Transcript of Rationally Speaking #137: Prof. Marc Lipsitch on, “Should scientists try to create dangerous viruses?”

Julia Galef: Welcome to Rationally Speaking, the podcast where we explore the borderlands between reason and nonsense. I'm your host, Julia Galef, and with me today is our guest, Professor Marc Lipsitch.

Marc is a professor of epidemiology and the Director of the Center for Communicable Disease Dynamics at the Harvard School of Public Health.

Marc, welcome to the show.

Marc Lipsitch: Thank you, it's nice to be here.

Julia Galef: Marc has been one of the leading voices warning about the dangers of a particular kind of research, which some people call gain-of-function research. We're going to be discussing, today in this episode, the potential risks of this kind of research, potential benefits as well, and whether or not the scientific community should in fact proceed with this research going forward.

Marc, maybe to kick things off you can just briefly explain what gain-of-function research is, and what has happened in the world in the last four years that makes this an issue now.

Marc Lipsitch: Gain-of-function is a term that is used very broadly in biology to describe an approach to biological experiments where one often uses genetic techniques, or natural selection or artificial selection techniques, to try to add some function to a living organism. Or in this case a virus.

What has been of concern in the last few years is the application of this very valuable, appropriate technique, to study a function that is quite concerning to many people. Which is to add the function of transmissibility to strains of influenza virus that are already very harmful to people that they infect.

Julia Galef: What do you mean by transmissibility?

Marc Lipsitch: I mean contagiousness. The ability to spread from one person to the next.

Of course, you don't do it in people, you do it in ferrets. You take a virus that is already very harmful when a person or a ferret gets infected, and you passage it from one ferret to the next, thereby teaching it genetically how to transmit through the air.

The idea is those are the sorts of changes that would occur if such a virus became able to transmit from person to person through the air.

Julia Galef: The virus -- before this experiment, what kind of transmissibility did the virus have? Not through the air, clearly.

Marc Lipsitch: The virus that has been the focus of most of the recent experiments has been H5N1 bird flu virus, which has infected at least several hundreds of people, basically by very close contact with infected animals.

There may have been occasional spread from one person to the next, but it was very inefficient and not enough to get the virus going as a full-fledged pandemic or epidemic. In its natural form, if it can spread from one person to another it's very inefficient.

Julia Galef: What is the justification for doing this kind of research? What motivated it?

Marc Lipsitch: The idea of this research is that one of the things that we would really like to know about flu viruses is: How it is that they jump from being viruses that transmit basically through the feces of birds through the water, to other birds, infecting the birds' gastrointestinal tracts, not their lungs? It starts out as a bird gastrointestinal virus, roughly speaking, and it occasionally becomes a human virus that transmits from lungs to lungs.

And when it does that it's extremely harmful to humans. We would like to know why it does that, how it does that, and whether we can predict the properties of viruses that are more likely to do that. And take counter measures in order to try to prevent that from happening.

That's the theory. And the concern on the other side is, first of all, that doing that may not be a simple as the proponents suggest. And that in the process we are doing an experiment that doesn't just put a few people at risk, the way other experiments with dangerous pathogens put the technicians in the lab at risk. This kind of experiment, if it went wrong, potentially puts the entire human population at risk. Because the strain of flu that's being created is potentially very transmissible and very harmful to people. The fear is of starting a new pandemic by mistake.

Julia Galef: Right. It sounds like you have concerns both about the potential benefits of this kind of research, whether those benefits are as strong as the proponents claim, and also concerns about the risks.

If we could break down the kind of risks involved here a little bit more, it seems to me like there's at least two kinds. There's the kind of risk where the pathogen, after it has been made more transmissible or more virulent, it escapes the lab. Either accidentally or, in theory, one of the lab workers could intentionally release it, I guess.

On the other hand there's the kind of risk where this sort of research, after it's been shared and published, disseminated, helps people, potentially terrorists, intentionally create more transmissible or virulent pathogens.

Does that seem like the right breakdown? And if so, which one are you pointing to, or both?

