[MUSIC PLAYING]

MICHAEL A. LIVERMORE: Welcome to the *Free Range* podcast. I'm your host Mike Livermore. This episode is sponsored by the Program in Law, Communities, and the Environment at the University of Virginia School of Law.

With me today is Lisa Robinson, a senior research scientist and the deputy director at the Center for Health Decision Science at the Harvard School of Public Health. Lisa is a leading expert in the use of cost-benefit analysis to evaluate public policy, which is what we'll be talking about today. Hi, Lisa. Thanks for joining me today.

LISA A. ROBINSON: Hey, Mike.

MICHAEL A. LIVERMORE: So you've spent a good portion of your career working on and thinking about formal cost-benefit analysis using formal cost-benefit analysis to evaluate lots of different kinds of public policies. Just to give a sense, what do we mean by public policy here? Like things like whether we should require a new safety feature on automobiles or whether we should cut certain kinds of air pollution or whether we want to ban or control a food additive that has some negative health risk, right? That's the kind of stuff we might do a cost-benefit analysis on.

And as you well know, as an expert in this area, not everyone likes this kind of cost-benefit analysis. It's very well entrenched in our public policy system, but there have been critics for as long as it's been around. So I maybe thought we would just start with your kind of 92nd elevator pitch for why we should be doing this.

Say, you had a senator in the senate office building. And in the elevator, you had 90 seconds. And she asked you knowing who you are, why are we doing this cost-benefit stuff? Why is this valuable?

LISA A. ROBINSON: I have to tell you, I don't need to wait for some senator in the elevator. I get these sorts of questions all the time.

[CHUCKLING]

So in the scholarly literature, a lot of times you see-- almost all the time you see people describe benefit-cost analysis as a way to see what policy is most economically efficient in the sense that the benefits exceed the cost by a significant amount. I don't think that's why we do benefit-cost, however. Usually, there is enough that we have difficulty quantifying, and there's enough uncertainty in things that we can't quantify that it doesn't give us a definitive answer.

What it does do is give us a well-established systematic framework for investigating policy impacts. I think this is just forcing us to poke at the problem. I've done a lot of analyses over the years that were actually never completed because somewhere along the way we discovered something that nobody knew that made the policy just not make sense, not feasible.

For example, I worked on one regulation where in the very early stages, that would have released some materials that were potentially hazardous. And we found very early in the analysis that the cost of testing to see whether or not the materials were hazardous were so high that it didn't make sense to even get started on doing some sort of regulation or policy. And I found that over and over again.

For people who are really deeply into these frameworks, there's all kinds of controversies that I'd be happy to talk about, but I think for the general public, there's really just two things. And one is that we made the unfortunate mistake many, many decades ago, I think, of labeling the way that we value mortality risk reductions as the value [? for ?] statistical life. If I could get in my little time machine, go back to whenever it was that somebody first coined that term, I would tell them to please, please, please come up with something else.

People think that we are placing a value on their life, and they find that very offensive. But that's not really what we're doing. What we're doing is asking the question, how much would you be willing to pay for a very small change in your risk of dying this year? Because most policies that we're evaluating only change your risk of dying by a very small amount. But it's very hard to get that across.

The second is I think a little bit more subtle, and I think we as analysts, need to do a better job of communicating it. But benefit-cost analysis is intentionally not paternalistic. What we're trying to get at is the preferences of people for spending on different goods or services that will benefit them somehow. So we're trying to compare their own evaluation of the cost to their own evaluation of the benefits that they accrue.

There are many other forms of economic evaluation, including the type of work that's done for health and medicine where they're using [INAUDIBLE] that is more paternalistic because there's this idea that health is what economists often call a merit good. So you should get lots and lots of health improvements. Regardless of your preferences, it doesn't matter if you'd rather spend your money on other things.

And benefit-cost analysis at least is-- conventionally applied isn't paternalistic. It attempts to reflect the preferences of the people who are affected.

MICHAEL A. LIVERMORE: Good. OK, great. Just to recapsulate that a little bit-- so on the benefits, so to speak, of cost-benefit analysis or the value that you think is most worth emphasizing is this kind of structured analysis, right?

The idea here isn't to come up with some magical number that's going to tell the decision-maker whether a policy is a good idea or not. But it does force analysts and ultimately policymakers to, in a systematic fashion, go through and try to understand the consequences of their actions. And so that's kind of on the pro side, and then you mentioned a couple of controversies.

So one is this notion of, quote unquote, "placing a value on life" and then questions around kind of paternalism versus nonpaternalism and how that might generate controversy. So I think we can return to all these things but maybe just focusing on the mortality risk one. It's been discussed a lot. It's something that you've done a lot of work on the value of statistical life, continuing to use that phrase. [CHUCKLES]

It's just so current. It'd be so hard to change the words there. But maybe we could just unpack a little bit. As you said, what that is often about is a small change in your mortality in a given year. This is kind of the fundamental unit that we're often dealing with. So maybe let's just again-- just to unpack all of this, what is a context where-- like a concrete context either that you've worked on?

Or we can even talk hypothetically where the value of statistical life or the value of mortality risk reduction is kind of an important part of a benefit-cost analysis. Like, when does this arise?

LISA A. ROBINSON: That's a more complicated question than it might seem to be on the surface, but let me start with the simple answer and then give you the more complicated one. Most policies have some impact on health. It's been interesting to see-- certainly, climate change, air pollution, transportation, food safety, Homeland Security, all of those policies have some effect on mortality risk.

