

General Equilibrium Effects of Cash Transfers in Kenya (Frisch Medal 2024 Winner) Podcast Transcript

Stefan Dercon

Welcome to the CSAE Research Podcasts, a series of conversations about projects taking place at the Centre for the Study of African Economies at the University of Oxford. I'm Stefan Dercon, I'm the director of the CSAE, and I'm a professor of economic policy at the Department of Economics in the Blavatnik School of Government.

This year, the Econometric Society awarded the 2024 Frisch Medal Award for the best applied (empirical or theoretical) paper published in *Econometrica* in the last four calendar years, to Dennis Egger, Johannes Haushofer, Ted Miguel, Paul Niehaus and Michael Walker, for their paper 'General Equilibrium Effects of Cash Transfers: Experimental Evidence from Kenya'. It's a great pleasure for me to be able to talk today to Dennis, who is an associate professor at the Department of Economics here at Oxford, and also linked to the CSAE, and we're going to discuss the research behind the award.

So welcome, Dennis.

Dennis Egger

Yeah, thanks very much. I'm looking forward to this.

Stefan Dercon

So, Dennis, this is quite an amazing achievement to get an award like this, awarded only every four years for the best paper in *Econometrica*. It's going to be really interesting for our listeners to understand a bit more about this research and maybe just to simply start with this is. What was the question you and your team set out to answer?

Dennis Egger

Yeah. Thanks, Stefan. So we started at a time when research on cash transfers, on conditional cash transfers in particular – which is basically cash given to poor households, with no strings attached, so they can decide what they want to do with it, different households will have different priorities – had sort of been growing and there had been a set of Randomised control trials (RCTs) looking at the impact of these cash transfers on recipient households. So how do they spend the money? Do they spend it on wise things? Do they invest it in things that will generate income down the line?

There were all kinds of worries around potentially poor households not being able to use the money effectively, and those had sort of been dismissed at the time when we entered the field.

But what was not clear is sort of whether these effects stop at the people that you give the money to, or whether these effects ripple out and sort of affect their neighbours, the local communities, local prices and the economy more broadly, this is what we call the general equilibrium effects, which are in the title of the paper. And in a way, it's a sort of natural question as an economist to ask, because cash, by nature, only has positive benefits for the recipients if they spend it, and they need a counterparty to spend that cash, and presumably the counterparty will benefit in some way, shape or form from receiving that cash. And we went out to try to estimate rigorously the spillover effects on neighbours, local economy prices, enterprises, and firms in the local area. And, and this is important because overall, we care not only what happens to recipients, but we care what happens to welfare of everybody in the community to local economies, both in the short and in the long run. And that's why we started to address that.

Stefan Dercon

Right. Tell us a little bit about, how do you go about answering a big question like this? You know, normally micro people, they look at micro impacts on households and on firms and then macro people look at big things, so this sounds to me like a micro person trying to do something at the macro scale.

Dennis Egger

Yeah, exactly. And this is partly why I got super excited about this paper from the start is that, is that really we are trying to merge here some of the tools and methods that we use in microeconomics with the sort of concepts and effects and questions that macro economists are asking. We're really sort of doing macroeconomics, but with an RCT and with our own really rich micro data collection. And I'm going to sort of talk a little bit now about how exactly did we do we do this?

The first thing, and you've already alluded to this, is that you need really extensive measurement, right? You need data not just on recipient households, but also on households that were ineligible. You need a full set of representative household surveys that map or capture the entire economy in the area that you're studying. But you also need data on firms, enterprises, so a representative set of all the firms and enterprises operating in these villages and economies. And you need data on prices, prices of labour, such as wages, but also prices of output goods that you can find in the markets. And to be able to do this, we had to construct our own measures of GDP basically. So, we had to be like a national statistical office, but with our own survey data, reconstruct what they would do to sort of capture broad aggregate economic activity in these villages.

