Grantmaking methodology for high-uncertainty contexts: what we learned from climate

EA Global London 2026 — Johannes Ackva & Stan Pinsent, with other project members Luisa Sandkühler, Hubert Thieriot, and Hannah Yang – Founders Pledge

This transcript is mostly AI-generated with light human review. Mistakes are possible.

Introduction

Moderator: Good afternoon, and welcome to this talk, which is about grantmaking methodology for high-uncertainty contexts: what we — the talented team at Founders Pledge — have learned from working on climate.

Giving this talk will be Johannes and Stan from Founders Pledge. Johannes is the climate lead for the Founders Pledge team. He has been losing sleep over high-uncertainty grantmaking for many years, and says he hopefully learned something along the way. And our other speaker is Stan. Stan is a mathematician by training, focused on the modelling behind climate grantmaking at Founders Pledge.

Please can we give them a hand?

Slide 1 · 0:37

Johannes: Welcome to our talk on grantmaking methodology for high-uncertainty contexts: what we learned from climate. You can almost not read “climate” here, and that’s kind of intentional —

Slide 2 · 0:45

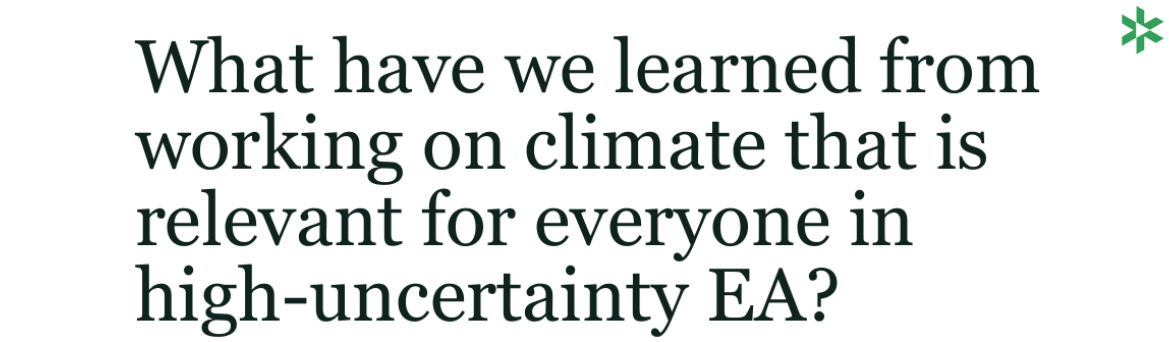

— because this talk is really about what we've learned from working on climate that is relevant for all of you, for everyone who's working in high-uncertainty effective altruism. And don't worry, we're going to define what "high uncertainty" means.

Three key points

Slide 3 · 1:00

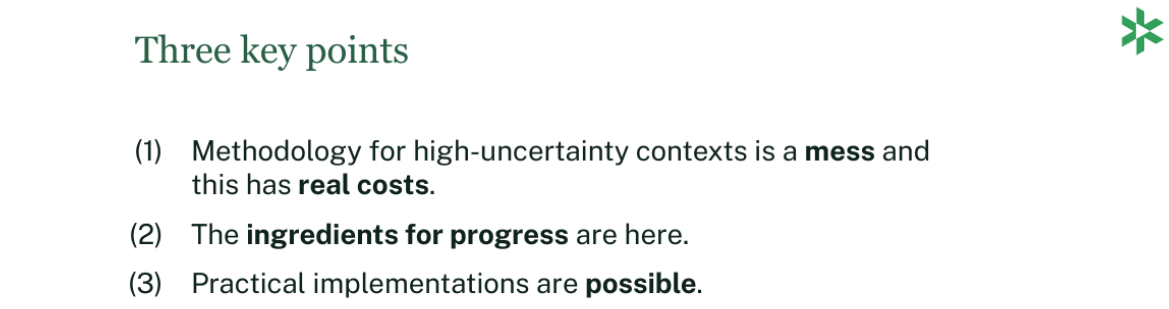

What we're going to try to do is make three key points.

First, that the methodology for high-uncertainty contexts is a mess, and that this has real costs. This is not free — that's our first point.

Second, that we have the ingredients for progress: there's actually a lot of stuff in the community, ideas and methods lying around, to make progress here.

And third, that practical implementations are possible. Stan is going to illustrate how we are thinking about this as a practical matter. So that's what we're going to cover.

Part I · Methodology for high-uncertainty contexts is a mess, and this has real costs

Slide 4 · 1:30

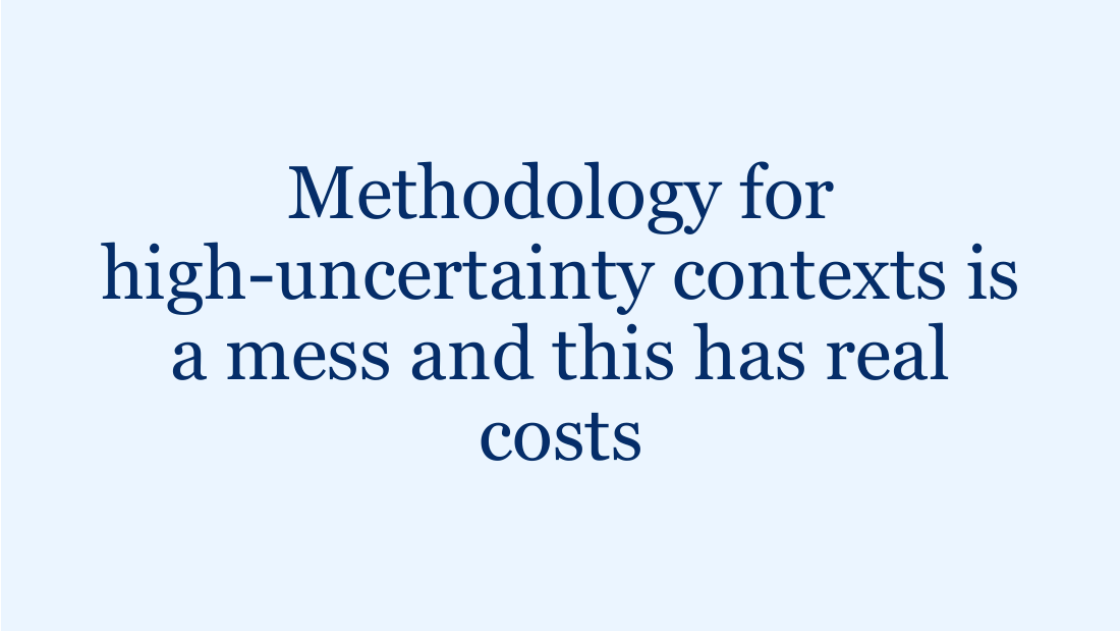

To start with: why do we think that charity evaluation in high-uncertainty contexts is a mess?

The gold standard: direct-delivery global health and development

Slide 5 · 1:35

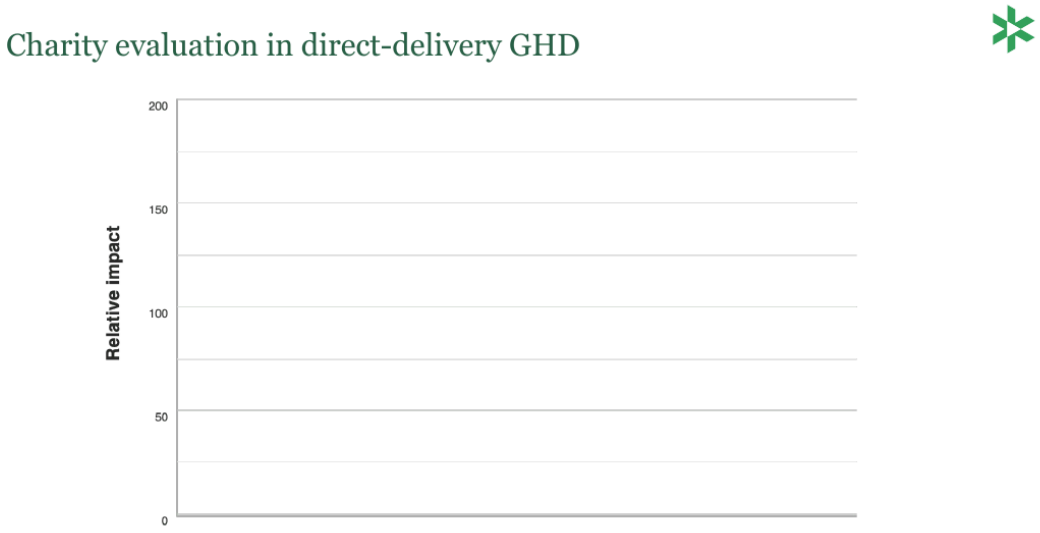

Let’s compare it to the gold standard. And the gold standard for charity evaluation really is direct-delivery global health and development.

Just to flag this: when I shared this with our global health and development lead, she gave me about seven footnotes on how what I’m going to say is basically correct but mildly wrong. So I forgot all of those footnotes — just take this as a stylised example, and don’t yell at her or at me about minor differences.