And, in fact, interventions that we often think of as being entirely outside of the health sector like education often have fairly significant effects on mortality. But I think that to some extent that's a little bit misleading. Because if you think about it, health includes how long you survive. It also includes the likelihood that you're going to get sick or you're going to get injured.

And we do not have enough research on nonfatal health effects, injuries, and illnesses to have a very good idea of their value. So I think in a lot of the work that we do, mortality risk reductions dominate the benefits estimate partially because we care a lot about living longer but partially because we're not very good at assessing the value of other effects. You might have a health effect that has a relatively small value compared to the value of saving a life.

I shouldn't use that language, should I? Even I get tripped up. The value of reducing mortality risks. And lots and lots and lots of cases. I mean, all I need to do is think about COVID. I've forgotten the numbers now, but there are so many people who've had nonfatal cases that even if the value is a lot lower than the value of avoiding the risk of dying, it could add up to something very large.

MICHAEL A. LIVERMORE: Yeah, that's an interesting point right there. A lot of the rules that we're talking about, a lot of the policies have lots of different kinds of effects, right? They can have effects on health. Your quality of life is what we're kind of talking about. And then there's effects on how long you're going to live.

Basically, those are the two general things that we're talking about. And then there could be other kinds of effects, what your consumption is or whatever else. We've been kind of ragging on the value of statistical life and the notion of saving life, but I think it is worth maybe just offering, if not a defense, a little kind of explanation of how we get to that kind of language.

We talk about saving a life or rule saving a life. And certainly, agencies talk that way, right? They don't say-- at least sometimes, they'll say, yeah, we're going to reduce mortality risk. But over a given population, we will be saving lives. That's often the idea.

So if you have a population of 200 million people and you eliminate a 1 in 100,000 risk for everybody in that-- what'd I say?-- 200 million population, you can be expected to save 2,000 lives, right? That's the idea, right? And so that's-- so there's a peculiar kind of feature of policymaking in a large country like the US, which is that small things add up, right?

Small reductions in risks of different sorts-- we're talking about life, risks to life and limb. Small reductions add up. And so the risk that any individual child will die by drowning is quite low, but over every single year, I believe hundreds, if not a couple thousand, kids die in drowning accidents. And the risk that any individual will slip in the tub and have a fatality is low, but it happens every year in the US.

So yeah. So there's this-- there's almost like a peculiarity. I think a moral peculiarity, something that we're not used to in our evolutionary history or even our cultural history. We think in terms of, oh, a small risk to me, or that very likely will not get realized. But a small risk to many people over a big population, essentially, it's inevitable that it's going to be realized.

LISA A. ROBINSON: Well, it is going to add up, but I think a lot of [? us ?] need to be very careful about our language. And I've played around with different ways of explaining this in both what I've written and what I've been teaching over the years. And it's challenging. But I think the key in what you said is we expect deaths to decrease by some amount.

We don't know who would die if we hadn't implemented the policy. We don't whose deaths are being averted. We also don't-- all we're doing is delaying those deaths. We're all going to die sometime perhaps, unfortunately.

MICHAEL A. LIVERMORE: For better or worse. I don't if the alternative of living forever would be all that wonderful either so Yes something we have to accept we're all going to die

LISA A. ROBINSON: Yes. Yeah. Maybe leave that one off with the living forever one off the table, at least for this conversation, because-- but let's unpack the term for a second because I think it relates exactly to what you were just saying.

It's got two pieces. Well, three pieces. Maybe value is the dollar value because we're using money as our medium of exchange because-- we could use apples or oranges or pick your favorite thing, but money's the easiest thing to use because we're used to spending money on different things in the marketplace.

The second is statistical. And that to me is really important because what that's talking about is the probability of dying. This policy, for example, is decreasing each person in the population's likelihood of dying by 1 in 10,000 in a defined year. And once we sum that up across everybody in that population, we get to a number of expected deaths. So it's a probability of dying. It's not death with certainty.

And then the final piece, of course, is just life. And the idea there is that goes back to what I said earlier that we're delaying deaths. And sometimes we're only delaying them for a year. It might not be that long. We're not saving somebody's life with certainty.

If we have a risk-reducing policy that's implemented this year, it doesn't mean that you're going to live to whatever average life expectancy is for somebody your age. It just means that you're less likely to die in this particular year.

MICHAEL A. LIVERMORE: Right. Yeah. So this is all true, right? So that's what we're talking about.

One is it's statistical, right? That's an important characteristic of it. And then the other is we're only ever extending anybody's life. That's the only thing that we can do.

And there's interesting counting questions that come up like I think-- let's do a little philosophical question. I'm curious what you think the answer to this is. Let's take away the statistical thing for a second just so that we can make it a bit more clear. And let's say I'm a crossing guard, and there's someone-- just person A who is just terrible at crossing the street.

They just don't look both ways before they cross the street. And so once a week, the crossing guard has to grab hold of person A and block them from walking right in front of a bus. And it just happens with-- it's like clockwork. Now, what's gone on every time that the crosswalk guard has or the crossing guard has saved this person's life clearly at least.

Let's just say that every time-- but for the crossing guard, the person A would have been flattened by this bus, OK? But the crossing guard has saved this person's life let's say 50 times a year. Should that count as one life saved or 50 lifesaving events? And I think that kind of is a peculiar question, but it's something that actually kind of comes up in regulatory cost-benefit analysis because it might be that we have all these different interventions.

