing additional molecular confinement strategies that are not vulnerable to human error. This is concerning as research with pathogens has shown that barrier confinement is fallible, which in the case of gene drives could readily result in the accidental alteration of a wild population. Discussions between relevant groups of scientists seeking to establish and widely publicize interim guidelines for the field while the NAS report is completed are currently ongoing. Regulation will be important to mitigate this sort of risk and to spread knowledge about safe applications of gene drives.

Public attitudes

Because it is a collective technology and any use of gene drives will affect the common environment, it will be important to include the public in discussions about gene drive applications. There should be mechanisms that allow for public directed research, as well as feedback and criticism from citizens.

If there was an accidental release of a gene drive, the public backlash could prevent any future applications, even if the release didn't have negative ecological effects.

Environmental groups may oppose the use of gene drives. These often oppose technologies that involve genetic modification. Dr. Esvelt is optimistic that these groups might not oppose gene drives given appropriate education and outreach efforts. Gene drives can also be used to promote conservation efforts, e.g. by using a gene drive to prevent a species from going extinct. Dr. Esvelt is attending a conference next month on new genetic technologies for conservation.

The public mounted a strong response against genetically modified (GM) food, but it hasn't had the same negative response to using synthetic biology to produce

medicines or to create fluorescent transgenic fish (which are now available commercially).

Some people are also concerned that gene drives rely on the CRISPR/cas9 system, which uses prokaryotic DNA to edit genes. Most eukaryotic organisms haven't had common access to these genes, so it is possible that this could alter evolutionary trajectories (though this would likely take thousands or tens of thousands of years). Species that now have access to the new prokaryotic genes may become more evolvable.

Potential benefits of gene drives

Beyond conservation efforts, there is the potential to use gene drives to eradicate disease (especially mosquito-borne diseases like malaria), combat locust swarming, and replace toxic pesticides. Dr. Esvelt would like to build a global community of citizens from different walks of life who are interested in gene drive research. He hopes to promote these potential positive applications and bring synthetic biology into museums in order to share it with the public. Synthetic biology will only be accepted if there are high levels of transparency and public engagement.

Building this sort of community and ensuring that gene drives are applied in as safe a manner as possible will require more than scientific expertise. Ideally within the next year, Dr. Esvelt would like to found a non-profit with a team of three to five people to create an open platform for community-level discussion, participation, and guidance of technology development and application in the area of gene drives.

Dr. Esvelt is also applying for faculty jobs now and plans to ensure that his future lab will be fully transparent with its research agenda. Ideally, advocating for full transparency in gene drive research and other synthetic biology applications will encourage broader reform in scientific research practices. Across all scientific research, experts are becoming more and more removed from the common citizen. Bureaucratic and political obstacles, like the delays of gatekeeper journals and the pressure to keep results secret, are impeding scientific progress.

Potential ecological risks from bioengineering

If an organism could be bioengineered to obtain a large fitness advantage, it could potentially become an invasive species. However, conferring a large fitness advantage to an engineered organism is challenging. There is only one commercial GM trait that might confer a fitness benefit in the wild. In general, nature is better at developing herbicide resistance, faster growth, and drought tolerance. Human engineered traits are unlikely to spread in the wild, but they could potentially spread in agricultural ecosystems, which are monocultures. A few possible mechanisms for increasing the fitness of organisms are discussed below.

Genomically recoded organisms

Progress on recoded E. coli

The laboratory of Dr. George Church, Professor of Genetics at Harvard Medical School, has produced a recoded strain of E. coli. The Church Lab changed a single codon in the E. coli genome, which will grant it trivial resistance to bacteriophages (viruses that infect bacteria). Thus far this is the most advanced successful recoding project.

In order to recode an organism, scientists need to:

1. **Free up some of the existing codons** – The Church Lab changed all 314 instances of the TAG stop codon to another stop codon, TAA. This freed up one of the 64 codons in the E. coli genome.
2. **Add new functionality** – The Church Lab then put in necessary genes to code for an unnatural amino acid. Following this recoded genome, the cell will incorporate this new amino acid into its proteins. This strategy will also keep the new organism contained, because it will be dependent on the laboratory media that provides the new amino acid.

In order to add the new genetic material, the Church Lab used a technique called recombineering.

- They produced a DNA oligonucleotide (a short DNA molecule) that binds to the DNA reverse strain during the synthesis process.
- Mutations in the DNA oligonucleotide are incorporated into the new strain of DNA.
- The Church Lab produced one DNA oligonucleotide for all 314 instances of the TAG stop codon and completed many rounds of synthesis with these oligonucleotides and combined and shuffled the results until they produced one E. coli strain that had all 314 TAG stop codons replaced with TAC.

Because it is relatively easy to make new organisms dependent on artificial, laboratory provided amino acids, scientists are not that concerned with these recoded organisms getting out into the broader environment.