Slide 6 · 2:15

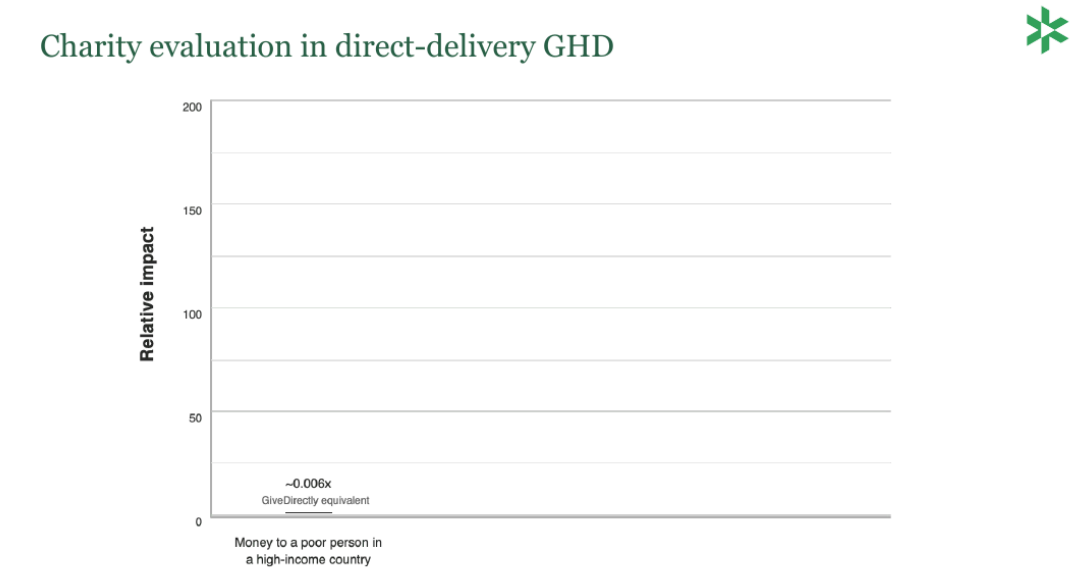

If you think about direct-delivery global health and development, you can think about GiveDirectly equivalents as a measure of impact, and then think about different interventions. Here you have giving money to a poor person in a high-income country. That's very low impact.

Slide 7 · 2:26

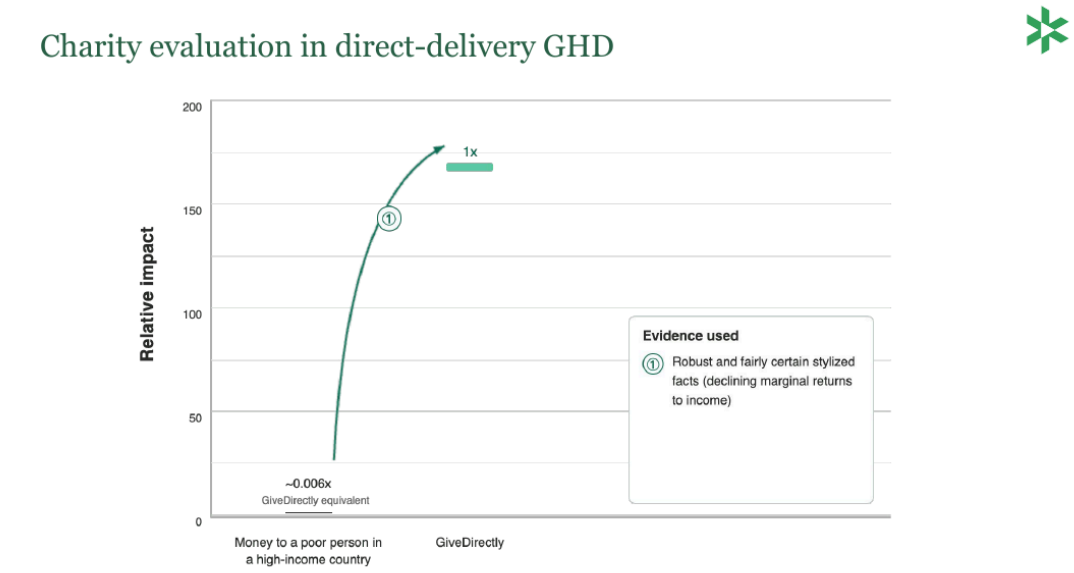

You get almost a 200x multiplier from focusing on the global poor, and that gets you to GiveDirectly at 1x GiveDirectly in terms of rating or now it is three.

And importantly: how do you get there? There's a robust stylised fact about the way of the world — declining marginal returns to income is the key dynamic here.

Slide 8 · 2:43

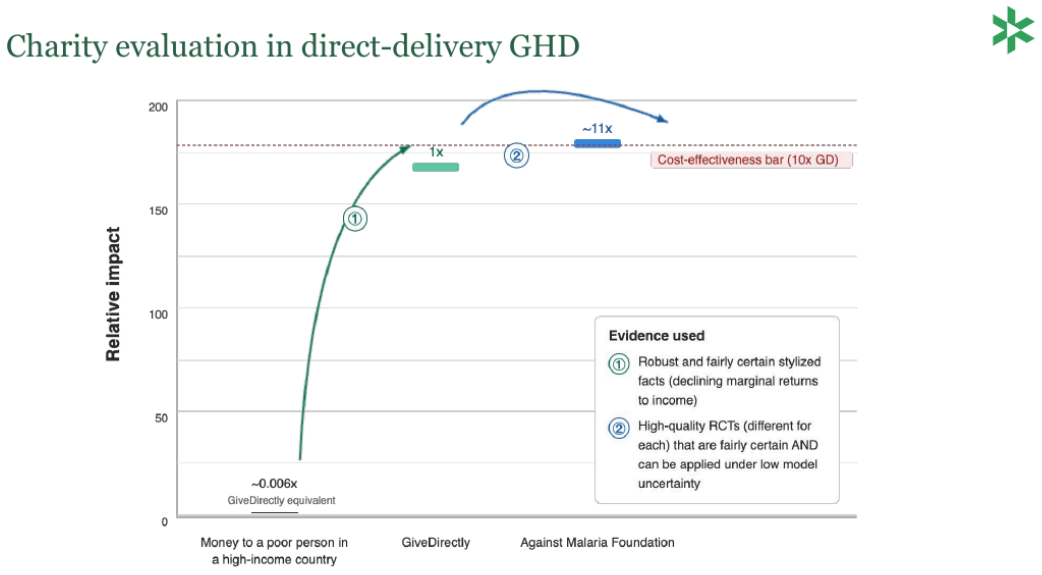

And then you get everything that GiveWell and others are doing, which gives you additional multipliers — for something like the Against Malaria Foundation.

Slide 9 · 2:51

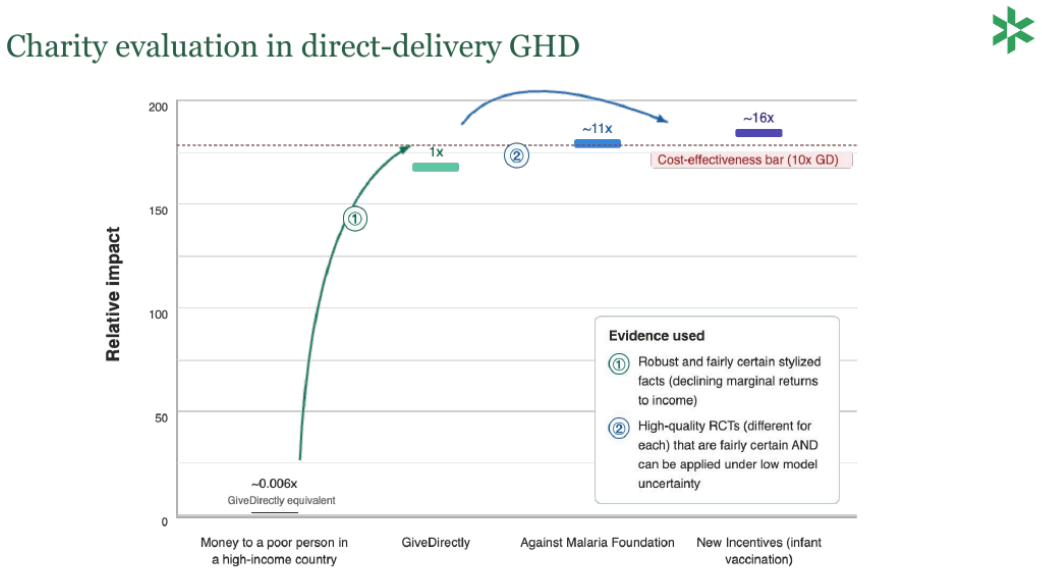

Or New Incentives. And that's very high-quality evidence: these are randomised controlled trials. So that's one key aspect of this evidence.

And there's another one, which is that you know fairly closely how this evidence applies to the interventions that you're trying to understand. So there's low model uncertainty.