Marc Lipsitch: I'm pointing to the first. It's an interesting fact about the way this debate has evolved. Really, the debate centered around the second, the so-called “biosecurity” concern of whether it was a problem to publish the data from any of these experiments. Because it didn't really come to anyone's attention until the work had already been done, so it was too late to ask the question, "Should we do these experiments?"

There was a debate about that. Eventually it was decided to publish the data from the two studies that had been done in 2011, published in 2012, for a variety of reasons.

As those decisions were made, several colleagues and I wrote one paper, and then several other people followed with similar concerns. Stating that while we don't know whether there's a risk from bio-terrorism or not from use of the published sequence, we were quite concerned that accidents happen in even the most respected high-containment labs. On a fairly regular basis.

They don't usually result in human infections -- and most importantly, when they do result in human infections, those infections don't go anywhere typically, because they are working with viruses or bacteria that are not easily transmitted.

The concern is that we're now entering an era where people can make very easily transmitted virulent pathogens, where there's not a lot of immunity in the population. And where the risk really goes well beyond the kinds of risks we've been tolerant of when they apply to one or two people in a lab.

Julia Galef: You, and I think your co-author Alison Galvani have, I believe, tried to estimate -- put some numbers on these potential risks. Can you give us a rough sense of what kind of risk we're talking about, in terms of number of lives? And probability?

Marc Lipsitch: I think the important thing to state at the outset is that we think that the risk of an accident is very small, but that the magnitude is very large. And that the combination of that is something to worry about.

We've been looking at these estimates in a series of different ways. But it seems that from available data on laboratory accidents in high-containment labs in the United States with select agents, which are the more heavily controlled infectious agents that are studied in research labs, about for every 1,000 laboratories working for one year, there are about two accidental infections of laboratory workers. That would be the first step in a chain of events that might lead to a pandemic.

An accidental infection wouldn't necessarily lead to a pandemic, because it might go nowhere or might be contained. But based on mathematical models of how infectious diseases like flus spread, and set parameters relevant to flu, we think that there's somewhere between a 5% and a 60% chance that one of those accidental infections might spread widely.

That's the probability. And when you multiply those together you get somewhere between 1 in 1,000 and 1 in 10,000 probability for every year that's spent in a high-containment laboratory, that there might be an accidental pandemic started.

Julia Galef: Then we multiply that by the number of labs doing this kind of research?

Marc Lipsitch: That's right. And that of course is what's up for discussion. It's in the western world very small right now, because the United States has put a temporary moratorium on funding. And we were the major funder.

But the question is whether it should be allowed to resume. And it's also probably happening elsewhere that we are less aware of. Although some papers have been published from China.

Julia Galef: Can you say a little bit more about this moratorium, just to give people the social context for this debate? This is an unusual moratorium, right, an unusual step to take, for the government to just step in and say, "Please pause these experiments that you're doing, scientific community, until we can figure out how risky this is."

Marc Lipsitch: That's right. Since this is a rationally speaking podcast, rationally speaking, the way to think about risks is to assess them, make a decision about whether they should be taken, and then either take them or not take them. Rather than to take them and *then* question the decision.

But historically, that's how it went. The sequence of events that led up to it really started with the publication of these papers. And it was brought back into the spotlight by a series of accidents and discoveries of protocol violations at major federal laboratories in the summer of 2014. There were three events, the discovery of smallpox at NIH and events involving anthrax and highly pathogenic bird flu at CDC. These are some of the leading labs in the country.

Julia Galef: When you say the discovery of smallpox, you mean that some sample --

Marc Lipsitch: There was a stock of smallpox which was supposedly destroyed in all laboratories worldwide, except two. Many years ago it was discovered that there was a vial of viable smallpox sitting, forgotten, in a cold room at NIH. Which was a protocol violation, because they should have destroyed it. But it was not a ... nobody was at any risk.

