#263: Is cash the best way to help the poor? (Michael Faye)

Julia: Welcome to Rationally Speaking, the podcast where we explore the borderlands between reason and nonsense. I'm your host, Julia Galef, and my guest for this episode is Michael Faye. He's an economist and the co-founder of several companies and organizations, including GiveDirectly, which is a nonprofit that helps you send cash directly to people living in poverty around the world.

And as you might have noticed, Give Directly has been a sponsor of several Rationally Speaking episodes this year. But this is not a sponsored interview; I wanted to talk to Michael because I think their mission is really interesting, both from an empirical perspective of, what happens when you just give poor people cash? And also from a more philosophical perspective of, why should you give people cash? Which is kind of an unusual approach compared to most philanthropy where you're giving people specific things like a cow or job training.

I find these questions really interesting and important and I'm excited to talk to Michael about them. So without further ado, here is my conversation with Michael Faye from GiveDirectly:

Julia: Hey, Michael. Welcome to Rationally Speaking.

Michael: Hey, Julia. Glad to be here.

Julia: So could you start by giving my audience just the basic overview of what GiveDirectly does and why you chose that mission?

Michael: Yeah. GiveDirectly does the simplest of things: it lets donors send money to those in extreme poverty to help get them above the poverty line.

Julia: And why did you choose that mission instead of the many other ways that people try to alleviate poverty?

Michael: Yeah. So to go to back a bit, I was doing a PhD program at what turned out to be a pretty important time for the field. It was when-

Julia: When was that?

Michael: Oh, gosh. Early 2000s. You're making me go way back. So early 2000s. I think I started grad school 2004.

Julia: And why was that an important time for the field?

Michael: So economists, and specifically development economists — these are folks that study the alleviation of poverty — started to apply some of the same evaluation techniques that are used to test whether a drug is effective to whether these poverty programs were effective. What we learned was that a lot of what we had been doing in the sector was not as effective as we would've hoped.

Julia: Like what?

Michael: Oh gosh, at the time… training programs. So as a sector, we spent so much on training programs and it turns out that we're not particularly good at training people in the emerging markets.

Which of course brings up the old aphorism, "Teach a man to fish, feed him for life. Give him a fish, feed him for a day."

And if you unpack that, there are a lot of assumptions in it, right? The biggest one is that we're good at teaching people to fish.

Julia: Right.

Michael: It also assumes that the person has enough money to buy a fishing rod. It also assumes that he wants to fish and feed his family fish. There are a lot of assumptions baked into that very simple aphorism. A lot of that turn out to be wrong when you looked at the numbers.

Julia: So what you do at GiveDirectly is give people just cash, instead of to give them goods or services that are supposed to help them become less poor.

Michael: That's right. It's almost the anti-paternalistic approach. So much of this sector has been based on the concept of the donor chooses what you need. And you can think of the most extreme example of that as the gift catalogs. You open up the foundation or charity gift catalog and you pick cow, goat, school uniform.

We say, it's a really hard choice. I can barely buy gifts for my family for the holidays, let alone someone in a refugee settlement. Why don't we just give them the cash and let them choose what they need?

Julia: How much cash are you giving people? Does it vary depending on the recipient?

Michael: It depends. So we're working in 12 countries. It varies a lot. In general, we're giving about the equivalent of one year's consumption, but to put that in perspective for what that means, the extreme poverty line is $1.90 purchasing power parity. People don't start with zero. So in a lot of the places we're working $1.50 or so per day is going to take somebody out of extreme poverty over the poverty line. Now, that's just not a lot of money, right?

It used to be the price of a cup of coffee, but now that's even less. So the idea and the opportunity to take another human out of poverty for $1.50 a day is a pretty remarkable one.

Julia: And how do you decide who gets the money? Is it just, anyone who says they're extremely poor gets the money?

Michael: Yeah. There's two bits of this. So this goes under the heading of targeting; How do you decide who gets the money? And even under that umbrella, there are two parts. There's the theoretical: Who would you want to get the money? And then operationally, how do you actually do it?

And both parts are hard. Who gets the money, we try to find the extreme poor. So those living below that $1.90 of extreme poverty. So these are folks that are in severe deprivation. Can't provide shelter, food, clothing for their families.

How you identify them is a lot harder. So when we started GiveDirectly, we were doing door-to-door surveys. So you'd go, you'd interview. You'd do something called a proxy means test to see how poor someone is and then base it off of that...

Julia: What is a proxy means test? Briefly.

Michael: So a proxy means test is you're going to do a survey of, how big is their house? What is their house made of? Do they have a TV? How many assets do they seem to have? And you're going to add that all up into some measure of how poor you think the person is. Then you're going to set a threshold, and below that threshold, you're going to give people the money.

Julia: Got it.

Michael: That was wave one. Wave two is we started targeting based on location, which is you find a location that's extremely poor, a village. And you're going to give to everybody in that village.

