Steven: Hello, this is Steven Durlauf, and welcome to the Inequality Podcast. Here at the Stone Center, we take pride in thinking of ourselves as fundamentally interdisciplinary in our approach to inequality. Today's guests epitomize that scholarly ethos. As you will hear, their approaches to the study of inequality span the use of tools from economics, to physics, to mathematics, and qualitative methods. Today, we will be focused on understanding the interdisciplinary insights that can be brought to bear to understand wealth inequality. Why is wealth inequality so important? In some ways, the answer is self-evident. Differences in wealth determine differences in the capacity to live flourishing lives, and wealth is very much a nexus that defines the ways that parents can influence children. Wealth inequality demands an interdisciplinary approach for several reasons. It is a case where mathematical models benefit immensely from interdisciplinary perspectives. Within physics, models have been developed primarily in the area called statistical mechanics, which has to do with the evolution of large populations of objects that are heterogeneous. So if physics might be interested in the collective behavior of a group of atoms, a social scientist is interested in the collective behavior of a group of people. Our first guest is Alberto Bisin. Alberto is a professor of economics at New York University. Alberto is remarkable for the breadth of research that he has done. He has made fundamental contributions to mathematical economics, to applied economic theory, to the study of culture, to the study of social economics, and to the study of economic history. For today, we'll discuss some of his work on wealth dynamics—in other words, understanding the evolution of the distribution of wealth over time. One of the interesting features of Alberto's work is the way that he has drawn from tools developed in the physics literature, and nobody is going to be better positioned to talk about the relationship between tools in physics and economics than Jean-Philippe Bouchaud. Jean-Philippe is one of the true polymaths in the world, and by that I mean he has made fundamental contributions spanning mathematics to physics to economics. We will hear from Jean-Philippe shortly, but first, here's my conversation with Alberto Bisin.

Steven: Alberto, I thought we'd spend most of the conversation today on wealth inequality, where you've been one of the leading scholars in thinking about the foundational theoretical underpinnings of inequalities. I thought for the audience it would be useful to start by you giving a sense of the stylized facts of wealth inequality that represent the background of your research.

Alberto: So in some sense, when I started working on this, I started from something that Paul Samuelson stated as the Pareto Law. Vilfredo Pareto, this Italian economist of the beginning of the 1900s, had done a lot of empirical work and theoretical work on wealth inequality. What Samuelson quoted as the Pareto Law is something that is obviously not true. It's not a law, first of all, and certainly it's empirically

not true. So that's really where I started—trying to understand what about this was wrong and what was actually going on. The law is something like: at all times, the distribution of income—and he means also wealth—remains the same, and neither institutional change nor egalitarian taxation can alter this fundamental constant of social science. This is clearly wrong. I was working with Jess Benhabib, a colleague of mine, and then with my student Luo on this, and we started really trying to understand what is actually true. How can we change the distribution of wealth? So it's not quite an empirical stylized fact, but it's almost like the opposite of an empirical stylized fact, and that's where we started. And then there's a bunch of other actual stylized facts that we thought we would try to understand better theoretically. One is, of course, the fact that however you measure it and really wherever you look, Pareto was right in one respect: the inequality of wealth is much larger than in the inequality of income or earnings. That's a fact. It's a very interesting fact, and in some sense, this is really what motivated a large part of our early work. Later on, once we got more into understanding the distribution of wealth, other stylized facts have been found interesting and sometimes difficult to understand and explain. They're both related to wealth mobility across generations. If you look at wealth mobility from parents to kids, you see more mobility than maybe somebody like me had as a prior—there is some mobility, thinking about the US, but not just the US. And then you look at this very long-run correlation between generations, five hundred years, and you see intergenerational effects that last this long. And that's hard to understand and fit together. So the second part of the work that we were doing was trying to get at this issue, which is still one we haven't understood completely, but are trying to.

Steven: I think one aspect of when people hear the name Pareto and wealth, they put them together and think about fat tails in distributions. Can you say a bit about that?

