

## Rationally Speaking #258: How to reason about COVID, and other hard things (Kelsey Piper)

Julia: Welcome to Rationally Speaking, the podcast where we explore the borderlands between reason and nonsense. I'm your host, Julia Galef. And for this episode I'm talking with Kelsey Piper, who is a journalist writing for Future Perfect, the section of Vox devoted to covering the issues that matter the most for global well-being.

I'm a big fan of Kelsey's careful and nuanced approach to thinking in general, and I've had her on the show a couple years ago — but since then, a lot has happened, such as a pandemic. And Kelsey's done a lot of great reporting on Covid. I've also enjoyed following her discussions on Twitter kind of explaining how she reasons about the evidence on masks, or social distancing, or vaccines.

So that's the impetus for our conversation, I wanted to get Kelsey's advice about how to reason about complicated topics, using Covid as kind of our test case. That is the focus of most of the episode. And then in the last fifteen, twenty minutes we talk about another piece Kelsey worked on recently that was interesting and controversial, about the de-growth movement.

And before we get started I just wanted to give you a heads up that a few days after we taped this episode, Kelsey emailed me to say that she had learned more and updated her opinion of one of the drugs we discussed in the episode. That's in the section near the midpoint of the episode when we're discussing Ivermectin. So I left in the original conversation but I've added Kelsey's update at the end of the episode.

Here is my conversation with Kelsey Piper:

Julia: Hey, Kelsey. Welcome back to Rationally Speaking.

Kelsey: Hey, it's good to talk with you again.

Julia: Yeah, likewise. So as I mentioned, one of the main things that I wanted to talk to you about was the intersection of COVID and epistemology. Basically, advice or lessons that you've learned from the last year and a half of thinking about and covering COVID, about how to think about hard things in general. Like how to think about different kinds of evidence, or disagreement among experts, or different levels of uncertainty. Things like that.

So I guess my first question is whether COVID has shifted your views at all on which kinds of sources you trust, or which particular sources you trust? And that could be which types of experts, or which particular institutions, you trust more or less than you did before, things like that.

Kelsey: Yeah. So obviously, I wish COVID had never happened, but it has been sort of incredible in terms of an opportunity to try to answer high-stakes questions with limited information, and see who else is doing a good job of that and who else is not doing a good job of that.

I don't know what I would've predicted in advance about the competence of our public health institutions, at whatever combination they're doing of figuring things out, and figuring messaging out, and making public announcements about it all. But I haven't felt like there's been a government agency where if they make an announcement, I'm sure of that announcement and don't feel the need to check their work.

And I have felt like it is possible as an individual, when you have the time to set aside, to look at a critical COVID-related question for 10, 20 hours and get a good understanding of what's going on, and sort of filter the limited evidence. And then when a couple of weeks later the picture is clearer, and reality has sort of come by to fact check you, get a sense of how well you did.

And I really recommend to everybody that they do that. Yes, there are people I think are worth listening to. Zeynep Tufekci came out really well as a journalist who was getting stuff right ahead of everybody else and explaining it well ahead of everybody else. But I also think it's just really valuable to try yourself, figure out what's going on while the evidence is limited, and then look back a month or two later when the evidence is less limited and say, "Was I right?" It's not very common that you have the chance to do something like that, and there's really no replacement for it.

Julia: And this would be on questions like, "Is it worth wearing masks," or on questions like, "Is it important for schools to be closed," or questions like that?

Kelsey: Schools to be closed, I think, is harder since that's a combination of so many different possible values, but... "Should you wear masks?" "How common is presymptomatic transmission?" "If your housemate has COVID, how likely are you to get it?" "How much more dangerous than an outdoor event is an indoor event?" "Should you get a booster?" All of those are of questions where I found it very productive to sit down and try and figure out your own answer. And then over time, as we've accumulated more evidence, go back and say, "Did I get this right? Why did I or didn't I get this one right?"

Julia: I often run into this tension, where I also tend to want to tell people to try to figure things out for themselves, and critically evaluate the consensus or the expert advice and so on. And the pushback I get, which I think is totally reasonable and I'm sympathetic to, is:

"Well, maybe people who are really careful and smart and knowledgeable about statistics can do this. And maybe they can do better than just a default simple heuristic of 'Just trust whatever the government says,' or something. But probably, a lot of people are not going to be doing their own... Their own work is not going to be better than that simple default heuristic, and so it's kind of dangerous advice to give."

And so... would you want to temper that advice by saying, I don't know, "Do your own work, but at the end of the day, if you want to play it safe, just do what the government tells you to do," or, "Do what the mainstream scientific opinion is"? Or would you not even give that advice?

Kelsey: So I think the thing I would say is doing it yourself is really, really hard. And part of why it's valuable with COVID is that you learn, often a pretty short time later, whether you were right when more evidence comes in. And it's absolutely critical to be able to listen to that feedback from the universe when more evidence comes in. If you try, and you come to a solid conclusion you feel sure of, and then a month later the evidence points a different way, you have learned something about how good you are at this. And that does change who you should be listening to and whether you can on your own improve on listening to the FDA, or whatever.

And if you're the kind of person where no matter what a month later, you're not going to concede that the evidence doesn't look as good as you thought, and you're not going to say, "Oh yeah, I got this one wrong," then doing it yourself isn't going to get you good results. But I don't really think your problem is that you're trying to do it yourself. I think your problem is that you can't admit that you're wrong.

I certainly do believe that lots of people are going to arrive at wrong conclusions and stick at wrong conclusions, and get worse results for themselves than if they listened to authorities that are okay at giving advice. But it feels to me like you are not in a helpless position with respect to whether that's true of you in particular.

You can ask yourself, "Am I more often right than the government," or, "When is the government more right than me?" And obviously, I should defer to people who are more right than me. And am I sometimes wrong? Am I capable of looking at the world and being like, "Yeah, I got that one wrong?" And if you're not, then you shouldn't be doing this, because you are just not capable of noticing when you got things wrong, and that's going to be a pretty crippling deficiency.

And I guess I don't really believe that your podcast or my Twitter feed or whatever are giving advice to completely randomly selected people. And I think that in terms of giving advice, you want to give advice where the

listener can figure out if they're the intended audience. And if the listener is trying to be honest with themselves, they can guess what they ought to be doing. I don't think you necessarily want to give advice that would be good if it was implemented by the median American with a median amount of competence or whatever. Firstly, because the median American is not actually listening to you. And secondly, because it feels really unfair to the people who are listening to you to be trying to give them advice that's not for them, and is for a different audience.

I am very careful not to recommend things in a context where I think that the person who's reading my writing is going to come away with a worse understanding than they had before. I do think that's really important. But "Think for yourself unless you notice that thinking for yourself isn't working" is advice that I feel comfortable giving, even if there's lots of people it wouldn't work for, because I think you can tell whether you're one of those people or not.

Julia: And what's the result of this process been for you, this process of trying to form answers and see if the universe proves you wrong, and learn about what kinds of things you get wrong, et cetera? Have there been any particular examples of things you realized after the fact...

Kelsey: Yeah, so I think one thing is you have to put in the time. Often, when I form an impression based on one hour of reading about something, I am often wrong. Often, it turns out when I have done a lot more reading or when more evidence comes in that I was totally underrating...

Julia: Are you often wrong in a systematic direction?

Kelsey: I think mostly I'm just neglecting a factor that's very important, and turns out to be the dominant one. Outdoor transmission is something where initially I kind of looked at it, and I was like, "Oh, it seems a little less than indoor transmission, and this is just not rising to the level where it's super important in my calculations." And then later, it turned out that actually, I think it basically dominates most calculations.

