0:01

DAVID KEITH: Welcome to the seventh of our Compton lectures. As for each of them, I'm going to start by trying to remind you where you are in the sequence of lectures and what we're going to do today. So just to remind you where you, where you've, where you are, where you've been. This is Tiffany Shaw and myself talking mostly about what is the consequences of sunlight reflection. So in the lecture that I gave two ago--the last one with me and Doug MacAyeal, was kind of something different--but the one lecture that I gave two ago, the lecture on November 1st is really the prelude to this lecture. And in that November 1st lecture, what I tried to do was explain how humans might make radiative forcing, how we might reflect sunlight, some of the kind of technical details, the different methods, some of the consequences. This lecture is more about what the consequences of reflecting sunlight would be, and especially about how uncertain those consequences are, or how certain they are, what our confidence is in predicting the outcomes. And we're going to do this with a combination of Tiffany and I, and I'll show you how it's going to work.

1:10

I'll start out by really trying to pose the question how uncertain it is, and talk a little bit about how we might address this question in models by kind of breaking the problem into two halves, as I'll describe. I'll say something about feedback controls, about answering the question how we can control something that we don't understand, where, at least in some circumstances, the answer is that that's quite easy and commonplace. Then Tiffany is going to give a few slides about a really specific research project that she's started that I'm completely thrilled about, some of the things that I've always thought needed to be done in this field about what we call perturbed physics ensembles. Then I'm going to say a little more about kind of the case for reliability research. And then the last 15 minutes or so, Tiffany are going to have a conversation which we really haven't scripted much, kind of talking about our different views about how certain or uncertain these predictions are. And so without further ado, here we go.

2:02

So I'm gonna start by showing you a slide I showed you last time. I won't go through it in detail again, but I want to use it as a kind of hook to ask the question, how much can we trust this? So I showed this slide before. It's results from a very fancy, high resolution climate model that I did with colleagues at Geophysical Fluid Dynamics Lab, Gabe Vecchi and Wenchang Yang and others. And the results of this are that apparently, in essentially every region of the world, reflecting some sunlight would reduce, or at least not change significantly, all the key variables that we are, many of the key variables that we care about in terms of driving climate impacts. So the results here was that for these four very important variables--average temperature; extreme temperature in the year; precipitation minus evaporation, which is the amount of water that's available to the surface, to plants, to rivers; and extreme precipitation, which is a big thing that drives risk of flooding--for all those variables, you see that in no region of the world in this particular model, are we making those variables, are we exacerbating them, are we making them farther from preindustrial, are we making them maybe worse. And that result from this climate model certainly a surprise in the sense that I and several other people going into this kind of 20 or 30 years ago were quite confident that that wouldn't be true, quite confident that, while there was no question that sunlight reflection could cool the whole planet, that it would do so unequally. And yet this seems to show it does it very equally.

3:31

But the question is, should we believe it? Should we believe this model? And that's really what we're going to focus on for the rest of today.

3:41

So the first thing to say, very generally, is when you think about how much you can believe or not, what sunlight reflection might do, if you ask this general question, how well can sunlight reflection reduce climate changes?, the answer is, it absolutely depends on the scenario. You'll hear people in public and everyday conversations, if you're in those conversations, have very strong opinions about sunlight reflection, doing something good or something bad. And it's important to say that there is no conclusion you could draw without some assumptions of how it's used. So in particular, any statement about the performance of sunlight reflection methods is contingent on assumptions about how it's used.

4:29

And there's a myriad of assumptions, but I'm going to divide them into three big buckets. One bucket is the amount that it's used, how much we're cooling the planet, for example. Another bucket is the spatial distribution of that reflection. Are we doing it just in one hemisphere, or just over one region, or evenly over the whole world? Those are all things that are possible, and the they would produce very different answers. And then the method used. In the lecture, the prelude to this, two lectures ago, I told you about all sorts of different methods, from putting sea salt aerosols that made some low lying clouds brighter to sun shields in Earth’s orbit. There are all these different methods that we might use, and so any statement about the performance depends on what methods we're using, are being used, how much is being done in the spatial distribution and a bunch of other things, but that's the top three.

5:25

And just to give you a sense, there's no question whatsoever that some choices, which I think is fair to label as bad choices, would produce some really bad outcomes for humans or for the natural environment. So just to give you some example, if I think one strong conclusion that we can draw from all the modeling work on this topic is if you think that you could use sunlight reflection as a substitute for cutting emissions, that is if we could just keep the world running with emissions going on forever and just keep reflecting more and more sun, sunlight, if you thought that was a solution, it is definitely not. There's no question that there are a set of risks that would grow essentially without bound, if that was what was done. And it's also true that if you only did this in one hemisphere, you shift this beta band of rainfall that we see around the planet roughly the equator, we call the intertropical conversion zone, you shift that, and that would certainly have some very severe impacts. It would make some places that are now really wet, very dry, and vice versa. And it's also absolutely true that if we, if humans, ramped up to do a whole lot of sunlight reflection and suddenly stop, there'd be bad consequences.

6:35

So there's no question that there's potentially bad ways to do it. So here's a scenario that might or might not be relatively good for lots of humans in the natural environment. That's what we can talk about, but it's a specific scenario. In the very last lecture--still one more lecture to come, we'll talk a little bit more about some of the politics of whether or not this might happen be feasible politically--but let's just say that we use sulfate aerosol--technically, it's precursors, not the actual aerosol--injected from high altitude aircraft in the stratosphere, and that we adjust, as we'll talk about, at the way those aerosols are adjusted to make things relatively even and to do a relatively small amount in a balanced way. And let's just assume that cooling starts in 2030, which I think is personally very unlikely to happen, but just for the sake of argument, assume that, and it ramps up so that we're adding a cooling force of a tenth of a degree of cooling for every decade. So at the beginning of 2030 you're not cooling the world at all. By 2040 you're cooling at a tenth of a degree. By 2050 you're cooling it two tenths of a degree, etc. And by cooling it, I don't make mean making it cooler than now. I mean making it cooler than it otherwise would have been. And to put this in your head right now or right now, over the last decade or two, warming is actually accelerated. Warming is now getting close to two tenths of a degree a decade. It had been more like one tenth of a degree, and now it seems like it's closer to two tenths of degree. Tiffany would be a better expert on the exact number, but that's roughly where things are. So this would be roughly cutting the rate of warming in half in very round numbers. So the world would still be warming, but it would be warming less than it would otherwise be doing under such a scenario.