But for any one of them, people might be dying, but it is the same population that we're saving over and over again. So yeah, I'm just curious what you think. Like, what is the right way to value something like that? Should we just have the value is this one person's life, or is it not whatever value we place on that? Or is it 50 lifesaving events?

LISA A. ROBINSON: So I have to say, first of all, I'm not that fond of the analogy. I hope that--

MICHAEL A. LIVERMORE: Just because of the statist-- you don't like the statistical part of it?

LISA A. ROBINSON: Well, no because these methods are not designed to deal with individuals. These methods--

MICHAEL A. LIVERMORE: OK, OK, I get it. So you're resisting the fact that I'm separating out the statistical from the repeat player kind of element.

LISA A. ROBINSON: Well, no. It's because you're talking about one person.

MICHAEL A. LIVERMORE: That's what I mean. That's what I'm talking about.

LISA A. ROBINSON: Yeah, I mean, I think the way that-- we talk about this in the literature is identified versus--

MICHAEL A. LIVERMORE: OK, OK, fine. You know what? I was trying to simplify, but I will deal with the population, if that makes you feel better.

LISA A. ROBINSON: Yeah. So we have a population of kids and a population of crossing guards and--

MICHAEL A. LIVERMORE: Something like that, but it's the same deal. Like, you've got the same kids. There's a group of them that will cross bad. And then one of them would get hit every time. And over a certain number of times, the same kid is going to end up getting saved several times, or the same group of kids is getting saved many times or whatever.

LISA A. ROBINSON: Yeah, I think this is a-- so I think the question you're asking is-- so let's say today we implement this new crossing guard program that keeps those bad streak crossers from getting hurt. And tomorrow we implement a safe driving program that also is going to keep kids who are crossing the crosswalk from getting hurt. So there's some aggregated effect across those two policies.

MICHAEL A. LIVERMORE: So say there's only 100 kids that you could possibly save, but you're going to save each of them 10 times.

LISA A. ROBINSON: Well, no, you're going to decrease the risk 10 times.

MICHAEL A. LIVERMORE: Well, there's a crossing guard that's going to actually reach out and stop a kid. Like, that kid will not have gotten hit. Right? The same way that like with a particulate matter reduction. I mean, I get the-- and actually, I think I'm going to resist the resistance to individuals at some point, some level because if like--

Ex ante, we cannot say. And we might not even be able to say ex post that like a air quality rule that reduced particulate matter exposure saved any individual's life. Like, we can't actually attribute at the level of an individual that something like this happened, that anybody was even benefited, right?

But if we think that reducing particulate matter pollution or if we think that particulate matter pollution causes people to die, that exposure to particulate matter pollution causes heart attacks, and we cut particulate matter exposure-- and we could even say let's even imagine, which would be very nice. Like, we have a nice clean experiment where we can say, OK, in this population here, particulate matter exposure went down by X amount.

And now instead of on average what we were seeing is a thousand heart attacks every year, which, of course, with some noise, now we're seeing 900 heart attacks every year in this population or fatal heart attacks. So we can say there's a hundred people. And we can't say what 100 people would have died, but we can say that there are people walking around, right?

And there really are. There's someone walking around in a counterfactual, where we had the same level of pollution, where we hadn't cut the pollution. And that counterfactual, there's someone walking around.

Sorry. In the real world, there's someone walking around who in the counterfactual would not be walking around. I think we have to think that. If we think that these policies are actually having effects in the world, we have to think that there's a counterfactual in which someone who is now alive is dead.

LISA A. ROBINSON: Yeah. The problem with a lot of this stuff is that simplifying. It takes away some of the clarity about what the concepts are because first of all, I think-- going back to the crossing guard for a moment, you've taken out the idea of risk. So let's say the crossing guard-- there's a 90% chance each time the crossing guard grabs that kid, right?

MICHAEL A. LIVERMORE: Yeah, there you go, something like that. Add that to it.

LISA A. ROBINSON: Or maybe there's only a 20% [INAUDIBLE].

MICHAEL A. LIVERMORE: Whatever it is. Right.

LISA A. ROBINSON: And there's a 90% or a 20% chance that that kid is actually going to do something dangerous. In the real world, these things happen in a very complicated way. The other thing is that the one key piece of the VSL definition that a lot of people miss is that it's in a defined time period.

So you would-- for your crossing guard example, we would want to talk about the chance that that crossing guard or that maybe this policy, this crossing guard policy, reduces the risk of X number of deaths within a month or within a year. And I think that that those two things take care of some of your concern about overlap or double counting.

MICHAEL A. LIVERMORE: I'm not even sure I'm concerned about it yet. [CHUCKLES] I was cur-- I'm not raising a concern. I was actually curious what-- I'm not sure whether that's double counting or not. I was curious whether you thought it was. I think it's an interesting question.

LISA A. ROBINSON: Yeah, I think it's a question of being very-- I mean, this is, I think, one of the benefits of doing benefit-cost analysis is it forces you to be clear about things. And for your examples, both the air pollution one and the crossing guard one, we need to be-- in order to do a good benefit-cost analysis, we need to be very precise about what we think would have happened in the absence of the policy, the baseline.

We need to be very clear about what we think would happen. And I'm using the word what we think intentionally because all of these are expectations. They're not things we know with certainty would happen with the policy. And we need to be very clear about the time frame and the population over which we're estimating those things.