The second thing, of course, is that you then need some identifying variation. So how do you actually estimate the effects of this? And we did a randomised control trial, which is a very established in the micro and development tradition, but with a twist. The problem with the normal

RCT is that you have a randomly assigned treatment group and a randomly assigned control group. But when you are worried about spillovers, about things that might be changing in this entire economy, then your control group could actually also be affected by the treatment. And so when you're now comparing treatment and control groups, if the control group is also affected, you don't have a valid counterfactual. You don't have a valid measure of what would have happened to the treated group in the absence of treatment. And in a study that specifically tries to get at the spillover effects, that's obviously no good. Right? And so the innovation methodologically there is to have two levels of randomization. So yes, you randomly select villages that receive cash, but then you also add a higher level of aggregation, which in our case was the sub-location. This is groups of 10 to 15 villages. You also have higher and lower saturation of cash. So in some regions, two thirds of villages were randomly selected to receive cash. In other regions, one third of villages only. And what that creates is now, at the local level, variation in the intensity or in the saturation of cash flowing into these local economies.

And now you can compare even control households in high versus low saturation areas where in some places a lot of their neighbours got cash, and in other places fewer of their neighbours got cash. And to the extent that you find any differences between those households because treatment was randomly assigned, those differences would have to be because of these spillovers. So the households with more neighbours treated within one, two, or three kilometres around them, if they're doing differently from the households with fewer neighbours treated, that would have to be as a result of these cash transfers. And lastly, you can only do this if you have an implementing partner. If you have somebody that actually does provide these cash transfers and is willing to conduct this study with you and learn from them. We were working with the non-government organisation (NGO) GiveDirectly. The business model of the NGO is to provide unconditional cash transfers to poor households. At the time they were working in Kenya, the region that we were working in is Siaya County, and that was sort of selected because it's a relatively poor part of Kenya and relatively densely populated, and therefore super important and also in some sense representative location to study these kinds of impact. And, yeah, those are sort of the ingredients that you start with. And then you go and do it.

Stefan Dercon

Yeah. I just want to follow up a little bit on how you did it. And before we get to these results, to pick up when cash transfer programs are normally happening, they give \$10, \$20 a month or relatively small amounts of money. With this design that you describe, you must be very rich, or at least the people you worked with. This must be a massive study. Can you just give me a sense of the scale and how big do you need to be able to do this and to be able to actually see impacts here?

Dennis Egger

Yeah. Thanks. That's one part I maybe didn't stress enough. You're right, to have not just effect at the household level, but effects at the economy level that you can actually measure and observe

and pick up with the statistical methods and tools that we have, the study needs a lot of variation and cash transfers. In our setting, that was true. First, households who were recipients of this, received \$1,000 or about \$2,000 in purchasing power parity (PPP) terms. That is something like 75 or 80% of average annual household income. So it's a very large transfer. And this is typical of how GiveDirectly operates. It's not small piecemeal transfers, it's a one time, over three tranches, paid out large lump sum cash transfer. That's quite transformative or expected to have quite big effects at the household level. But it was also large in terms of the saturation. Right? We just talked about this. In high saturation areas, two thirds of villages receive transfers. And in treated villages the transfers were something like 15% of village level gross domestic product (GDP). This is this is a massive stimulus. If you compare this to tax rebates or other types of study that people do in macroeconomics, these often are sort of 2, 3, 4%, of course, national GDP in their case, but they're locally not as intense as this study. And so yes, you're right. If you're going after big effects, you need a big treatment, you need a lot of variation in your treatment at the geographical scale that is relevant for your study. And fortunately, we were able to do this with GiveDirectly.

Stefan Dercon

This is quite interesting because now when you describe how big this intervention is, I can imagine that your priors would have been “this could have had quite big effects on prices, big effects on wages”. There were probably a lot of people expecting to find that this is really disrupting a lot of what's going on there. So, what did you find?

Dennis Egger

Yeah. I'll walk you through the results, as we do in the paper, starting with the recipients. We're going to look at effects first just on people who received the cash and there we followed in a long line of research, hundreds I think by now of RCTs have shown similar results. Recipients, of course, benefit. And in this setting, they basically spend all of it over the first year, year and a half. All of the cash is spent. And that's not necessarily because recipients aren't saving, even when they're saving, they're purchasing things like livestock or durables, which then will have benefits for a long term to come. But they're still market transactions. So, they're still purchases in the local market where the money then ends up in somebody else's hands. And this is going to turn out to be really, really important. So, we estimate here a marginal propensity to consume of about 0.9 over the first year and a half, two years, which basically means they spent everything.