What we've been doing more recently is shifting that more and more digital, which was almost required by COVID. So going to door-to-door becomes a lot harder when you're in a virtual no contact world.

So one of the examples which was led by the Togo government is to actually look at people's data on their phone usage and, applying some machine learning algorithms, you can tell how poor if someone is, and do the entire targeting remotely. And roll them into the program remotely, and then put money on their phone remotely.

Julia: How could you tell from someone's phone usage how poor they are?

Michael: This is what's interesting. So, you get a lot of data on phone usage. So I'll just give you some examples: Is it a monthly plan? Do they top up with 50 cents right before they make a call? Are they receiving money from abroad or are they only using the phone sporadically in emergencies?

These are the sorts of inputs that you put in, and then the algorithms you train against the proxy means test. So the algorithms say, "Okay, if you're receiving a thousand dollars from the UK every month, you're probably not the poorest of the poor."

Julia: Interesting.

Michael: Once you know, you tell people you're eligible, get a text, radio campaigns. They type in a code and shortly thereafter they'll get paid.

So the exciting thing about this — cash in general — is this is a very scalable solution to extreme poverty. There are a lot of important problems in the world, climate change, otherwise. Poverty is a problem that, at its simplest level, is solvable. You give someone $1.50, $1.90, they're over the poverty line.

Globally, the poverty gap right now is estimated by Brookings to be $95 billion. That's about 0.1% of global GDP. So if we all decided to go home and just give 0.1%, that's more money than the global poverty gap.

So that just starts to put in perspective how solvable this problem is.

Julia: I do want to talk later in our conversation about how long-lasting the effects of giving people cash are, and more about the evidence we have.

But for now, I wanted to just get a few more details, concrete details, in our listeners' minds. For example, what kinds of things do people spend the money on?

Michael: We've seen such a range. You'll see people replace their roof. You'll see people buy food. You'll see people buy motorbikes to start a business.

And then you'll see things that you didn't imagine, like someone that bought a speaker to start a band because he had a friend ages ago that taught him how to play keyboard and always wanted to start a band. And he's doing really well.

I've talked to people that have started planting eucalyptus trees, to farm eucalyptus.

So there're just so many different uses of money. And that's sort of the point of cash, which is that even within a village, people are going to spend money differently. Each person's needs are different than their neighbors.

Julia: When a donor gives money to GiveDirectly, are they giving it directly to a specific recipient? Or do they just give money to you, and then you choose the recipients?

Michael: Historically, donors gave to us, and we selected and then passed it on to recipients.

In some cases, donors would have a preference. They might want to target the urban poor. They might want to target a specific country, like Liberia, or women or something else.

What we are in the process of launching is actually letting donors see whom they're giving the money to before they donate…

Julia: I imagine that's satisfying for donors. I'd find it satisfying.

Michael: We're really excited. It's been in the ether for a while, the idea that you can do that. And there have been variants of "sponsor a child," and things that aren't quite as direct as this.

But to do it where you can see, "this person is getting money because you, Julia just donated,"… and then you can try the operational process. So you will get an email exactly when that person gets paid. You'll get an email when that person gets interviewed and asked "What did you do with the money?" You could follow along and see, in real time, what's happening with your money, the same way you might track a package.

Now, what we're not going to do, is we're not going to let donors choose.

Julia: Why is that?

Michael: When donors choose, the evidence is that people choose based on things like the color of someone's skin or how attractive they are, and that doesn't make for particularly good social policy or operations. So we're not going to give the choice. It'll be a predetermined queue of people, but you will see who you're giving to.

Julia: So a potential recipient will be essentially next in the queue until someone funds them, and then we go to the next recipient — like that?

Michael: Exactly. Imagine being able to get a group of friends together and say, "Actually, we're going to take an entire village out of poverty or you know what, maybe we'll try to take an entire county, or country."

Or, as the movement grows, the entire world.

Julia: I wanted to talk more about this basic premise of GiveDirectly — which I really love, but want to pick apart a little bit anyway — which is that people know what's best for them. So it's better to just give people cash, rather than assuming that we know how to make them less poor, and giving them training, or a cow, or something like that.

Do you think that this is just *always* true — that essentially all anti-poverty work should just be giving people cash? Or do you think that there is sometimes a reason to give people goods or services instead of cash? How would you decide?

Michael: No, it should not always be cash, is the simple answer.

So there are public goods and interventions with big externalities that I think are important for the social sector or governments to provide. Giving people cash is not going to discover a COVID vaccine. It's not going to build public roads. It's not going to solve a lot of issues of coordination where you need a government and the social sector.

But it is a very good default option to start with.

So a lot of people will ask us, why cash? And in some way, they're putting the burden of proof on the recipient. "Why Mr. or Mrs. Recipient, should I give you cash, as opposed to spending it, as a western NGO?"

I think we should flip that, and say, "Why *not* cash?" Let's start with the idea that these resources belong to the extreme poor. What do we feel comfortable asking them for those resources back for?