One issue that does come up, especially with large policies that affect lots of people-- so this idea that maybe policy A is affecting these hundred people in the population, and policy B is affecting a different hundred people in the population, assuming you have a large population, is whether we need to revise the baseline. So if my risk of dying-- I don't even want to think about this since I'm getting older, but if my risk of dying is 2 in 1,000 this year, we can't keep comparing each policy to a 2 [? and a ?] 1,000 baseline.

We need to think about how the incremental effect of preceding policies have changed my risk. Maybe my risk is now 1 instead of 2 because of everything that's come before it. But I think that's often a-- is that not what you're getting at, or is that what you're getting--

MICHAEL A. LIVERMORE: Well, it's funny because [CHUCKLES] we both know a lot about cost-benefit analysis. And I think a listener who's unfamiliar with some of these ideas would be shocked that we would make the claim that these techniques have a way of clarifying the questions that policy analysts are trying to figure out because it is-- I mean, it's actually quite technical. It's quite difficult. Requires very careful thought.

And so maybe we could kind of change gears a little bit. And I would be curious your thoughts about the kind of this meta question about the degree to which cost-benefit analysis can be quite difficult for a nonexpert to comprehend what's going on because as soon as-- we start talking about what is really a fairly straightforward policy that's crossing guard policy, where the mortality risks are kind of straightforward, like the causal pathways at least are straightforward and so on.

And it actually quite quickly turns into a complex question involving what are the suite of policies that we're talking about, what are the populations, what is the time frame, what are the base rate risks that we're talking about. And that's an imp-- those are all-- those are the right questions to ask. Absolutely.

One has to ask those questions in any kind of [INAUDIBLE] even of a simple pol-- almost like the simplest policy that you can think of still requires kind of a nuanced really understanding of these various features of the policy problem. And I think that can be part of what folks resist-- leads folks to resist cost-benefit analysis because then they start to say, well, this all looks like gobbledygook. This just looks like jargon. And I can't understand what's going on here. And it seems to--

I think some of the critics of cost-benefit analysis will say is it actually obscures the stakes rather than clarifies them. So what's your response to that kind of take? Because I actually think that's fairly common is that people start to hear things like baseline and population risk and et cetera, et cetera. And they just think, well, you're overcomplexifying something that we can think about in more straightforward ways.

LISA A. ROBINSON: Well, I think-- so let me-- I feel like there's several different issues there. And one goes back to what we were talking about at the very beginning, which is that my experience has been that if we don't do some sort of analysis, people always have an opinion or almost always have an opinion, but it's unexamined. I'll never forget working on an environmental rule in the US where somebody asked me, aren't there benefits for women? And--

MICHAEL A. LIVERMORE: To this environmental rule?

LISA A. ROBINSON: In the United States. Or I think it was actually more a statement that we'd left out the benefits to women. And if I wasn't-- if we weren't trying to do this benefit-cost analysis, I don't think that person would have ever even asked the question. It would have just been implicit in their thinking.

And after talking to this person for a while, I realized that what they were thinking about was actually low- and middle-income countries, where women walk to get water. So having a cleaner water source nearby has labor savings for the women. But we weren't talking about doing a policy in a low- and middle income country.

And the person I was talking to hadn't really grappled with that in the US. We just go-- well, most of us are fortunate enough to be able to just go and turn on our tap. So I think it's-- one of my colleagues likes to talk about as an aid to thought, but I think it's a way of just getting people to talk about what they think the policy is all about and ask questions and try and reach a deeper understanding of it.

In terms of complexity, I think that's a real problem for us as a profession. What happens to me and I think happens to all of us is we get deeply into the details of our analysis. We write up the analysis the way we think about it and the way we did it, and nobody except us understands it. And I think we need to get much better at communicating.

And I think the other problem is that we ourselves need to be better trained. And some of that's-- there's a lot of issues in terms of technical training that would be really helpful, but I think learning how to communicate what we're doing is incredibly important. It's striking to me that I've had the honor and privilege to teach many, many, many incredibly smart students.

And it's not unusual for them to tell me that I'm over their head, even though I think I'm giving a simple explanation and I've had to think hard about that over the years and try and communicate better in a more basic way. But it's a huge challenge for our profession and one that I think, as you said, it gets us in big trouble.

MICHAEL A. LIVERMORE: Yeah. I mean, I think part of it is-- just as we're thinking about this, I think part of it is that there is-- it's with any discipline, right? You start with some basic premises, and then you kind of start to build from there. And then you need to build from there and build from there.

And what you end up with can be a fairly complex edifice that's built out of fairly-- that could be built out of simple parts. And that's just the way it is. And it sometimes-- to a certain extent, it might just be that it's irreducible. Like, in order to understand complex biology, you need to know the building blocks, right? And I think part of what is tricky about cost-benefit analysis is that we're talking about public policy, right?

So people don't go into a lecture on recent findings in particle physics, or they don't go open a scientific journal like *Nature and Science* and open to the genetics section and expect to kind of understand what's going on, right? Whereas when we're talking about, should we hire more crossing guards? or even, should we require more stringent air quality standards? or should we require rear-facing camera in new automobiles? these are questions that in a democratic society we think people have opinions on, people feel that they ought to be able to have opinions on.