Julia: Wait, sorry, outdoor transmission dominates —

Kelsey: Sorry, the fact that outdoor transmission is so much rarer than indoor transmission. A lot of policies I sort of supported last spring I should have not supported, in favor of just move everything outdoors. Don't close things down, just put them outside. Which is obviously easier done here in the Bay Area than in many parts of the country.

But once I actually was like, "Okay, I have identified how much lower is outdoor than indoor transmission as a question that merits 30 hours of

research," then it wasn't hard. But one hour of looking at it didn't give me the impression that it was worth doing those 30 hours, that it was going to be one of the most decisive elements of good COVID precautions, and it should have. Ideally, your first hour does give you the sense, "Okay, I still don't know exactly what's going on here, but I do recognize that a reasonably plausible outcome of looking into this more is that it ends up being very decisive and very important, and is worth looking into more."

Julia: Got it. So a meta lesson is that it's not always obvious from the first hour that your opinion will change if you do more research?

Kelsey: Yeah. It's not always obvious. I think it should be. I think that there is a skill level that you can attain where your first hour of research successfully gets you to the answer to that question, but that's hard, and-

Julia: I don't suppose you can articulate what those skills are?

Kelsey: Oh, man. That's a really good question. I think you want to start with a pretty concrete model of why you're trying to answer this question, and what the range of possible answers to this question are, and how much they matter. And if your end conclusion is the range of possible answers is like this, and nothing in that range is very behavior-relevant, then that's a good reason to stop doing research.

Your initial look into a question should not necessarily be focused on answering the question. It should be focused on figuring out what is the range of possible answers to this question that look at all plausible... Try and figure out your 90% confidence interval before you try and figure out your point estimate, and then that will tell you how much time it's even worth investing in your point estimate.

Julia: So on the specific example of outdoor versus indoor transmission and how much that should dominate policy recommendations, what might that have looked like?

Kelsey: Hmm. So we're talking about March, right? By April, I think I had published things that were like "Outdoor transmission is much less common, and we should change policy accordingly." But in March, I think I was kind of thinking about lockdowns and the stay at home order issued in California as a sort of block measure. And I was thinking, "Okay, probably doing something like this is better than doing nothing and just letting 60% of the population get COVID."

And I think ideally, I instead break that down, and I'm like, "This is a combination of a bunch of policies. Which are going to be the most important ones?" If I spend an hour ranking all the things that went into the stay at

home order by how much I bet they matter, which ones come to the top of the list? And sort of doing that sort of work that's much more model-based and much more looking at specific elements and saying, "What is this specific element doing here," instead of trying to evaluate only on a very big-picture level.

Julia: I was just thinking about what my process looks like for evaluating questions like this. And I think when I'm really trying, I do something that's kind of like what you're describing, where I'm really tracking the nuts and bolts of the evidence, and I'm forming models in my head and so on.

But when I'm not trying really extra special hard, my process looks more like forming heuristics about which kinds of sources I should be considering trustworthy, and then putting more weight on those people. And so I might have heuristics like, "Well, the people who sound really shrill or snide, I trust them less, and the people who sound really measured and careful, I trust them more," or something.

And I don't know how good those heuristics are. I haven't really subjected them to much self-critical pushback. But I'm wondering what you think about how useful heuristics like that are?

Kelsey: Hmm. So I guess with the COVID discourse in particular, it feels to me like there've been incredibly toxic and snide and dismissive and horrible people advocating every position you can think of...

Julia: Yes, but there are very few reasonable people — so if you just focus on them, maybe they're not on all sides.

Kelsey: Yeah, that's fair. I guess I think "listen to reasonable people" is often a very good thing, because you can't put in the level of effort I was just describing for every question. You cannot. There's too many questions for it. And "listen to reasonable people" is often good.

I do think with COVID in particular, the reasonable people were often just as confused as we were back last February, last March. And so you have to check if any of them are doing the object-level thing. The whole thing needs to somewhere bottom out in somebody doing the object-level thing.

Julia: Right, that's true.

Kelsey: And as long as it does that, then I think it can be very reasonable to just mostly listen to people who seem to you like they're careful thinkers. I definitely do that a lot. But you have to think they're careful thinkers because they are themselves doing the careful thing, or that somebody is somewhere.

Julia: Right. As opposed to just... they sound reasonable, but really what they're doing is they're just judging who seems reasonable to them —

Kelsey: — and everybody's just trying to sound reasonable. Which is certainly a failure mode you get in some contexts.

Julia: Right, right. So, I interrupted you a couple of times, but to pop back up a couple levels... I had asked you if there were any specific examples of things, in this process of trying to figure out COVID-related stuff in this past year and a half, if there were specific things that you realized you'd gotten wrong. Or specific lessons you've learned. And one thing we talked about was the "outdoor versus indoor." Were there any other specific examples that came to mind?

Kelsey: Yeah. So another thing is last early February, kind of when everything was just getting started, people thought there was no circulation in the US yet, and I kind of bought into that. I kind of took that as a premise. I was like, "Okay, it's not in the US yet, but my model is that it is eventually going to get into the US." And how I thought that that was going to happen was I thought that rich countries would successfully suppress it for a while, with border closures and contact tracing and screening, and it would get into poor countries worldwide, and we would see really bad outbreaks in Indonesia and India and Kenya, and densely-populated places that didn't have the same public health infrastructure and contact tracing stuff.

And then once it was basically ubiquitous there, then something would slip past a screening in rich countries — even though I was imagining that that screening was pretty good — and it would get here too. And I don't know if I said that in public. I did say that in private messages, that my model was sort of, "It'll get here eventually, but it will get to poor countries first."

And that turned out wrong. By March, it was pretty clear that it had actually hit the rich countries much sooner than poor countries, just because of international travel. There was a lot more people flying back and forth. And also, there's some mysterious stuff there too. There's a bunch of areas where I'm like, "Why did COVID hit this place harder than this place? Their policies look similar to me. I don't get it."

And definitely, there's still a lot of, "Yeah, that looks weird, and I don't get it," but that's one where I've sort of gone back and thought, "Wait, why did I think that? Why did it seem to me like rich countries' contact tracing procedures were going to be good enough?"

And I think that's another level of, ideally, you get even more object-level. I was like, "Well, we're pretty competent. Probably, our contact tracing is pretty competent." And I think if you had actually called up a bunch of health

agencies and been like, "Hey, what is your contact tracing procedure? How many people are you tracing right now? How many people are you testing," you would have pretty quickly gotten the answer, "Not really any. The test doesn't work yet." And then it's really obvious that that model is just completely askew from reality.

And obviously, there are some great journalists who did make those calls and figured that out and reported it, and that was incredibly valuable work. And I was working at the time on some biosecurity stories that I also endorse sort of having focused on. COVID is not the only important story out there. But I'm really glad that there were those people who had the same maybe vague impression as me, but were like, "All right, I'm going to call and get a sense of exactly what our procedures are."

I think it's really important to sort of have in your mind, "Okay, I have this vague model that our procedures are good. What does that mean? What could I see that would convince me that our procedures are bad?"

Julia: That's kind of interesting, because you had a thread, I think it was last March, about how you wished you had been more proactive in kind of sounding the alarm that you thought COVID was going to be a big deal, before it actually became a big deal. You were like, "As of February, I was telling friends and family, 'I think this is going to be a big deal and we should prepare for it.' But I wasn't writing that in my Vox pieces, because I didn't want to sound alarmist, and I didn't want to get out in front of the mainstream government position or mainstream scientific position. And I wish I had."