This is the ideal case, and it is extremely different from what everyone in a high-uncertainty context is experiencing.

How this (doesn't) translate to a high-uncertainty context

Slide 10 · 3:19

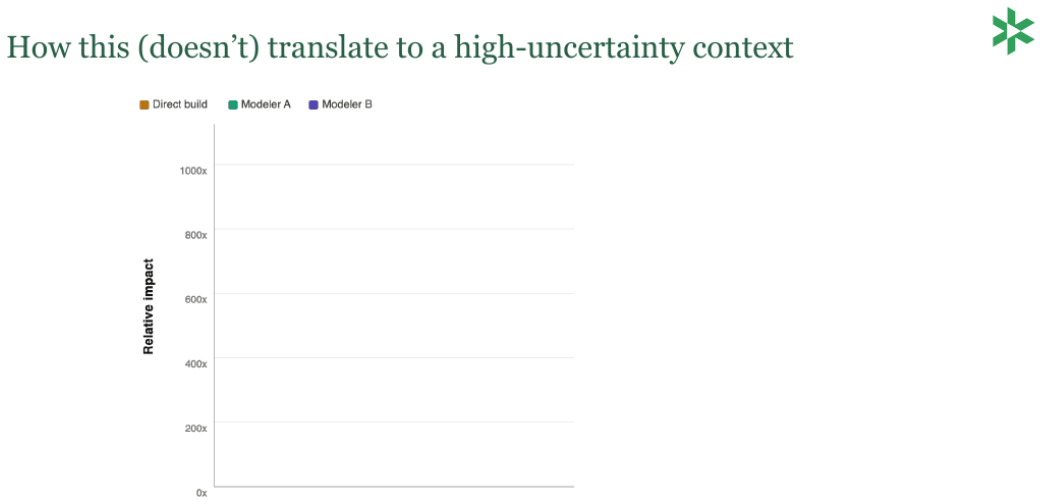

To give this intuition — and this is really how I got started on this journey of thinking about it, I think six years ago now — when you think about something like climate, you might want to think about the relative impact of building a new energy technology that's neglected.

Slide 11 · 3:32

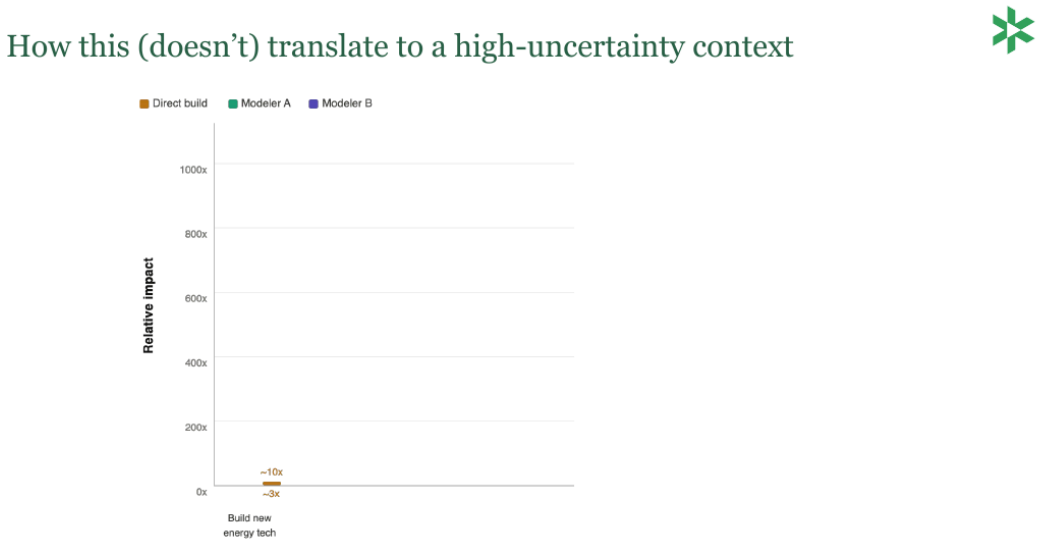

Directly building it: it's not very cost-effective. Here you have a kind of stylised fact.

Slide 12 · 3:40

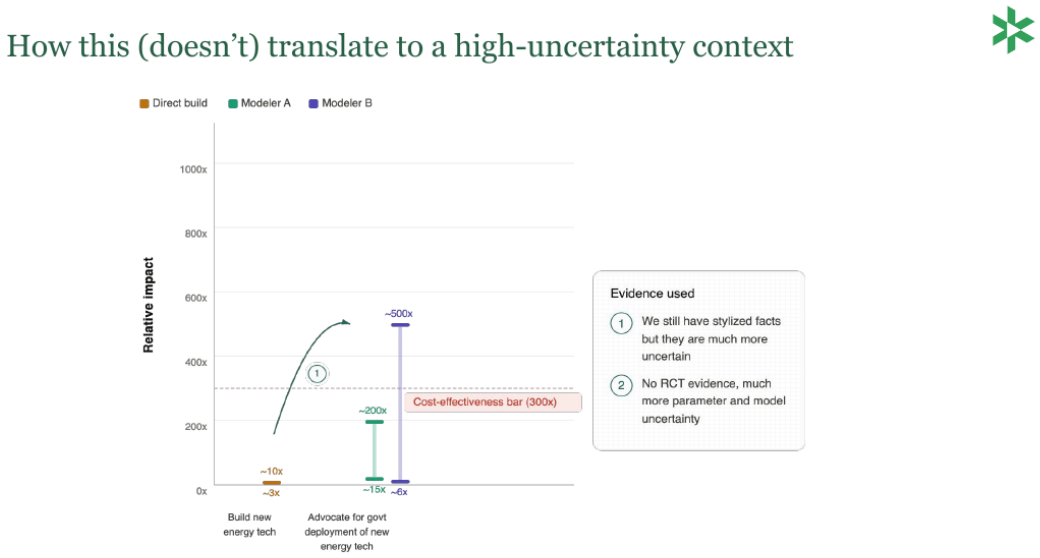

You do this through advocacy rather than building it directly — you advocate for the government to deploy a new energy technology. That's probably more cost-effective. No one knows exactly how much more cost-effective, and, crucially, two different modelers — I've called them A and B here — will have very different intuitions about this.

Slide 13 · 4:01

How this (doesn't) translate to a high-uncertainty context

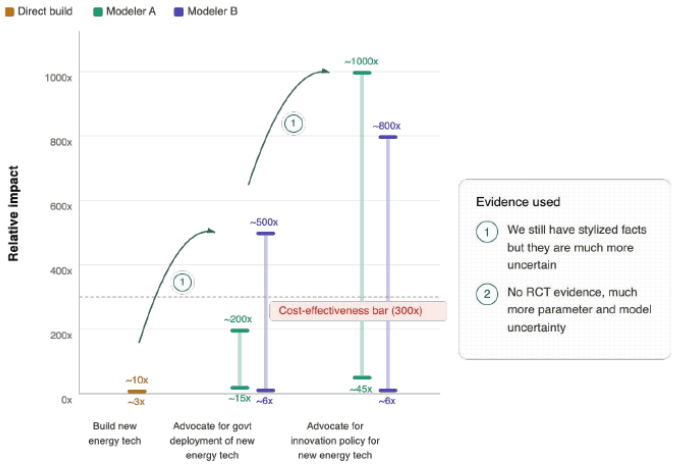

And then you add another uncertainty. Instead of only advocating for a technology, you advocate for innovation policy for the technologies — you try to lever technological change as an additional variable.

What you're seeing here is that things get extremely uncertain. It becomes an almost meaningless question whether something meets the bar.

And there's something more subtle, which is a bit harder to see: not only is the evidence much more uncertain, there's also a lot more model uncertainty. You don't really know how the evidence applies, and different people have very different intuitions even though they may agree about the broad structure.

So this is a really big problem.

Slide 14 · 4:43

I think this is the first major conclusion: essentially everyone agrees that copying the direct-delivery GHD methodology doesn't really translate to high-uncertainty contexts.

This affects a lot of fields

Slide 15 · 4:54

And this is not an obscure problem. This is not a problem that only climate has, because it affects a lot of fields.

The only low-uncertainty case that we're actually working with in EA is direct-delivery GHD — the example I showed at the beginning, the Against Malaria Foundation and the like. Here are all of the other fields that have varying degrees of high uncertainty. It's essentially all the other fields that EAs care about.

So we're in this really interesting situation where we have a really robust methodology for one cause, and then most other cause areas are characterised by high uncertainty. We don't really have an answer here. We don't have a unified methodology for those different fields. This is a really big issue.