8:16

So now let me talk about how we'd go about assessing the outcome of that, something like that, using the suite of models, the suite of underlying physical knowledge we have about the climate system. So one way to think about what we're talking about here is this kind of chain of, causal chain of what would happen, where you'd have aerosols injected--and again, technically, it's actually not aerosols it’s the aerosol precursors, those would be ejected from aircraft in some way one could specify--that would somehow shape the distribution of aerosols in the stratosphere, which, from the point of view of the climate system, you could think about as the distribution of reflectivity in the stratosphere, that is, where and how much and so on. And then a climate model responds to a specification of that of that distribution of reflectivity, and it predicts some climate changes.

9:09

So one way to think about it is divide this problem in half. And I'm going to talk a little bit about, sorry, one way to think about is to divide the problem in half. You say, part of the problem is figuring out how we can adjust stratospheric aerosols to make this kind of even layer that approximates turning down the sun. And the second thing we can do is look in climate models, and with this suite of different climate models and background knowledge, to assess what is the consequences of turning down the sun.

9:38

So first thing to say is that the uncertainty in predicting the amount of reflectivity, which is proportionally the amount of cooling, that we would get for a given injection of sulfate aerosol precursors using some given aircraft, that uncertainty is really big, at least a factor of two or three. And so your first reaction to that should be, well, forget about this then, it's hopeless. If that uncertainty is so big, how could you ever use this if we actually can't predict now, with the kind of science we have, how much cooling you get for a given amount of sulfate. And the reason for that, I kind of hinted at last time, it's because the amount of reflection depends very sensitively on the size of the resulting particles. If you think back to that last lecture, if you're injecting a ton of sulfur, you know you've injected a ton of sulfur. And we know from stratospheric circulation that ton of sulfur will stay the stratosphere for quite a while. But the reflectivity of a ton of sulfur, divided into particles that are half a micron, is very big. The reflectivity of that sulfur, if it's a particles are much smaller or much bigger, is much less. So the answer depends very sensitively on how big the particles are. And I hinted at some of that uncertainty last time.

10:54

So you might think, how could we manage that? The answer is feedback. So the answer is, you could adjust the way you were injecting. There's all sorts of adjustable parameters I won't get into here, we can answer some of that in questions--lots of interesting ways you could adjust it to achieve a given amount of reflectivity. And that adjustment is called feedback, and that's a pretty everyday phenomenon. So let me step back a little bit and sort of show you some slides about feedback.

11:22

So this gets back to a very general question we keep asking about this. How can you control a system that you don't understand very well? This is a comment that a colleague of ours, respected colleague, Ron Prinn at MIT, made about, to some reporter about--or to some reporter, to Oliver Morton, I forgot, who's written thoughtfully on this topic--this was Ron Prinn, kind of implying that we couldn't, plausibly ever do this, because we can't understand the system very well. And Ron is right. There are lots of things we don't understand about the system, but I think Ron is not right in the in the kind of conclusion that that therefore we can't do it. And the reason is that we routinely manage systems we don't understand by feedback. We do it all the time. We're sitting in this room right now, and the temperature is some relatively even temperature, 20 degrees or 21 degrees centigrade. The people who designed this room didn't know exactly what the heat loss from the room was, that is how much heat would be lost if you're adding a given amount of heat. And they didn't know exactly the performance of their heaters, of the, of the whatever is the heating and cooling system of the room. And if they tried to predict how, what the temperature of the room would be just by some setting of the heating, they would get that prediction quite wrong. But it doesn't really matter, because they use a feedback control from from this thing, probably this thermostat, and the thermostat just says, ah, if it's a little too warm, I turn down the heater a little bit. If it's a little too cold, I turn up the heater a little bit. And that thermostat has maybe some clever little control loop, engineers called them proportional and differential integral controllers, but, but that controller means that we can keep the temperature here pretty constant, even though we don't have a good model of exactly how the room performs or exactly how the heating system performs.

13:12

A related question that you might have is, is it, is it, in general, possible to control weather and climate? And I think that you might think, especially for weather, that control is impossible, because the system is chaotic. Chaotic--chaos is defined as, indeed, was invented by people thinking about weather prediction--chaos is defined as extreme sensitivity to initial conditions. So if you're just a little bit long, wrong about the initial condition, you get a very big change sometime in the future. That's this idea that if the butterfly flaps its wings, the weather will be different in a few months. And that's actually true in some sense. But that doesn't mean that it's impossible to do control. So you might assume it was impossible, but it's not. And the answer is, we routinely control different kinds of chaotic systems. And in general, what you need to do to control a chaotic system is you need some model that, with errors, predicts what the future state is from some initial conditions. You need some observations, sort of know what's happening. You need some appropriate lever where you're adjusting the system, and then you use feedback.

14:35

And an example of this thing is there are aircraft that are deliberately designed to be unstable. So conventional aircraft, it used to were, are stable in the sense that if an old fashioned aircraft, if the pilot takes her hands off the controls, the thing will just stably fly. But you can make aircraft, and we do make aircraft, that are dynamically unstable, so that without very fast feedback control, they'd spin out of control. And in fact, there's reasons to do that, because you can make the aircraft perform better in its unstable state. So it's an example of one of these unstable things.

14:58

So there's this interesting idea that we're actually, we, the kind of collective scientists learning about climate and weather forecasting, are essentially building a lot of the mechanisms we need for even weather control without the intention of doing weather control. We're just building them. And that's because we are building steadily over the last 50 years, better and better observations of the climate system. So we have better and better weather satellites and measurements. Arguably, they've got a little worse in the last while, as we haven't spent quite enough money on that. We've got better and better models, and we've got these better and better analysis forecast systems, and you can see that in the better and better weather forecasts. You may not realize that, but weather forecasts are much better now than they were 50 years ago. All those things together means that we could, in principle, do some weather control, or climate control, because we have that model, that's a forward model, and so smaller and smaller levers would allow us to do prediction.