And maybe under some theories of democracy, what ultimately we do as a government should reflect what people think about [INAUDIBLE] whether these are good ideas or not. So there's this really big tension because-- you know what? Potentially, I'd be curious what you think about this.

On the one hand, the fact is we're dealing with public policy questions in a democracy, where people have often strongly held views. On the other hand, these questions can be approached with an edifice. Basically, that is this cost-benefit analysis that we're talking about. That is a little bit more like a sophisticated scientific enterprise.

And I think we need to be careful about that to not-- there's a lot of values questions. There's a scientific element to it too. But it's also a sophisticated moral language. I think we can say that it's a highly articulated language about trade-offs and how you make those trade-offs and so on.

And so whatever it is, it's a highly developed discourse. Let's call it that. I don't want to imply that it's a purely scientific discourse and make too many analogies to fields like biology, but it's a highly articulated discourse. It's a highly developed discourse.

That it's going to be difficult for most people to access without a pretty substantial amount of investment of time to understand it. So we have this dilemma. It's almost like a dilemma. On the one hand, the highly articulated discourse, that can be difficult for the uninitiated. On the other hand, we're talking about public policy in a democracy, where folks are going to have opinions.

Yeah. So what do we do about that? What do you-- you've confronted this issue many times over the years. How do you kind of think about this problem of negotiating these two realities?

LISA A. ROBINSON: [INAUDIBLE] several points starting with-- well, I'll get to your core question in a minute, but let me step back through a few other things first. So any form of research-- I don't care whether it's particle physics, chemistry, or benefit-cost analysis-- has a lot of uncertainties. All you need to do is look at the news reports about scientific findings, and you'll see they get reversed. They get amended or changed in different ways.

Uncertainty is a fact of life, but it's also something that people don't like very much. They would much rather have somebody say to us, if you don't want to get COVID and do this, you definitely won't get COVID.

MICHAEL A. LIVERMORE: That would be nice.

LISA A. ROBINSON: You don't want to hear all these messages about, well, you could decrease your risk if you do this and that, aside from the question of whether-- I probably shouldn't be using COVID as an example.

MICHAEL A. LIVERMORE: So complicated. Such a complicated case but whatever. People don't like uncertainty. That's clear, right?

LISA A. ROBINSON: Yeah, people don't like uncertainty, yet there's uncertainty everywhere. And that's a challenge for those of us who recognize that. The second is we always talk about benefit-cost analysis as a way of informing decisions. I don't think there's anybody who's a practitioner who thinks that benefit-cost analysis should be used to make the decision because there's a lot of normative value judgments that underlie the framework.

It shouldn't be the only source of information because of those value judgments, but there's also a number of things that it doesn't address very well. It doesn't address legal and political issues. It doesn't do a very good job of looking at the distribution of effects across people who are advantaged and disadvantaged.

So I think that one thing that we need to be making more clearer to a general audience is, hey, we're providing you with some information, but there's lots of other things you should be thinking about, including your own moral beliefs and other things. I also think, though, that we really need to think carefully about audience because-- I know when I'm talking to a scholarly audience there's a lot of misperceptions about benefit-cost analysis.

But it's a very different type of discourse than when I'm talking to the general public. One of the most effective things I've seen as I was working on a radiation protection standard years and years ago-- and because you can't see or feel radiation [? unless ?] it gets really, really high, we are having real trouble communicating the amount of risk reduction that you would get from this regulation. And one of the other people in the project came up with this idea of having an illustration of how many hours you'd have to spend.

It wasn't even hours. I think minutes. You have to spend in an airplane flying at 30,000 feet to get the same dose of radiation. And that's what resonated.

And I don't think-- many people don't really care about the things that we debate about academically. They just want to know how much of-- how bad is this going to be, how good is this going to be. And we need to figure out ways of explaining that that airplane exhibit could be reduced to a bunch of numbers and formula.

But the airplane worked a lot better in helping people understand what the benefits of this particular policy were. Does that make sense? Because I think that you and I live in an academic world where there's lots of detailed technical debates about these things. But when I talk to people who don't do this type of work, what they really want me to tell them is, is this a good thing? Is it a bad thing?

MICHAEL A. LIVERMORE: Well, this is the tricky problem. This is exactly the tricky problem because that's exactly what folks want to hear. And in some sense, I think this was kind of what one of your early points was, which is cost-benefit analysis often-- or analysts often don't want to provide that answer in part because they still don't know, right?

There's uncertainty, right? So there's uncertainty about what regulatory consequences are going to be on the ground. What is the real risk reduction? There's controversy about how to value those risk reductions. There's potentially other values in play that could conflict, or there's normative choices that go into cost-benefit analysis.

So like kind of an all things considered, is this a good idea, or is this a bad idea? Arguably is something-- now I think that in some times anyway, people say-- I think what you said was that analysts don't want to provide that. They don't see that as what's going on. It's a tool to inform analysis. But if the tool to inform analysis just says, yeah, this is a good idea, or no, this isn't a good idea, then it's kind of-- it is ultimately making quite a forceful recommendation.

LISA A. ROBINSON: Well, I think there's actually two points there. The first is that we all-- those of us who are in-- the people who do this sort of work are people who like analysis. They like having to sort through all these thorny things. And not everybody's like us, but we get too much caught up in our own heads.

So I think a lot of the oh, I can't tell you anything hand-waving is because we are so caught up in all the details of what we just did that we can't step back and say, oh, wait a minute. I can tell you that maybe my message is there's so much uncertainty here that I don't really know if the benefits of this policy are going to exceed the cost.