So I would feel very comfortable explaining to someone, I'm going to take, whatever, 10% of this cash transfer and apply it to vaccine research. Or I'm going to actually build a road into the village. I could explain that.

But there are a lot of things that are a lot harder for me to explain. I'm going to take, whatever, 500 of your dollars to buy some grain in the US, pay to ship it over and then give it to you, even though you could have bought the grain locally, more cheaply, helped a local farmer in the process… And by the way, you might not have even wanted the grain.

That's a lot harder of a case. And that's how we talk about it internally.

Julia: At GiveDirectly?

Michael: Yeah. Our first value is “recipients first,” and the way we talk about that is: Could you explain a decision to a recipient?

Julia: So I was thinking about this question before I asked it to you. And I also came up with the idea that positive externalities can justify giving a good or service instead of a cash transfer.

I was also thinking that maybe there are some goods or services that we just wouldn't expect people to know about or to see the value in, because they just have no reason to know about that — but we have good reason to believe those things would actually make someone less poor.

Like... I don't know, suppose scientists invented a new seed that is more productive. Better at producing whatever crop people are growing. Or more robust to, I don't know, bad weather, or something like that.

I don't think we could assume, “Well, of course, if people cared about that, they would buy it themselves.” Because they've never heard about it and would have no reason to believe it was good.

So maybe something like that, there would be a case for giving people goods in kind, instead of cash. Does that sound right?

Michael: I think you want to test it and be careful. Because I hear that argument used to potentially give recipients things that they may not need.

So, markets in my experience respond pretty quickly. I'll give you an example from literally the first trip I took, is I went to a refugee settlement in Kenya, or an “internally displaced person” camp in Kenya. This was after the election violence. And there were a bunch of local stores that were completely empty.

A lot of the criticism that we came under during that first trip was, "Why are you giving people money? There's nothing for them to buy. Go look at the shops. They're completely empty."

Which was true. But what happened is once people had money, there was an incentive for a shopkeeper to come and stock the stores. And very quickly those stores weren't empty.

I've seen this time and time again. When you talk to recipients about seed varieties or types of cows, they have a really good sense of what makes money and what doesn't make money. The incentives are really aligned for them to make the right choice. So I do think the markets respond.

And we can give them that choice, right? Give them the different types of seeds and give them the money and say, "Buy which one you want."

I think where it becomes problematic is we say, "Well, recipients don't seem to buy this cow type or of the seed type, so they must be dumb. They must not understand how valuable it is. So we're just going to subsidize it and have some Western donor pay to give it to them." I think that's a slippery, dangerous slope.

Julia: Right. Okay. So I guess the modified argument would be:

If there's something that seems to us like we have good evidence would really help people, that they just don't know about… Find a way to make that available on the market in their region. And give them money. And if they agree that that's good for them, then they'll spend the money to buy it. But if there's something that seems better to them, like maybe a new roof to replace their leaking one, then they'll buy that instead.

And we should expect that to be better than giving the good in kind.

Michael: Exactly. To say the same thing slightly differently… there's the supply side, which to me is the market. So start a business selling the super seeds. So start that business, which by the way, if it's going to be a good market, you're going to get capital. You're going to be able to build that business.

And then make sure the demand side is there. Make sure people actually have capital to buy it.

We did a very small project on index insurance. So index insurance, it's weather insurance, which basically says if it doesn't rain, it pays out to recipients. This has been shown to be quite good for people, but in a lot of the studies, people don't buy it as much as you might expect.

So we did a project where we just offered it to people out of the cash transfer. Would you like to get whatever it was, $20 less and also get weather insurance? It turns out a lot of people said yes.

Julia: I see. So they weren't buying it before because they didn't have the money, but when you give them money, they're willing to buy it.

So in theory, I guess that would predict that if you just gave people money without the choice to purchase the insurance right then, they would still end up buying it, just on the market like a normal product…

Michael: Correct.

Julia: But it's kind of hard. I guess there's a bit of a chicken and egg thing, where you don't want to start the business when people don't have the money to buy it. And you also can't test to see if people buy it, if the product doesn't exist yet.

So it helps, I guess, to package those two together — where you're giving people the money, *and* the option to buy the good or the service, at the same time. To get around the chicken and egg problem.

Michael: That's right. And this chicken and egg problem comes in different varieties for us.

The other chicken and egg problem that we've run into is mobile money. So we pay people on mobile money, which is their phone. So they get a text saying you've received $100 from GiveDirectly and they can go pick up that money at any local agent. So like, a local shop. They can convert that digital to physical money.

Now, we've worked in places where there is no mobile agent within a 10, 15-hour drive. So the first question we hear is: "How is it possibly going to work? There's no mobile money agent."

What we've tried… there was a place in Northern Uganda, which was up a mountain — and you hike up the mountain and you can hear the baboons, and there were the snakes up the mountain. You barely have cellphone reception at the top, and there's certainly no mobile money at the top of the mountain.