Existing responses

Slide 16 · 5:43

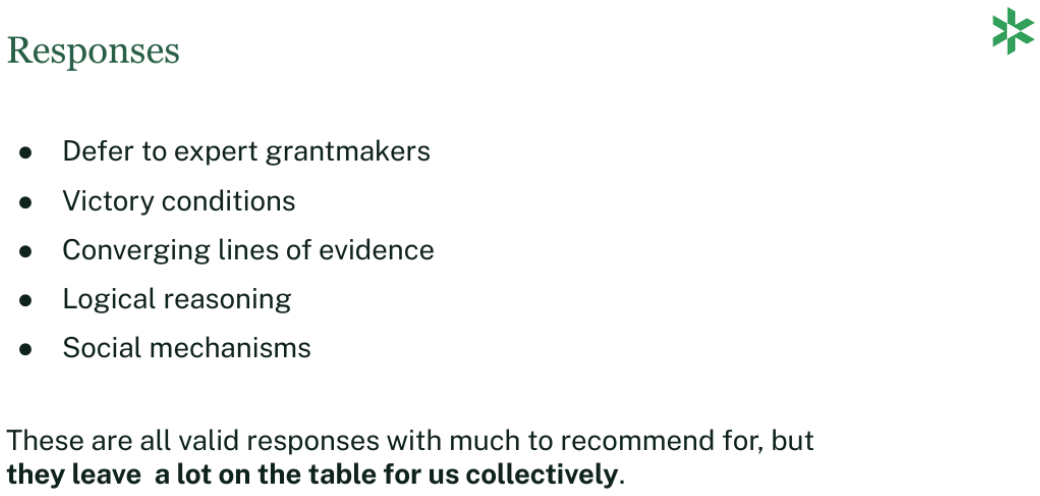

And to be clear, that doesn't mean there's no methodology — don't get me wrong. So, to walk very briefly through the different responses to this problem:

One response is that you do a lot of cause prioritisation and then you defer to an expert grantmaker on what the different opportunities are that should be funded.

The second one is slightly more reasoning-transparent than that: you're not only deferring to an expert grantmaker, but you're laying out a very detailed theory of victory. How do you solve a given problem? You define all of the parts that are necessary for doing this.

Another approach you often see is thinking about converging lines of evidence — saying, we're not believing a single cost-effectiveness estimate, but here are all of the arguments why you should believe that this is much better than other things.

So that's another approach. Logical reasoning — often: if you believe this, then you should believe that this is high impact. Or social mechanisms: for example red-teaming, or group-based grantmaking as we see with some of the funds, like the Long-Term Future Fund.

Those are all valid responses. These are not bad responses; we use many of them ourselves, and they have much to recommend them. But they really leave a lot on the table for us as a community, because they don't encourage us to make progress on this problem of high-uncertainty methodology in general, and they're also often not really justified in a very clear way.

The costs of having no agreed methodology

Slide 17 · 7:05

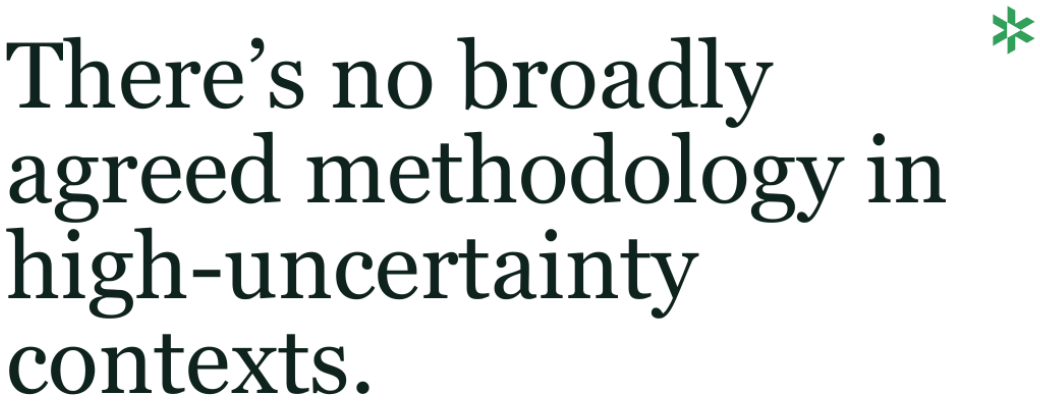

So this is, I think, the first major conclusion. There's no broadly agreed methodology in high-uncertainty contexts. And this is actually a problem.

Slide 18 · 7:13

There's no broadly agreed methodology in high-uncertainty contexts. This has significant costs.

This has significant costs. Let's talk a little bit about what those costs are.

Slide 19 · 7:18

Collective Costs

- **Lack of Reasoning Transparency**
  - Credibility
  - Legibility → harder fundraising
  - Lack of cross-learning
- **Lack of Comparative Theory of Impact**
  - Arbitrary Decisions
  - Lack of transfer across structurally similar opportunities

Over time and across the community, these **costs accumulate.**

Broadly speaking — there are more costs, but in the interest of time — we're going to focus on two sources of costs.

One implication of this is a lack of reasoning transparency. A lack of reasoning transparency has many different costs. It's a cost to credibility: when there's no clear methodology, it's difficult to understand why something is recommended.

There's also the related concept of legibility. It's hard to understand, and this makes it harder to fundraise for a given problem. This is a weird situation, because many of us believe that work that is highly uncertain is often better in expectation, but it's harder to fundraise for, because the case is harder to make.

And there's a lack of cross-learning across different high-uncertainty causes. Right now all the different high-uncertainty causes have different and bespoke methodologies that are not clearly justified. So there's less learning than there could be.

And then there's something related to this: there's a lack of a comparative theory of impact. In many high-uncertainty contexts, grants are understood as — you just go grant by grant, and you're focusing on specific grants rather than comparing sets of grants. This can lead to arbitrary decisions. I showed this at the beginning: depending on which of the modelers you trust, you might recommend something or you might not.

It also leads to a lack of transfer across structurally similar opportunities. One thing you will often find is that something is changing in the world and it should affect many different grants, or you're learning something new — but if you don't have a systematic methodology, you're not able to propagate those changes.

Sorry, this is a little bit abstract, but Stan is going to make all of these points much more concrete in just a few minutes. So we're going to have examples. But really our claim is that over time and across the community those costs are accumulating, and that that's actually a big problem. It's a bit of a commons problem.

Part II · Three reasons why this could be different

Slide 20 · 9:14

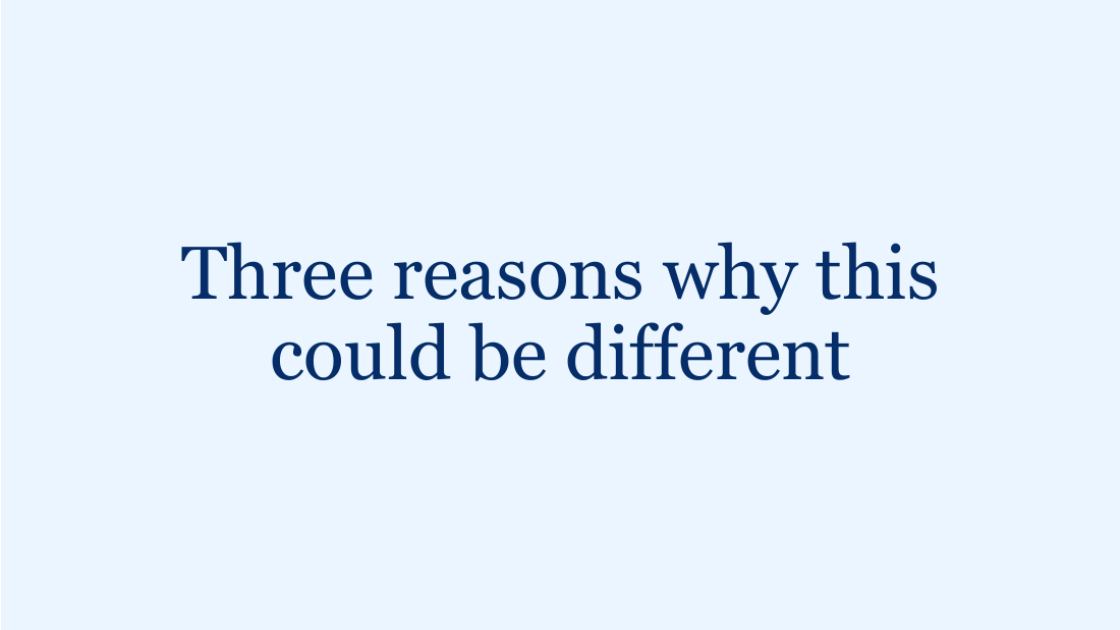

Okay. So before I hand over to Stan, I want to give you three reasons why this could be different — the optimistic case that the world could be much better.

