Steven Durlauf: Welcome everyone to the Inequality Podcast. I'm Steven Durlauf. Today I'm absolutely delighted to introduce Francisco Ferreira as today's guest. Chico is the Amartya Sen Professor of Inequality Studies at the London School of Economics, and he is also the director of the International Inequalities Institute.

He has done extraordinarily broad work in studying inequality across the entire globe and is one of the leading figures now in rethinking the ways in which intergenerational mobility, for example, ought to be reconceptualized. Chico, it's just a delight to have you here and thank you so much for joining me.

Francisco "Chico" Ferreira: Thank you for having me and for that overly kind introduction. It's a pleasure to be here.

Steven: Chico, I thought that we would start with a discussion of your current perspectives on the ways in which inequality and mobility ought to be measured. I emphasize that because even though it's obvious that intergenerational mobility has profound moral salience, that doesn't necessarily answer the question of how we ought to analyze data and extract evidence about degrees of mobility and degrees of persistence. So I thought we would start with an overview of your thinking on inherited inequality.

Chico: Thank you for that, Steven. My current thinking on inherited inequality, rather than necessarily representing some original innovation, really is about recognizing the commonality between two different literatures that I came to separately.

One was the better-known literature on intergenerational mobility to which you contributed a lot, along with many other people. Then there's this literature, perhaps a little less well known, on inequality of opportunity, which draws on John Roemer, Doug van de Gaer, Marc Fleurbaey, and many others.

For the longest time, we understood very well that intergenerational mobility is about understanding the association between a particular outcome—income or education or occupation—across two generations, possibly more. In inequality of opportunity, we often thought in terms of social justice, of decomposing inequality between inequality that is morally acceptable because it's due to responsibility or effort, and inequality which is less morally acceptable or unacceptable because it's due to circumstances beyond people's control.

What has struck me, and may have struck many other people, is that there's a common ground—a very big intersection between those two—which is the idea of how much of the inequality we observe today can be predicted by inherited circumstances. So I tend to think of it as lying somewhere in between, or perhaps encompassing at the two extremes, both of those two other concepts.

Clearly parental income is something that is part of what we inherit in the family. But it's arguably much more than that. Parental education may have separate effects. Race is something that we were given at birth, biological sex, communities where we live—a variety of other

variables that may very well be correlated with parental income, but aren't fully captured by parental income.

We can expand on that by adding some of these other things to see if we're interested in the overall effect of what people inherit on their outcomes. Then I think we can move a little bit beyond mobility and enlarge it with this additional set of variables.

From the other extreme, inequality of opportunity had for the longest time big debates about what is a circumstance versus what is an effort. In some sense, you simplify things if you say, "Let's just look at the subset of circumstances that are actually inherited either at birth or by some age of consent—some early age in life—and say those things are inherited. They are a subset of the circumstances that are outside people's control." Let's look at the extent to which those variables can predict future outcomes.

The benchmark case, and if you like, the moral anchor to the idea, is we'd like these circumstances not to be predictive of outcomes. So the extent to which they are is the measure of inequality of opportunity or inherited inequality in this case that I think of as unfair.

Steven: A couple of things I would highlight in the approach that you're taking. One of them is it moves beyond the conventional economist notion, which is asking a particular question: if I raise parental income by 10%, what's the expected change in the income of children? Instead, you're asking the more comprehensive question: what fraction of the child's income or adult outcomes are being determined by circumstances that are measured at age 18 or younger?

I think that distinction is actually very important because you can end up with environments in which you get a high measure of persistence from the intergenerational elasticity of income. But in fact, it's not really particularly predictive because there are so many other things going on. So I would emphasize that I think your work does have this comprehensiveness.

The other thing is that the way you describe the circumstances also allows for, I'll say, moral disagreement as well as moral consensus. What I mean by that is if we were to go through a conversation on which circumstances we think are most morally compelling as sources of unjust inequality, I think we would start with race and gender. In other words, we have notions of discrimination that are beyond any reasonable disputation.

Then a second question is what is it about parents that we think is particularly troubling with reference to inheritance? I think you and I would agree that the effects of income and buying tutors is different than the effects of parents making different choices in how much they read to their kids, conditional on the same income. Then we could append to that thinking about social circumstances—do we think about neighborhoods and schools in a different way than the direct family?

I could go down the list, but I wanted to emphasize that the richness of your conception has to do with the fact that it's moving beyond regression coefficients and statistics to asking comprehensive implications of circumstances. And second, allowing us to then tease out what I'll

call a hierarchy of circumstances where we can identify the consensus morally versus the ones that are in disagreement.