But if your model was that this is going to be a big deal first in poor countries, and then in rich countries, then I feel like your strategy, given that premise, was actually quite reasonable. Because it would become clear to the US that this was going to be a big deal with enough lead time for people to actually prepare.

Kelsey: Yeah, I thought we'd have more time. I still feel like there's something culpable there, because I shouldn't have thought we had more time. That was wrong.

Julia: I guess, but it's a different cause. Then it's not so much a matter of you being too timid, it's a matter of you just having made wrong assumptions about how well the US's state capacity would work.

Kelsey: So in very early February... I think I pitched the piece to my editor January 20-something, and it published February 6th. I wrote a piece that was like, "Stop dismissing people who are worried about COVID. They are right to be worried about COVID. We don't know yet how bad it's going to be, but the



upper end here is quite bad." And then I kind of wrote articles about other stuff that was on my docket and important.

And I think since I thought that COVID was going to spread in poor countries and then get into rich countries, the ideal thing to be doing, from that flawed model at the time, would have been trying to figure out, "Okay, at what point should the US be taking measures about this? If we know for a fact that this is going to be here in four months, what does that mean? What should we be doing?"

And when I was talking to people, I was like, "I think it's going to get here eventually, so we could have lots of masks, and we should probably stock up on food in case there's a period where we don't want to go to the store."

And it feels very alarmist to write a Vox article that's like, "I think this is going to happen at some point, and it'll probably be more obvious before it happens." But in hindsight, I think trying to write an article with my reasoning would also have clarified my reasoning a little bit. If I had asked 20 people, "When do you think COVID will become community-circulating in the US, if that's going to happen," maybe I would have run across somebody who said, "I think it might be happening now." And maybe that gets you to call up places that are supposedly doing contact tracing, and learn that they're really not.

All of this is very hard. And when I discuss ways I think I could have done it better, especially on Twitter, which is not very amenable to complexity in any of its forms, I always kind of run into, "Okay, where exactly was the mistake?" And this is often a pretty high-level mistake. It's a standard to which I wouldn't even hold other people, because I don't really expect them to get questions like this right, necessarily.

But I do think that if I had been like, "All right, since I expect this big, bizarre thing to happen, I should write an article," that process would probably have helped me notice some of the ways in which my expectations about how it was going to go down were too optimistic, or assumed we'd have more time than we really would.

Julia: Is there an official or unofficial norm in the media, to be really reluctant to be more worried than the mainstream scientific consensus? Or the mainstream government position, or something? Is it really taboo to express more concern than that?

Kelsey: So I think there is an obvious incentive gradient towards alarmism and "The sky is falling"-ism, and everybody is very aware of that. You will do quite well with a story that's like, "The pandemic is coming, and it's going to kill

millions of us," because people absolutely click on stories about how a pandemic is coming and it's going to kill millions of us.

But whenever you write that story and it doesn't happen, you lose credibility with people. People start thinking of claims about big, scary events as basically just nonsense designed to get them to click. You're really spending down institutional credibility.

So among outlets that value their institutional credibility a lot, which Vox definitely is, there's a lot of... If you're like, "Hey, I want to write this story about how I think this is incredibly dangerous," people are going to be like, "Well, do we have a lot of experts saying that? Is it just you saying that? Is this kind of more of an op-ed?" which Vox doesn't even run.

There's certainly something there. But part of what Future Perfect is about is trying to do journalism differently, so I think Future Perfect has very much had more leeway, too. And later in the pandemic, sort of learning from some of the lessons in the early pandemic, I would say to my editor things like, "All right, I don't have experts who are super willing to venture numbers yet on how much the vaccines reduce transmission, but I do think that the evidence is there, and I want to write the story." And he was like, "Great, write the story." And a month before the CDC acknowledged that the vaccines reduced transmission, we have the story out saying, "We expect that they reduce transmission by about 90%," which I think held up pretty well.

So certainly, if you're right — and I think being consistently right is important to the longevity of your career as a journalist, if you're trying to do this — there's some space to go out on a limb. But you're going out on a limb, and you're spending institutional credibility, and it's not something to do lightly.

Julia: Right. I think in that mea culpa thread from last March, you linked to a paragraph by Zeynep Tufekci, where she says that the reason that we were so slow to react was that we — I forget if it was we or our institutions — we're bad at thinking about complex systems. I might be misparaphrasing her. But it seemed more to me like the problem was each actor was kind of assuming that other actors would be the one to sound the alarm, that someone else was going to sound the alarm.

Kelsey: Yeah, I think there was a lot of waiting for somebody at the CDC or FDA to say, "This is an emergency," at which point all of the news outlets would have dutifully reported that it was an emergency.

I think there's a sort of interesting dynamic, where health and science reporters relate to the FDA and CDC very differently than, say, Capitol Hill reporters relate to their sources on Capitol Hill, or military reporters relate to the military. I think if the US Defense Department is saying, "This situation in

another country is not an emergency," and somebody who's a national security reporter wanted to write an article that was like, "The military says this is fine, but actually it's a disaster waiting to happen," I think they would feel pretty able to do that. Because it's kind of understood to be part of their job description that they challenge the official sources.

And I don't think in health and science there's a similar ethos. "The CDC says that this is fine, but actually, this is very dangerous and scary and bad" is just not a story people were prepared to write.

And then there was both, I think, incompetence at the CDC and internal political pressure not to make a big deal out of something that the president was sort of coming around to the position wasn't a big deal. And it was kind of a bunch of things that went wrong at once.

But still, I think it is possible for a sufficiently careful and competent individual who was working full-time on understanding what is going on to have figured out what was going on sooner than we did, and to have been that first person to say, "This seems really bad." And I did some of that. I am proud of the early February piece that was like, "Hey, the bad case here is sufficiently bad that worrying is extremely warranted," but I wish I had followed that up with three more, you know?

Julia: Yeah. I mean, I would give you basically a 98% for the February piece. But I'm kind of grading on a curve!

So the very first question I asked you was about if your trust in different kinds of sources had shifted at all, and you said, "Even before COVID, I don't think I would have just heard a statement from government agency like the CDC and been like, 'Well, okay. I can just trust that without any questioning.'"

But I was wondering if you could get a little more specific about, to the extent that claims from the CDC seem mistaken or incomplete to you, are they mistaken in systematic ways? Are they mistaken in an "underweighting evidence that isn't RCTs" kind of way? Or in a "not wanting to sound alarmist" kind of way? Or in a "overly risk-averse" kind of way, et cetera?

Kelsey: I think there's a lot of fighting the last battle, whatever it was. In the last couple pandemics a big problem was like... With Ebola, a big problem was people overreacting and utterly panicking about the idea of bringing Americans who needed treatment back to America for the higher-quality treatment, when Ebola is not that transmissible, and there was not really a significant risk posed by that. And so I think there was a lot of sort of correcting for that, and trying to emphasize things that would have been right to emphasize in the last thing.

The other thing I've noticed, if you look at the FDA approval documents for the vaccines, or the CDC analysis of different treatment options, it's usually quite good. It reasons about uncertainty in a basically reasonable way. It tends to have most of the studies that are studies that I find when I'm independently trying to research the question. It tends to present the information you'd want.

The broad conclusions, in terms of recommendations, I often disagree with. I think they're doing cost benefit analysis wrong. I think they're too conservative. I think they're fighting the wrong battles or whatever. But I find them a very good source of, "Here's what we know on this question." And then you just have to look through what we know on this question yourself and be like, "Given what we know on this question, do I agree with their recommendation about what to do with what we know?"