The other two incidents were at CDC and involved the exposure of about 80 CDC employees, possible exposure, to anthrax, because of inadequate decontamination. Something we just heard more about was a series of accidents, of inadequate decontamination of army labs in the last few weeks.

Then there was another incident involving sending out the wrong strain of bird flu, supposedly a mild strain but actually a very severe strain, because some vials got switched at CDC.

There was this convergence of multiple events involving human error, circumventing the very high levels of containment that were available in the well-designed labs at CDC. But then undoing all the benefits of that because the agent was handled in a way that it shouldn't have been, because people didn't realize what it was.

Julia Galef: So, none of these incidents were involving the gain-of-function research specifically, but they did increase the probability that we should put on a similar accident happening with the gain-of-function research?

Marc Lipsitch: Exactly right. They did not involve gain-of-function, they didn't even involve, in most cases, the same organisms.

What they did was to focus public attention on something you could learn if you read obscure papers in the American Biosafety Journal, but was not something people knew about -- which was that accidents happen in high-containment labs at a quite high rate, as I described. None of these accidents involved human infections. But two per 1,000 laboratories a year is due.

It focused people on the fact that these pathogens are dangerous and we need to improve our efforts to contain them. But also on the idea that, as I've tried to phrase it, risks we might be willing to accept when they involve one person, or a few people getting sick in the laboratory… We don't like them, but we might be willing to accept them for the sake of biomedical science if they're rare. We might not be willing to accept it if the consequences are for the entire globe instead of a few people.

Julia Galef: As I've understood it, one of the counterpoints is that the risks are just not one-sided. Deciding to be risk averse does not necessarily point to *not* doing gain-of-function research. In that there is already a risk that there will be naturally occurring mutations, or maliciously induced mutations, in some strain of flu virus, that will cause it to be simply more transmissible between humans, and can put us at risk of a pandemic.

And that the gain-of-function research helps us stay ahead of that game, and do various things like develop vaccines, or monitor strains of flu developing around the world, et cetera, to see which ones could be a threat. And that is actually reducing risk. It's not really clear that the risk is lower by not doing the research. What do you think about that?

Marc Lipsitch: That's right. And that's another way of asking the question, what are the potential benefits of this kind of research? It's a complicated question and it depends particularly on what we're comparing this work to.

A very hard question to answer is, what might we forego in terms of scientific knowledge if, instead of doing this work, we did nothing? Or we put the money towards deficit reduction, or towards a bomber or something? It wouldn't buy very much of a bomber.

Then the question is, should we do science or should we not? We know that many scientific discoveries lead to totally un-anticipatable benefits and really great things for human well-being, including health. If the question were, “Should we just ban this research and thereby make a loss to science?”, I think it would be a hard question to predict what the benefits are.

But what would actually happen is that we would do other research. Probably on flu, maybe on other infectious diseases with the relatively small amount of money that's at stake. And so it's really a question of whether we want to do *this* research on flu or other research on flu. Let's just keep it on flu for now.

There the question is whether the marginal benefits of doing gain-of-function research compared to other completely safe, alternative kinds of flu research, are really compelling.

If we frame it as, “What are the unique benefits of gain-of-function research that we can't really hope to gain any other way?” then I think it's a little bit easier to answer the question. I think there are some scientific questions that can only be answered by gain-of-function research. Such as, “If you take the Vietnam strain of H5N1 and put it in ferrets, what is required to make it transmissible between ferrets?” I think the only way to answer that is to make it transmissible between ferrets. And that's been done. That's what the one of the first studies was.

Julia Galef: But surely we weren't interested in that question specifically. We were interested in that as part of the broader question of whether avian flu could mutate into something more dangerous for humans, right? You don't think that question is uniquely answerable by gain-of-function research?

Marc Lipsitch: I think that the question of whether the avian flu can mutate into something that's dangerous for humans, in principle could only be answered in humans, and that's an unethical study to do. Doing it in ferrets perhaps gets us closer to answering the question of how easily transmissibility in ferrets can develop.