Chico: Thank you for that. Two brief comments in response. On the second point, I do say in this recent paper where we propose this idea with Paolo Brunori of inherited inequality as this middle ground that there is a choice of what factors should society want income to be orthogonal to—or whatever advantage we are talking about. What is it that we don't want to predict?

Some people might say we don't want genetic ability to be predictive. Other people may disagree. So there is scope, therefore, as you say, for moral debate, political debate, political discussion about what ought we to try including in that set that we would like not to be predictive, and that society should compensate for when there is a dependence.

And then just very briefly on your first point, I think most people—obviously yourself, but most other people who work with intergenerational mobility—do see the measures, the IGE or the correlation coefficient or what have you, as just a measure of association, a measure of correlation. But I think some people might still see it the way you describe it: what happens if I raise parental income by 10 percent?

The IGE is not the answer to that because if it were, it would be a causal model. We know it's not a causal model—it's hopelessly misspecified as a causal model. It's a measure of association. I think most people who work in that literature are perfectly okay with that idea and familiar with that idea, but perhaps some of the readers are not. Emphasizing the statistical prediction nature of our exercise rather than falling prey to the temptation of interpreting regression coefficients as causality is another possible advantage.

Steven: A second area that I thought we might chat about is the data work you've done. You've really been involved in an extraordinary effort to expand the discussion and analysis of intergenerational mobility to the entire planet. So I was hoping you could describe for the audience the data collection exercise—the GEOM project—and say something about that.

Chico: Thank you for that. First I have to say it's a huge team project. A lot of the data collection—none of the data is primary in the sense we didn't go to any household and collect data ourselves on their incomes or their parents' characteristics or anything. It's all based on household surveys, but they are nationally representative household surveys.

We have 196 of them at the moment for 72 countries around the world, which cover about two thirds of the world's population. So there is a ways to go still, but we are covering a good amount of the population in the world.

With my collaborators—Vito Peragine and Paolo Brunori and Pedro Salas-Rojo and a lot of other people working with us—we were able to amass these data, then treat it and clean it in a harmonized way and be able to construct a vector of circumstance variables or inherited characteristic variables, with income defined in a similar way. We're able to apply different paths to decompose inequality into the predicted element—the part that's predicted by the inherited circumstances—and the residual to that.

It's been fascinating for me. I've been learning with my co-authors to see the variety of outcomes. Just one thing I'll mention—not to spend too long on this—is the techniques that we use, the statistical techniques that we use for that decomposition, are prediction-based as I say, because the spirit of the exercise is to understand to what extent these circumstances and inherited characteristics predict incomes in our case.

To do the best possible job with prediction, we rely on some of these machine learning techniques that are now becoming quite standard, like conditional inference trees and random forests, but also some newer ones like transformation trees. The interesting thing is that they generate little pictures of inequality in a country by basically picking which characteristic first splits the population, and then split the other nodes and so on and so forth.

You get this picture which is purely descriptive—again there's no causality of any kind—but it's almost like a sociological description using some standard statistical techniques, which are to me as interesting, if not more interesting, than the headline numbers.

Steven: I think that's actually one of the frontiers in mobility research: using these powerful machine learning methods. There are obviously two dimensions where they make advances. One of them is they allow for non-linearities, which I would say typically, if not universally, are part of the economist work on mobility.

The second is, as you said, the notion of visualizing the data in terms of splits into objects that are becoming locally linear facilitates thinking about interactions across the different mechanisms. The genius of the techniques is finding ways to find the interactions without running into terrible curse of dimensionality problems. So I think that's really just a very exciting area.

Steven: In thinking about or describing the findings from this truly global effort to study mobility, I thought it would be more sufficient if we decomposed it into questions. If I asked you to describe the entire planet, there's so much heterogeneity it beggars a dimension reduction. So I thought maybe we'd start with the United States and Europe. What do you think are the best-documented differences and similarities in mobility patterns?

Chico: I think that comparison has been the focus of most of the empirical attention in the literature. As we know, going back to Alberto Alesina and Angeletos and other people, there's been a lot of theory written about this with Bénabou and Tirole and many people working on this area. They typically found that, as we see in the Great Gatsby curve, the US tends to have lower mobility by most measures than most places in Europe, certainly these days.

That actually contradicts this interesting perception from the American dream—the idea that the US may have been more unequal in terms of outcomes, but the tradeoff was that it was very equal in terms of opportunities. It turns out if we measure opportunities in these ways, that's not true.