Reason I: there is a broadly shared intuition about the structure of impact

Slide 21 · 9:21

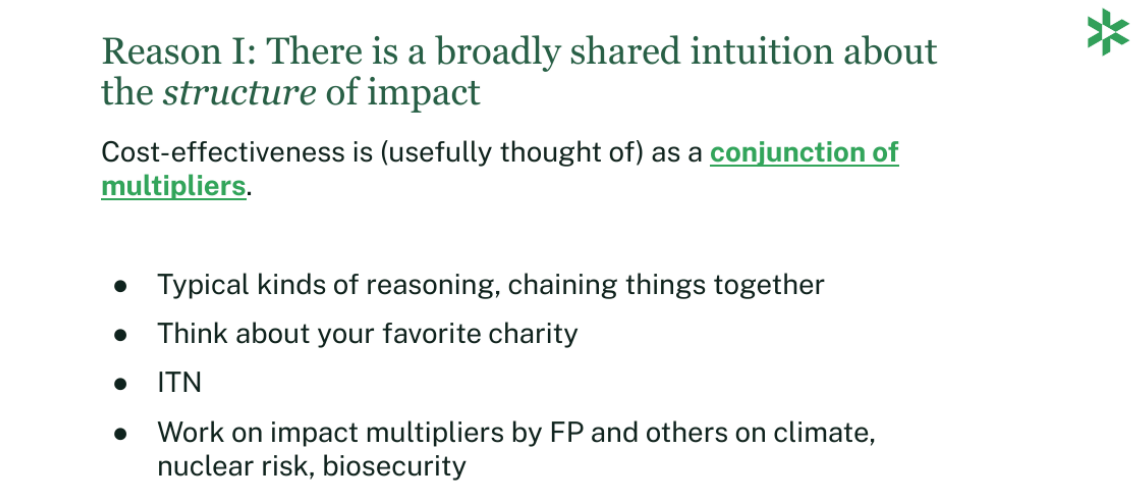

Reason number one is that most of us actually agree, or have a shared intuition, about the structure of impact.

There's a very powerful idea here, which is that cost-effectiveness can be thought of as a conjunction of multipliers — my favourite EA Forum blog post ever goes into more detail there. But a couple of ways to motivate this idea:

If you go back to where I started — the different ways of going from giving to a poor person in a rich country, to GiveDirectly, to the next intervention — that kind of reasoning is very typical.

If you think about your favourite charity for a moment, you will probably also think: why do I think this is good? You can probably reduce it to a couple of reasons that relate to each other in a multiplicative fashion. For example, when I think about my favourite animal charity, GFI: why do I believe they're good? Well, they're working on a neglected moral patient — farmed animals. They're leveraging advocacy, so they're leveraging public resources. And they're leveraging a mechanism like induced technological change. Bruce may have given you a lot of other detailed reasons why he thinks they're great, or why one can think they're great. But it comes down to this: 80% of why this is great can be explained in three sentences, and these reasons relate to each other in a multiplicative fashion.

Also, the most important intellectual contribution of our community — the importance, tractability and neglectedness framework — is a multiplicative framework. You have different forces, importance, tractability, neglectedness, together jointly giving us an expectation of impact. So this is another way in which we already believe in this, but we're not really carrying it over to grantmaking at a lower level than just cause prioritisation.

And then, we've been doing a lot of work on this, both in the climate team, but also our former colleague Christian Ruhl has brought this approach to other areas that are high uncertainty: nuclear risk and biosecurity. So there's a lot of nascent work and thinking, and actually shared understanding, that you can leverage there, and it's underutilised right now.

Reason II: thinking this way is a cost-effective way to reduce uncertainty

Slide 22 · 11:35

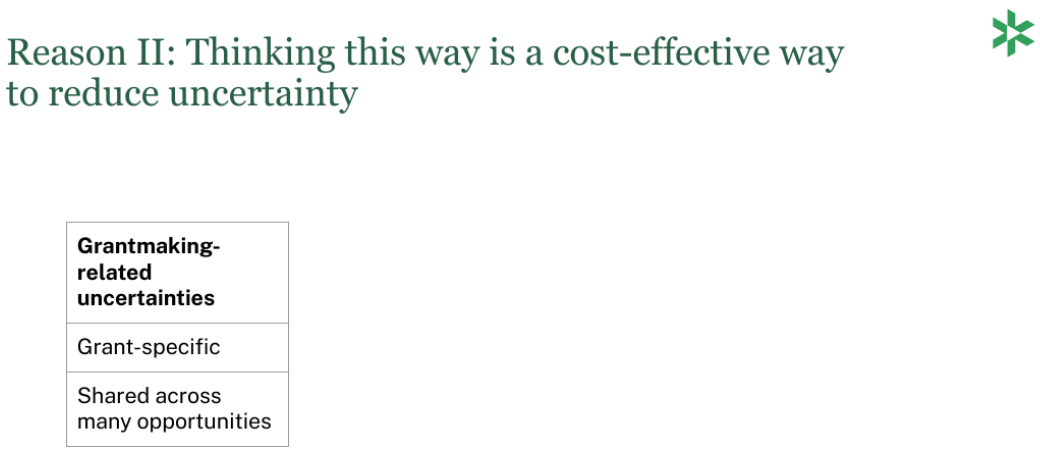

The second reason why this could be different is that this is actually a cost-effective way to think about how to reduce uncertainty.

What do I mean by that? If you talk about different grants and the uncertainties that you face, they're essentially always of two sorts. There are grant-specific or idiosyncratic uncertainties — what is the specific plan of this charity, how good is the team, and so on. And then there are uncertainties about this grant that are shared across many opportunities.

Slide 23 · 12:05

Reason II: Thinking this way is a cost-effective way to reduce uncertainty

Grantmaking-related uncertainties	Cost to reduce uncertainty
Grant-specific	High
Shared across many opportunities	Low

If you think about cost: where is it costly to reduce uncertainty? It's often very costly to reduce uncertainty that's highly specific to a given charity, because the charity always has a massive information advantage — asymmetric information. The cost to reduce that uncertainty is high.

Whereas uncertainty that's shared across many opportunities often lies in the environment — for example, how good is innovation policy, or how effective is advocacy. This is something you can observe in the environment.

Slide 24 · 12:32

Reason II: Thinking this way is a cost-effective way to reduce uncertainty

Grantmaking-related uncertainties	Cost to reduce uncertainty	Benefit of reduction
Grant-specific	High	Narrow
Shared across many opportunities	Low	Broad

And then, the benefit of reduction. If you reduce grant-specific uncertainties, it helps you with a single grant decision. If you reduce uncertainty about something that's shared across many, it's broad.

Slide 25 · 12:42

Reason II: Thinking this way is a cost-effective way to reduce uncertainty

Grantmaking-related uncertainties	Cost to reduce uncertainty	Benefit of reduction	Cost-effectiveness
Grant-specific	High	Narrow	Low
Shared across many opportunities	Low	Broad	High

This really leads to the situation that the cost-effectiveness of focusing on shared uncertainties is actually much, much higher. So this is a cost-effective way to think about uncertainty reduction.

Reason III: our community is already obsessed with quantifying and modeling uncertainty

Slide 26 · 12:50

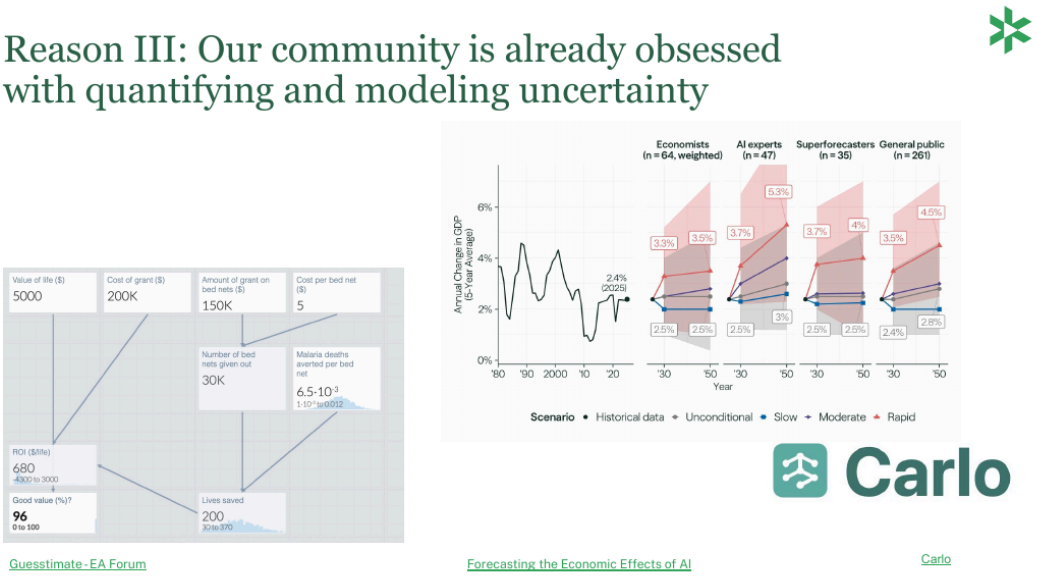

Guesstimate - EA Forum

Forecasting the Economic Effects of AI

Carlo

The third and last piece is that our community is already obsessed with quantifying and modelling uncertainty. We're the only community that has created two different tools for Monte Carlo modelling.[1] This is one thing that we really like to do. And obviously we're also massive fans of forecasting.[2] So we're very obsessed with this already.

But — and that's the key claim — we're not really bringing this together in a systematic fashion to work on methodology for high-uncertainty contexts.

---

[1] Guesstimate and Carlo, both shown on the slide.