And what also turns out to be important is that about 75% of that spending was local. So it's on goods that are locally produced and locally sourced. A lot of it is food. Others are local services, local manufacturing and so on. So, a lot of that money in the next round ends up in the hands of somebody else, or some business owner who locally sells these products.

I just want to say one quick thing again to flag. We didn't find any of the worrying potential negative impacts on purchases of tobacco or alcohol or gambling that some people had been worried about. We don't find that at all. People spend more on food, they improve their homes, they send their kids to school for a bit longer, they start businesses. As you would expect, different

households have different priorities, but generally it looks like a relatively meaningful, sensible, positive change in recipient people's lives.

We followed this money as it flows through the economy. We're now going to turn to what we found on the enterprise side. So reassuringly, a year and a half later, recipients had 15% higher consumption. The \$350 per year. We find exactly the same increase in local revenue, \$350, increases in revenue on average per household for local businesses. And that's just to say that our measurement worked in the sense that what we're able to capture in terms of increased spending from a household service, we're also able to capture as increased income for local firms. So local firms output expands. And one thing that we're going to come back to and that's really interesting, is it expanded without them really hiring any additional workers or purchasing all that much more in terms of machines sort of inventories. And what seemed to happen is that these firms were able to increase production by reducing their slack or idle capacity, and therefore sell more.

Corresponding to that, we also don't really find them increasing their prices. On average there was inflation. And we again, just to highlight the richness of our data, we collected data for something like 100 products in 60 markets over 30 years and up to three price quotes. So there are like 500,000 measures of prices underlying this. And on average, over this period, prices rose and treated areas by 0.2%. Compare this to the 15% increase in spending, you can already see that basically, almost all of the gains that we see in terms of spending were real gains. People were actually doing better. They didn't spend more just because prices had been going up. So, we find very little inflation, overall.

Now how does that now affect non recipient households? Many of these businesses who now have higher revenues and higher profits are not owned by recipient households. Right? They are owned by neighbours or people who are in a control group. Yet they still benefit from having this increased demand for their products. They're able to pay higher wages. We find a pretty sizable increase in wages, and we find a pretty sizable increase in profits. And some of these increases in wages and profits will go to non-recipients. strikingly, if you look at a year and a half after the transfers, we had recipients consumption go up by 13 or 15%. Non-recipients increase in consumption was exactly the same, 13 or 15%, a year and a half later. And that really highlights this idea of multiplier effect. So, recipients spend the money, local businesses see an increase in their demand. They're able to pay higher profits and higher wages to their workers and their owners. These owners now have higher incomes, which they spend again in the local market raising further local demand, which they then, which then trickles back to the owners and so on. And as the money starts flowing around the economy, you have this typical Keynesian idea of a multiplier where the tide lifts all boats. A year and a half later, everybody seems to have benefited roughly equally, from this large injection of cash.

What we then try to do is tie this all together and summarise this in one multiplier number. Because this is the classic Keynesian idea of how much extra economic output you get for every \$1 that you transferred into this economy. And what we find is that this multiplier in our setting is 2.5 on average, and we measure this in two different ways. On the expenditure side, you can see what the expenditures are that are going on into the economy. And on the income side, what are all the sources of income. I've already mentioned wages and profits but there's also rental payments. And how much does GDP, measured in these two different ways, go up as a result of sort of injecting \$1

into these economies. And the result is 2.5. So, for every dollar transferred, there's two and a half times more than that generated in local economic activity. And that is net of any price increases, so in real terms.

That's a pretty large multiplier. And it's larger in particular than many of these studies in high income, richer countries, such as the US or Europe, find in response to cash rebates or military spending, for example. And it's certainly something that surprised us. But it's super encouraging because it sort of shows that if you just took the direct benefits of cash that we had already measured in RCTs before, you would significantly underestimate the total benefit of these programmes because you're missing another one and a half times as much benefits that accrue to people, who are the neighbours, who are local business owners and who also shop in these markets.