Alberto: Basically, what Pareto did was look at some data—very little data that he had on wealth and incomes across countries—and notice something, which is that basically the tail of the distribution is relatively fat or thick. We can try to be more specific about what this means. What he did then was come up with a functional form for distributions that has this property, which is now called the Pareto distribution. The idea has to do with how fast the distribution declines in the tail for the very rich. A fat-tailed distribution is one that declines not exponentially, but geometrically, much slower than exponential. This means that, of course, there is a lot of wealth in the tail. The fraction of wealth owned by the relatively rich is larger than you'd expect from a normal distribution. This is a property of the tail, a limit property of the distribution. Of course, nobody is really richer than some number; the distributions we see are bounded. So this is a theoretical property that is very interesting to

think about generally. It's never something you can test empirically very precisely, but it's a property of distributions that's important theoretically because you want your model to generate wealth distributions with a lot of wealth in the tail.

Steven: Can you outline your main theoretical contributions? I know there are a lot of them, but in terms of distilling them into perspectives and mechanisms, what would you want to put on the table?

Alberto: We started from this point, which is that distributions in this sense—once again Pareto was right—it's not true that the distribution of wealth is the same and untouchable, but it's true that it's got a bunch of properties that are very general across space and time. One of these is the fact that the tails are relatively fat. How do you explain that? When we started thinking about this, the consensus in economics, or the class of models that people were using in economics, were models where the fatness of the tail in the distribution of wealth was coming from the fatness of the tail of the distribution of income. People were writing down models where basically labor earnings had a distribution which was very fat, so there were a lot of people making a lot of money income-wise, and this translated into fat tails in the distribution of wealth. We thought these models had problems because to fit the distribution of wealth, it looked like earnings had to be very, very fat. We started writing down models and tried to understand this, and of course, physicists had done a lot about these things. They're obsessed by power laws, not necessarily in a negative way, but they've studied this a lot. Basically, there's a bunch of theorems out there, and the theorems suggested that abstracting from all sorts of stuff, those models which had stochastic earnings inducing stochastic wealth implied that the fatness of the tail of earnings was the same as the fatness of the tail of wealth, which is counterfactual. When you were going to look at these models empirically, people were making up extreme income shocks that they were not observing, that were needed to fit the theory. Because implicitly what they needed was a fatness in earnings which was the same as the fatness in wealth, but the fatness in earnings is not the same in the data. So you kind of needed to make them up. In fairness, a lot of the data on the distribution came after Piketty resuscitating this and working on the data. A lot of new data came from there and back then, the data wasn't there. In any case, we started thinking that these models where wealth was driven basically by earnings were not satisfactory. And we started working with models where there was a stochastic element to the rate of return on wealth. The idea is that the rate of return that people have on their own wealth is stochastic. Some people get higher returns than others—not necessarily very persistently, but they do get higher returns. Think of entrepreneurship if you want, but we were working on very abstract theoretical models. And we've shown in a series of papers that these models generate fatness of the tail in wealth, which is much higher than the fatness of the tail in earnings. In fact, it's independent of

the fatness of the tail in earnings. Basically, in the limit, the fatness of the tail in wealth only depends on the rate of return on wealth.

Steven: I think for the audience, part of what is nice to understand in this work is the blending of ideas that came from the probability theory literature and from the physics literature with economics. What I mean by that is there was an idea that if you think about the level of something, and there's a shock—the rate of return, so some random return on my wealth today is one plus that thing next period—what's happening sequentially is you're getting the products, the multiplication of the shocks over time. What had been understood from the mathematics literature is that multiplicative shocks generate thick tails. Where social science comes in, and what's so essential about the Benhabib–Bisin work, is it shows that the mathematical logic isn't by assumption. It is derivative from thinking about the decisions of individual actors in the socioeconomic environment. I emphasize this because I think it's one of the examples where you have very constructive engagement between the probability theory literature, work by physicists who have thought about power laws in physical contexts, and economists.