So you Mr. And Mrs. Decision-maker, either need to figure out some other basis for the decision, or maybe you decide not to do anything at all. Or I think most things that I work on, we can pretty clearly say that the benefits can exceed the costs or that the costs exceed the benefits. We just can't tell you sort of right out to the decimal point with any certainty. So I think it's a question of needing to get out of our own heads and think about in simple terms what's the bottom line of what we've done.

But also, I think when I say informed decisions, what I mean by that is-- I could say to you, Michael, this crossing guard policy that you're thinking about, my benefit-cost analysis says that the net benefits will be positive and could be significant. But the thing that I haven't looked at carefully or that I don't know that you, Michael, should take into account is the fact that most of these kids are getting hurt are very low-income kids. So we should add that into your decision-making along with the fact that-- I don't know-- maybe it's illegal to grab kids in the crosswalk.

So I think that we need to tell you-- tell a story about what's in the analysis and what the analysis tells you, but we also need to be able to tell a story about the other things that you might want to think about.

MICHAEL A. LIVERMORE: Right. Well, the legal part, that's-- that agencies always have to deal with their constraints and what they can and can't do, right? So we could say, well, this policy would be great if you could do it, but you don't have the power to do it, right?

In some sense, that's fairly straightforward. The other issue maybe we could get into a little bit because this is a more of a-- I don't know if to call it a cutting-edge issue. It's an issue that has been long discussed in policy circles, right? Should we be-- how should we account in for the distribution of costs and benefits?

And so in the example you gave, that was the easy case. That would be like, look, there's more costs than benefits. And the benefits will mostly go to folks who are less well-off. Well, that's fantastic. The real issue comes up when there's a conflict, right?

So we say, OK, well, we do our cost-benefit analysis, which still at this point, even the simple case, there's something a little mysterious about it for the average person, like, we look at the populations. We look at years. We look at these risks. [VOCALIZING] We have this value of statistical life that we're using.

We could explain where that comes from, but it's complicated. But anyway, we do our analyst thing, and we come up with the bottom line which says the benefits are going to be greater than the cost. And we could even say that we looked at a suite of policies, and this is the one that has the highest net benefits will be adding crossing guards. And that's what's going to make most sense as opposed to changing around the routing for the kids or changing the school day or putting in a stoplight or whatever else.

So anyway, we looked at a suite of policies. This is the one that has the highest net benefits. But it's mostly going to help advantaged, highly advantaged kids.

The alternative policy that we looked at actually would have lower net benefits. It would cost a little bit more, but it would also change who's going to be advantaged by the policy. And it's actually the case that the kids who would be helped by the second policy, which let's say is rerouting traffic.

The rerouting traffic policy, again, which is costlier, but it's going to actually help more disadvantaged kids. So that's where you have a conflict between what might be appealing on cost-benefit grounds, the one that has the highest net benefits versus one that might be appealing on distributional grounds, the one that helps the least well-off. So this is a hard kind of problem.

So one possibility is that we just kick that to the decision-maker, and we just say, look, we've done the kind of standard cost-benefit analysis. We've also looked at these distributional questions. And so we can tell you kind of who benefits and who bears different kinds of costs. And then you policymaker, you need to decide.

Another alternative would be to incorporate the distributional analysis somehow into the cost-benefit analysis, which there's different technical ways of trying to do that. I'm just curious what, broadly at that level where you have a conflict between what we might think of as normatively desirable distributional or undesirable distributional effects versus efficiency effects, just the total cost and benefits. How do you think that should--

given the state of knowledge now, what should we be doing? Should we trying to-- should we do-- we just lay it all out for the decision-maker and that person has to do the calculus, or is there something different we should be doing?

LISA A. ROBINSON: So this is an area that I've been working a lot on recently. And I don't think you phrased that as what should we be doing, and I think-- we're still trying to figure that one out. You and I both have a strong background in regulatory analysis, but I think it's important for us both to keep in mind that many benefit-cost analyses are not done for regulations.

A lot of the work that I've been doing lately has to do with social programs or with policies that don't require regulation. They just require somebody to fund the policy or government to decide that they're going to--

MICHAEL A. LIVERMORE: Spending programs, basically.

LISA A. ROBINSON: Yeah. Yeah, well, it could be a foundation, or it could be something that a private firm could do. And I think that that changes the role of-- the legal issue's quite a bit. It's not just that we have requirements for what we need to do for regulatory analysis. You're in a different world.

If it's a foundation, you don't have Congress and the courts breathing down.

MICHAEL A. LIVERMORE: Yeah, do whatever they want as long as it's not illegal, right?

LISA A. ROBINSON: Yeah, almost whatever they want. They still have some constraints. But I think that the-- as you know from having seen this stuff that we've written, people do not do a good job of distributional analysis even if all they're trying to do is describe the distribution.

MICHAEL A. LIVERMORE: Yeah just-- sorry to intrude here, but you've done a lot of work on this very question of just, how is distributional analysis done at the federal level amongst agencies? and just this practical question of, what are we currently doing? And so yeah, as you were saying, some of the findings there is that we're not doing all that good of a job

LISA A. ROBINSON: Yeah. I mean, I think that-- and this isn't just true for the regulatory stuff. It's also true in-- I've been doing a lot of work in global health and development, same thing. Every single guidance document-- there's probably some exception out there that if anybody listens to this, they're going to tell [? me it ?] exist.