15:57

I enjoy showing you this slide because it's actually one of the oldest slides that I have. I first showed this slide probably for a conference 25 years ago or so, roughly 20 years ago, for a physicist called Steve Schneider's 60th birthday, a climate scientist. And this actually became the start of a whole set of work with collaborators about feedback control, because a guy called Doug McMartin was in the audience, and he was a professor of control and dynamical systems, and he came up to me afterwards and said, well, let's work on this topic. And now he's done an enormous amount actually doing this in climate models, and I'll show you a bit of it.

16:32

So the very basics of any feedback control system are something like this: that you have some process, the process produces some measurable output, you look at the output, you adjust the feedback so that you change the input to that process, and you do that in a way that closes the loop, like the thermostat works. So that's the most generic version of it. And this is a version from a paper that this guy, Doug McMartin, I mentioned to you, wrote, I was part of it as well. And this has got a bunch of fancy formulae, but it's the same idea. What we were doing was using a big climate model, and we had noise and so on, and we were looking at the different outputs, like the global temperature or precipitation or some, you could do it one area, then we're assuming that this is the climate system here on the top and the second half is the geoengineering, and we have some adjustment. And so we did this in real climate models. This shows you what the climate model prediction was for the precipitation averaged over all land under some particular scenario, as the black line, and then the red line is what happens if, starting in 2020, we have a target to keep precipitation at some fixed amount.

17:46

We, even in these models, tried tuning our feedback loop like our thermostat on one climate model and then using it on a different climate model to kind of simulate the fact that we're never going to in the real world know exactly how the world works, we'll have the wrong model. And we wanted to kind of see what the consequences of having the wrong model were.

18:03

So back to this simple idea. I think it is true, there's a lot of uncertainty between the injection of aerosol precursors and the distribution of reflectivity in the stratosphere, but I'm not too concerned about that uncertainty, because it's relatively easy to apply feedback there. And the reason is, it's quick. So you see the, you can see or measure the distribution of aerosol sizes within a week or two after you've released materials from an aircraft, and each individual aircraft is doing essentially nothing to change the climate, just, just incrementally. So you can afford to experiment a lot and figure out what the distribution of aerosols are, and do that at a larger scale across the stratosphere. So I'm kind of optimistic that that could work. I think they defend that optimism.

18:50

But I've sort of implied in this idea that you could divide the problem in half. I've made it look simpler than it is. So now I want to show you some complexity and some limits to what I just said.

18:59

So in this earlier slide where I said, I divide the problem in half, and we, on the one half, just control the stratospheric aerosols, and the other half we give it to the climate modelers. Real world is actually a lot more complicated than that. The question is, how useful is what I just said. I won't take questions right now.

19:15

So one of the examples that I'm leaving out is that the climate will change in response to the added reflectivity, or whatever, the climate will change the circulation of the stratosphere, the way the stratosphere flows. And when the stratosphere flows in different way, then you'll get a different distribution of particles for a kind of input. So that's a way in which you can't just draw a line between these two things.

19:39

And there are a bunch of other ways where you can't draw a line between these things. So aerosols that you put in the stratosphere don't stay there. The fundamental thing. They're going to fall back to the lower atmosphere. They're going to cause acid rain on the ground. They're going to cause air pollution on the ground. In the upper troposphere, the high part of the troposphere, below the stratosphere, we find that air is very clean of particles, has very few particles in it, and so it's very sensitive to additional particles that might, for example, alter the way cirrus clouds, these thin, high, wispy clouds that heat the planet, might alter the way those work, in ways that might, might be heating or might be cooling, actually can be either. But the point is, you can't divide the problem in half as cleanly as you'd like, because putting aerosols into the stratosphere will change the way the climate of the lower atmosphere works.

20:26

And also I just wrote as if the aerosols were just reflecting sunlight, but they're not. Aerosols are also absorbing some sunlight, and that heats the stratosphere, and they're also interfering with stratospheric chemistry, and that changes the amount of ozone, and that changes stratospheric heating. So there's a bunch of different interactions here that I'm not showing. So what I've done is show you three, but there's really, like 27 or whatever, a large number of interactions that I'm not dealing with.

20:50

Now you might say, Ah, if feedback is so good, then what I could do is just do the feedback directly on the climate variables, and then I don't even know anything if I've got enough feedback. I think in practice, that really doesn't work, and that's because the climate system is noisy, and it's hard to see responses. It takes a long time to see responses. And while we have published papers, as I just showed you, that do this kind of feedback control the whole thing, and I think that is relevant to it, there are limits to feedback. I don't think you can do this in a way that ignores all the uncertainties.

21:25

So now let me get a little more careful about what I think the limits might be, and then I'll turn this over to Tiffany in a minute, and then happy to take a question right as I do that, and then then we'll go on. So--trying to be a little more specific, I think if this stratospheric aerosol injection, radiative forcing, that's the amount that we're doing it, is approximately uniform, and it's comparatively small, where I'm not defining what approximately comparatively means--and the problem is, we as a community haven't done that, and this is the kind of cutting edge of research, I think we haven't quantified well enough how even it needs to be and how much is too much before we can't use these kind of simple methods to effectively predict what would happen--but if this is true, then I think it's fair to say the climate response is well approximated by this kind of turn down the sun method of predicting the response.

22:21

And we know a lot about the way the climate system responds to turning down the sun. That's actually been a thing that climate modelers have looked at from the very beginning. So if you think of these early papers--I think in Tiffany's lecture, she mentioned Syukuro Manabe, who won the Nobel Prize for getting some of the predictions of climate right back in the 1960s and early 70s--very early versions of Manabe’s papers looked at the consequences of turning down the sun, not because he was thinking about climate engineering, but because it's a basic thing that, that is a, it's a basic alteration to climate system that we care about a lot.