Stefan Dercon

So, okay, those are cool results, clearly. Tell us a little bit about what you think this means. Are you saying we're now going to do this over the whole of Kenya, and we would see these effects, we should be doing this? What do you draw as conclusions for broadly defined policy of NGOs or philanthropists, of governments? What are you confident about in terms of what it tells you?

Dennis Egger

I think the first point is just to reiterate what I just said. What seems clear to me is that if you want to assess how effective cash transfers are, looking at the direct effects on recipients only is going to significantly underestimate the total positive benefits.

Stefan Dercon

Yeah.

Dennis Egger

So in a setting like ours, you're going to only have 40% of the full effect if you're doing that. And you're going to want to reassess your evaluation of how good cash is relative to other types of interventions that you could do upwards, because these spillovers are positive and concerns around negative spillovers, including on prices - one thing I haven't mentioned is on sort of, potential negative psychological spillovers that people now perceive that there's inequality in the economy and they're unhappy about this- we don't see these at all. So even the non-recipients, if anything, have no change in sort of reported psychological well-being, if anything slightly positive. All these potential negative concerns around negative spillovers aren't borne out. On the contrary, the sort of maybe the positive spillovers are larger than we may have expected. So it's going to presumably sort of assess your evaluation of cash upwards by a little bit. I'd want to be careful,

though, to say, look, the multiplier is going to be 2.5 everywhere. And these negative effects are never going to show up.

And now I'm going to get a little bit nerdy about what's really underlying this. And really that the aggregate impact of this cash depends on this low observed inflation, even though output in the local economy increased by a lot. It's what we would call a supply elasticity. It just seems to be true that in this type of setting, the supply in the local economy is able to match that increased demand. And that's a substantial increase in demand of like 15-20% of GDP, without the local firms having to raise prices. And that's really the big puzzle that that was raised by the study is, why is it that economies like this are able to do that? And when I see economies like this, I mean, a vast share or probably a majority of economies in this world look relatively similar to this. Relatively rural, dominated by agriculture, with local weekly markets. That is prevalent all over sub-Saharan Africa, South Asia and many other developing country settings. And presumably, some of these findings that we have would also translate to there. And this sort of leads us a little bit into sort of ongoing work of how can we explain why that happened.

One idea that we have is to look at how firms actually did it. And as I already mentioned, they didn't really hire any more workers. They didn't really invest and sort of purchase many new machines. Yet they were able to increase output by 30 or 40% without doing so. Clearly a part of it is increased intermediate. So, they purchase the inputs that they need for their products. These are imported typically from Mombasa Port or from Uganda in our setting. But we show that this is a small part of the total effect. A big part of the fact is just coming from increased utilisation or increased output with the same number of workers and machines. And this in economics terms, this is a productivity increase. So, what we think is going on is that there are sort significant underutilisation or slack in these economies in normal times.

And one reason why this might be is, just to give you a statistic, 90% of firms in our markets, and this is the universe of all firms, are one person firms. And anyone who's been in these kind of settings or works in development economics knows that this is super prevalent everywhere. You have a lot of these micro-enterprises. And what we think might be going on is that basically these firms face a sort of integer constraint. And that is the fact that ideally, maybe, they only have enough demand to be there 3 or 4 hours in the day, but they have to go to the market and "man the shop", for the whole market day. And now if you have 20% more customers, you work five hours instead of four hours and you wait a little less. And you're able to serve those customers without having to sort of hire a second worker.

A similar sort of logic applies with the most predominant form of manufacturing here, which is, grain milling. Every village has a grain mill. That grain mill may run four hours a day. Now, with more people bringing more grain to be milled, it runs five, six hours a day. But we didn't really need a new machine to service the additional demand. And that will tell you now a little bit about how generalisable these effects are.

In other settings where you don't have so much under utilisation, where firms are much closer to capacity, where there are larger firms. Say for example, urban, highly densely populated Nairobi, potentially you'd find smaller multiplier effects.