But I don't think I've ever read a guidance document that doesn't say, along with the estimate of net benefits, you need to tell us the distribution across the disadvantaged. Almost never see that. I have many theories why not, but I think one piece of it is that it's often fairly easy to think about-- or at least estimate how benefits are distributed.

You know what the population is who gets heart disease. So if you have a policy that reduces heart disease, you can just estimate that it'll be distributed the same way those base cases are, unless there's something about your policy that makes one subgroup affected more than the other. Costs are very difficult because often costs are incurred by a government agency or by private industry.

And you need to think about how is that agency or that industry going to react to [? a ?] position of additional costs. Are they going to absorb it by decreasing their profits. Are they going to lay people off, reduce people's salaries, increase the prices of their product. And we know very little about that.

And if you can't estimate who's going to be affected by the costs and the benefits? It's very hard to do any of these sort of more sophisticated approaches that look at, for example, people's preferences for distribution. As I think you know, there are approaches for weighting cost and benefits by--

We always talk about them by the marginal utility of income, which is a fancy way of saying the fact that-- saying that a poor person-- an additional dollar to a poor person means a lot more than an additional dollar to a millionaire. But there are also variants on that theme like what gets called prioritarianism that try to provide greater weight to people who are disadvantaged across a number of different dimensions. It might be health. It might be something else. But we can't apply.

So all those approaches are being developed. People are experimenting with them. But in order to apply them, we need to get much better at this sort of starting point of describing who bears the benefits-- I mean, who bears the cost, and who receives the benefits? I think there's ways of doing that, but we need to think hard about that.

And we need to start implementing the ways of doing it rather than just waving our hands around and saying, oh, distribution is important. And then not addressing it in our analysis.

But I think those sticky questions, the sticky questions about, OK, you've got something that has net costs [? but ?] [? would ?] [? advantage ?] [? poor ?] would be really good for poor people. I think no matter whether you do some sort of weighting inside the analysis itself or whether you just display the results without doing anything to weight the effects on the advantage versus the disadvantage, it still comes down to normative judgment that needs to be made by somebody and--

MICHAEL A. LIVERMORE: Right. That's interesting, right? Somebody needs to make this normative judgment.

LISA A. ROBINSON: Yeah, and if it's government policy, it's the government agency or maybe the legislature. If it's a foundation-spending program, it's whoever makes their decisions. And I don't know-- I think that's outside the benefit-cost analysis. My job is fortunately maybe not to make these decisions. My job is to make sure that the person who is making this decision has some good information in front of them.

MICHAEL A. LIVERMORE: Right. OK, good. But just to maybe clarify or reiterate a point you were making there, which is just it's actually quite difficult. So just to maybe quickly summarize how this debate has unfolded over the past 20-odd years, cost-benefit analysis tends basically, and it's kind of traditional form, focuses on aggregate costs and benefits.

So if we take our crossing guard example, we're not looking at whether the kids are advantaged or disadvantaged. We're not looking at who pays the cost, whether it's the city or individual drivers, if it's like traffic rerouting is the policy. We're not interested in that kind of question. We're solely focused on just summing-- everything just gets summed up.

And you look at the aggregate number, and you say, OK, what are the benefits and costs look like in total? And a consistent critique that has been offered over the years is well, that is-- there's something wrong with that that we ought to, as a society, be sensitive to who bears the costs and benefits. And in particular, we should care about costs and benefits that are born more by people who are less well-off exactly for the reason that you described at least, Lisa, which is the diminishing marginal utility of consumption, which is just exactly that a cost that's imposed on someone who doesn't have a lot of money.

That's a much bigger hit to their well-being than a cost-- than the same cost that's imposed on someone who has a lot of dough. And so that's a very straightforward normative idea. And cost-benefit analysis again in its most traditional form is not sensitive to that basic normative impulse that many people share.

And so the line then says as you note in most guidance documents in most of the dictates that agencies follow for doing cost-benefit analysis, it will say something like-- and you should do distributional analysis too, right? But the problem is it's really tough. It's really tough. And I think maybe we can just spend a couple of minutes investigating the difficulties there.

And again, you've really got the experience with real-world analyses. And you've gone through, and you've looked at agencies analyses and what they're doing. So an example that I use sometimes is a very classic one. I believe it's in the A-4 Circular actually. The existing government-wide guidelines on cost-benefit analysis is-- say you improve air quality in a low-income neighborhood, then you might say, well, that is a benefit that accrues to low-income people.

Well, not necessarily. If everyone there is a renter and their rents go up as a consequence of where they're living, the property is now more-- it's more valuable. There's more demand for that real estate as a rental unit, so the landlords charge more money. And they end up extracting some of that benefit.

And so actually the benefit goes to wealthy landowners. So you might think in the first instance that the benefit goes to low-income people, but it actually-- there's a kind of a transfer there that happens in the marketplace where the landlords extract some of the benefit.

LISA A. ROBINSON: Well, it's even worse than that. If you clean up the air in that neighborhood, wealthier people are going to move in [INAUDIBLE].

MICHAEL A. LIVERMORE: Exactly. Then there's a gentrification issue, right? So then you've got over time. So there's all these complex dynamics that affects, so we might have a benefit or a cost that in the first instance is the benefit goes to low-income people because they're living where there's the air quality.