22:55

So I think that we know a lot about the consequences of turning down the sun, of adjusting the solar constant, to use the fancy words, and I think the key thing to say is that the uncertainty in predictions of the climate response turning down the sun are, I think, and in this conversation Tiffany, I'll have Tiffany, I'll have at the end, I think, no more uncertain than our predictions about the response to CO2. They're both pretty even, well-mixed things. We've looked at both of them in different ways for decades and decades. There are still uncertainties, lots of uncertainties about predicting the response to well mixed greenhouse gasses, but the uncertainties are of similar order.

23:32

I'm saying that to essentially disagree with an underlying idea many people have, which is that somehow predicting the consequences of sunlight reflection is much more uncertain than predicting the consequences of CO_2. I think that's probably not true for these assumptions about sunlight reflection. It certainly would be true for some kinds of sunlight reflection.

23:53

And let me even--push that a little bit further. It's generally true that our uncertainty about predicting the response to sunlight, to carbon dioxide, to greenhouse gasses. That uncertainty gets bigger the more greenhouse gasses there are, or the more warming there is, that as the uncertainty about climate predictions, about the response to climate it's bigger if we're talking about a climate response that's say, three degrees warmer than preindustrial than if we're talking about a climate response at one and a half degrees bigger than preindustrial. So it might be interesting to hear what Tiffany says. It might be that, in some respects, the uncertainty would be smaller if we were doing sunlight reflection, because the net amount of cooling is less. So you've got kind of a balance of two effects. On the one hand, if you're doing sunlight reflection, you've added new uncertainty from the fact you're doing sunlight reflection, this new thing that hasn't been done before in quite this way. On the other hand, if you're doing sunlight reflection, the overall warming is less than it otherwise would have been. And to the extent that uncertainty grows with the amount of warming, with the total amount of what we call radiative forcing, the amount we're pushing the climate system, then it might actually be that uncertainties are in some respects less in a world with sunlight reflection than a world without.

25:06

Then there's the uncertainty in the direct effect of aerosols. So what I'm saying is the overall impact is a combination of the kind of turn down the sun response and the direct effect of aerosols, meaning the chemical effect of the aerosols that come into the atmosphere and they form, they interact in complicated ways to with atmospheric chemistry and ultimately make aerosols on the surface. But that turns out to be something we know a lot about, because we've got this huge knowledge that I'll come to later about the way that sulfur interacts in sulfur interacts in the atmosphere.

25:35

So this is, I see, kind of a cutting edge of the extent to which we do or do not know about how confidently we can predict response to climate engineering. And I'll turn it over to Tiffany in a second. What we want to say here is that that I think you can make claims that in some limited set of circumstances, doing this not too much and evenly, the uncertainties might be pretty low, but you'll notice there's a lot of weasel words in there, “might be” and “not too much,” and I can't quantify exactly what those are. So over to Tiffany to give a really specific example of how we might reduce uncertainties.

26:11

I guess I’ll take one question now. Yep.

AUDIENCE MEMBER:

What I'm not getting is, how do you evaluate the effectiveness of the-the changes, the engineering changes that you are advocating? Engineering—in engineering, we usually deal with a system--like you mentioned, this room, this room is a closed system, but you're talking about a system which is the globe. How do you evaluate? What are the limits in which you make an evaluation? Just for someone who's not doing that every day? I don't see that you get my question.

26:55

KEITH: Yeah, I think I do. I think we should take that at the end. I mean, some ways, your question is absolutely the central question I'm trying to answer, and I don't have a crisp answer for what are the limits of how well we can do it. I think a key thing that maybe we haven't said here is that, and this maybe is a nice transition to Tiffany, is that we're not talking about doing something totally new. We're talking about doing something intentional, new.

So the world is already being cooled today roughly a third of a degree centigrade by sulfur aerosols that humans are putting in the atmosphere. That's aerosols from pollution. They're not deliberately there to cool, but they are cooling, and we've studied them immensely. And Tiffany's one of the leaders in understanding that. And so there are literally, not an exaggeration, 10,000 papers in the last 100 years that deal somehow with aerosols and climate. And while climate engineering seems really striking and new, and it is new with respect to intention and all the politics and philosophy that comes from intention, the actual thing we're talking about is, I think in some respects, well within the realm of what we've already studied. But now I'll turn it over to Tiffany.

28:07

TIFFANY SHAW: Yeah, thanks for the question. I think this will motivate things as we go forward, and the open discussion that we will welcome to come back to that point. But I think David's right. You know, I'm a climate scientist, so I have studied the climate as it has evolved that have—including both natural and human influence unintended, what we might call unintended human influence, okay? And so we have developed a set of knowledge and understanding of the climate system and the uncertainties of that understanding by watching the system evolve, by making predictions, like I talked about in my lecture, last lecture in October. Also was a very beautiful day that day, so maybe I got lucky.

28:53

But what I want to talk today about is the uncertainty of those predictions, and in particular, a particular source of uncertainty that we've learned a lot about as climate scientists that needs to be brought into the fold of these questions that we have when we turn our attention from looking at unintended human influence to intended human influence.

29:16

So one of those key sources of uncertainty comes into this climate modeling aspect. So we have an intention to make a change, and we're using these physics-based models to get an answer of what that change would imply. And I talked in detail last time about these climate models, which are physics based, okay, so they are based on partial differential equations, nonlinear partial differential equations, essentially Newtonian mechanics applied to the atmosphere, the ocean system. And so we can't solve them by hand. They're very, very complicated, but we can solve them numerically. So this is illustrating the numerical grid of such a climate model. And so you've got this well understood large scale physics. And each of these grid points, or vertices of the triangles, have variables that we are integrating in time using these equations. But as you'll also notice, the second big part of being able to move forward in time to get that predicted climate response is all of the physics that's too small to be represented within this numerical discretization, some very, very important physical processes, as shown here, convection clouds, radiation, ocean mixing. There's the radiation, you know, sea ice, aspects of the land surface. All of this is small scale physics that needs to be represented. And essentially as a closure problem, we represent this small-scale physics as a function of the large-scale physical variables, wind, temperature, humidity.