The people doing this research recently have begun to say that if the strain that came out of their ferrets was released on the subway, it would not lead to extensive transmission. They've begun saying that it in fact is adapted to ferrets, not to humans.

So there's a bit of a disconnect between the claims of why this is supposed to be beneficial, which is that it's a model for humans, and the claims in response to concerns about risk, which is that it's not actually going to be harmful for humans. Both claims have been made so it's a little bit difficult to disentangle.

Julia Galef: I see. Isn't the fact that the virus was shown to be able to mutate into something transmissible between ferrets -- whereas that had not previously known to be possible -- isn't that at least Bayesian evidence that the strain of the flu could mutate into something transmissible between humans?

Marc Lipsitch: I would say it probably is. But I think that incremental Bayesian evidence is of limited value for making decisions. It does increase our posterior on the idea that we might have a threat from H5N1.

But I think that before that experiment was done, the prudent decision was to put a certain amount of resources into preparations for H5N1. I would say more resources into preparations, that would be useful against any flu pandemic. Because we don't really know which one the next one is going to be, and if you're uncertain of what it's going to be you put more resources towards general purpose solutions.

After that study, the prudent decision is the same decision. I don't think that it's updated our information enough to make any different decision.

Julia Galef: Interesting.

Marc Lipsitch: What the proponents of this work further claim is that as we survey the landscape of the hundreds to thousands of known outbreaks of flu in birds -- and there are obviously many other outbreaks of flu in birds that we don't know about, because we don't have enough surveillance, and in other animals… As we survey those, they say if we know what mutations to look for in the viral genomes, we might be able to prioritize better which flu strains we take action against and which ones we don't.

That's where the question of general purpose, versus specific actions against certain strains, comes into play. The sorts of things we could do against specific strains if we see a strain that we think is really a pandemic risk, like some of the H5N1 strains in Asia have seemed to be over the last decade… is that we can go and kill the chickens that we know of that are infected with those strains. We can develop vaccines, seed stocks against those strains, which gets us somewhat closer to having a vaccine if we need to develop one. Those are the main two kinds of activities.

Whereas general purpose actions would be stockpiling antivirals, working to develop a vaccine that works against all strains of flu, which is a major research program underway in many labs in the world. Or making some

headway we can try… surveillance so that we can deal better with the epidemic when it comes. Those sorts of things.

Of course, we would like to know which strains are most threatening and try to be responsive to those. But given the large numbers of strains that we never see, like the Mexican strain that caused the last pandemic -- We never saw that coming. It wasn't until hundreds of people in Mexico had pneumonia that we knew we had a pandemic on our hands. We didn't have some kind of advanced warning because we weren't looking in pigs in Mexico.

The question is, do we really want to make an even brighter lamppost to search under for our lost keys, or do we want to invest in something that will make us more prepared for whatever it is?

Julia Galef: For those listeners who haven't heard the parable of the lost keys, do you want to tell it, Marc, or should I?

Marc Lipsitch: Yeah, sorry -- a guy was searching under a lamp post for his keys that he had dropped, and someone said, "Why are you looking under the lamp post for your keys? Didn't you drop them over here on the other side of the street?" Ad he says, "This is where the light is. That's why I'm looking here."

So the question is do we want to figure out a better way to interpret the little bit of data that we have? Or do we want to focus our efforts on the very likely outcome that we will not see the strain coming? Therefore, having the best tools in the world to predict its risk level isn't much help. Or do we want to rather focus on strategies for public health that are robust to our being wrong about predictions?

This is a general idea that is out there. Richard Danzig has written about in his article, "Driving beyond the headlights." He's written about the idea that humans have a tendency to try to make predictions, almost a compulsion to try to make predictions. And a tendency, unfortunately, to over-believe those predictions.

And what we should be doing, in his view, is making our decisions much more robust against the possibility that our predictions are wrong. Keep trying to make them, because we can't help it -- but set up our decision making so that the predictable level of being wrong, very often, isn't catastrophic for our decisions.