[2] The chart is from the Forecasting Research Institute's Forecasting the Economic Effects of AI.

Hypothesis

Slide 27 · 13:17

Hypothesis: Integrating and formalizing these ideas is promising for progress.

And this is really the hypothesis that we have: integrating and formalising these ideas is promising for making progress.

And to integrate and formalise, it's great to have a mathematician on board — which is why I'm handing over to our in-house mathematician, who will talk about this more.

Part III · Our approach

Slide 28 · 13:34

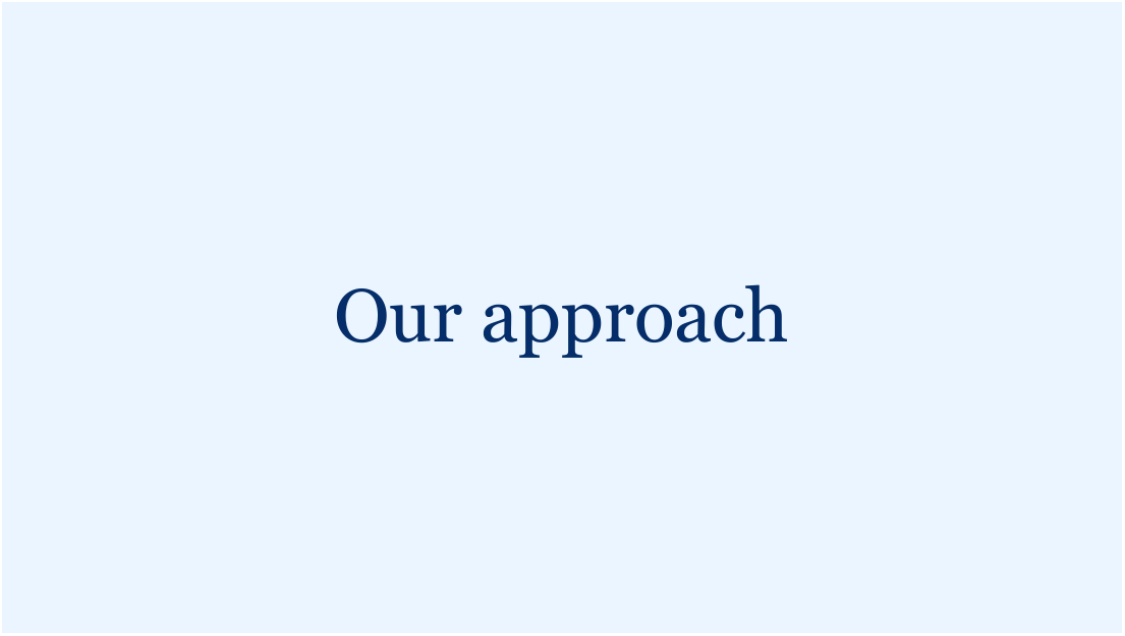

Stan: Thanks, Johannes. So, yes, I'm Stan, the fabled mathematician. I've been with Founders Pledge for almost two years, working on our mathematical modelling and on how we take a quantitative approach to making climate grants.

A general framework

Slide 29 · 13:48

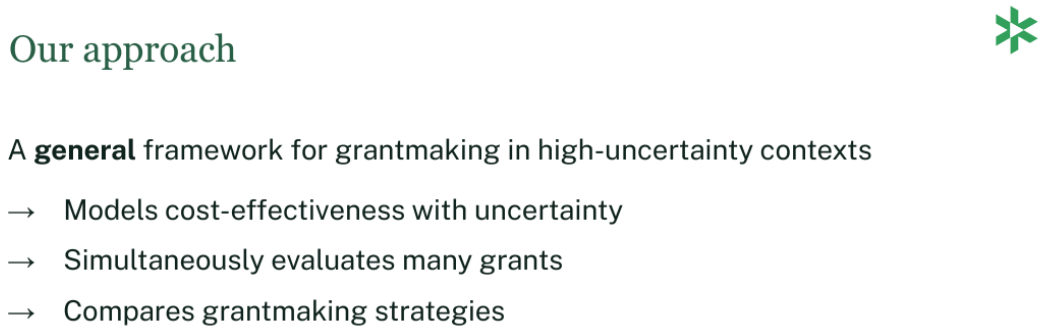

I'm going to show you our approach and the framework that we use to make grants in a high-uncertainty context. And I'm going to focus on a software tool that we've made with help from our engineer, Hubert, which brings together all these features and hopefully provides a scaffolding to use some of this EA wisdom that we've built around grantmaking under uncertainty.

So I'm going to show you how we model cost-effectiveness under uncertainty with Monte Carlo simulation. We're doing this simultaneously — we're doing this analysis for many grants at the same time. And this approach allows us to compare not just grants with each other, but whole grantmaking strategies with each other, which brings a whole host of benefits.

Cost-effectiveness with uncertainty

Slide 30 · 14:38

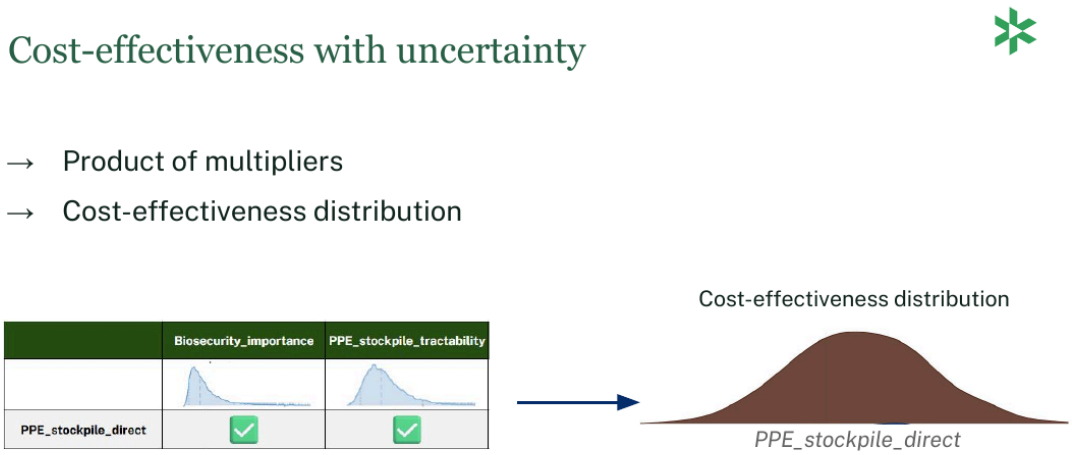

The base unit of our approach, our framework, is that we're doing cost-effectiveness analysis, and we're using this multiplicative impact model which was mentioned earlier. Importantly, we're retaining the uncertainty: we're in a high-uncertainty cause area, we're uncertain about the variables.

And so in this simple biosecurity example, the cost-effectiveness of the grant is the product of two variables. Super, super simple. And that means we get an estimate for the cost-effectiveness of this grant which is itself a random variable, a distribution. So far, so familiar — this is basically Squiggle, Guesstimate.

Simultaneous analysis

Slide 31 · 15:14

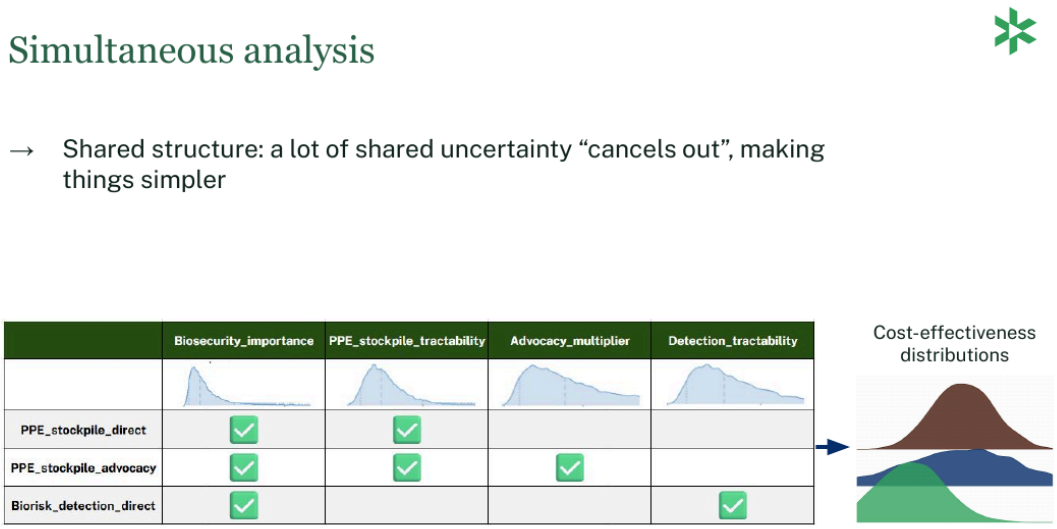

So we’re doing this simultaneously. We’re looking at all of the grants that are under consideration at the same time. I’ve extended the toy example to three grants — those are the three rows. And there are four sources of uncertainty — those are the four columns.