But we did two things. One, we told some of the agents in the city we're going to be delivering cash there. So some went. But even more, what happened is people in the village just coordinated. They said, "Okay. Well, I'm going to take the trip to the village."

So the first person takes the trip, the 10-hour trip. The person says, "Well, why are you coming?" He says, "Well, we just have all this cash now in the village." And they say, "Oh, great. I'll come out." I, the agent, will come out there and just cash people out.

Now, why does he do that? He does that because there's money to be made and it's a market.

So things respond pretty amazingly. When you're living on the brink of poverty, you're forced to be extremely resourceful and creative. And that's what we see time and time again.

Julia: I've seen some people argue… I think they're going even farther than you are, if I'm not mistaken, in advocating for just trusting people to do what they think is best with the money.

So I've seen them argue that if someone is begging on the street, say… you shouldn't ask yourself, "Well, are they going to spend my dollars on drugs and alcohol? If I think that's likely, then I won't give to them." You should just give them the money, *even if* they're going to spend it on drugs and alcohol. Because it's not your place to judge how they spend the money.

So I'm curious if you would go that far. Is your support for cash transfers conditional on the evidence that people seem to not be spending the money in terrible ways? Or would you support it *even if* it seemed to us like they're spending the money in terrible ways?

Michael: Yeah. That's such a good question…

There is a threshold for me. So if you ran the cash studies and it was all going towards drugs and alcohol, and violence increased in the village and that… I would not be doing what I'm doing, is the short answer.

That goes back to the "why not cash?" question. If that was the case, I'd say, "Well, because everyone just spends it on alcohol and violence, and all these other bad things go up." Thankfully, that's just not the case.

Julia: I think that is an important point to hit. Because I think the stronger case appeals to some people philosophically or ethically, but it really turns off a lot of people who are like, "No, if you know people are spending it on terrible things that make their lives worse, you shouldn't be giving them cash."

And so to be clear that this isn't good by definition, tautologically — it's good *empirically*. Because it does turn out that people seem to not be spending the money in terrible ways.

Michael: And it's interesting, because it's a bit of both. So there's some empirical evidence threshold by which I would not do cash…

Julia: Right. I guess the philosophical argument is just about the burden of proof. That you should assume people are spending it in ways that are best for them, *unless* you have good reason to think otherwise. Which is different from what many people say, which is like, "Well, people should have to prove to my satisfaction that they're spending the money in good ways, otherwise we shouldn't give them money."

So you're flipping the burden of proof, but there's still some burden of proof. It's not just an automatic "people should definitely get money no matter what."

Michael: That's exactly right. And it's also avoiding the subjectivity of moral weights. So once we've got past this threshold, the "don't do a lot of bad" threshold, let's call it… When people evaluate interventions against each other, they need to aggregate different outcomes.

So any intervention, it might impact education, health, income, consumption, mental health, and so on. Someone needs to take all of those outcomes and say, "Which ones matter? How many points am I going to give for one unit of consumption? How many points am I going to give for one unit of education?" And that's inherently a subjective exercise.

The philosophy of direct giving says, "Don't do that exercise as a donor. Let the recipient do it." One, because there is no right universal answer. There's no single source of truth of what moral weights matter. People are individuals and they're going to know about their own subjective weights better than us. There's no way for me to possibly know. I hardly know my own subjective moral weights, let alone those for someone else.

Julia: Shifting tracks a little bit… One thing that is interesting to me after reading a bit more about GiveDirectly, and your history and the specifics of what you do and how you do… it is that it had seemed to me like a very simple intervention to just give people cash. And yet there are actually so many different ways to do that, and the differences can matter a lot.

Like, how do you decide who to give cash to, which we talked about? How do you decide how much cash to give them? Should you attach strings to the cash? Should you give a one time lump sum gift, or a recurring income that they can count on, et cetera?

So I just wanted to talk a little more about how you made those decisions. Maybe let's start with the gift amount, which I think was about $1,000. How did you pick that number and has it changed over time?

Michael: It depends. So cash is not monolithic globally and GiveDirectly doesn't have a single model of cash. We've given people $30 a month for 12 years. And we've also given families a one time gift of a thousand dollars. Then of course, if you do it in the US, it's even more than that. So —

Julia: How do you decide?

Michael: So, how do you decide? That's one of the hardest questions of designing a cash program. Where we started was: If we gave people this amount and they invested it, it would bring them up to the level of people that were not eligible for the program. So, get them up a level.

Where we want to shift more and more is to actually ask recipients themselves. So you could imagine the veil of ignorance, telling a recipient, "You can either get a thousand dollars, but it's a 50-50 that you get the money. Or with certainty, we can give you $500. What do you think we should do for the next village, or what would you prefer?"

Julia: Oh, interesting. It kind of is a veil of ignorance. It's not just a thought experiment, right?

Michael: No, it's not. We can do this.

Julia: Yeah. It's for real. You can ask them.