What's sort of nice is that since we published this paper, a lot of other papers have come out using different methodologies, trying to estimate multipliers across different countries. And the pattern seems to hold that you do find larger multipliers in developing countries settings more broadly, and this is true in Mexico and in Brazil and in the Philippines and in other sub-Saharan African countries. So, we are capturing here something really, truly, different about developing country settings that, maybe is more generalisable than we initially thought. And so that was a long winded answer to say you should also think that direct cash transfers, social protection programmes, things like that would have bigger multipliers in developing countries more broadly, even though at some point potentially, they'll start to become a little bit smaller as in our city and inflation will or potentially might kick in if you make them large enough.

Stefan Dercon

Yeah. So it's interesting thing that you get here and you emphasise the relative rural nature of the settings and the under underutilisation of capacity. You alluded to what is effectively a bit of the trade balance in these, so that there's a little bit more inputs of intermediate inputs coming in, but you say it's relatively small. It could be very tempting for a government in Kenya to say, well, Keynesianism is back. Let's do a lot of consumption, government consumption, expenditure. But the concern clearly has to be, at some point, something to do with your trade balance. So maybe the effects are not inflationary directly, but there's something here, if you do it in an urban area, you'll have to get these inputs from somewhere. Then, there's may be something about other effects here as well. I don't know, does that seem to make sense?

Dennis Egger

Yeah. I mean, I think you're really pushing the research and the envelope in a really interesting direction. Even though this was, by many measures, one of the largest, if not the largest cash transfer RCT, it was still a small area in rural Kenya. Less than 1% of national Kenyan GDP. And it's certainly not comparable to a large billion dollar aid programme, or the central bank just printing money to pay for the fiscal debt of the country. I'm not going to try to say that you should translate what we found to these types of programs, which much other research has shown may not necessarily make sense.

Stefan Dercon

Yeah.

Dennis Egger

But I agree that a big difference between the locally big, but nationally small programme, and a nationally big programme, is this idea of the trade balance. If you print money or get, not just rural households, but urban households, or everybody in Kenya to want to consume more mattresses, this is sort of an example, and there's no mattress production in Kenya. Then whoever is importing mattresses is going to feel super happy. And you clearly want to start looking at what happens to the exchange rate. Might this lead to a Dutch disease idea where the exchange rate appreciates, and then all imports are becoming more expensive. At some point there must be a supply constraint, right? The fact that local production here is relatively unconstrained because of this idle capacity leads to these large multipliers. But this is clearly not true for all types of goods, imports, or particular type of good, but also local manufacturing. At some point, if everybody buys more of the limited manufacturing that Kenya does, prices may have to rise. And so when we take this and think about the future and scaling this up even bigger, those are the questions that are top of my mind. What happens in larger firms? What happens to supply chains, imports, exchange rates, tax revenues? These types of things that clearly governments and larger policy would care about. And maybe we get a chance to talk a little bit about how we would tackle that, but that's the answer to the next the next step in this bringing micro to the macro.

Stefan Dercon

This is really interesting because what you are saying here is that the small can be bigger than maybe we sometimes have assumed. That this is quite a big intervention. Yes, in a relatively local area, but probably bigger than many of these interventions we can do with transfer schemes. And so, it's about where is that sweet spot? How big can you go? It probably talks to something that you are working in a pretty poor area. Consumption patterns of relatively poor people are definitely less import dependent, less requirements for some of these large farms and the frontier type of things. Their basic needs are still very basic.

Dennis Egger

Yeah.

Stefan Dercon

So there's something quite encouraging here that you can also probably think about maybe scaling these things for really poor people may not have these problematic effects, but it is clear that you want to be careful with this Dutch disease kind of effects. Once you're going to go on to a national scale, and so on. There are a lot of interesting angles here.