Julia Galef: There was an interesting point that you made -- I forget where, maybe in the CSER debate -- that I want to talk about now. You said that the debate over whether gain-of-function research should proceed, the answer that you give to that question, involves both your estimate of what the potential benefits are and also your estimate of what the potential risks are. And in theory the answers that someone would give to those two questions are, a priori, independent. The risks could be high, the benefits could be high; the risks could be low, the benefits could be low. Or, the risks could be high, the benefits low; or vice versa. There are those four possibilities. And in theory there should be people in all four quadrants.

But in practice, it seems that the people who think the risks are high also think the benefits are low. And the people who think the risks are low are also the people who think the benefits are high. The overall answer for most people is sort of clear, because there are two points in the pro column and two points in the con column.

This is interesting, that this is actually the pattern of risk and benefit calculus that we see. You sort of mentioned this point in passing, and didn't really go into an explanation of why you think that is, or what we should conclude from that observation.

It reminded me of some research in the field of biases and heuristics, in cognitive science, about this phenomenon. That when people tend to think that the risks of something are low, they tend to think the benefits are high, and vice versa. Even when that's objectively not the case.

I was wondering if you were trying to point to that potential bias there? Or why do you think we see that pattern?

Marc Lipsitch: I think the nature of this kind of bias is that it's very hard to analyze it from within the debate, once you have a position. Even this answer, obviously, should be taken with a grain of salt.

Julia Galef: Sure.

Marc Lipsitch: I think that part of the explanation may be that we are very unused to, and we should be unused to, in science, trying to demand a very clear direct benefit from research. That's not what most science is about. Sensible science policy does not demand immediate or predictable benefits for every project. There probably should be some projects like that, but not all.

Also, most science is essentially risk-less or very close to risk-less, with a few exceptions. I think that to even come to the benefit question at the level that I and others have been pushing it requires that you assume that you be concerned about a risk. Risky research, in other words, should have a much higher bar for benefits than risk-free research.

I think that the people who started the debate, and I was one of them, came at it from noticing that there was a large risk -- and then, at least my own evolution was, I started looking at the benefits and thinking, wow, these seem to be significantly over-claimed. Because they're not as generalizable as people think, as people claimed. And all sorts of other reasons.

At least in my own case, it was a matter of: the threshold condition for even entering the thought process was noticing the risk, and that the benefit then becomes subject to much more rigorous treatment than science normally should be.

As a practicing scientist, I run a lab with bacteria and do a lot of epidemiologic work. I would not want every study that I was proposing to do to get rigorously analyzed for whether it was going to have a life saving benefit in the short term. I don't think most science should have that. Most flu research certainly shouldn't have that.

But I think that when you propose research that puts large numbers of people at risk, the ethical and societal constraints should change. And there should be a much stronger presumption against doing it, until you really have an overwhelming reason to do that.

Julia Galef: It seems to me that you're pointing at a selection effect, where the debate is mostly populated by people who think the answer is relatively clear cut, those being the people who think the benefits are low and the risks are high, are relative to the common wisdom. Because those are the people who think the issue is important enough to be worth discussing publicly.

Marc Lipsitch: I think something like that is probably at work.

Julia Galef: Interesting.

Marc Lipsitch: This morning I actually just thought of an area where I fall into one of those off-diagonal categories, and I was very pleased with that, which is antibiotic use in animals. Which many people think is an important cause of anti-microbial resistance, and it is in the bacteria in animals.

The industry has argued, although they're kind of softening now, that it's important to making food cheap that we can use lots of antibiotics in animals, and it increases productivity and all that. The “anti” side says it causes tremendous drug resistance.

I actually think it's low risk, low benefit. And would probably say that it's more risk than benefit, and be against it.

But almost all my friends in infectious disease think it's high risk, low benefit, which makes the decision easy. I think the risk is pretty low. It does make resistant organisms. But those are not organisms that typically infect and kill people. Sometimes they infect people and don't kill them, and sometimes they don't get into people. But the evidence that people have died from resistant organisms that got resistant because of animal use and antibiotics, I think, is very small.