Michael: It's for real. So I think it's just pushing the philosophy of recipient choice further and further. Same thing with, "What's better, monthly transactions or one time transactions?" Why can't we give that choice to recipients?

Julia: What do people usually choose?

Michael: It's a mix. When we did it, it was very, very mixed. There's some small suggestive evidence, that I don't want to overstate, but it says that if you actually give people what they want — so if I pick monthly and I get monthly — I do slightly better than if you made that decision for me.

Julia: I guess that makes sense. People know how good they are at saving or not.

Michael: Exactly.

Julia: And my impression is that in the past you used to give money to individuals irrespective of what village they lived in. Whereas now you choose a village to give to, and then give to anyone who is eligible within that village.

Am I right that's the direction you're trending in? And if so, why did you make that shift?

Michael: Yeah, this is one of the biggest operational changes we've made over time, and it's something we've debated a lot internally.

So the reason why you would pick within a village is because you wanted to find the poorest of the poor within the village. But anytime you do that targeting operationally, you're going to make two types of errors. Sometimes you're going to include people that shouldn't have been included — so, they're slightly wealthier than you would've liked. And you're going to exclude people that should have been included.

So to give you an example, to make this real, we looked at housing material and construction as one of the inputs. Now, it's a reasonably good predictor of wealth, but sometimes you'll have something like someone whose husband died, or whose family died, and is still in the old housing. So it's a cement house, but they have no income. They're taking care of a few children. So they're really extremely poor. They just happen to have inherited a cement house.

If you're using that to exclude people, they're not going to get money.

So you wind up spending a lot of money doing this individual-level targeting, making both these exclusion and inclusion errors. As we got more and more data, it got to the point where we thought the additional cost of doing that targeting wasn't worth the slightly better targeting that we might get out of it, and that it made sense to give to everyone.

We're now trying a universal basic income, which is over large geographies. We're giving everyone the same amount over a long period of time. And we don't have the results yet, but anecdotally… I was in Kenya, and one of my favorite questions to ask was, "Do you think this is fair? Is it fair that someone that's slightly wealthier in the village gets the exact same amount of money as someone that's on the poor end in the village?"

Julia: Interesting. What'd people say?

Michael: To set the expectation, this is not people living on a dollar a day, and billionaires. This is people living on 60 cents a day, and maybe people living on a few dollars a day at most.

Julia: But presumably, those differences, even though they might seem small to an outsider, will feel significant to the people living that life.

Michael: Absolutely.

And people will say that they really like it. Because it one, destigmatizes them. So it's not the poor and the not poor.

But two, it's allowed them to actually start to have conversations about how to spend the money. So the comment that stuck in my mind was: "Before we used to have richer and poorer, and now we just have everyone getting $30. So we can go into the market and talk about what are you doing with your money? What am I doing? And it allows us to actually trade insights, maybe a new seed to adopt, and have better conversations than we did historically."

Julia: And that would be harder if people were getting different amounts of money because… what one person could do with their $30 is different from what another person could do with their $300?

Michael: And nobody would know who is getting what…

Julia: And it would be awkward.

Michael: And if you said "I'm a hundred-dollar person," you say, well, oh, you must be really poor. Or the opposite. "Why'd you get a hundred dollars? You're not that poor." So you have all these tense conversations, so people tend to close up and not talk about it as much.

We had something similar in Houston after Hurricane Harvey, one of the earlier places we went. We were going to target based on income level. And in the US, it's a bit easier. So you can look at people that are on welfare programs and target accordingly.

And what people said was, "We don't want that. We're already stigmatized for being on welfare. It is better to just draw names out of a hat. And that way it's fair, and we're in it together as a community."

Julia: That actually reminds me of another question. Why do you do some of your work in the US, given that a dollar can go farther in Africa than it can here?

Michael: This was another much-debated internal decision.

So, you're absolutely right. A dollar goes further in the context of extreme poverty, right? So just to put that in context, universal basic income in a place like Kenya will cost about $30 a month to get people over that threshold. In the US, people talk about that as about a thousand dollars a month.

So it's a meaningfully different amount of money to get over the threshold.

What we realized is two things. One is that a lot of people think about their giving in separate mental accounts. They're going to give to their local community because it's the right thing to do. They want to be part of the community. They feel some obligation to the community that's also supporting them. And two, they want to solve some of the extreme global problems, like poverty.

And they've already done that allocation. In that context, the best that we can help them with is doing both as effectively as possible. If you want to give locally, and you want to give cash, be less paternalistic and do it as efficiently as possible, we want to provide that opportunity.

And the second thing that turned out to be true is the US gets more attention than global poverty. There was a great piece in Quartz comparing the number of media articles on Harvey, Hurricane Harvey, in Texas, and Maria in Puerto Rico and it was meaningfully different.

So if we want people to be thinking about giving directly and giving cash, I think working in the US is really helpful to that. And we saw exactly that, which is a lot of people that started giving directly, through the US programs, have actually started to also do that internationally and globally.