Julia: That's a really interesting distinction to make. And it reminds me a little bit of something that I sometimes do when I'm talking to doctors, who have similarly kind of warped incentives when it comes to talking about risk to their patients. Like, they understandably worry that if they say something "probably" won't happen and then it does, then their patient will be like, "You promised!" and sue them. Or they worry that people aren't able to think about risk. Or a lot of doctors themselves aren't that great at thinking about risk.

And so if I just ask them for the recommendation, I'm not sure how useful their recommendation would be — but you can kind of extract the raw data from them, to some extent, if you ask the right questions. You can ask, "How many of these surgeries have you done?" And you can ask, "Well, how many times has there been a serious complication?"

You can get that info, and then you can kind of evaluate it for yourself about how high the risk is, and whether it's worth taking. But that's different from just asking for the advice from them.

Kelsey: Yeah, I think that feels very analogous to my current attitude towards official recommendations, is they tend to have a good sense of what's going on, and to be doing a good job of collecting evidence. And then in terms of what to do with that evidence, it's often worth reading it yourself and thinking through your own incentives and your own needs and what problem you are trying to solve.

Julia: Yeah. So I have a couple more specific epistemological case studies that I wanted to ask you about, for how to think about evidence. One of them is related to an article you wrote very recently for Vox, about a drug that it has just recently started to seem like might be useful in treating COVID. It's called fluvoxamine. Is that the right pronunciation?

Kelsey: Yes, that's right.

Julia: Great. Okay. Fluvoxamine. So it's an antidepressant, originally. That's what was intended for. But it looks from some preliminary evidence, like it might reduce inflammation and thereby reduce the severity of people's reaction to COVID.

Could you talk a little bit about like what the state of the evidence looks like, and how optimistic you are about it?

Kelsey: Yeah, so I am pretty optimistic.

The way I try to ground all of my COVID treatments and vaccine booster questions type stuff is: My grandparents are 85. They're very confused by all of the public pronouncements and all of the media pronouncements. They'll email me and be like, "Kelsey. What do you think we should actually do?" So whenever I'm looking into this stuff, I try to ask myself, "These people who I love so much who are at very heightened risk, if they got COVID, would I want them to take fluvoxamine?"

And the answer right now is, "Yes. I think that I would." It's very clarifying to stop trying to think about, "What can I defend in a Vox article?" and try to think about, you know...

Julia: It's very much reminiscent of this distinction between the "press secretary" mode of our minds and the "board of directors" mode, which I talk a little bit about in my book. But yeah, if you're in the press secretary mode, you're just trying to come up with justifications for something. But if you're in the mode of a board of directors that actually has to make bets, where the company will live or die based on how good the bet was, you're really thinking, "What do I actually think is the right choice here? Do I actually think that our company can compete with this newcomer or do I..." etcetera.

And so asking yourself, "Would I want my 85-year-old grandparents, who I love dearly, to take this drug?" it's a great way to turn this question into a bet that you're actually thinking about taking.

Kelsey: Yes. So fluvoxamine's situation is that there were some very promising early studies — some observational, some randomized, but quite small — where you just saw a huge difference in hospitalization-related end points between the fluvoxamine group and the control group.

But a couple of things made those iffy to me, and at that point I would not have told my grandparents to take the drug. One thing is that anytime you see a weird end point... Like one of the early RCTs for fluvoxamine had not hospitalization, but a level of respiratory dysfunction that they measured. Even if they preregistered that, and announced it in advance, it's like, "Why

are we using that end point?" I'm a lot happier about studies with standardized end points for the effects of the thing.

Julia: Sorry, are you suspicious because it might not be the case that we can extrapolate from that proxy end point to the things we really care about? Or is it like, you're suspicious that they chose that end point because they thought they'd be more likely to be able to find an effect for that?

Kelsey: Yeah. Both of those. And also just that in general, it seems like maybe they don't have a clear mechanistic model yet. And sometimes drugs work before we have a clear mechanistic model, and that's how it is... But it definitely makes me hesitate, whenever I see an end point that doesn't seem to clearly follow from how we expect it to work, or it doesn't seem to like clearly cut between "This works" and "If you look for enough things, you'll find something."

And obviously all of that is much worse if the weird end point isn't preregistered. If the weird endpoint isn't preregistered, I think you can basically just ignore it.

So anyway, these early studies, they look good, the effect size is huge. In some cases it's like a hundred percent prevention of hospitalization, and I never believe that. I'm always like, "Yeah, sure." And also, yeah, just not bad studies.

And I've talked since then to the people conducting the studies, and I think they did a really good job with the resources they had at the time, while looking into a hypothesis that, at the time, didn't have very much to differentiate itself from a vast sea of hypotheses. So kudos to them for that, but it wasn't enough.

But it was enough for Fast Grants, this is the Patrick Collison and Tyler Cowen thing, to, issue a call last December where they were like, "Let's do a good study on fluvoxamine." And the team at McMaster University, which has been conducting adaptive platform trials for COVID treatment and done a really good job, I think, ran a fluvoxamine arm. The data from that is recently in, up as a pre-print now.

And there's a bunch of things about it. It's not definitive, but there's a bunch of things about it that tell my gut, "This looks good."

In terms of statistical significance, a lot of people said, "Okay, the primary end points, which is ER observation or hospitalization, are statistically significant. But ventilation and death aren't statistically significant." And that's true. But the estimated effect size is the same for the amount it reduces ER observation, the amount it reduces hospitalization, the amount it reduces ventilation, the amount it reduces death. It's looking like a 30, 35% decrease

across the board. It's just that our error bars get wider and wider as we look at rarer and rarer outcomes, so that we should be less sure of the death stuff than of the hospitalization stuff.

But anyway, the primary end points reached statistical significance and look good. The secondary end points don't all reach statistical significance, but they tend to have about the same median as the other ones and just wider error bars than the way you would expect, because death is very hard to study in clinical trials for COVID because most COVID people recover.

And so that looks good to my gut. That looks like not what you'd find with fishing, and what you'd find with a real result. The end points obviously were set before the trial. Very important.

The mechanism makes sense. And there's observational evidence from Europe suggesting that many antidepressants have this effect. But fluvoxamine, because of the way it interacts with Sigma-1 receptors, has it the most — which is consistent with it being the best drug of antidepressants for this, but other ones working less well. Which just feels like what I would expect to see if there's something real there and not just what you'd expect to see if you're finding random noise.

So there's still a lot of questions it would be cool to see answered, but the whole thing looks to me like there's something real there, and it significantly reduces the risk of severe symptoms and likely also the risk of death from COVID. That's my current read.

Julia: I'm just curious to investigate this general question of how should we react to results that are not statistically significant. In general... suppose there was a social psychology study that got some result at  $P$  equals... What was the  $p$ -value for death in the fluvoxamine study? Was it —

Kelsey: 0.19?

Julia: Yeah. So, above the agreed upon threshold of 0.05. Would you be inclined to look at that and be like, "Well, it's not significant, but these results would still be more likely to happen in a world in which their theory was true than a world in which their theory was false. So I'm updating"? Or would you just be like, "Eh, not significant. I'm just going to ignore it"?

Kelsey: It depends a lot on whether it's the primary outcome or a secondary outcome. If they're like:

"Hey, our primary theory is that behavioral priming works, and we primarily measured how much people donated through our little donation game. We also measured how much they answered a questionnaire about their donation behavior. And we had a reason in advance to expect the

questionnaire effect size to be a lot smaller than the donation game effect size. And that's why we declared in advance one of them, the primary end point. The other, a secondary endpoint. And the secondary end point looked like our study was underpowered to detect an effect there. But what we see is consistent with our hypothesis."