Alberto: Absolutely right. I was going around when I explained this, and in some sense, I didn't say the key, which is exactly what you said—that the physics and mathematics literature is teaching us that the reason why it's the rate of return to wealth that is important is that it enters multiplicatively, whereas earnings enter additively in the stochastic process of wealth. There's only so much you can do by adding stuff; there's so much more you can do by multiplying. That's exactly the point on the relationship with physics and mathematics. I think that's interesting because we went both ways. We started thinking, OK, so physicists have written down a model of how wealth moves—equations like "wealth tomorrow equals something times earnings plus interest rate times wealth," so basically a stochastic process of wealth where there is no micro-foundation, no choices of individuals. Of course, wealth is the outcome of a process where people consume and save and invest the wealth into a portfolio. Economists, of course, like to do this. So we started thinking we need to put economics into these equations. So we wrote down a relatively standard macroeconomic model that any economist can understand—a very standard one where agents choose how much to consume and how much to save, what they consume and what they save depends on their wealth, and then they look at the process, they invest, and so on. And we showed that in the end, the mechanism was in a way the same as, or very related to, the one that the physicists had in mind. But of course, there was a fundamental contribution of what saving means, what does it mean when do you save, all these things that economists like. So we started thinking that the economics on top of this would change in an interesting and dramatic way the conclusions that the physicists had. And we concluded that it didn't change that much, in a way. The

theoretical core is right. The theoretical core is that linear equation that the physicist writes down. What economists put on it is very interesting—at some point when you start playing with this, for instance trying to understand the effect of taxes on the distribution of wealth, from an economist's point of view, it's crucial to go back and understand decisions and the access to the world. The question is how much of the accumulation comes out of saving, how much comes out of bequests, and all these issues. The core of the argument in the end goes back to: it's multiplicative and not additive, and that's where the key point is. The taxes act in different ways when they act on the multiplicative component or the additive component. So the economics is fundamental—I'm not minimizing it—but in the end, a lot of these results come from this basic intuition in physics, which is very important.

Steven: And it also ties together the two Paretos you identified. The successful empirical claim is the thickness of tails, and there's a certain commonality in mathematical structures because, after all, it's natural to think that a multiplicative accumulation of rates of return is going to be a source of universal thickness. On the other hand, the idea that the shapes themselves are invariant with respect to policy—which is empirically false—is exactly where you have to have the economics to do the counterfactuals or to think through what the consequences are. The work that you have done is one of the most successful examples of the integration of different perspectives, but also, I think, speaks to the importance of constructive collaborations between physics and economists. The extent to which such collaborations have not been successful, I think, is when there's been a view by physicists that the economics is not important and that simply the stochastic processes suffice. But what matters is when it has worked, and this is exactly when it's worked, in other words, the integration of what our advantage is which is modeling behavior, with the knowledge of certain generic properties of particular stochastic systems.

Alberto: I agree absolutely. I think it worked very well. I'm really very happy about it—the collaboration worked very well.

Steven: What would you say are competitive explanations?

Alberto: A couple of things. As I said, this distinction between whether it's earnings or stochastic returns, my understanding, though I'm biased because I'm on the stochastic return part, is that stochastic returns kind of won that one, though I don't think it's at all agreed upon by everybody. But I want to go back to what I was saying before, which is that the stochastic return story is really about the tail, and the tail is very relevant. We very often discuss the top one percent—that's what we care about, and we should. The inequality is there, and it's important. But there's also all the rest of the distribution, and in all the rest of the distribution, the