31:05

So that's the setup of these models. The challenge of representing small scale physics as a function of large-scale physics is the classic closure problem. Okay, it is not an easy problem, and over the last 50 years, many very smart people have worked on this problem and have come up with some good closures, but they're very hard to test, and they often, because of the small spatial scale, are very hard to observe. Ok? So they're often thought of as empirical closures with some physical insights, but they depend on parameters that are not well observed.

31:47

Okay, so this ultimately then represents a form of uncertainty. Okay, so there's parameters that we need to set in these climate models. A strong function here will be represented by clouds. That's why they are the major source of uncertainty in climate models. Very hard to observe clouds, to set the values of those parameters, in order to, you know, make our predictions more certain. So we call this structural uncertainty. And amongst all the different climate modeling centers around the world, which I think in my lecture I told you, there's at least 40, there's a table such as this one, where they will—it's become more commonplace now for modelers to talk about these parameters and to talk about the choices that they make, because there's a very large list.

32:37

So here's an example from a model at the Max Planck Institute for Meteorology, which is in Germany, and a collaborator of mine that I will work—collaborators of mine there will work with me to develop an ensemble based on perturbing parameters for solar—stratospheric aerosol injection. So here are some examples. So we have the Prandtl number, which is a very important number to do with small scale turbulence, that of course, involves small worlds that are not represented when you have a global grid like we have for a climate model. We also have things that you might recognize, the conversion of cloud water to precipitation. What are the temperature and dependence of that? What is the humidity dependence of that? That's something that we have to set in these models. Vertical mixing in the ocean, among you know, terminal fall velocity of ice, things that are very hard to measure.

33:34

And so the thing that we need to do in order to quantify the uncertainty associated with these small-scale physical parameters is to do a perturbed parameter ensemble. So the large-scale physics is set. We know the equations of motion. It's the small-scale physics that's uncertain. We have choices for these parameters, and what we're going to do, or the approach that the community has taken, is to do and construct perturbed parameters ensemble.

34:04

So what do these represent? Well, you can imagine choosing three parameters shown here, three axes, all related to clouds, the accretion, the auto conversion from liquid to other forms of water, and the ice fall speed factor. This is, in fact, in a different climate model at the National Center for Atmospheric Research in Boulder, Colorado. And what they've done here is take an approach of a Latin Hypercube, which is a very efficient sampling of this three-dimensional space that ensures that within each group of parameters you're calculating—and these dots represent each of the ensemble members that is essentially efficiently sampling this high dimensional space that represents this physical science uncertainty of small-scale convective parameters. And so you can think of each one of these as a single climate model simulation, okay, where everything is fixed about these, you know, Newtonian equations, and all you've done is then changed the value of these cloud parameters. And we do this again because we don't know the value. There's some constraints for some of these parameters, so we can sample within the observations we do have, but for many of them, it is very difficult to constrain them.

35:27

So we can make these choices, and then what is done then is you run these climate model simulations for each of these parameter values, and then you can show how this fundamentally affects the response of this climate model to a particular change. And so the change that this has really been applied to is an increase of carbon dioxide. So this standard problem of thinking about the climate response, in particular, on the x-axis here, the warming response to a doubling of the carbon dioxide concentration in the global mean. So this is the global warming, and it has a unit of Kelvin. And what's shown, this is a nice paper of one of the earliest such perturbed parameter ensembles from climateprediction.net is the distribution of the warming for the global mean in their kind of standard set up in black, which, of course, has a very high tail, you know, which is the very concerning kind of high climate sensitivity, as we talk about, but very low probability in terms of this distribution; and then the blue, representing a set of perturbed parameter ensembles where they change a parameter called entrainment. This is a convection parameter. It relates to the mixing between the convection, the cloud and the dry environmental air. So right at the edge of that cloud, it can entrain dry air, and the magnitude of that was varied to create this blue distribution. And you can see how it slightly changes relative to the standard setup. And then there's another distribution in red that is shown when they change this conversion from cloud, you know, cloud to rain conversion threshold.

37:12

So this is, you know, a very standard approach for looking at the sensitivity of our uncertainty of our predictions to the increase of carbon dioxide. Of course, global mean is one number that we often talk about, but the regional response to such changes are also of large importance.

37:31

So in that same paper, they showed three different members of this ensemble, the standard, what they called unperturbed, so the model off the shelf, okay, where the warming with a doubling of carbon dioxide, its climate sensitivity in the global mean is 3.4 Kelvin. So what you can see is it's warming everywhere, more warming over land than ocean. This is a signature of a response I talked about during my first lecture. Manabe and his colleagues predicted this well back in 1960s and 70s. And then the changes in precipitation as a fractional change here on the right. Okay, so you're seeing some shifts of the rain bands, some changes here, increases in the monsoonal regions and so on and so forth.

38:19

And so what's shown below are the same two variables, temperature and rainfall. Temperature rainfall, but now when you've changed these convective parameters, in one instance, such that the global mean warming is slightly less, so two and a half Kelvin global mean warming. And in one case, one of these very extreme warming ensemble members, where you get up to 10 and a half Kelvin warming. And so, you know, in general, the patterns show similarities, right? More warming over land than ocean, more warming over land than ocean, but some of the hydrological cycle changes can get, you know, quite a bit more dramatic as you, you know, move into this world that has a lot more warming. And, you know, ideally we would want to constrain this, and as we more, you know, observe the system more we do, you know, use that additional observations to better refine and understand, you know, the distribution, or the likelihood of where we will fall as we continue this unintended experiment.

39:19

But what it really folds into nicely is the idea that climate models have uncertainty, and that uncertainty has to be quantified and reported. Okay, so we use our models to under—to quantify the uncertainty through perturbed parameter ensembles, which is shown as the structural uncertainty in blue. But there's other forms of uncertainty, and so when we do perturbed parameter ensembles, we get to quantify this portion of the fractional total variance. You can think of this as the total uncertainty, or the total change of, this is the global mean temperature with different lead times from the year 2000. And then there's scenario uncertainty, which comes down to how much you know CO2 will be emitted. So it's a very different form of uncertainty. That depends on our human choices. It's not uncertainty in our understanding of the clouds and the climate system. So that's a different type of uncertainty. And then we have noise, because the climate system is quite uncertain.