Julia: So maybe this is a good segue into talking about the evidence.

Like, how long-lasting are those effects? If we look at people who receive a gift from GiveDirectly, five years later, are they still better off — or is this a more temporary benefit to them? Et cetera.

So what do we know?

Michael: Yeah. So there are papers that show long-term evidence, the longest one actually coming from the US on something called the Mother's Pension Program, where they looked 50 years after the program at how the children of the recipients did. As it turned out the children live longer and had higher income during their lifetime. So that's by far the longest study and it's published in the American Economic Review. It's a great paper.

Julia: What about evidence in the areas where GiveDirectly [mostly operates]?

Michael: In the developing world? There are two things. So, in some cases in Sri Lanka, you saw the gains persisting at seven years. Other places, you've seen them persist as well, and then in other areas —

Julia: And the gains we're talking about here are… they have more money?

Michael: Income and consumption. Yeah.

Julia: Okay.

Michael: Then in other contexts, you see the people that didn't receive money catching up to those that did.

Julia: Within about how long?

Michael: I think that was also the 7 to 10 year time [frame]. So people catch up.

Julia: Is that surprising? Is that what you expected?

Michael: It gets at this “poverty trap" literature. And this will get a bit wonky.

So, the poverty trap is the idea of that if people are below a certain line they're stuck, and all we need to do is give them one shot in the arm and get them over that line and everything will be great.

The other view is the classic view, which is that the poorer you are the higher the return to capital is. So that over time you're going to catch up and converge.

And this theory is talked about both in the context of individuals, but also countries. The concept of a poverty trap casually underlines a lot of comments you'll hear, “Oh, we just need to help people once and get them out of poverty and that'll be it.”

But the reality is when you look at the evidence, there's not really strong evidence that poverty traps exist at the individual level. And there's a lot more work to be done on this, but it may be — which I think would be great news — which is that people will catch up over time.

And I think we should take that as a positive, not a negative. It doesn't mean that accelerating the end of poverty is not an important objective or that giving people money to get them over the poverty line doesn't do tremendous good in the short term.

Julia: Yeah. One thing I've noticed with COVID is that people, in general, don't seem to be that interested in accelerating the good that we're doing. Like accelerating the deployment of vaccines.

I don't know, my friends are very interested in that — and you can do the back of the envelope calculation and notice, “Oh if we just sped up the vaccines by this many days, we'd save this many lives on average or something.”

But I feel like, for most people, they look at that and they're like, “Well, we get the vaccines; we don't get the vaccines. That's what I care about.” And they're not thinking about how much benefit we get by doing it earlier, instead of later.

So there's probably a similar thing happening for people looking at anti-poverty interventions, where they're like, “Well, if people were going to catch up anyway eventually then we didn't do anything.”

Michael: That’s exactly right. And if you flipped the question and say, there are two worlds, a world where people would've caught up anyways, it's just going to take a very long time… or a world where people wouldn't catch up and would be stuck in poverty forever? Certainly, we would all prefer the world where people caught up.

The other thing I don't want to miss about cash, which is important, and is important for interpreting the evidence, is there's now evidence that not only do the people who receive cash benefit, but the people who don't receive cash, that live near the people that do, benefit.

Julia: Why would that be?

Michael: So if I get cash, I need to spend it. It is going to be spent likely at a local business or somewhere. So I'm literally taking the dollar that I received and giving it to someone else. So there's another form of cash transfer in the market and that person is then going to give it to someone else and so on and so forth.

So when you look at the impact of cash or any development intervention, you need to ask, not just what the impact is on the person that receives but on everybody else.

And as time goes, this cash is going to start sloshing around. The effect is going to diffuse in the community. And that's a really important fact.

Julia: Got it. I was curious about the positive spillovers that this paper found, because I would have thought that there would be at least the negative spillover of inflation. Where, if all these people get say a thousand dollars to buy new tin roofs and they all want to buy new tin roofs, then the guys selling the tin roofs are like, “Oh wow, I can afford to charge more now, because there's all these people who can suddenly afford it.” So the price of tin roofs goes up.

Why wouldn't we find something like that?

Michael: So, theoretically, that's exactly the right question. When you give money to a place, either prices could go up and there's no real change in the economy… or real things — output, consumption, can change. And that becomes an empirical question, which is one of the things this study was looking to address.

In that study, they actually gave a really large amount of money. So I think it was about 15% of GDP that they dropped on the village… and I forget the exact price impact, but it was tiny. I want to say 0.1%.

Julia: So that's surprising, right?

Michael: I think it depends. If you told me to design a cash project to create inflation, I'm sure I could do it. I would go to a very isolated island that wasn't to the global economy, I would drop an enormous amount of money on that economy. And I'm pretty sure I could create inflation.

But the reality is, the places where we're working, even though they're extremely poor, do seem to be connected to the economy. And do seem to have some slack capacity, where you don't see that inflation.