earnings structure is really very important, especially at the bottom, but not just at the bottom—also in the center. The model with stochastic shocks really is a model of the tail. It works well in the tail, and it's important for the tail—people who have wealth to invest are in the tail. So the competitive explanation regarding earnings is really very important. We should go back to that and think about it seriously to understand the whole distribution. That's the first point. The second point is on the determination of the rate of return on wealth. Does it depend on wealth? For instance, is it really true that the rich get higher returns on their wealth? It's most probably true. There's a lot of evidence, both direct and indirect, in the sense that when you model returns that depend on wealth, you do better in explaining the tail. So that's most probably true, but we need to understand where that comes from. Is it just increasing returns to having a lot of wealth, which lets them invest better? Or is it just that, of course, they're less worried about risk, so they can invest in more risky assets and get higher average returns? There are all sorts of questions on the determination of the rate of return to wealth. In the work that I've done with Jess, the rate of return is exogenous—it's something we take, we put there, and we estimate the properties. But there's a lot of modeling to do there. That's one of the theoretical issues that's open and very interesting.

Steven: Could we focus on one of your papers, "Heterogeneous Dynasties and Long-Run Mobility" both because I think it's a wonderful paper and because it draws very interesting relations between wealth dynamics and intergenerational mobility?

Alberto: The motivation behind that paper is that when you look at social mobility, you find some substantial mobility from parents to kids when you look at the wealth elasticity between parents and kids, but then you find a lot of persistence of wealth over generations. There is some work on very long horizons—Greg Clark, who studied England from the 1800s up to now, using rare names. There's also work on Florence. Florence was the only place in the Renaissance that did what the Italians call the catasto, which is a registry of land ownership, basically a measure of wealth. Most of what historians look at when they think about mobility and wealth is Florence—that's the data we have. Lots of people have gone and looked at that, and they really find intergenerational persistence of wealth that lasts four hundred to five hundred years. What I and other people have always thought about when thinking about long-run persistence is that stochastic returns on wealth are the driver. If I am a person that has high returns to wealth and I am able to pass this on to my kid, then he'll have high returns on average, and his kid, and so on. That's the model we have in mind. But that model doesn't work because the passing-on multiplicatively declines financially very fast. You're not going to fit the fact that there is some elasticity—my son is in fact going to be richer if I'm richer, but not by that much—and then when you multiply this for various generations, it goes down. You should

have no effect over five hundred generations, but in fact, you still do. So that's what we try to do in this paper. We go back into stuff that physicists know well and write down models where there's a latent factor that explains the persistence of differences in rates of return on wealth. We don't model that—that's going back to what we were saying before about where we want to do more modeling. Aristocrats somehow must have had an institutional mechanism that protected them from losing their wealth.

Steven: There's a very broad set of research areas in economics that look for "deep roots," which I'll define as events in the relatively distant past whose footprints you can see today. Italy is actually one of the standard examples: the institutions of different parts of the peninsula five hundred years ago are reflective of differences today. A different example would be arguing that the slave trade in Africa has contemporary consequences for the economic performance of different countries. Nathan Nunn, to drop a name, would be a leading scholar there. I wanted to ask about your thinking with respect to deep roots and wealth inequality. You gave a little bit of it already with the Greg Clark idea that there seems to be something about dynasties in certain contexts that are extraordinarily persistent.

Alberto: I have been thinking a lot about these issues. Of course, as an Italian, the issue of what we call the "questione meridionale"—the fact that the south and the north have forever had very different GDP per capita levels and other socioeconomic activity—is really a fundamental question. I can't say that I got to it through the study of wealth distribution; I got to it in different ways. But you're absolutely right that they link very much. There are all these examples of divergences in growth between countries, across regions, that are really fascinating as a topic. This is, of course, something that historians study—the great divergence between China and the West, the long divergence between Islamic countries and the West, the north-south of Italy, and Haiti and the Dominican Republic, and Argentina and the US—all these situations where you have countries or regions that were at the same level at some point in time. That stuff is interesting. This is one of the reasons why an Italian becomes an economist, at least in my age—the north and south issue. What I've tried to do is connect with historians, because this is really hard. It's not hard to work with physicists, but in my personal experience, it's much harder to work without historians—you really need historians to understand this stuff. My issue is trying to write down growth models that can understand these divergences. The kind of model I've been working on, this is work with Thierry Verdier at the Paris School of Economics, these are models of the interaction between elites, different groups of elites, and the civic societies of workers in some cases. This goes back to what you were saying before—these phenomena seem to be linked to the distribution of power between groups in the polity. A natural way is to distinguish elites from the rest of us, or different kinds of elites from the rest of us. There's a lot of work