Genetic recoding comes at a fitness cost. The Church Lab worked with the TAG codon because it was the least used codon in the entire genome. However, in the process of changing only that single codon, scientists created around 350 accidental mutations across the genome. The recoded E. coli's doubling time is 80 minutes, four times that of wild E. coli. This decrease in fitness is likely the result of the accidental mutations and the altered TAG codons. Many important processes that influence gene expression and regulation are dependent on the gene sequence, such as:

- Rate of transcription and translation
- Transcription factor binding
- Secondary structure of mRNA
- RNA binding proteins

Recoding a gene disrupts these processes. Engineering extensive resistance to bacteriophages will require scientists to free up multiple codons and shuffle them extensively. For example, a codon that once encoded the amino acid leucine might

be changed to encode arginine. When scientists attempt to change codons that code for amino acids, rather than stop codons, the negative fitness effects could be even worse. A large fraction of the ensuing mRNA will have a different sequence. It will also be more technically difficult to change multiple codons across the genome (and scientists will still have to alter every instance of the chosen codons across the genome).

Scientists are only beginning to understand codon optimization and the cellular rules that direct gene expression. For example, it was recently discovered that the first 50 amino acids have a large effect on gene expression because they control the structural characteristics of the resulting mRNA. As these processes are better understood, scientists can begin to replace codons that are not very good at coding for specific amino acids.

Recoding whole organisms

If a very large fraction of an organism were sufficiently genomically recoded, it would become incompatible with natural viruses, and therefore resistant to phages. Bacteriophages are a major factor constraining the growth of some ocean bacteria, such as cyanobacteria. Phage resistance could, in theory, greatly enhance the fitness of such an organism, potentially resulting in an invasive species. However, bacteriophages are not the only factor constraining the population growth of cyanobacteria because they still need to find resources. Moreover, genomically recoding a single codon in the *E. coli* genome was challenging and carried substantial fitness costs. Therefore, it would likely be very technically challenging to genomically recode a large portion of the DNA of any variety of ocean bacteria, and there would be major challenges in addressing the fitness losses associated with the recoding process. For this reason, it will likely take at least 10 years—and perhaps much longer—before we have the technical capacity needed to create genomically recoded bacteria that would pose significant risks as invasive species.

Artificial photosynthesis

Beyond adding resistance, it might also be possible to provide a fitness benefit by improving an organism's ability to acquire resources. As a general rule, scientists can't beat natural evolution at this, but photosynthesis may be the exception. It is an incredibly inefficient process. Many laboratories are working on designing artificial photosynthesis processes that would improve upon it.

For example:

- c4 carbon fixation is more efficient than the more common c3 carbon fixation process. It avoids wasting water.
- RuBisCO (ribulose-1,5-bisphosphate carboxylase/oxygenase), an enzyme involved in the first step of carbon fixation process, is not very efficient.

Evolution has likely not come up with a solution because it requires too drastic a change in the genomic sequence. Improving upon photosynthesis will be a

challenging project, but with the ability to recode organisms to utilize artificial amino acids, it is not unsolvable.

If someone does successfully engineer a more efficient photosynthesis, it could cause serious ecological disruptions. Cyanobacteria with enhanced photosynthetic capabilities could replace existing bacteria, with the increase in productive capacity substantially altering associated ecosystems. Artificial photosynthesis would be useful in agriculture, but it would be difficult to ensure that engineered organisms didn't spread to the wild or that crops don't become incredibly invasive. If altered genes were able to transfer horizontally between species, artificial photosynthesis could become a major ecological disrupter across multiple species.

It is difficult to estimate when artificial photosynthesis might become a reality. This won't occur in the immediate future, but perhaps sometime in the next 10-50 years. More sophisticated control strategies capable of ensuring the trait remains confined to crops are under development.

Mirror organisms

All molecules possess a chirality (i.e., they are either "left handed" or "right handed" – one part of a mirror image pair). If scientists were able to engineer organisms of the opposite chirality, this would make them immune to existing viruses. However, it might also affect their ability to utilize food sources effectively. In general, organisms will prefer molecules of one chirality. For example, most sugars are "D" sugars. Organisms can metabolize sugars of the opposite chirality ("L" sugars), but not as effectively. However, because the photosynthesis apparatus should work just as well in an opposite chirality organism, this metabolic issue would not affect phototrophs.

Scientists have yet to successfully build a mirror organism. In order to do so, scientists need to:

- Synthesize DNA of the opposite chirality
- Synthesize DNA polymerase of the opposite chirality
- Synthesize ribosomes and ribosomal proteins of the opposite chirality
- Create a cell with these engineered elements that will "boot up" effectively

Scientists have successfully synthesized opposite chirality DNA and DNA polymerase. No one has successfully made a ribosome of the opposite chirality – this will be difficult because ribosomes are especially large and complicated. Scientists have not successfully engineered a completely new cell (even without altering its chirality). This requires purifying ribosomes, polymerases, and metabolites from a cell, building a ghost cell that has nothing in it, putting in the engineered materials, and starting it up so that the cell will carry on the necessary processes. It will likely take at least an additional 10 years until we are capable of creating opposite chirality organisms.