40:20

So being able to to understand the partition of these three things for the CO2 response has been, you know, one of the important advances in climate science. This paper came out, very famous paper came out in 2009. What we don't currently have right now is this kind of partitioning for the type of interventions that David has been talking about. So research that is ongoing as part of the Climate Systems Engineering initiative is funding new modeling to quantify the structural uncertainty of these cloud—small scale cloud processes and what that would how that would impact some of the temperature, hydrological cycle changes that we would then expect, as a consequence of modifying the sunlight. So I think that's what I wanted to say. And then we'll hand it back to David before the discussion.

41:13

KETIH: Yeah, I mean, I think what, what Tiffany didn't quite say there is, is that I'm so excited that these methods of these perturbed physics ensembles that Tiffany talked about—Tiffany and Sarah Kang, one of the climate science leaders at Max Planck, are now working together to use this method, which has been used before for other topics, as you see, to explicitly go after what is the uncertainty in some versions of sunlight reflection. And the reason I'm excited about that is, I think so far, work on sunlight reflection has mostly just used a few off the shelf models and kind of rerun the scenarios in kind of similar ways, and it hasn't done that good a job to really answer the “how certain are we?” And it's not as if Tiffany's research will answer that question once and for all, but it will make an important new contribution, I hope to going after that question.

42:00

Now I'm going to say a few really sort of general words about why I think that broadly, the reliability of research on this topic might be bigger than you initially think. And then Tiffany and I'll have a conversation about that.

42:14

So let me start with a guess. Ansatz is just a old German physicist word for what your guess is. So, forgetting everything about the details, thinking only about big questions of, how do you individually decide whether or not you trust some body of science? I don't have a right answer. What I'm gonna tell you is my answer, my kind of general heuristic for how I decide whether or not I trust something in science when I don't know the specifics.

42:44

So in general, when I don't know the specifics, I tend to trust research more if the research has been going on for a long time, if the scientists have been getting the same answer for a long time, I tend to take it more seriously. If there's a lot more research, I tend to take it more seriously. Volume. If the research is from a kind of intellectually and globally diverse community, I take it more seriously. And finally, if there's not commercial interests, I take it more seriously.

43:07

So let me give you an example. You may or may not agree with this. I can't prove this is the right way to judge reliability research. I'm just telling you this is how I judge reliability research when I don't know the answer. You know, for example, I've been trying to decide how better I should get certain kinds of knee treatments for some minor arthroscopic problem I have in my knee. And I'm not an expert on this stuff, but this is how I go about judging it.

43:29

So I'll give you two different examples. So one—these are real papers. One is a paper that came out in Nature a few years ago on something I really don't understand at all, this interleukin—I don't even know how to pronounce it—11, some signaling pathway that might show that you could extend the lives of mammals. It's pretty exciting. It’s a Nature paper. If it's true, it might really be a way to help people like me who are aging age a little less. It's totally, awesomely exciting. But how much do I trust it? Yeah, not so much. I don’t know anything specific about these authors. I'm not criticizing them directly, but the reasons I don't trust this too much, the reason I'm not going out to, like, buy stocks in this company, or to start taking IL-11 myself, is that this is really new. It's one group. The group does have at least some conflicts of interest, in the sense that they're, there's connections to people who are trying to commercialize it. And while there is, it's true, a kind of a, turns out, 10 or 15 year history of understanding that this IL-11 is potentially important, this is really the first experiment that's explicitly tested whether it can extend lifetime of mammals. So I'm excited, but I'm pretty skeptical.

44:38

Another paper, actually by a good personal friend of mine, is a physics paper, not expecting you to know anything about what this is, but it's—he was—my friend Eric was trying to measure the electron dipole moment, which is an important thing, and the answer is very likely zero, but we trying to show exactly how zero the answer is. And Eric did it in a particularly clever way. Added some more zeros to the—knowing how zero “zero” is, and I'm really very confident that Eric's number is reasonable, because this grows out of this broad set of physics on electron dipole moment that's sort of global spanning. There's people like all over the world physicists who work on this. Physicist--there's no commercial application anybody can think of at all, so there's really no issue of commerciality. So there's big intellectual diversity, at least world spanning. Maybe you can all argue that all us physicists who care about the electron dipole moment have some kind of strange mental problems, but in any way, there's duration, there's intellectual diversity, there's no commercial aspect. And this is a part of the mainstream of physics where there's immense volume. So I'm—I trust Eric's result quite a lot.

45:46

So how does this apply to SRM? I think let me first give you the argument that might make you think you should not trust research on SRM. So here are the reasons you shouldn't trust research on SRM using this algorithm. So first of all, most research is done in the last decade. More than half of all the papers published about sunlight reflection are published in the last decade. So that's, all else equal, a reason to not trust it that well. Second, much of the work has been done by a relatively small clique of researchers, of whom I'm one, and all else equal, that's a reason to be distrustful. And in some kind of sense, the volume is low, in the sense that the total funding for research, and you know, physicists, I mean, how do you herd cats? You put out cat food. Physicists or climate scientists do work that's funded, and the funding for work on sunlight reflection is, at most 0.1% or so of the funding for work on climate science. So not much work, in some sense. So those are reasons to be skeptical about any kind of general answers about SRM.

46:46

But here's the counterargument. So while it's true that most of this has been published in the last decade, this idea is not new at all. The original ideas here are, in some respects, as old, or even slightly older, than ideas about measuring or predicting the response to CO2. So the very, very first suggestions and analysis of what we now call silent reflection basically arise at exactly the same time as people first started to worry about CO2 in the atmosphere. So the idea is not new.