to be done there, some of which I'm doing, though it's a lot for my ability to work. There are very many projects I'm involved in, maybe too many, but it's really fascinating. It gets to the issue of the roots of growth phenomena and to the issue of the roots of inequality, because these are fundamentally examples of the creation of inequalities.

Steven: Alberto, I can't thank you enough; just a wonderful, delightful conversation as always whenever we have a chance to talk.

Alberto: Thank you, Stephen. This was a lot of fun. Thanks.

Steven: As hinted in the conversation with Alberto Bisin, there seem to be some deep links between models that have been developed in the physics literature and the dynamics of wealth inequality. Few understand that better than Jean-Philippe Bouchaud, a mathematician-physicist by training who has also been a pioneer in econophysics, which refers to the broad effort to use tools from physical contexts to understand social phenomena. Jean-Philippe is the chairman of Capital Fund Management. He is also a professor holding the Bettencourt Innovation Chair at the College de France, and he is a member of the French Academy of Sciences. I firmly believe that economists have much to learn from physicists, and I really can think of no better person to make those connections than Jean-Philippe Bouchaud. Given the subject matter of the podcast, I was hoping that you could give us an outline of the work that you've done in wealth dynamics.

Jean-Philippe: Yes, sure. It's something that I've been thinking about for many years. My first work on that subject goes back to around 2000, with another physicist colleague of mine, Marc Mezard. From then, we've always tried to think about inequalities in a physicsy way—that is, trying to identify mechanisms, not think too much about rational agents and equilibrium theories, but just have descriptive models of how the inequalities that we observe in the world can arise from minimal models. Stochastic models, models where we can parameterize skill versus luck if you want, models where we can test the influence of taxes, models that help thinking about how much inequality is fair, which is something I've been pretty obsessed with since the beginning. It's clear that inequalities seem to be at a level beyond which it becomes unethical or unsustainable or unfair, but what is this level exactly, and can we try to pin down what makes inequalities unacceptable and what makes inequalities acceptable? That's in a sense the background of what I've been trying to understand.

Steven: What stylized facts about inequality do you think your models have been particularly successful in elucidating?

Jean-Philippe: Well, I think that since Pareto in 1900, where the

distribution of wealth was found to be first of all universal, pretty much universal with a high degree of stability across countries, which is something that Pareto himself was surprised by, but also very fat-tailed, with these cute distributions that have become one of the main topics in statistical physics in the last thirty to forty years. Traditionally, in physics, people were more naturally attracted to distributions like the Gaussian, with very narrow tails. And it's surprising that Pareto was in some sense the first scientist to find uses for these fat-tailed distributions. So I think that really captured the fascination of the physicist—this appearance of fat tails in many social phenomena, actually wealth but also others, like crashes in financial markets or economic crises, or distributions of firm sizes or city sizes. And these sound very much like things we now encounter very often in physics, in complex physical systems as well. So to answer your question, I think this stylized fact of the very cute fat-tailed distributions is perhaps the thing that physicists are well-trained to try to understand and interpret.

Steven: And so your work, and I'm thinking partially the work with Roger Farmer but of course a broader body of work, what mechanisms would you identify that are understandable to the layman in terms of generating the fat tail?