As with recoded organisms, scientists will be able engineer mirror organisms to be dependent on artificial amino acids in order to keep them contained in the laboratory or other controlled environments.

Other functions

It is possible that providing an organism with the ability to synthesize a new unnatural amino acid might prove to be a fitness benefit in the wild, but it hard to imagine what that fitness advantage might be. Because evolution is hard to predict, there is a risk that this new chemical functionality could lead to unexpected results.

Theoretically, it is also possible to genetically engineer new organisms from scratch. However, it will probably be easier to recode and replace the genome in parts. It's unlikely that cell engineered from scratch will "boot up" immediately. Engineering new organisms from scratch or extensively recoding organisms is a big project that would probably take at least five to ten years.

Potential threats from pathogens

Intentionally evolved pathogens

Because crop systems are monocultures, they are much more vulnerable to new pathogens (whether naturally evolved or engineered) than to gene drives. It would be much more difficult to engineer or evolve a pathogen from scratch that would affect humans. With crops, it would be fairly easy to get a hold of some of the dominant monoculture seeds and a few greenhouses and start testing various pathogens. Newly evolved pathogens pose more of a risk than newly engineered pathogens. It remains unclear how DNA synthesis technologies will interact with natural evolutionary strategies.

Naturally evolved pathogens

Naturally evolved pathogens pose as much as a threat as those developed by humans. If a pathogen destroyed a large fraction of any given crop in successive years, it could pose an existential risk.

For example, 40-60% of wheat crops used to be lost to wheat rust (a fungal disease) annually. In the 1950s, Dr. Norman Borlaug figured out how to combine different resistance genes from wheat strains into a single high-yielding variety. His techniques have worked for the last 50 to 60 years, but farmers are now beginning to see rust that breaks through those defenses.

Given modern transport, it is possible that a wheat rust could spread worldwide. The Irish potato famine was caused by a fungal blight similar to a wheat rust. In principle, this could happen again to any grain crop. Researchers need to keep an eye on potential attackers and new resistance genes that can be added to future generations.

Smallpox threat

Synthesized smallpox poses the biggest threat of all engineered pathogens from a synthetic biology perspective. It is relatively easy to synthesize a virus because it only requires inserting the virus genome into a cell. The smallpox sequence is publically available.

It is unlikely that anyone would synthesize smallpox from scratch, though as technology improves, this is increasingly feasible. Someone could contract or take over a synthesis company and redirect its synthetic capabilities. Someone could hack into a company, sneak in a synthesis order and simultaneously erase any record of it. DNA synthesis companies in China and elsewhere don't run screening algorithms to check for dangerous viruses. There needs to be more attention and diplomatic pressure in this area.

Creating a central repository where records of all DNA synthesis orders are automatically deposited and screened would help mitigate these risks. Even if the orders remained anonymous, maintaining a central database (or multiple databases) with independent security would be a big improvement. Currently, responsibility for screening for dangerous pathogens is distributed across synthesis companies and they are already too vulnerable.

An outbreak of smallpox could pose an existential risk to civilization at large. On an unprotected population, smallpox has a 30-80% mortality rate. Smallpox may not pose the largest threat to the United States. The United States already has 350 million vaccines ready to go and mortality rate in a vaccinated population is fairly low. However, if someone was able to synthesize smallpox, he or she could conceivably engineer or evolve a vaccine-resistant version.

Gain of function research on pathogens with pandemic potential

Gain-of-function research on pathogens with pandemic potential (such as H5N1) has created a schism in the scientific community and debates around it have become very ideological. The release of pathogens created in this way could result in a global catastrophe. Dr. Esvelt strongly believes that scientists should not work with potentially civilization-ending pathogens, but there is an ongoing debate in the scientific community about the ethics of this kind of research.

Synthetic biology community

For the most part, the scientists who are currently working on recoded organisms, gene drives, and mirror organisms are in George Church's lab or among his former students.

Many scientists are interested in mirror organisms and may have a side project in their labs devoted to it, but there are few labs focused on this project. It has gotten a lot of attention in the popular science press.

There are very few people examining possible counter measures. For example, there are no dedicated funding mechanisms for scientists who want to understand how

rust pathogens infect monocultures. It is unlikely that agricultural companies will fund this research on their own.

Changing the way the government thinks about science would help to improve regulation and funding for smart and safe applications of synthetic biology. An independent think tank with a lobbying arm focused on science policy could achieve a lot of good.

Other people to talk to

- **Dr. Nicola Patron** – Dr. Patron is the head of Synthetic Biology at the Sainsbury Laboratory. She is a plant-engineering expert with a background in plant pathogens.

*All Open Philanthropy Project conversations are available at
<http://www.givewell.org/conversations>*