Now, you’ll notice that each grant leverages its own subset of the uncertainties. So these grants have some things in common: they are all leveraging biosecurity importance. Or, saying that in a different way: the importance of biosecurity is important for the cost-effectiveness of all of these grants.

And that’s actually a really good example of simplicity popping out of this. In this particular grantmaking situation, the importance of biosecurity is not decision-relevant, because it applies to everything. So that’s cancelled out, that’s simplified.

And we can also see that you’ve got one of these uncertainties which is leveraged by multiple grants. This is the kind of many-opportunity information which is often cheaper to observe but really valuable to have.

Slide 32 · 16:19

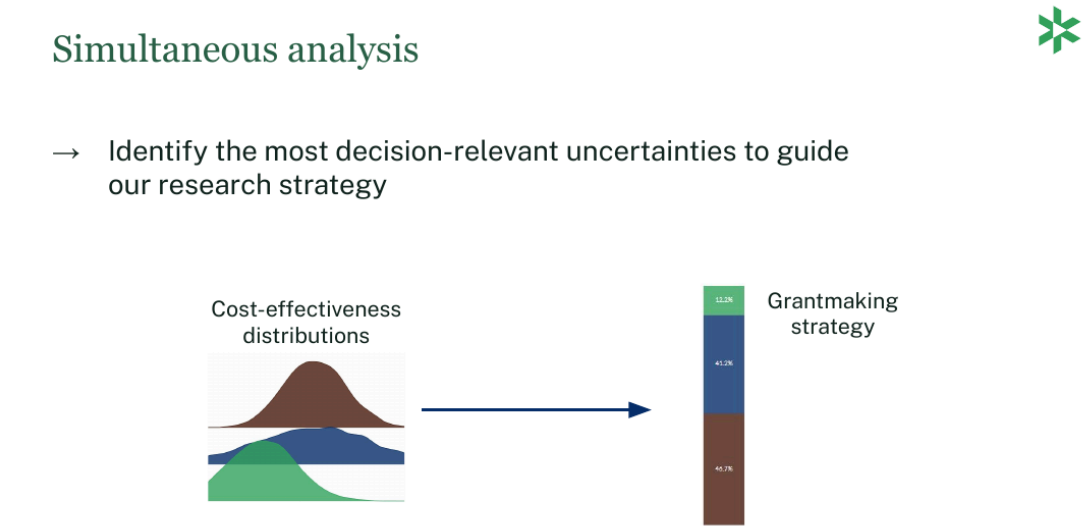

So we're doing the same thing as on the previous slide, but for three grants. We're going to get three cost-effectiveness distributions. In the Monte Carlo simulation we repeat the experiment many times with the random variables, and we get different values each time, to get something that looks like a distribution.

Then we can use a simple algorithm. In each simulation we make a portfolio of all the best grants: we'd fund the best grant, then the second-best grant, then the third, until our budget has been exhausted. And that gives us a portfolio, or a grantmaking strategy, which was the best one in that simulation. This is really powerful, and helps us look at strategy from above, top down, instead of piecemeal by looking at individual grants.

Evaluating grantmaking strategies

Slide 33 · 17:10

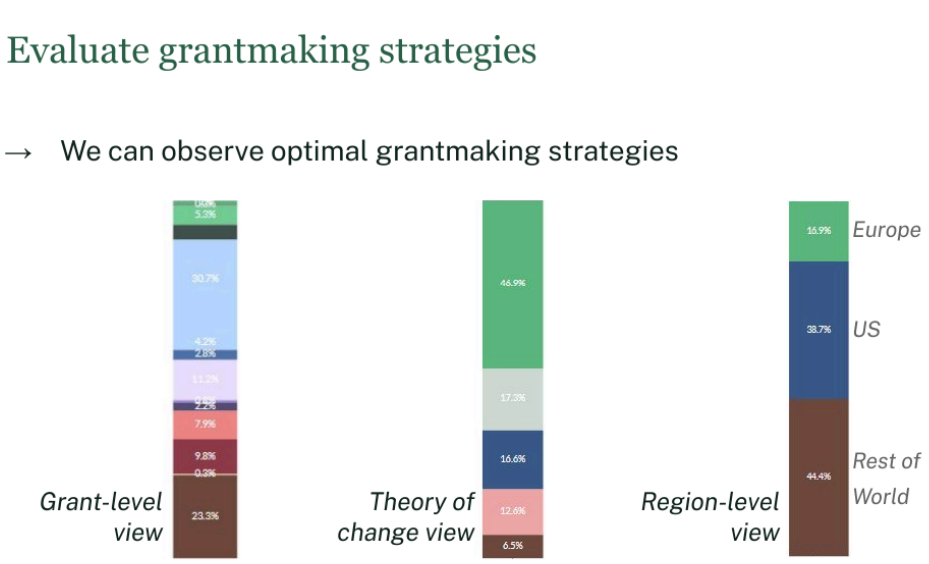

Now, when we do this it's not a toy example with three grants. We might be looking at dozens of potential grants, or over a hundred potential grants we could make. And when you try to visualise these portfolios, they're pretty messy — it's a mixed bag of grants of different sizes.

So we can categorise these to try and understand what they mean and what consequences they have for our grantmaking. We might categorise the portfolios by theory of change and see how money should be distributed that way. We could cut them up by region to see how things should be distributed geographically.

This is really helpful. We've gone all the way bottom-up here: we've put into the model assumptions and variables about individual orgs, or things that are more general, and we've built up a whole grantmaking portfolio to see what's most valuable.

We can also do sensitivity analysis, to understand which of the uncertainties in the model actually change our decisions. Some of them are not decision-relevant, and some of them make massive differences to how we would allocate our money. This is really useful as a research strategy tool: by identifying the most decision-relevant uncertainties, we're essentially spotlighting the things which we should put our research effort on — the uncertainties which, if we reduce them, mean we're going to do a better job.

Updating quickly when the world changes

Slide 34 · 18:25

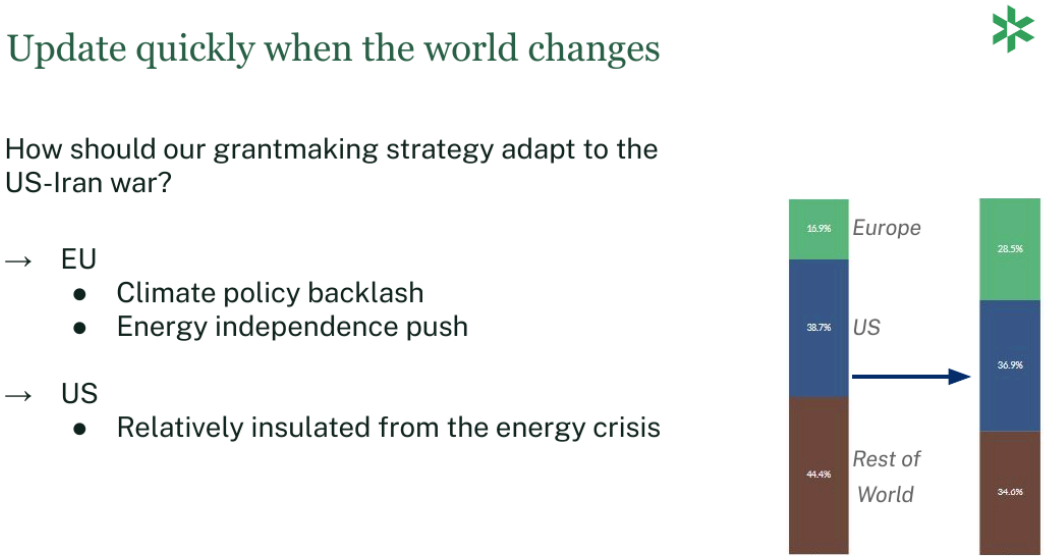

I'm now going to highlight one big advantage of having this ability to look at portfolios, and having everything in a simultaneous analysis.

You may have noticed, or you may have heard, that there's been a war between the United States and Iran. This has big implications for a lot of people, but especially in energy. And so we're faced with this challenge: okay, there's been a big and quite sudden change in the world. How should our grantmaking respond?

To simplify things a lot: in the EU there's perhaps an increased risk of a policy backlash against climate policy, because now we're all obsessed with the cost of energy. But also, perhaps longer term, there's going to be more political capital around energy independence, which might make it easier to support, say, clean energy innovation funding.

We might expect similar dynamics elsewhere, like in the US, but in the example of the US it's relatively independent — it's got its own energy sources, and so those effects are more muted.

We can reflect this change in the world in the variables of the model, in the inputs. And when we do that, we can not only see how the cost-effectiveness of every grant changes, but we can see how those optimal portfolios change, and we can actually put a number, or an estimate, on how much our grantmaking should now shift geographically to other regions. So you can see in this example that the allocation to Europe has increased in light of this update.