Our work is certainly not pioneering in that regard in any way, but it confirms it. When we apply our techniques and methods within this GEOM project to those countries, we find the United

States in a very similar place to what it does in the Great Gatsby mobility curve. It's less unequal and has less persistence than developing countries—a majority of developing countries, certainly in Latin America and many countries in Africa—but it is more unequal and less mobile than pretty much every European country, I think, in our data.

Steven: Stepping back to the Great Gatsby curve, I think the reason that it became so salient in public policy debates in the United States and in fact led Alan Krueger to call it the Great Gatsby curve was because of its inconsistency with America's self-image. I can remember quite distinctly when I was an undergraduate in the late 1970s that this was an idea that America tolerated more inequality of outcomes because of the equality of opportunity.

Chico: Right. I think we were all taught that, weren't we? We took that as a stylized fact, I think, for a long time without necessarily having the data to support it, but that was the view, that was the idea.

Steven: You're maybe once again being a little too modest in terms of your confirming that fact. The reason I say that is that using the more powerful methods gives you a richer conception of what mobility means, and consequently, implicit in what you're doing is accounting for non-linearities and interactions.

Turning to this idea that we want to look for configurations of circumstances that are a perfect storm to inhibit mobility—that's intrinsic to thinking about things such as poverty traps. Another, of course, is I could look at another configuration which locks in success, and I'll call that an affluence trap. The fact that once you use methods that allow for poverty traps and affluence traps you're getting a consistent message—in my judgment, this is a significant addition to the received wisdom.

Chico: I think that's absolutely right. In your work and other people's work on the non-linearity that shows up in fact in the conditional expectation function, if you like, of parental income to current income, you see those non-linearities there as these traps.

Sometimes in our trees you see interesting things as well. For example, there's often a first branch that separates a smaller group of people—an elite—from the majority of the population. Now who that elite is in terms of inherited characteristics varies quite informatively and interestingly from country to country.

In South Africa, it's just the whites. In Brazil, it's both parents went to university. In Bolivia, it's just your father having gone to university was enough. And then you see that these groups' children have very substantially higher average incomes compared to the other groups. So it's another way of seeing those non-linearities.

Steven: Might we turn to Latin America, which is one of your areas of expertise, and ask you to describe broadly the main facts and dimensions that you've written about and what the project has revealed?

Chico: In terms of this project, I've worked on many different things in Latin America. Some of my much earlier work was not so much related to inequality of opportunity—it was more trying to understand the dynamics and the factors, the proximate factors, behind changes in inequality in the region. In particular, the interesting relationship between the educational distribution and the income distribution.

Brazil, but many other countries in Latin America, had massive educational expansions in the 80s, which translated into big shifts in the educational composition of the labor force in the 90s and the 2000s. And yet we didn't see a corresponding increase—until later—in income inequality. With François Bourguignon and Nora Lustig, we would describe this as the paradox of progress. It had to do with something very simple, really, which was the fact that the returns to schooling are convex. So as you're shifting this mass of education to where returns are higher and higher, you're making the mean higher, but you're actually increasing the spread on the projection, in some sense. Obviously the returns move, but if they don't move enough, then you get this effect, which is interesting.

More recently we've been looking at inherited inequality in Latin America, which is, I think I should say, the region with the least mobility or the most inequality of opportunity. The reason I hesitate is it's always difficult to compare Latin America with Africa, in part because most countries that we have information on for well-being in Africa collect data on consumption rather than income. Obviously consumption expenditures and income are related, but they're not related linearly, as we know. So it's difficult to compare the two.

Typically measures of inequality in consumption are lower than measures of inequality in income when you have the same variables for two populations, so it's a little difficult to compare it with Africa. Also, the other reason is Africa is—at least in our data—more heterogeneous in terms of inequality and inherited inequality than Latin America. There are some countries which do better, have lower degrees of persistence, but then at the other end you have South Africa, which is in our data the world champion of inherited inequality in terms of shares.

In our latest estimate, in South Africa, 81% of inequality observed in 2017 could be predicted by inherited circumstances in the nationally representative survey—81%. There's nowhere else where you get to that.

Steven: In thinking about these differences across regions, where do you see the role of deep roots? Certainly in the economic growth and development literatures there's a very active literature that looks at very long-term causes. So what would you identify, let's say starting with a comparison of Latin America and the United States and Canada?

Chico: That particular comparison immediately, I think, sends us back to Engerman and Sokoloff's work, and also to some extent Acemoglu, Robinson, and sometimes with Johnson—the recent Nobel Prize winners—have also worked on that, though they weren't as explicitly motivated by the Latin America versus North America comparison which Engerman and Sokoloff were.