Dennis Egger

Can I add to that I think it's a really interesting question around how to target these programmes. I think about targeting in two different ways. One is the individual level, let's try to find the most

deserving, the poorest, the people who will make the most use of this cash at the individual level. Should we target pregnant women because this might reduce child mortality. So that's the targeting at the individual level. When micro economists have thought very, very deeply about this and there's lots of great papers around this, there are clearly heterogeneous impacts for different types of households. And there's a lot to be gained from this type of targeting. But if you're a macro economist and if you're sitting at a budget office or at the treasury of whichever country, that's part of what you think about. But then the other thing you think about is targeting at the macro level. Typically, you want to target sectors where there is a lot of under utilisations that have a lot of slack. Maybe you want, after our findings, to target rural areas where a lot of slack is present. So just like with a stimulus package in a recession, maybe in developing countries, they will have to also start thinking about where do we have this underutilised capacity? There are many, many, many small entrepreneurs who can just basically, costlessly, without having to raise prices, increase the services that they sell. One direction that I'm super interested in as well is how can we operationalise this? How can we make this something that policymakers can take into account when they design stimulus packages or tax rebates or social protection programmes that target not only at the individual, but also maybe at the regional or economy level, both over time and over space.

Stefan Dercon

Yeah. But it's interesting what you've now just said, because in some sense, what you could probably say is that these poverty traps that these people may be in are spatial, and it doesn't really matter who gets the stimulus here because they are spatial. If we don't need to target the poorest people in that particular area because the multipliers work through the system. That's a way of looking at things here. But it goes back to the spatial nature of it. And, of course, that links back to the other growth theories that talk about the spatial externalities and that growth is all about these spatial externalities. And so in some sense, with this kind of size of multiplier it doesn't matter too much who you actually targeted here. Is that a fair interpretation?

Dennis Egger

I wouldn't go so far, but yeah, certainly it's like it's a counterpoint to saying the people that you don't target are going to end up with nothing or even be aggrieved. No, they'll certainly benefit from this large aggregate economic effect. It is still true that some people spend it on different things and households are different in the way they benefit. One thing, for example, and this is the only slight potential downside of what we found is that the spillovers are a bit more likely to go to relatively richer, and when I say relatively richer, these are still very poor households. But it's the local business owners, it's people who are active in the market that benefit the most. And that makes sense because the spillovers do operate as we document through these market interactions, purchases and local sales of firms. They don't operate through things that we might have expected, like people sharing the money with their relatives directly. We don't find any evidence of that. So the spillovers are really market based. And whoever is able to be there and have a business is going to benefit a little bit more from that. But it's certainly true, yes, it's coming back to this tide that lifts

all boats. When we do this welfare analysis and new, more structural work, we find that absolutely every village on average is better off. So even the control villages benefited. And there's not a single one that that might have been hurt by the programme.

Stefan Dercon

So with all you know about these villages now, is there anything that you now wish you had done a bit differently or had more data on? I'm just curious, how would you improve your own study if you got the chance?

Dennis Egger

Oh, that's a great question. I mean, it was already a heroic data collection effort, but I think one thing that I wish we had done differently, and I'm actually now doing differently in some follow on studies, is to really carefully measure not just baseline and then a year and a half later, end line, but to do monthly – even if they're shorter, even if they're with a subsample – surveys of what people do as the money rolls out. Because the bulk of the decisions in the spending are happening right when they get the big transfer and a year later people might not recall it as well, much of our recall periods in these consumption surveys are one month or seven days. And a lot of this has already been, has already happened and is outside of our measurement. What we do in the paper, in the end, is we use comparable studies to estimate what would have happened over this three to six month period that we missed and fill that in. We think that's relatively plausible because the GiveDirectly programme has not changed very much. And these areas are very similar. But in the future, I might have built that into this study. Maybe doing a shorter survey, but more frequently, and not just the one big end line. And this is, I think, something that everybody should think about who is doing RCTs. What is the right time to measure outcomes? Does it make sense to measure the dynamics of it, not just a one time, this is now, effect. Clearly, we care about the whole trajectory.

Just to flag, we're in the field again. We just started a ten year follow up last month. We did a five to seven year follow up as well. So really to understand what this is doing, not just at the individual level, but does this transform local economies in the longer term, which is what we're ultimately interested in. You also want to be doing these long run follow ups. This matters, not just in terms of our understanding of these economies and the research, but also in terms of our evaluation of these types of programmes. If effects persist for ten years, that's a very different type of positive impact than if it fades away after two. So, stay tuned for updates here on what's happening.