...I think in behavioral sciences, I would also find that convincing. And my issue is that often what I'm seeing is the primary... If your study is underpowered to detect its primary end point, what are you doing? Don't run that study.

But if you're correctly powered to detect the thing you're primarily looking for, and then you're like, "While we're at it, we're also going to look at these other things, but we're underpowered to detect an effect there." And you're like, "Yep. Our result is that we're underpowered to detect an effect there." Yeah. That's a more promising pattern of evidence to my eyes.

Julia: I was thinking about this  $P = 0.05$  cutoff for statistical significance that we, collectively, ended up using. And it seems to me that what it's doing, even though it's kind of arbitrary, it's serving as this crude fix for the problem that... if everyone's allowed to do what you, Kelsey, are doing... where you just think about the specifics of the experiment, and you think about how likely those results would have been to obtain if the hypothesis were true versus false, etcetera, you look for the good signs versus the bad signs about the study...

If everyone's allowed to just do that, then it gives everyone just a ton of wiggle room to say, like, "Well, my results weren't significant. But I still think that we should update on them, because blah, blah, blah, blah, blah reasons."

And so in theory, if everyone was being really careful in doing what you're doing, we would be better off abandoning the 0.05 cutoff. But given that everyone's not going to do that, maybe the "You're only allowed to update on this on the study if it gets  $P$  less than 0.05," maybe that's just a crude solution, but better than the alternative.

What do you think about that?

Kelsey: Yeah, I'm pretty sympathetic to that. In that as a mechanism for controlling which things we're allowed to promote to each other's attention as hypotheses, it's reasonable to expect a certain threshold of evidence before you even consider a hypothesis.

I guess what I would say is... say you run a study and it does pass your statistical significance threshold. And then somebody runs a replication, and their replication finds almost the same effect size, but it just barely fails to



achieve statistical significance. It's 0.06. That's not a failed replication, right? You've got to use the p-values when you're trying to figure out what hypotheses you're considering in the first place, but then once you have a hypothesis and once you have an effect size and you're trying to nail it down, it makes a lot more sense to think incrementally, and think more Bayes-ically and be like, "Okay. How close is this result to the result we would have expected to observe?"

Julia: Yeah. That makes sense. Okay. So to get a little back closer to the object level question I asked... for fluvoxamine, how much better do you think the state of evidence is for it compared to other drugs that have seemed like, "Oh, maybe this drug is a promising approach to treating COVID," like ivermectin? How much better does the evidence look for fluvoxamine?

Kelsey: The fluvoxamine evidence is a lot more concordant. With ivermectin in particular, there's studies that look really good, there's studies that are finding nothing at all. There's a lot of trying to explain that by saying, "Oh yeah. You have to use this dosage and method of administration and supplement with zinc, as well." And maybe that's what's going on, we could get more into the specifics of ivermectin later, because there's a lot going on there. But you see like a bunch of things pointing you in different directions.

And with fluvoxamine, you see some stuff that's stronger than others, and some stuff that's like, "I would love this to have had more statistical power and given us a little more information," but you don't see this "pointing in different directions" thing. You sort of see what you would expect to see if it works.

And its effect size is not huge. It's not as good as the vaccines were, at like 30%, but it's working. So just in terms of, you look at fluvoxamine for 10 hours, I think you're like, "Yeah, it looks like there's something here. There's still some more stuff I'm uncertain on. And it wouldn't be a huge shocker if the effect size turned out to be smaller than this study has found, or smaller for death, or something like that. But I see a picture here of what's going on, and the evidence I have broadly is all evidence that I would expect to see if that picture is true."

Julia: Yeah, let's do talk about ivermectin. Because I'm sure a lot of my listeners who have been following the discourse on ivermectin recently were probably kind of surprised that the way you talked about it didn't sound as clear cut, "This is really dumb," as a lot of other people talk about it as being.

Kelsey: Yeah. So to be clear, you can hurt yourself quite badly if you take concentrations of a drug that were not designed to be safe for human consumption. You can also hurt yourself quite badly with drugs that are a

good idea and safe. Tylenol is incredibly dangerous. But certainly if you're DIYing ivermectin, you are running a very serious risk.

And I don't want to... To the extent that all of the, "Don't do the dumb thing" stuff is a response to people who have been hospitalized with serious health problems because they dose themselves badly, I do want those people to know that they should be really careful, because this is not something to do lightly. And it is something where you can cause yourself extremely serious health problems, at a point where hospitals in many areas are overwhelmed and going to give you less good treatment for your extremely serious health problems.

Julia: Yeah, I didn't think you were saying, "People aren't wrong to take ivermectin." I thought you were saying, "They're less obviously wrong than some people think they are," or something.

Kelsey: Yeah. So my impression is that the ivermectin fandom, such as it exists... Okay, essentially, you can sort of divide a drug into:

"You should take it prophylactically," you should take it before you have any reason to believe that you're sick, and it will prevent infection.

There's, "You should take it in early disease." That's where fluvoximine is good. As soon as you start showing symptoms, you should start taking this drug, and it will keep your symptoms mild. That's also where monoclonal antibodies, which work very well, are good.

And then there's, "You should take this in late stage disease." You are in the hospital. You are potentially dying. And dexamethazone is the gold standard there. It helps a lot with people who are on the verge of death.

My read is that, ivermectin, there is not really any reason to think it helps people who are hospitalized with serious disease, like you are in the hospital on a ventilator. I do not see any reason to believe that ivermectin would increase your odds of recovery. I haven't seen any studies that I think are high quality enough that they should update us from the prior that taking most random substances won't help you when you're in the hospital on a ventilator.

I have a good friend who's an ICU nurse and talks a lot about what ICU nursing actually is. And I've seen some people in the ivermectin community say, "Well, if you're dying, you may as well take it. What can it possibly hurt?" And the answer is a lot of things. Because in the ICU, a bunch of different bodily systems are going wrong in different directions at different times. And being an ICU nurse is like basically trying to keep all of your body's many different electrolyte and blood pressure and heart rate and bodily systems

within the ranges that won't kill you, while giving you enough drugs to try and treat whatever got you into that situation in the first place. And adding a drug that doesn't look like it helps to that mix is really bad. You shouldn't do that. ICUs aren't doing that. Again, don't.

Early disease, my read is that ivermectin might have a very small effect size that's positive. It might reduce severity of disease if you take it as soon as you get sick. It might not. My error bars here very, very much encompass zero. A lot of the studies that I think are high quality find, "We can't find a statistically significant result." But when you do a meta analysis, then it does look significant because there were a lot of studies that did find positive results. They were just small, positive results.

There's also been some outright fraudulent studies that just made up their data. And that's sort of messing with the ivermectin conversation a lot. It's such a sabotage to everybody's ability to think about stuff. You have to go back and rerun all these meta-analyses yourself, taking out the studies that were just... We're not talking about some sketchy behavior. We're talking about like outright, duplicated the data to make the sample size look four times bigger than it was. And then changed a bunch of numbers to make the drug look good. Just stuff that's completely burning the epistemic commons here. I'm really mad about it.

Julia: Yeah.

Kelsey: Prophylactically, the ivermectin evidence actually looks significantly better. There's some decent quality studies that seem to suggest it is decent at preventing getting COVID. I'm not quite sure what to think about this. Like, on a mechanistic level, I don't understand it. There are also other studies that are much more ambiguous or finding nothing. So I'm not thrilled about the mix of evidence there... but yeah, there might be something there. The people who think there's something there are on reasonably solid ground.