47:21

And it was, in fact, part of the big climate assessments of like the 1977 Energy and Climate Assessment of the National Academy, the ‘83 National Academy report, the ‘92 report. They were all built in and my colleague, Ken Caldeira and Bala, they, for the first time, used a modern climate model to assess how well sunlight reflection might work in a paper that was frankly motivated by somebody who was overclaiming how well it would work, and motivated with the goal of showing that sunlight reflection wouldn't work very well. And they found, to their surprise, that it did work pretty well in this model. And that model is already more than a quarter of a century ago. It was published in, I forget, ‘99 I think was publication date, so.

48:06

And I think there's been, I think, a very healthy and appropriate bias in the broad climate science community to be very skeptical about research on sunlight reflection and try and show why it wouldn't work. And so one of the reasons that I think we could argue for reliability is it’s now been a quarter of a century, and there's been lots of people motivated in the climate science community to publish a paper showing this is the proof that SRM doesn't work, that it really inherently has some big flaw. And there is no such paper. A quarter of a century that, since this, Bala and Ken's paper, and there's no such thing. And not to say there aren't lots of uncertainties, but there's no fundamental problem that's been resolved, that's been shown.

48:45

And I think, and this really, actually, is it picks up directly on Tiffany's work as an example. I think the reason that I'm saying that that there might be a good case for a liability of research in SRM is that SRM is not thought, should not be thought of is some entirely new thing that is like a new technology that needs to be tested and gone through experiments because it's a new thing. Really, the idea of doing this intentionally is new, but the underlying science isn't really new at all. That's a way in which I think it's not right to think of sunlight reflection as a field. Sunlight reflection is an application of an existing set of knowledge about aerosols in the atmosphere, about the way the atmosphere circulates, about certain kinds of aerospace, about aerosol science and physics. These are all existing disciplines that have been around for a long, long time. So to give you a sense, there have been more than 10,000 papers published that deal with aerosols and climate. So it's not a new topic. And there have been more than 2500 papers that deal with stratospheric aerosols. And again, thinking of duration, the first of those papers was published in 1961 or actually the first review paper.

50:00

And I actually, in preparing an earlier version of this talk, I went back and read that paper. I'd never read it before. And I was kind of stunned by how there aren't obvious things are that clearly wrong about it. In some fields, you read papers from, you know, more than 60 years ago, and they just look like these people are living on a different planet. But this paper is quite consistent with our current knowledge of aerosols, and what that tells me is our knowledge of stratospheric aerosols has been quite consistent over the decades. We still learn things. There's still surprises. But it's been quite consistent.

50:27

And for that reason, I think you can draw two kinds of conclusions. One is that I think there are good reasons to believe that with those caveats I gave you before, if it's even, if it's not too much, where I can't quite define what “even” and “too much” is, that existing stratospheric science and existing climate science is a very strong basis with which to make predictions about the consequences of SRM. This is not a statement that SRM won't have bad consequences. It will have bad consequences. It's a statement that we have a deep scientific basis to quantitatively predict what those bad and good consequences might be, because this isn't like some totally new thing out of left field. It's building on existing knowledge. It is different that we're doing it deliberately, but with respect to the scientific uncertainties, they don't care about whether it's deliberate or not. The scientific uncertainties just care, if you like, about what is being done. And in lots of respects, what is being done is relatively small, and there's evidence—or what might be done, what I'm talking about under that scenario would be done is relatively small—and there's this bunch of evidence that that it could have low impacts.

51:37

So what I want to do now is bring Tiffany back up and and I'd like to hear, and we haven't pre-done this, kind of her reaction to these general claims about kind of basis of confidence, because I think this is really the cutting edge. This question this gentleman asked about how much we should believe this is the essential question that I think our community hasn't answered clearly.

51:57

SHAW: Yeah, I hope this is still working. Okay, great. Yeah, I appreciate the freeform nature of this. So I had the benefit of having you talk more than me. I got to write a few things down. I think I collected my thoughts as follows. I mean, this is the question even facing climate prediction as a whole. Okay, so in my lecture in October, I talked a significant amount about how we started making these predictions in the 1960s and 1970s based on some physical ingredients, based on some assumptions, but we were largely predicting worlds we couldn't see, we couldn't observe, so we didn't know if this, you know, we couldn't test our understanding as represented by these climate models. Okay, so what has evolved in the intervening 50 years is—emissions have not been reduced. CO_2 has continued to increase, and we have started to see the signal emerge from the noise. So this has enabled us to now assess, you know, what is the success of our predictions? And so I think that you know the challenges, or the successes that we see, if we go back to Manabe, you know, if you want to make a successful prediction, what Manabe, who won the Nobel Prize, showed us, is you have to have a certain, or there is a pathway, and it is of a certain approach, which is to break the problem down. Okay? It's a reductionist approach.

53:26

I talked in much detail in my lecture about a climate model hierarchy, where you take the global mean, you increase CO_2, what do you see? Now you add the complexity of land versus ocean, what do you see? Then you add the ocean circulation, what do you see? And by doing so, what he showed is that you would expect global warming at the surface, stratosphere cooling, land warming more than the ocean, that there could be transient cooling at the surface, particularly over the Southern Ocean. These have all emerged. That is a successful prediction. So I think that if we're trying to say, if that's the baseline now for climate science and what it has successfully achieved, a lot of where we see discrepancies or uncertainties or potential surprises is where we haven't laid that groundwork.

54:14

And so I think with SRM or intentional modification or intentional engineering, what we need to do is build the scaffolding and build the knowledge, particularly with this hierarchical approach, since that's my preferred research path, to really understand the system in the context of solar dimming. So indeed, many simulations have been run with solar dimming, Manabe being one endpoint, and then very complicated models, with all these uncertainties being the other. But there's a huge knowledge gap in between. There's much that can be done to run simulations where clouds are not represented in these empirical ways, where clouds are actually fully simulated, because we know the equations. And so I think what—you know—we need to do is invest in that as at least a means of kind of backstopping some unintended consequences.