And the cost might be imposed on, say, a local polluter. But then the costs can get shifted to consumers. Maybe it affects the taxes that are paid and the like. So what do you think of that?

I mean, someone might say, OK, this is just too complicated. This is impossible. And so we should really focus cost-benefit analysis on aggregate effects just looking at wherever the first instance is and not trying to tease through all these distributional effects.

Or is your sense, given that you've really have gotten into this in a bunch of different contexts that actually it is worth the candle, that we can do this? It might be complicated, but we can do complicated things. And it'll be imperfect and uncertain, but the information that we would get out of that would have at least enough value to justify the analysis.

LISA A. ROBINSON: Mike, you just described why it is that doing this sort of work is so fun. It's funny talking-- I've been working for long enough. A lot of the people I'm surrounded by are retiring because they feel like they've been doing the same thing over and over again for years. But there's no danger of ever doing that if you do benefit-cost analysis because there's so many interesting issues to explore.

But I think that-- this really gets back back to what brings people like me to benefit-cost analysis. I fundamentally believe that benefit-cost analysis is an important contributor to good policy decisions and good policy decisions, policies that will increase the welfare of the people affected. Maybe I can tell you a short story. When I first started working, I worked at a-- maybe I shouldn't tell this story. How many people listen to this podcast?

MICHAEL A. LIVERMORE: Well--

LISA A. ROBINSON: I worked in someplace. No, I worked in a-- let me keep this vague. I worked in a federal agency, where we had the experience over and over again of political appointees making what seemed like arbitrary decisions or legislative folks.

They'd say, Oh, I'm just going to-- I'm going to do X because my buddy thinks that's the right thing to do. Or I'm going to do X because I feel like disagreeing with you. I mean, that literally was what I had somebody tell me once.

No rationale. And this is a-- benefit-cost analysis is a way of pushing back against that of saying, here's some information that I'm hoping you will take into account in your decision-making instead of making these arbitrary decisions. But if my goal in terms of doing this type of work, which I think it's not just my goal. I think it's the goal of most of us who do this sort of analysis is to improve the welfare of the people who are affected.

Distribution is an important piece of that, and I really don't think we can ignore it. So you're right. It's hard. That's what makes it fun and interesting. But it's also important, and I think we need to work [? at ?] [? it. ?]

MICHAEL A. LIVERMORE: Now are there cases that you've come across where you've seen that you think that an agency has done a good job at some kind of distributional analysis?

LISA A. ROBINSON: It's interesting that you ask that because I think that recently, at least in the US, there's been a lot of pressure to improve distributional analysis. And I have not gone through everything that federal agencies are doing to try and see whether some of them have gotten far enough along to deal with the cost side because that's where the challenges are, I think, there's lots of people, lots of agencies and scholars who've been able to think about the distribution of benefits, mortality and morbidity risk reductions and other types of benefits.

The one place where people have done a lot of this is with macroeconomic modeling, where you have really large policies that you can feed into these big computer-- well, general equilibrium models because they allow you to-- in different ways, often model households. They have mechanisms that allow you to estimate the effects on the rich and the poor.

But even there, they're making underneath all the complexity of these models, there are some very strong assumptions about things like whether costs are passed on as changes in prices. And these models tend to be useful only for very large policies because there's enough uncertainty in them that smaller policies sort of get lost in the error term. You don't really end up with a good understanding of the effects.

But I think that's an interesting question. And one of your next shows, you should pull on a bunch of the federal economists who are working on this sort of stuff and see if they've-- they've made some progress on this because I really don't know and the cost side because I really don't know.

MICHAEL A. LIVERMORE: Yeah, it's tough. I mean, it's a tough issue. I think that maybe the kind of the final just note on all of this is it's easy to talk in very general terms about, well, we should look at cost and benefits. We should kind of examine the pros and cons of our policies.

We should do things like account for the distribution of those costs and benefits. And I think if there's a lesson for folks who are deeply steeped in all of these questions, if there's a lesson from our podcast today it's that it turns out that there's just-- it's actually-- even though conceptually fairly easy to state the goals of this research program, it's very, very complex to carry out in practice. It's conceptually complex. There's enormous data requirements.

There's careful thought that's required. There are assumptions that are going to be made. And so on. And so you've stuck with it notwithstanding alk of that complexity over the years.

And so yeah, I'm just curious if you have any kind of parting thoughts on your general takeaway from the fact that this is a very sensible thing to do, at least-- arguably, there's a lot of smart people who think that it's a sensible thing to do. But it is also a very difficult and complex thing to do.

If there's people out there listening to this who are looking for something that's incredibly interesting and challenging to do and who are very interested in increasing the social welfare of their country, state, region, world is a great sort of thing to be working on. I think it's really exciting. It's really challenging.

It keeps your brain young. Can I say that now that I'm over 65? And I think it's incredibly important. And there is so much that we can do both to extend the framework into different policy areas but also to improve it.

MICHAEL A. LIVERMORE: Yeah, great. Well, thanks so much for a really fun and interesting conversation today, Lisa. It's always fun chatting with you. And thanks for all the great work that you've done on these issues over the years.

LISA A. ROBINSON: Well, thank you for inviting me on. It's always great to talk to you, Michael.

MICHAEL A. LIVERMORE: And listeners, if you enjoyed this episode, please let us know. You can give us a like, a rating, subscribe to the podcast and follow us on social media. It'd be great to hear from you. Till next time.

[MUSIC PLAYING]