So I think it is possible to have an off-diagonal view. But it would take an awful lot of activation energy to get me going into the public space saying that, because there's not a good op-ed to write about it… No one wants to read that, it's not very interesting.

Julia Galef: Not seeing the page views skyrocketing for that one, indeed!

We have a few minutes left. And I think what I want to cover in our remaining time is: The object level question about gain-of-function research, and the risks versus benefits, is very interesting and important in its own right. But there's also this interesting meta question about the way that this issue has been discussed and handled, by not just scientists, but governmental bodies and the press. We could widen the sphere of actors here.

I'm wondering whether you think the scientific community and the government, et cetera, have handled this well or not. There's different ways that you could approach that. Like, should they have done a risk/benefit calculus before the research preceded, instead of halting it in the middle? Also there are smart, well-intentioned, very accomplished scientists on both sides of this debate. How well do you think they have handled the debate? Productively or no?

Marc Lipsitch: A few things to be said. I think that if it had been flagged properly it would have been very appropriate to do the risk/benefit assessment early. But in practice, for whatever reason, it was not appropriately flagged as a danger.

Even once the research had been done, it took a while for people to decide what it was that was really concerning about it. You can't fault people too badly about the retrospective nature of the debate.

In terms of why it was not flagged early, I have to remind people that information on laboratory accidents is extremely hard to pry out of the hands of the authorities. USA Today has been trying valiantly to get a Freedom of Information request answered by the CDC, on laboratory accidents, and has been told it will take three years. That was about a month ago.

Julia Galef: Wow.

Marc Lipsitch: There's all sorts of secrecy about laboratory accidents, and that's bad for everyone. It makes decision making very hard, and it makes it hard to figure out the rates at which these things happen.

In terms of the scientific community, I actually think the debate has been reasonably high level and cordial. With the exception of one other podcast -- not this one -- where it sometimes gets a little bit ad hominem.

I'd say overall that the public debate, and even the private discussions that I've had, has been nothing but polite and even respectful. There are definitely some friendships across this divide that were formed in the course of this discussion. That's a nice surprise, especially surprising for people in Washington who aren't used to bipartisan friendship anymore.

There is a lot of very careful work being done now within the government to try and get this right. And I think that's crucial, because I think this is the first of a number of problems that are going to come up as biology becomes more powerful, and the scope of what we can do to organisms becomes greater.

We've already heard the debates over gene editing. A little taste of other kinds of discussions where society and science intersect. And there will be many more of those. I think having a system, a process for discussing risks and benefits and ethics, in a context where we're not used to it, is going to be very important going forward.

Julia Galef: Good. That gives me a little glimmer of hope about the future of technology and science and humanity. Thank you for that, I don't often get those.

Marc Lipsitch: Good.

Julia Galef: We are just about out of time for this section of the podcast so I'm going to wrap up this conversation, and we'll move on to the Rationally Speaking Pick.

[musical interlude]

Julia Galef: Welcome back. Every episode on Rationally Speaking, we invite our guest to recommend the Rationally Speaking Pick of the episode. This is a book or website or movie, or something else that tickles his or her rational fancy. Marc, what's your pick of the episode?

Marc Lipsitch: My pick is a policy brief. I'm working in the really exciting area of policy briefs! … But this one was really inspiring for me.

It was written by Richard Danzig. It's from the Center for A New American Security, and it's called "Driving in the Dark: 10 Propositions about Prediction and National Security." I read it this past winter and found it one of the most compelling descriptions of how to think rationally about rare events and the problems of prediction. Not just in the national security context, which is his specialty, but in many other contexts. It's an addition to rational thinking.

Julia Galef: Excellent. We'll put a link to that on the podcast web site alongside this episode.

We are all out of time, Marc. Thank you so much for joining us on the show, it's been a pleasure.

Marc Lipsitch: Thank you. My pleasure. Bye bye.

Julia Galef: This concludes another episode of Rationally Speaking. Join us next time for more explorations on the borderlands between reason and nonsense.