55:09

KETIH: Where Tiffany's work, her sort of actual professional work, is at the cutting edge of science has been trying to understand some of the discrepancies in the current climate predictions, and particularly looking at examples like as we have reduced aerosol pollution in Europe, because we these aerosols were killing people from air pollution, and we've reduced them. You're seeing climate changes in in Europe and down into Asia that are significant and can be at least partially attributed to the fact that we've, we were doing this cooling from pollution, and now we've reduced this cooling for pollution, and some of the examples of European heat waves, I think, are attributable to that with uncertainties. I think that the question is, how much can we learn from that attribution, and how much easier is it if you're talking about relatively even stratospheric distribution? If, basically, if I'm right, and we can argue about whether that's true or not. If it's true that you that you could achieve a pretty even stratospheric distribution of aerosols, to what extent is that an actually easier prediction problem with respect to predicting the results of aerosols and the than the aerosol prediction problem we've been dealing with, which is this patchiness, where lower atmospheric aerosol pollution is patchy, and the cleanup of pollution has been patchy? You know, we in North America and Europe have cleaned it up more, quicker than before the Chinese did, and now we're seeing all these cleanups and kind of patchy changes. So let me speculate on that.

56:32

SHAW: Yeah. I mean, that's a fair question. I mean, this gets into how this is done. And, you know, there are engineering issues that I am not an expert to speak about, but I do agree that we if we can condition the statements on having distributed the stratospheric aerosols across both hemispheres, then we do have less of a knowledge gap, let's say. We have very strong theory that tells us that if you can achieve that, you will not shift the—you won't have any unintended consequences in shifting rainfall patterns, which are the, you know, the surprises we don't want. We want to avoid large shifts in rainfall patterns, such as the monsoons. And so I think it is a good idea to condition on a scenario, on a particular, you know, way of which we would distribute these aerosols to reflect the sunlight. Achieving that scenario is definitely another thing we should talk about. But in that particular case, I do see theory, idealized modeling, global climate modeling, that would support not shifting these hydrological patterns, cycle patterns, as much.

57:44

I think what you're also touching on, though, is natural experiments. Okay, so the recent changes in the Clean Air Act to remove sulfate aerosol for our own you know, human health and air quality has been an experiment that we can put models, our models to the test. Do they represent what has happened, you know, since the Clean Air Act? And so work of mine, among many, many colleagues, has shown that in those regions over North America and Europe, where this was done by intention, not for the climate, but for the human health aspects—I mean, okay, also, acid rain—models perform pretty well. And I think it does surprise a lot of people, because of what it, we see from that is that clouds don't seem to be involved a lot. So David mentioned on his slides the direct effect of aerosols. And that little jargon says it's just how the aerosols interact with the sun. The sunlight comes in, it reflects. If you take this, the aerosol away, the sunlight just comes down to the ground. So that, I think was a good test of our modeling and understanding. But aerosols do interact with clouds, and so ideally, we would have another natural experiment that could test that.

59:05

Another natural experiment, though, just as an aside, is volcanic eruptions. Okay, these are very important. And Pinatubo has been a very rich source of understanding of natural experiment for stratospheric aerosols and the surface changes of the surface climate. And so there has been work that has really, a lot of work that has looked at Pinatubo and has shown, you know, the temperature response, the changes in the humidity, and models, you know, tend to do a good job with that. Of course then there's more volcanic eruptions like Hunga Tonga, this really crazy one, and it injected a whole bunch of water into the stratosphere. So continuing to use these natural experiments as tests for our models continues to advance the science forward and makes the models better.

59:56

KEITH: Maybe just one last thought and then turn over to questions. So in the lecture two before, and in the opening lecture, we talked about the fact there's a range of sunlight reflection methods from marine cloud brightening. I talked about specifically stratospheric aerosols, mirrors in space, in principle. So marine cloud brightening really only works where there's, if it works at all, in certain areas where there's the right kind of marine cloud. And those areas are in particular places that we actually understand pretty well. Basically, think of kind of off the coast, off the coast of Seattle, or off the coast of San Francisco, where there's a certain kind of marine cloud deck. And if we could make those brighter, we could reflect away sunlight. There's no question. We could cool the Earth, but it would inherently be more patchy, because marine cloud brightening only works in these certain set of areas where it works. But I think one of the things that you can take from what Tiffany said is, in fact, the models did a sort of surprisingly good job predicting the response to the patchy cooling from sulfur pollution. And so they might do a pretty good job predicting response to the patchy cooling from MCB, but it's probably even easier for the models to a pretty good job predicting the response to stratospheric aerosols, because stratospheric aerosols could be, if we wanted it to be, much less patchy and much more even. But there's still the uncertainties.

1:01:15

SHAW: Okay, but I would have to come back at you, David.

KEITH: Please.

SHAW: These low marine clouds, you know, in those, in those distributions I showed, yeah, where it goes really warm to like 10 and a half degrees—what happens in those climate models, in those worlds is that those marine clouds go away. They dissipate, they break up. And that's a huge source of uncertainty and huge potential for warming. So whether we understand marine clouds well enough to go in there and get them to do what we want them to do, I would say there's certainly some uncertainty that needs to be quantified

1:01:58

KETIH: 100% agree. Maybe let's close with that last question then can take questions. So far been sort of talking about actually doing sunlight reflection at scale. But of course, there's an intermediate thing where we do it in this experimental mode. And I think there are kinds of experiments you could do, even if you never were going to do sunlight reflection, if you were committed to never, ever, ever doing it. There's still actually experiments you do that would teach you stuff about the climate system you want to know.

1:02:24

So this question of exactly what is the effect of aerosols on these marine boundary layer clouds is really important, even completely forgetting marine cloud brightening. It's important to understand in climate predictions. And you could imagine experiments that put deliberately sea salt aerosols into these marine boundary layer clouds in some little patches for a week here and a week there, that would teach us something about the way those clouds work that we need to know to make the prediction about global warming, completely independent of whether we ever want to do marine cloud brightening. And there are other examples of that with cirrus clouds that kind of connect the question of the long term actually doing this to the near term question of what the benefits or harms of research would be. But I think with that, I'd like to take questions. Thank you.