It's something to keep an eye on, and we should be evaluating it in many contexts… but that's at least what they found in this paper, and several other papers that have been done on the GiveDirectly programs.

Julia: Why does it matter whether the region where you gave the infusion of cash is connected to the larger economy? Or how connected it is?

Michael: This is great. This is back to my macro oral exams… If it's globally connected — say, cars — the price is going to be set on the international market. The fact that a village in Kenya… or motorbikes. A village in Kenya has slightly more money. In the big picture, that's not going to move the needle.

Now, if it's a completely closed economy, there's only one person giving haircuts or something on this island, you would expect differently.

The channel by which it changes in this general equilibrium paper is kind of interesting. So, well, how does output change? And it's suggestive evidence at this point, but it goes as follows.

If I'm cutting hair and I need to go to the shop every day, I'm going to be there for eight hours, no matter whether one person comes in or two people. So it's eight hours worked.

And in the old world, where there wasn't much demand, these people didn't have the money, so maybe I cut one or two heads of hair. Now, for the same eight hours, I might be cutting 10 heads of hair.

So you have this integer constraint, which is I have to go to the shop no matter what. Now, I'm actually doing more and producing more in that same quantity of time than I was previously.

Julia: So this is about the amount of slack in the economy, basically… and so it suggests that infusing money into a region can take advantage of this slack?

Michael: That's right, and that's how you see this multiplier effect, which is how much more good gets done in the economy than just the impact on an individual. And they estimate in that paper, that it's about 2.6. [Edit: Michael later corrected himself, the multiplier found was 2.4]

Julia: I see. Well, so if inflation wasn't a big problem in this case, because these particular regions that were getting the cash were quite connected to the broader economy, et cetera… Does that suggest that it might be a problem if you guys were to scale up, and maybe do something like the universal basic income you're doing in Kenya? Because there you're having a much bigger input into the demand side.

Michael: So, you would absolutely want to keep an eye on this and keep testing it. I suspect even if you were doing universal basic income in a county or a district of Kenya, at these levels, you would be okay, because a fair bit of the economy will be traded. But it's something you want to keep an eye on.

Julia: I've got this list of observed or potential or theoretical downsides of giving people money…

We've hit on the question of, "Does it last?"

We've hit on the question of, "Could this make their neighbors feel worse about themselves?"

We've hit on the question of inflation.

Another one on the list is: "Do we have any evidence that receiving these cash gifts makes people less likely to work?"

Michael: No. So, there's a great paper called "Debunking the Myth of the Lazy Welfare Recipient," by Abhijit Banerjee and colleagues, and they look across papers at exactly this question and find that there is not evidence of people stopping to work.

And in some cases, there's evidence of people working more. And the intuition of that is pretty straightforward. It's, "If I have more capital, the return to my time is greater. So, if I have a motorbike and can make $20 a day, versus if I didn't have a motorbike and maybe what I'd be doing earns me $5 a day, I should work more, because I have the chance of making more money because of this capital." And that's what you tend to see in the data. I don't want to overstate the working more, but you don't see evidence of working less.

Julia: Are there any other observed or theoretical downsides to cash transfers that you might be concerned about?

Michael: Oh, I think you've hit the big ones. I think there's inflation, there's how neighbors feel, there's people stop working, there's people drinking. All of these things, and they should all be evaluated empirically. Right? You need to look at this. And there's certainly going to be cases, right? Some people, there are no question, people that we have given money to that have had a beer, that have probably decided to work less.

So, we're really looking at aggregate impacts, and we're looking at aggregate impacts in specific context of the projects we're working.

Julia: Have you had people turn down the offer of free money? That could indicate a potential downside.

Michael: Oh, lots of people turn it down. It's one of the sadder parts of what we do.

Julia: Why do they turn it down?

Michael: Imagine if someone knocked on your door and said, "Hi, Julia. We're here from an organization called Give Directly, and we just needed your phone number, and we'll send you $30,000."

Julia: Uh-huh.

Michael: You would probably close the door.

Julia: Yeah. That's true.

Michael: And then you'd probably think to yourself, what do these people want from me? What's the catch?

Julia: "What's the small print?" Sure, yeah.

Michael: Especially when a lot of people that have given to recipients have asked for something, right? "We'll give you this, but we'd really like you to convert to this religion."

Julia: Right.

Michael: Or we'll give you this, but we expect you to do the following behavior.

So, "We'll give you this, see you later"? It's pretty unusual.

And we actually had a problem in the early days, where we were using the housing type to decide who got and who didn't get. And people are so easily persuaded that they thought we were asking them to improve their houses with the money.

Julia: Oh. I see.

Michael: And they did it, but [they said], "But I thought you wanted me to do that. I thought that's what made me eligible." We said, "No, no, no. This is really no strings attached."

So, we had to revamp our communications and figure out, how do you persuade people that we're really not asking for anything?