My big issue there is that taking a drug prophylactically has a very different cost benefit trade-off situation than taking it when you get sick. If you're taking it prophylactically, that means you're taking it potentially indefinitely, potentially just for months and months. If somebody was like, "I'm an ER physician in a country where vaccines aren't available yet, and I'm taking it prophylactically," I would be like, "Yeah, that seems like a pretty warranted thing to do."

But if you're in the US and you're vaccinated, which is one of the most, evidence-based, smartest things you can possibly do to protect yourself from COVID, taking a drug that kills all your gut bacteria, and in some people causes various other side effects, and that is quite dangerous if you dose it wrong — and taking that months and months, or potentially indefinitely,

because COVID is going to stay around — that just doesn't seem worth it to me. Maybe it does work reasonably well, but I just don't think that the cost benefit pans out there. So, in no case should you be taking veterinary medicine that you don't know how to dose safely.

Julia: That is an important point that we didn't really emphasize yet, which is that there's the dose for horses, and there's the dose for humans, and a lot of people are taking the former and not the latter.

Kelsey: Yeah. That's sort of the problem with any complicated evidence situation. There are some people who I know who are like, "Yeah, I found the prophylaxis stuff convincing, and I'm taking it at the dose recommended by Front Line COVID-19 Critical Care, which is the big pro ivermectin advocacy group. And I don't think they're making the right cost-benefit trade off, but I'm not worried about them.

But then there's also a lot of people who just... somebody goes on Joe Rogan and is like, "Ivermectin is the miracle cure that the government is hiding from you." And then they don't have the information about how to dose it safely, how to measure it safely, how to like react if they are getting side effects. And yeah, I don't think you could do that.

I was saying earlier I think it's okay to give advice that wouldn't be good for everybody if you're thoughtful about how people can tell if it's good advice for them. But I don't think you should go on Joe Rogan's show and declare a drug, even a drug that I think is great, like fluvoxamine or monoclonal antibodies, "the miracle cure that the government is hiding from you."

Julia: I mean, there are a lot of problems with that. Not just one, right?

Kelsey: Yeah.

Julia: I mean, I was a little surprised by the intensity of the reactions against the ivermectin crowd. If the evidence is, like... not that compelling, but not completely made up out of whole cloth, then why are so many people just so incredulous and furious that this is even a thing?

And I guess a lot of it is because... Well, probably some of it is because people are taking ivermectin who won't get vaccinated. Which is crazy if you're... It's hard to find a consistent way to think about the evidence that would say ivermectin is worth taking and vaccines are not.

Kelsey: Yes. Strongly agree.

Julia: So I can understand getting angry about that. And yeah, then probably also "the government is hiding the truth from you," or "pharmaceuticals are hiding the truth from you" stuff just drives people up a wall, very understandably.

But nevertheless, I was kind of surprised that the evidence isn't completely as non-existent as it seemed from the reactions of the anti-ivermectin crowd.

Kelsey: Yeah. The FDA ran a public service announcement. It was like, "You are not a horse. You are not a pig. Don't take ivermectin."

Julia: "Seriously, just stop," or something. That was the tone.

Kelsey: Yeah. Yeah. "Seriously, just stop." And I think a lot of people were responding to that and participating in that.

And you know, I am not an expert on public health messaging. It would be convenient for my worldview to make a claim like, "Mocking people doesn't get them to change their minds. Reasoned engagement gets people to change their minds." But I don't actually know that that's true in this case.

Julia: Definitely. Yes, that is convenient. I've said the same thing. Yeah, I agree.

Kelsey: Maybe mocking people gets them to change their minds, but I don't like it. It's not how I like approaching the world.

I think I am interested in trying to say to people, "Yeah, I don't know. I did do a lot of reading about this, and the evidence is really, really messy. And I get where you're coming from, but I think that this is not one of the treatments that I would recommend to my grandparents. This is not one of the treatments that I would personally choose to take. And I think that you should agree with me about that, if you read through this and think about this."

And ivermectin is recommended for humans in other contexts, like as an anti-parasite drug. So the, "You're not a horse. You're not a cow," that's not really responsive, you know?

Julia: Right. And the people who were mocking the people taking ivermectin responded to that point by saying, "Well, they are taking the dosage for horses. So yes, it is technically a horse de-wormer. You can't deny that."

Which is true. It's just that the connotation of saying, "They're taking horse de-wormer," is... it overstates the case. Like it's not... A typical case of someone taking a medicine for animals is more idiotic than this case, given that there's also a human form, and there's some, albeit not great, evidence that the form for humans can be helpful. And so this is like a special, unrepresentative case of "humans taking veterinary medicine" that is less ridiculous than just calling it "humans taking veterinary medicine" makes it sound.

Kelsey: Yeah, and I guess I feel it's really important for like the long-term epistemic health of society or whatever that we try talking about what's actually going on and not just going for cheap shots.

Julia: Right. Yeah. Well, I have a few other questions unrelated to COVID that I was hoping to talk to you about, so I think I'll switch tracks.

Kelsey: Yeah. Absolutely.

Julia: One of them is about a very widely-shared article you wrote for Vox about a month ago on the de-growth movement. So would you mind just summarizing what de-growth is? Then you can give your high-level take on it.

Kelsey: So de-growth is a movement I mostly ran across in leftist, combating climate change spaces, although it's not totally specific to climate change.

And it basically argues that since humans are ecologically overloading the earth and driving other species extinct and causing all these cascading ecological problems, we can't solve those problems while we're focused on economic growth, and productivity growth in economic and GDP terms. And the only way to solve those problems is a transition to a philosophy where we don't care about growth. We still care about human well-being, but we don't go, "Oh yeah. We hope for 30 percent GDP growth next year," because we're not aiming for that anymore.

And I think that this is very, very much underrating just how much human flourishing wealth and abundance makes possible. And I think that, "Yeah, we'll keep the human flourishing and get rid of the wealth and abundance," is not really achievable.

And the other thing that I focused on a lot in the piece is I don't think it solves climate change at all. We need to get to net zero in order to address climate change, but we can get to net zero emissions in a growing society. There are lots of countries that are already successfully growing while shrinking their carbon emissions.

And degrowthers can kind of say, "Yeah, but we're not doing it fast enough to avert the worst effects of climate change." But if your complaint is we need to do something fast, then "We need to completely rethink the way that our society sees economics and civil society and the fundamental aims of the economy" is not something that's going to happen fast.

Julia: It's not a quick fix.

Kelsey: Right. There is no avenue by which degrowth is the thing we do in the next 10 years. So the more you think that we need to act soon on climate change, I

think the less compelling you could find, "We need to fundamentally change the way that the entire economy works." Because it's going to take a while.

Julia: I mean, to me this seems very obviously, politically, just a non-starter. But I'm more interested in the question of: Suppose you could somehow get buy-in for de-growth among all the wealthy countries. Would it actually work?

And so I wanted to go back to the thing you were saying about how you think they're underrating the connection between GDP growth, and human flourishing, on the other hand. Do you understand what the crux of disagreement is there, between you and the degrowthers? Why they don't think that's as connected as you do?

Kelsey: I think so. If I sort of put on my degrowth hat, what I would say is GDP obviously measures tons of stuff that has nothing to do with human well-being.

Like building lots of tear gas is good for GDP, building tons of carbon emitting things is good for GDP. Leaded gasoline, even, sales contributed to GDP, even though we now estimate that actually lead gasoline was dragging down GDP enormously by causing everybody brain damage. If you're optimizing for a metric that's not human well-being, such as GDP, then you're going to get lots of instances where GDP is being prioritized at the expense of human wellbeing. That's just how metrics work, right?