As you know, the argument there was that some combination of initial factor endowments in these countries led to some countries which had suitable land and climate to be able to import lots of slaves, whereas other countries didn't and attracted other kinds of immigrants. Sometimes, in order to attract those immigrants, they had to give them land, as through the Homestead Act in the United States, for example, generating communities where the distribution of land was more equal than it was in the hacienda or plantation structure of much of South America.

Still at a very high level of simplification, in South America, of course, slavery is the key institution on the eastern side, particularly Brazil, the Caribbean, and to some extent Colombia. But on the other side, in the Andes, you have extractive institutions that are not quite slavery but are indentured labor of the resident natives—the indigenous populations that lived in those areas and that were very numerous and well developed, of course, in the Inca empire, the Aztec empire, the Mayans and so on.

Anyway, either through the institutions of exploitation of indigenous labor or through slavery, Latin America generated, at a very stylized level, these incredibly unequal distributions of land and wealth in the beginning, with a mass of dispossessed people that had basically nothing—and in fact, not even a right to the fruits of their labor.

Whereas in northern North America, people were given land, they weren't as exploited, their labor wasn't coerced, and as a consequence, Engerman and Sokoloff write, suffrage expanded more quickly. People voted, and when people voted, people voted for schools, and so education and literacy spread more quickly. So they paint this broad picture where initial factor endowments and composition of the distribution of wealth lead on to different institutions and then to development.

Of course, the role of the United States is an interesting one because the United States did have two halves in this regard. It did have slavery with the plantation systems in the south that were very similar to those of northeastern Brazil, but of course it had the north, and the north prevailed. But it's interesting, sometimes particularly these days, to think of the US being in this hybrid position somewhere in between in the Engerman and Sokoloff model.

Just to say, bringing it back to Latin America, we did another project I've been involved in recently—this Latin American and Caribbean Inequality Review—where we brought together a large panel of scholars to look at very different aspects of inequality in the region. There we had one chapter by Felipe Valencia and Francisco Eslava on the origins, the historical origin, the colonial origins of Latin American inequality.

They use modern econometric techniques to try and identify causality by looking, like Melissa Dell did, for example, at the borders of where certain regimes of slavery or other extractive institutions worked, and trying to identify modern consequences of those differences. Through a variety of mechanisms—this is not my work, so let me not try to misrepresent it—but they do find, the bottom line I do know, is that they find very strong consequences, particularly of slavery, in terms of predicting both worse outcomes on average and more inequality today.

Steven: I think that work is obviously an important complement to understanding American inequality between blacks and whites. Its salience possibly couldn't be higher. There is a change in the narrative of the American government with reference to the consequences of slavery and Jim Crow in contemporary inequalities. So the work that you describe is a perfect example of why social scientists have to weigh in and measure these things, because the implications for public policy of the narrative are so important. So I appreciate you bringing that up.

Steven: Might we turn to Africa? And I should mention that you were the chief economist for the African region at the World Bank for a number of years, so it's an area that is very close to you both intellectually and personally. What would be the broad stylized facts? You've already mentioned two of them, and I was hoping you could elaborate on the heterogeneity within Africa and what may be known about that, as well as the uniqueness of South Africa.

Chico: Obviously being the poorest region in the world, it also is the most data-deprived region in the world. So the number of countries, if you look at this GEOM—GEOM for the people listening is this Global Estimates of Opportunity and Mobility database that we've set up and put out on the web—the African map is much more sparsely populated than the Latin American map, even though we do have a number of countries there which we got largely with help from the World Bank.

But to the extent that we can say something, it seems that West Africa is less unequal and less opportunity-unequal than either Eastern or Southern Africa. Eastern Africa lies somewhere in between, and Southern Africa is at this extreme.

The most interesting thing, actually, about South Africa, which I'd like to mention briefly, which shows to some extent the limitations also of what we do, is that in 2017 it just so happens that the South African statistical institute did a great job in a piece of innovation, which was to provide the regular sample they have but also to oversample the rich in an attempt to correct the problems with coverage of household surveys at the top, which we are all familiar with.

When they oversampled the rich and got arguably a more precise estimate of the upper tail of the distribution, inequality went up, but inherited inequality went down as a share. And why did it go down as a share? Well, there are many possibilities, but one of them is we use very coarse variables. We have race, so for all the rich people—the vast majority of these rich people were white—so that's not differentiating between them.

We use "have you gone to university?" which is a huge difference in the broad population. Amongst these people, the question is, did they go to an excellent university in Cape Town or Johannesburg, or did they go to a much worse one? This granularity and fineness of the characteristics is something we're not picking up in our data. In that regard, parental income, when you have it, is actually better because you do have the granularity a little bit.