Julia: Actually, that does make me think of another potential downside, which is: Maybe if people know that you're using the quality of their roof as a proxy for whether they're eligible for receiving money, that might give people incentive to not improve their roof. Potentially for years. Because they're like, "Well, I don't want to price myself out of this potential gift I might get in the future."

Would that happen?

Michael: Yeah, we've had all these concerns. I think as much as we've grown, for the random Kenyan household to say, "There's a one out of — whatever it is, million chance, that I become a Give Directly household, and therefore I'll continue to let my children get rained on…"

Julia: Yeah. Just seems unlikely, a priori?

Michael: We haven't really run across it yet.

And it's one of those interesting cases where we all come in with biases. I've been doing this for a while, and I still catch myself with those biases. You go in... The roof one was the perfect example. "Why are you putting on a nice roof? Do you really need a nice roof? Is that the equivalent of the marble bathroom? Is that just a status symbol?"

And people say... They'll give you this very rational argument:

"One, it's cheaper over the long term than replacing thatched roofs. Two, I can collect clean water off the roof. Three, there are less mosquitoes when I have a tin roof." And then they say the thing that we don't measure, which is, "Do you know hard it is to watch my children get rained on every night and know that I can't literally put a roof over their head, and how painful that is for me?"

And then I pause, and just step back and say… there's just so much I bring as a Westerner, in terms of judgment and bias that is so unfair. When someone says, "Do you know what it means to watch your kid get rained on?" And the answer is, "No I don't."

Julia: Yeah. Is there anything else that you've changed your mind about since you founded Give Directly about poverty or how to approach it? Or just anything else that was kind of surprising to you as you really dug into the details?

Michael: I think our thinking has evolved. When we started this, I was doing a PhD. The co-founders were doing PhDs. It was very much rooted in the evidence, and the new technology that let us do this. But it was still part of that question of what is best for the extreme poor. And over time, we realized how impossible that question is to answer, for the reasons that we've described.

Julia: What about mistakes that you feel you've made in retrospect, things you would do differently if you were starting from scratch?

Michael: So, I think we would not have done targeting within the village and randomization within the village. I think it's really hard to tell two people within one village that one of you will get — and one of you, that's essentially the same, won't get, so that we can learn what the impact is.

Now, the types of study we're doing now, where you randomize at a village level, I think are fairer and better. But they're also meaningfully more costly.

Julia: Just because you're giving money to so many more people?

Michael: Exactly. And it's a trade off that I think we need to be really thoughtful about, and we as an organization have moved away — and have had the luxury, in many ways, because we have enough funding that we could move away from the individual. But I think you really need to think about those dynamics.

Julia: There's a common critique that I'm sure you've heard many times, and anyone doing any kind of antipoverty interventions, not just cash transfers, probably hears this a lot. Certainly EAs hear it.

And the critique is that what we really need is systemic change that reduces poverty on a large scale, in a lasting way. And just giving people money in the meantime is just this bandaid. And we should be working to cure the disease instead of put bandaids on the symptoms.

Do you have any thoughts about that?

Michael: Cash is a very unappealing intervention as a donor. You'd be hard-pressed to come up with a less satisfying intervention than just giving someone money to make them less poor. Right? It goes against all of our human instincts.

So, one, "So, wait a second. I have no control? I just give someone all of the control and power?" Yes. You have no control.

Two, "But I'm really smart. I'm creative. Surely I can find something better." Right? So, you don't get to do that with cash.

Three, it's not soft and cuddly like a goat or a sheep.

So, there's a lot about cash that's really unappealing from the donor perspective. And I think it takes a fair bit of humility to step back and say, I don't know what's best. I'm just going to hand it over.

To your broader question on systemic change, we would all love systemic change, but ultimately lofty goals need to be tracked back to specific actions and interventions. How are we going to get that change? And people will say, "Governance matters." Governance does matter, but what can we do to improve governance, right?

And I'm not sure we have great interventions to do those things. Now, it doesn't mean we shouldn't try, and we shouldn't keep trying. But in the meantime, we have something that works and it works today.

And rather than worry about what's *best*, I think taking the lens of, "Look, we have a lot of money globally right now. Is it worth spending a bit of that 0.1% to take people over the poverty line, while we work on systemic change and some of the other problems?"

My answer is yes. I don't think we should be waiting.

Julia: Well Michael, thank you so much for coming on Rationally Speaking, it's been great talking to you.

Michael: Awesome. Thanks so much, Julia, super fun.

[musical interlude]

Julia: That was Michael Faye, president and co-founder of Give Directly, and if you want to learn more, I encourage you to check out their website, givedirectly.org. I especially like the section called GD Live where you can see updates from the people who receive cash, explaining what they're doing with it. And I'll put links in the show notes to the papers that we talked about, studying the impacts of giving cash on, for example, inflation.

And of course if you want to donate to Give Directly, you can do that as well, at givedirectly.org / rationallyspeaking.

That's all for this episode, and I hope you'll join me next time for more exploration on the borderlands between reason and nonsense.