And they point to lots of instances of countries that have much better health care than the US despite spending less on it, or that have better food security, lower infant mortality, stuff like that, despite being significantly poorer. And they're like, "See? Maybe you have to get up to a certain level of wealth, but at some point the wealth and flourishing break apart. And when they do, go for flourishing, not for wealth."

That's sort of, with my de-growth hat on, that's how I would pitch it.

And I think my big cruxes with that are: Wealth is not just a proxy for human flourishing in the sense that we will try to optimize wealth and hope that human flourishing happens. People have the wealth — also inequality matters, but the ideal situation is that people have wealth, which they then spend on the stuff that they want. And that's not a proxy for human flourishing. That is the thing itself, is people being able to get the things that they want.

And a lot of proposals to stop using wealth for human flourishing are kind of like, "But we'll provide people with the things that are most important for them to have, through a state system." Which... I'm probably in favor of

universal health care through the state. I think like our existing healthcare system is worse than a universal healthcare system would be.

But for most other things, I'm pretty nervous about that not being provided through people choosing to spend their money on it, and being provided in some other way. I would expect that to produce a lot less human flourishing. And it's kind of like... no matter what, you're using proxies for the thing you really care about, right? There's no perfect measure of the thing that actually matters here. But I think "How many resources can people put to work on getting what they want" is just much, much better as a proxy for human flourishing than almost anything else. And I think that's our crux.

Julia: Okay. So then the crux is just like, the degrowthers think that we can...

Kelsey: That there's something better, at getting people what's good for their flourishing, than giving them more resources to get what they want. I would say, yeah, it's not that giving them more resources to get what they want is perfect in theory, but I suspect there's not anything better.

Or maybe there's something better in some edge cases. Like if somebody invented a perfectly addictive drug, that once you took it, it made you want nothing but the drug, I'd probably be in favor of banning the production and marketing and sale of the drug. But those are sort of edge cases, rather than it being systematically true that people will be better off if they have fewer resources to get what they want, and instead more external forces providing what it thinks that is good for their flourishing.

Julia: I guess I also was never quite clear... I mean, I've read maybe a couple of Jason Hickel's essays. He's one of the most prominent degrowthers. But I haven't read everything they've written, of course. So maybe they talk about this and I haven't encountered it, but... Usually these conversations are at this very kind of high level of abstraction, of like, "We" — sort of this collective we — "We can pursue flourishing in such and such ways, and not in other ways."

It's never really clear to me, what is the actual intervention that they think we should do? Is it like, we should ban certain kinds of economic activity because central planners decide they're not conducive to flourishing enough? Or is it a kind of social... are we trying to change people's hearts and minds, and convince them that "No, we shouldn't be producing SUV's. We should be producing classical music"?

I don't know. I don't actually know what the intervention is that they're suggesting, but both of those sound not feasible.

Kelsey: That's a very good question, and it was one of my frustrations with Hickel's book *Less is More*. When it got to the section that was on solutions, it had lots



of stuff about “End planned obsolescence by making there be less financial incentive for companies to sell laptops that will break down in three years,” or whatever.

And it's like, okay, I guess you could come up with a structure of laws or incentives or tax breaks or whatever, that change the incentives there. And this seems like maybe a positive thing. Certainly I find it annoying having to buy a new computer every couple of years, and a new dishwasher every couple of years.

But that doesn't solve climate change. It just feels like a very small and sort of neat thing compared to the actual situation here. And directing any resources from climate change to it seems like it would be a bad thing, from a climate change perspective.

And there's like... “Cancel student debt” is in there.

Julia: What's the logic for how that helps with degrowth?

Kelsey: That people have to take jobs where they earn a lot of money, because they start out their careers in debt. And the whole thing is meant to create conditions under which all you can do is try and get more resources to pay off the debt that you accumulated when you were like 18. And again, this feels very small next to like, “We need to get to net zero carbon emissions in the next couple of decades.”

I do think it's pretty important that we burdened a bunch of 18 year olds with hundreds of thousands of dollars of debt, often under some false pretenses about how the college system is going to help them. And I'm down for various reforms, although “cancel student debt” isn't one of the reforms I like. I have some sympathy here, but it just feels like not the most... it feels sort of non-responsive.

Julia: Like if you sat down to think about, okay, what is the best way to fight climate change, this is not what you would come up with. It feels more, kind of like, reverse-engineered or something?

Kelsey: Yeah, like what are the degrowth policies that are most popular?

Which, if you're trying to get people to do more degrowth things, then finding the most popular degrowth policies are fine. But a lot of the justification for degrowth is the claim that we're in a state of ecological emergency, that the planet cannot bear up under the burden of our continued economic growth. And if you believe that, then you have to also believe in solutions that are a lot more transformative than that.

Julia: Right. Well, that brings me to another question I really wanted to ask you, which is that earlier this year I had Matt Yglesias on my podcast, and we were talking about his proposal for 1 billion Americans, i.e., raise the population of the US to, ideally, around a billion. Both through increasing the birth rate in the US, and also through letting in way more immigrants from poorer countries.

And so Matt has his arguments for why this is a good thing, but he acknowledges one of the common objections he hears, is which is that hundreds of millions more people living in the US, where energy usage per capita is relatively high — that's going to make climate change worse.

And he acknowledges, "Yes, that is true, I just think that —" and I'm paraphrasing him here, "...I just think that the value of all those hundreds of millions of people, living good lives, who otherwise wouldn't have existed, or living better lives than they were living in their poorer countries, that outweighs [my guess about] how big the impact on climate change would be."

And partly, he thinks that trade-off is more worth it than other people do, because I think he thinks that climate change, while serious and extremely bad, is not as existentially bad as a lot of people, maybe including the degrowthers, think it is.

Kelsey: Yeah. I think that winds up being a very fundamental disagreement there, is that there are a lot of people who believe that climate change is running a significant risk of making the planet uninhabitable. And if that's true, then "millions of people will be happier" is just small compared to that.

I think I disagree that climate change is running a significant risk of making the planet uninhabitable. But I think that, from that premise, it makes sense to be like, "None of those things matter, next to keeping our planet a place where humans can live." I basically do agree that keeping our planet a place where humans can live is the most important task ahead of us this century.

Julia: I mean that's one reason the de-growth plans confuse me. Given that they're supposedly premised on the urgency of solving climate change and the supreme importance of solving climate change, and then when the proposals seem more optimized for solving other things like, I don't know, social justice or economic justice than they are for saving the planet. I'm like, "I could see why you might want that policy, but how would you arrive at that if your goal was to save the planet from extinction in the next 10 years? That just doesn't follow to me.

But the question I wanted to ask you about Matt Yglesias' argument was, I saw you had a thread on Twitter recently about how it is valuable, in your opinion, to have new people existing with positive expected value lives, and

that creating more happy people is a good thing. And so I was wondering... that's definitely a premise that you and Matt share, and that not everyone shares. Would you still say that it's worth doing, even given the impact on climate change? Would you go as far as Matt goes?

Kelsey: Yeah. I mean, I think that the impact of a individual person on climate change, the negative impact caused by their emissions, is not actually very large. You see a lot of really big numbers which come from counting not just that person, but also all over their prospective descendants as well, assuming that like 2021 levels of carbon emissions continued. There's just a lot of very irresponsible data there.

But if I could just create another person in the US today — or if I could just press a button and in some non-evil way, there's 400 million of us — I think that's a society that is better and healthier and makes more progress on important challenges of the century, probably including climate change.

This isn't something I'm highly certain of, the “makes more progress” part. It seems probably true to me, but also there's definitely some areas where I'm not sure that there'd be differential progress where we solve problems faster than we create them.