We also ex post started measuring a bit more of what happens to fertility decisions, child mortality, children's outcomes. This is the intergenerational transmission of the benefits, which turn out to be substantial. So, an ongoing work, we find substantial reductions in child mortality, for households which had a pregnancy in the in utero period or two years after the cash transfers. And that, of course, is another massive benefit in terms of lives saved that you cannot measure in terms of a GDP multiplier, but that may be just as important or even more important if we think about the net benefit. And so, in hindsight, I may have wanted to do a little bit more in real time. But the good

thing is that we could go back and did a full census of all households, all of the birth histories over the last ten years, capturing over 120,000 births and their outcomes. And we did verbal autopsies, to measure, in great detail, what changed and how cash lead to these sort of transformative child health outcomes. So, stay tuned, a working paper should be coming out soon here. But there are more angles to this than just the local macro economy angle. And we're still following up on those.

Stefan Dercon

Some of the people listening, besides being generally jealous of the amazing work you've been doing, and you fully deserve all the accolades and the awards that you got for it. Some of the people listening may not think they would have access to the kind of resources, the kind of facilities, the kind of teams. If you're interested in this kind of work, what are the kinds of questions that they should be looking at in this whole field that maybe shed light on side questions and so on, when you don't have all these access to resources. What are the kinds of things that people could work on and explore? What kind of things are thrown up by this kind of research that maybe with simpler, cheaper studies could be followed up on by others?

Dennis Egger

For full transparency, I joined this project as a graduate student, myself and did not have all these resources available as well. These are very large studies that require a lot of financial and time resources and big teams to collect all this type of data. They require really good implementing partners that are trustworthy, that implement the design as intended. And they're not easy studies to pull off and just as a word of caution, also, I see a lot of people doing smaller RCTs wanting to also look at spillover effects. And I just would like to say that having thought about this quite a bit and wanted to do this in follow on studies too, do your power calculations really carefully, you do need a lot of variation in treatment intensity and a really big treatment effect. So just to give you an example, suppose the multiplier is two and you give cash to 10% of people in the local economy, the 90% of people who did not get it would have to have a nine times smaller effect than the direct effect. And to be able to be powered for that you need 81 times as much sample among the non-recipients as you do for the for the recipients. So just to say, think through it carefully, if you can't, it doesn't mean that every study has to be this big.

All of our data is online and we're super happy to make that publicly available, and other people can dig in and look into that data. The other thing is you could combine multiple other studies with online data repositories and combine multiple studies into meta-analysis to get further insights at potentially a larger scale where each individual study is underpowered. The third thing is not everything needs to be an RCT. There are other great strategies, such as difference and differences, that people have been using convincingly for many years. So what I'm saying is that some of the follow up studies, for example, on multipliers in Brazil, used your traditional shift chair approach of increases in Social Security spending at the national level, which differentially impacted different regions because they had different poverty levels to start with and use that to look at than regional level GDP. And those studies don't require all these big resources. So, think creatively about using

natural variation that happens from changes in government policy and big influxes of donor money, for example. And you could look at these things too. And lastly, I've tried to give you a little bit of a flavour of the open questions of why supply so elastic? Why is there this under utilisation? Why are there so many small firms? Answers to all of these questions will help understand and put into context, why was the multiplier so big and will also help inform policy in terms of how we can overcome or how we can make the multiplier as big as it could be. And those are mechanisms questions which you can get at with smaller scale data collections or smaller RCTs. And I would love to see more work in this direction.

Stefan Dercon

This is great. This has been a great conversation. And I just want to thank you very much, Dennis, for your time and your enthusiastic explanations of all the things related to the research. I just want to refer back to the paper we were discussing: 'General Equilibrium Effects of Cash Transfers: Experimental Evidence from Kenya.' Published in *Econometrica* by Dennis, by Johannes Haushofer, Ted Miguel, Paul Niehaus, and Michael Walker. It's been a great pleasure to talk to you and I look forward to seeing the next papers related to all this research. Thank you very much, Dennis.

Dennis Egger

Thanks, Stefan, for all the great questions and the discussion.