The FPRS Portal: an AI interface

Slide 35 · 19:55

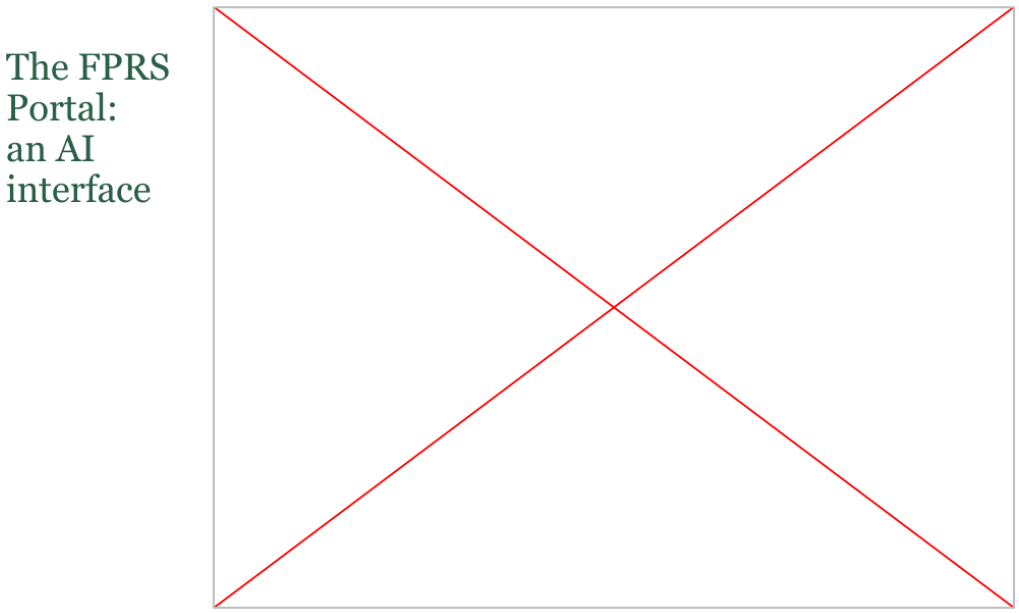

So this is all quite technical-sounding, and believe me, when we actually do our grantmaking work we're looking at much bigger numbers of grants. There are perhaps hundreds of variables. It's all got to go into our software tool, in the spreadsheet. And it's quite confusing — you've got loads of rows, and if you make a small mistake the model maybe won't run. We realised this would limit its potential as a tool that other people might want to use and learn from.

So we have a new, exciting project, which is essentially an AI interface through which you can use the model without touching a spreadsheet, without logging into our software. You're just interacting with an LLM chat interface.

So we're going to risk a live video demo here, which I will narrate for anyone who can't read this at the back. What we're doing is uploading a Google Sheets URL. This will be the spreadsheet which encodes all the assumptions we're making in this model run — I've used the biosecurity example from earlier.

The model has now got that, and I've given it some information, some new evidence which I want it to incorporate in the model. So I'm basically saying: change the variables, add a variable, change the structure of the model, to try and reflect this new update that we have about the world. And so now the portal has made a suggestion about what it can change. I've said "great, go ahead with it, run the model and show me what the new optimal portfolios look like" — and it can also go ahead and do that. It's made the changes to the model internally, it's run the process, and we can see this visualisation shows us some of the optimal grant portfolios. The top row shows giving $2 million to the second grant and $1 million to

the third. And we've got different ways of evaluating what the best portfolios are, depending on how we think about expected value.

But the main illustration here is that what we're working on is something which is going to radically lower the bar, technically, for using a high-uncertainty framework for grantmaking. So the grantmakers just focus on what they do best, which is their expertise in their cause area. They've got a lot of hunches, they've got a lot of intuitions about grants and about the structure of impact, and we're hopefully going to make it easier for them to convert those intuitions into cold, hard numbers in our modelling framework.

It's also possible that we can grow the knowledge base of this LLM over time as we get more experience, and so its performance should gradually improve.

Conclusion

Slide 36 · 22:14

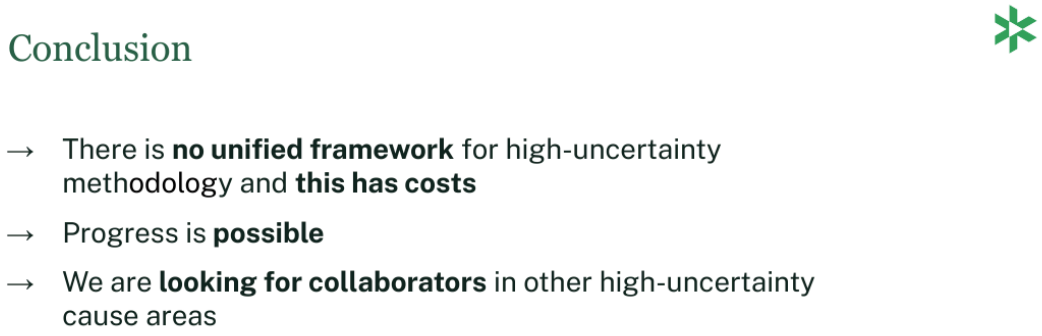

So, thanks for sticking with us this long. We've covered a lot of ground here.

Johannes has shown you that although we do have an established methodology in low-uncertainty areas — which is basically just direct-delivery GHD — there's nothing equivalent that we use in the high-uncertainty areas, which is everything else. And this has some pretty bad costs for all of us. Between cause areas, we have very little visibility of how other people are doing things. It's difficult for us to assess and understand their work; when they make improvements and when they learn something about methodology, it's difficult for us to adopt them. Within cause areas, we have similar problems: not being able to directly compare options which are structurally different, and not really having a learning process where we're all improving because we're using similar structures.

But Johannes has also pointed out that there's a lot of wisdom already out there. The seeds of progress are available. We're obsessed as a movement with modelling uncertainty, and we have some really cool tools for doing that. We understand that impact is often multiplicative, and we also understand that there are different types of evidence — some of it is quite easy to increase and quite valuable to use.

What we're trying to do is develop a framework which is on those rails, using that wisdom about modelling with high uncertainty, but giving the modeller freedom to actually communicate and express the situation that they're in.

So, importantly, we're here to collaborate with other people. If you're in a situation like us — you're in a high-uncertainty cause area and you're interested in making good grantmaking decisions — we really would love to collaborate with you. We don't think we've solved the problem, made the solution, but we want to work towards a shared methodology framework which could be common across different high-uncertainty areas.

Learn more

Slide 37 · 24:04

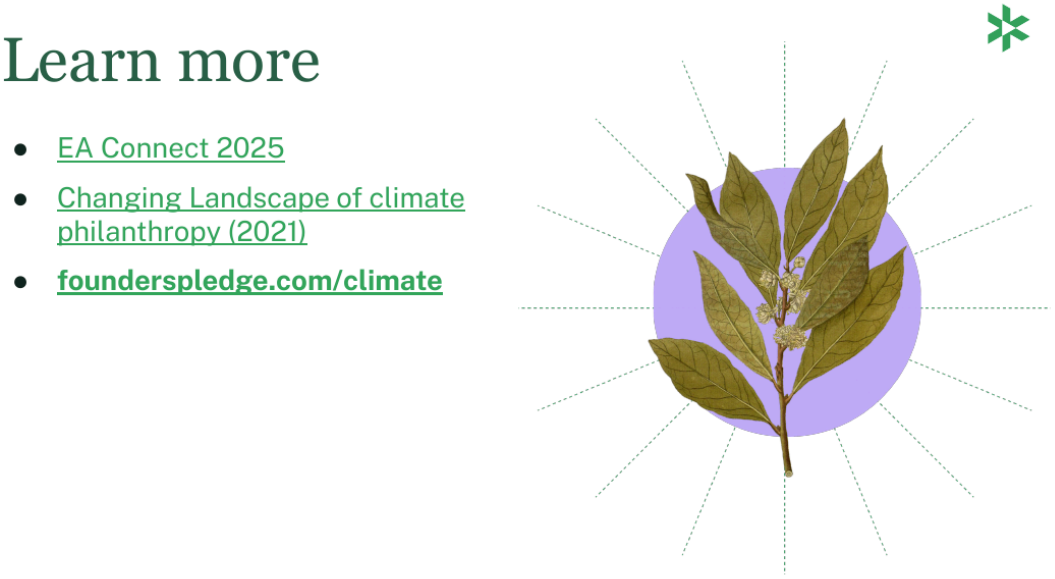

I encourage everyone here, whether you're a grantmaker or not, to read more about our work. We've spoken at previous EA conferences — you can find videos.[3] We've

[3] EA Connect 2025 (YouTube).

got very high-quality reports, as I said myself; they're on the internet.[4] And you can keep in touch with our recent grantmaking and our updates online.[5]

We have pretty much exhausted our time, so we're not going to be doing Q&A on stage, but please do come and join us at our office hours. They're going to be on this floor, in the room I can't remember the name of — in the room called Docklands.

So, thank you all for sticking through perhaps the nerdiest talk of this EAG. We have been Stan Pinsent and Johannes Ackva of Founders Pledge. Thank you.

[4] The Changing Landscape of Climate Philanthropy (2021).

[5] founderspledge.com/climate.