Julia: Just to clarify, you're saying the thing you're not sure about is whether having a hundred million more people would-

Kelsey: Actively make things better.

Julia: Overall? You're including both the positive benefit of increased chance of finding good technological solutions with all those new people working on them, and also the downside of their consumption?

Kelsey: Yeah. I'm uncertain about that. I think I lean on the side of “it would be good,” but I could be convinced that that's wrong.

But there's also a deeper thing of, I feel like a coherent philosophy of how “lots of human flourishing is good,” is really, really important. I'm really reluctant to see people retreat from that, or start thinking of people as basically a problem to be managed, rather than the source of value and good things and everything that I care about.

Julia: Meaning, like —

Kelsey: Given uncertainty — given uncertainty about whether the 400 million America is better than the 300 million America for climate change, considering both all of the benefits and all of the costs — I prefer to err on the side of more people. Because of a sense that it's good for humanity in the long run to be conceptualizing ourselves as a good thing, that we need to

create the space for more good things in the future, instead of as a cancer that we try and mitigate.

Julia: But do you think that the issue of timing is relevant, where we need a solution in the next 10 years, and the benefits to technological progress would take much longer than that?

Kelsey: That seems true for having children, although they also don't produce very much carbon when they're little babies. It seems less true for letting in immigrants who can contribute to innovation right away.

The other thing is that the 10 year framing comes up a lot in degrowth and it comes up a lot in climate circles. I think I don't expect it's discontinuous like that. We probably do have less than 10 years to stay beneath the IPCC 1.5 Celsius benchmark or whatever, but I don't think of that as discontinuous. Climate change will be bad, and cause a lot of harm, and the longer we let it go on the more harm it will cause, and the more bad it will be — but “we have 10 years to solve it” isn't very true to my model.

There are some speculative feedback loop type of things where if we don't keep it below a specific level it gets worse, but I find those pretty unconvincing, so I don't tend to think of it as like a “we have 10 years” thing, so much as a “we need to be moving as soon as possible and the quicker, the better” thing. And it is still worth doing things that won't help for 20 years, or 30.

Julia: I guess my last question for you is just about this philosophical intuition that having more happy people existing, who otherwise wouldn't have existed, is a good thing. A lot of people disagree with you about that. And this disagreement, over what's essentially population ethics, philosophers would call it... it seems like it's at the heart of a lot of very policy-relevant disagreements.

And I have such a hard time negotiating these disagreements, or making progress on these disagreements. How do you, or how would you, approach a disagreement with someone who's like, “I don't see why a new happy person existing, who otherwise wouldn't have existed, is a good thing. I want whoever does exist to be happy and not suffer, but I don't see the point in creating new happy people.”

Kelsey: Yeah. Changing minds around population ethics is very hard, because it's very much something where we're all just kind of consulting our intuitions and it's often hard to figure out where they even bottom out. And it's also something where I think, no matter what, you're forced into some counterintuitive conclusions. There's no beliefs you can have about population ethics that

don't require some bullet biting somewhere. Aside from like, "Well, my beliefs are inconsistent." I guess you can always do that.

So it's very hard to convince people. But I guess a thought experiment I found valuable when I was thinking about this is:

Imagine that we found out that the nearest sun has another planet around it. And the planet is very much like Earth. It's full of people who, very much like people on Earth, have good days and bad days, and write beautiful music, and commit horrible crimes, and have joyous experiences and horrible ones.

How do we feel about that world existing?

Personally, I care a lot about factory farming, so part of how I feel is like, "Okay, are they torturing billions of animals?" Like, that's important. Are they torturing billions of animals? But setting that aside, for non-moral reasons, they are not torturing billions of animals... I'm excited that they're out there. I'm like: "Oh cool! More people!"

They've got some problems. We've got some problems. But one, there's two points of light in the universe now. Two worlds from which we might spread out into the galaxy. And secondly, their symphonies are different than our symphonies. Their great plays are different than our great plays. Their movies are different than our movies. And I'm so excited that those exist, too. And then also just like... All of those people, if their sun went supernova and destroyed them, it would be really, really tragic. And not just in a "they would die" way, but all of those stories and all of those lives that could exist, not existing, feels tragic to me.

And so, similarly, if we think that there could be more people here, or that the people here could eventually go live on other stars and we don't, then that feels tragic to me. There's all those plays and all those experiences and all those romances and all those things that never happen. And I think they should.

It's complicated. It's certainly not one of my policy priorities right now, because I think there's other problems that are even more pressing. But other things equal, if I can make that other planet exist, and there wasn't any factory farming there... and who knows what's going on with wild animal suffering... but speaking about the people, I would.

Julia: It's so weird and interesting, how imagining a planet full of novel and unique and happy people who are different from us — imagining that planet existing, versus not — it feels more meaningful. It budes my intuitions more than just imagining, "Well, take our planet, but imagine 20% more people," or something. That feels so much less important. But you put them on a separate

planet, and you're like, "Oh, I want that planet to exist. It's better if that planet exists." Why should it matter which rock they're on?

Kelsey: Some of it is probably legitimate scarcity intuitions, right? Like trying to imagine where the people fit in. There's something legitimate about the difference between my thought experiment and Matt Yglesias' proposal.

But I do think in terms of a thought experiment for the pure impulse of positivity about humans existing, it's useful to say, "Okay, they get their own rock. We don't have to process our feelings about resource scarcity or anything like that." We want that rock, right? It would be good if they were out there.

Julia: All right, well, Kelsey, I will let you go. This has been a fascinating and an enlightening conversation, thank you so much.

Kelsey: Yeah, absolutely. It was great talking with you, too. I really enjoyed it.

Julia: Bye, Kelsey.

[musical interlude]

Julia: That was Kelsey Piper, you should follow her work at Vox's Future Perfect section. You should also follow Kelsey on Twitter, her handle is @kelseyTUOC. That's KELSEY, TUOC.

And as I mentioned at the beginning of the episode, Kelsey wanted me to let you guys know that her opinion of the evidence on ivermectin shifted after we taped this. I'll just read you the email she sent me:

"Hey Julia! When we talked for your podcast, I said the good evidence suggested ivermectin's benefits were modest to non-existent - but that those benefits looked best for prophylaxis, and that while I wouldn't recommend it I could see the case, especially for unvaccinated people at high exposure risk.

Since then, as I've worked on a Vox article about this, I've been doing a closer read of the studies that suggested large benefits for ivermectin prophylaxis, and I've come to the conclusion that many of them cannot have actually occurred as claimed by their methods. I've been emailing back and forth with Hector Carvello, lead author on one of the big ivermectin prophylaxis trials (one of the hospitals he listed as involved says no trial was ever conducted there), and I honestly don't entirely buy that it happened at the other hospitals either.

I was talking with Gideon Meyerowitz-Katz about this and he mentioned that the team of science-fraud researchers who've been doing data forensics (ie checking for unnatural patterns of repeating digits, unlikely distributions of

traits, etc) have found that with nearly every huge-effect-size ivermectin paper, they've either found markers of fraud or the researchers have declined to share their data. They're going to publish their forensic results one by one, paper by paper (as appropriate, for huge allegations like this) but in the Bayesian spirit, I'm going to try to update once rather than ten times in the same predictable direction: I don't think any of the published studies on ivermectin prophylaxis ought to move us.

I really enjoyed our conversation and I hope this is a useful supplement to your listeners. While I try to read studies with a skeptical eye, I'm usually looking for flaws, not fraud - but both are out there!"

So there you have it. This concludes another episode of Rationally Speaking. I hope you'll join me next time for more explorations on the borderlands between reason and nonsense.