So there's still, I think, a long way to go and a lot of work to do to combine the greater variety of variables that we have and the information with parental income when it's available. I think you get an even better picture then.

Steven: Might we turn to your work—it's almost meta-science. In other words, you have a very interesting paper talking about the evolution of thinking about inequality at the IMF and the World Bank. So what were the main themes?

Chico: That's interesting. I never thought I'd be asked about that paper because it's a marginal and slightly navel-gazing paper that I only wrote because I was asked to write it for a book project on multilateralism, and I thought that was important. It was a project that actually began supported by the World Bank and was begun by two colleagues, Vijayendra Rao and Michael Woolcock, at the World Bank at that time.

The ask was: how do the multilateral institutions, and in particular, I narrowed it down to the Bretton Woods institutions—the World Bank and the IMF—how do they think about and interact with the notion of inequality?

I think the two main findings from looking at some bibliometrics, but also just a little bit of the history of the influence of different people and the nature of debate in these two institutions, and the extent to which the topic was permitted to influence headline reports like the World Development Report at the Bank or the World Economic Outlook at the Fund—my conclusion, probably unsurprising, was there were two major factors.

One is actually the salience of the concept in the economics profession broadly, because a lot of the people at these two institutions get their PhDs in leading Western universities, for better or worse. If they studied in the 1980s when nobody was talking about inequality much, except perhaps Steven and a few others—but it was not a central topic, and macroeconomics was completely going somewhere else completely—then people didn't think of that when they came to work at these multilateral organizations.

Before that, when there was more of an emphasis on poverty, at least, and inequality, and then following the 90s, Tony Atkinson had that nice paper in the Economic Journal called "Bringing Income Distribution in From the Cold," which you may remember, which was written and published in 1997. It talks about the emergence of a whole literature on that, and that was becoming more salient and more important. Then, of course, all the way to Piketty and Stiglitz writing bestsellers that capture the imagination of a public much beyond academia.

In those times, people in these institutions think it's respectable and okay to worry about those things. So there's really an influence that flows from academia to the thinking of the Bank and the IMF, and it should not be underestimated by academics. That's whether or not the academics have anything to do with—many people go and visit the World Bank or the IMF, but it's nothing to do with that necessarily. It's really to do with dominant themes in the profession.

And I say the economics profession because, again for better or worse, it is the dominant intellectual—it provides the dominant intellectual framework in those institutions. There are a few sociologists here and there, anthropologists and so on, but they are a minority, and it is the economics profession that dominates in shaping the view of those institutions.

The qualifier to that is that, though they are staffed by a technocracy, they are politically driven and politically influenced institutions. Speaking more about the Bank, which I know a lot better than the Fund—I've worked at the Bank for many years—there's a very strong power of the president. It's a very presidential institution, and the president appoints his or her—only been his so far—vice presidents and chief economist and so on.

So there is some amount of censorship occasionally here and there, but that's actually very light. Most of it works through self-censorship—another question that's relevant these days. People feel the way the wind is blowing and adjust what they say and what they work on and what they do. These are very measurable effects. If you look at the bibliometrics, for example, as we did in that paper, there are clear cycles that correspond to different presidents of the Bank.

And of course, the presidents of the World Bank are, as the listeners may or may not know, effectively always appointed by the US, even though that is not written anywhere. But it is an agreement between the US and the Europeans which, since 1944, has held that basically the Europeans get to pick the head of the IMF, the Americans get to pick the head of the World Bank. So flows, ebbs and flows, on how or whether inequality is something that you can talk about in the US do affect what happens at the Bank and the Fund, albeit with a lag.

Steven: Chico, I cannot thank you enough for this absolutely delightful conversation.

Chico: No, my pleasure. It was really wonderful. Thank you so much for having me and for spending this time with me.

Steven: The Inequality Podcast is a production of the Stone Center for Research on Wealth Inequality and Mobility at the University of Chicago. I want to end the podcast with thanks to the people who really make it happen.

First, I want to express deep appreciation to our producer and engineer, Shane McKeon, who oversees every aspect of the process of creating these podcasts and really does just a splendid job.

Second, I'd like to thank our Assistant Director, Nina Gray, for production oversight and the role she plays in bringing the podcast to fruition.

Finally, I'd like to thank Grace Kolovo, who's the Executive Director of the Stone Center, who basically does everything in terms of making the center work.

You may get in touch with us at StoneCenter@uchicago.edu. Thank you so much for listening.

This transcript was prepared with the assistance of AI.