Jean-Philippe: I think that the main driver is multiplicative growth. Most economic phenomena lead to an increase of wealth that is multiplicative—we always speak about interest rates or rates of return. This mechanism leads to an exponential growth, albeit with fluctuations, and these fluctuations, I think, are absolutely relevant to understanding, in particular, the far tail. So I think the mechanism by which these inequalities appear is this multiplicative nature of growth. And actually, I spoke about city sizes before, and it's the same with population size, which increases exponentially as well, with fluctuations. This is the basic mechanism for fat tails. And the mitigating mechanism, if you want, is redistribution; it's how once you get wealthy, part of your wealth gets redistributed through various channels. Some of these channels can be taxes; other channels can be exchanges, and the way the exchange network is set up is going to be important in determining the type or the amplitude of inequalities. But something also that we discovered with Roger is that financial markets are a way to kind of elicit this exponential mechanism, this growth rate, and this randomness. So we have a model where agents disagree, and they disagree on, if you want, the fundamental value of an asset, and some are temporarily smarter, but there's no real skill in our mind. So the model—and that's what makes it interesting—is that just by luck in a sense, some people seem to get it better than others during a long period of time, and during that long period of time they get extremely wealthy. And then they lose track, and other people get wealthy, and so there's a kind of constant renewal of the wealthy versus the poor, but this is really mediated by this exponential growth. Our toy model of disagreement in

markets that allow people to express their disagreements is a kind of way to incarnate this exponential growth that I was speaking about before.

Steven: Is it appropriate to think that people with higher wealth have access to rates of return that are on average higher than the poor, or is it better to think about a non-ergodicity in the system, in other words, that the sequences of shocks or the realizations of the shocks are not averaging out over time, and so you're ending up with very different, permanently different trajectories?

Jean-Philippe: Well, I think both can be correct. In the very first model that we proposed, there was no real non-ergodicity or long time scale especially involved, just some amount of noise. In that first model there was really no skill at all, and in the model with Roger Farmer, it's also a model with no skill, just luck. Indeed, if you add a level of skill-of permanent rate of return that is intrinsically higher for the wealthy, which there is probably an element of, then the inequalities tend to be even higher than the type of inequalities that you get through sheer luck. But the two are very hard to disentangle. To come back to your question of non-ergodicity: in the case where skill becomes dominant compared to luck, then you can really enter in a truly non-ergodic phase where wealth gets more and more concentrated over time in the hands of a very few people. Whereas in the luck-only model, after a while the wheel turns, and although there is a stationary level of high inequalities, it's not the same people generation after generation that remain the rich people. So there's a kind of variety of different possible behaviors that these types of models can generate, and I think it's quite interesting to get that at the cost of a minimal investment in mathematics, if you want, and also of the sophistication of the models.

Steven: One of the key major contributions that you have made in thinking about wealth inequality, from my vantage point, has to do with almost a philosophical difference in modeling from the conventional economic one. Let me give you my stereotypes of the approaches and see what you think of them. Economists typically start by specifying an environment: these are the preferences of actors, this is how they form beliefs, these are the constraints they face, this is how wedefine an equilibrium, and then see whether or not it produces properties that are similar to certain facts. The approach you've described, I would say, starts with facts-fat tails, certain types of discontinuities almost between booms and crashes-then asks what sort of stochastic processes, what sort of mathematical systems, generate these and generate them in a robust fashion-in other words, intrinsic to the logic of the systems as opposed to an artifact of the model's specification. I also think that's one step back from asking what are the behavioral interpretations of the stochastic processes. I put that on the table as a way to think about the differences in vantage points, and I emphasize that because of the idea-central in

your philosophical writings about social science—of moving from clear definitions of measurements of interesting dimensions to systems and mathematical processes that have the universality property you describe, and then going to the micro foundations, whereas we economists tend to start with the micro foundations.

Jean-Philippe: Absolutely, absolutely. And this is very much also related, I think, to a more general difference between economists and physicists. Physicists start from observation and build what's called phenomenological models—models that are not necessarily micro-founded, but that allow, little by little, step by step, trying to capture mechanisms, going from simple toy models that can then be expanded and more sophisticated, and at the end calibrated to data. So I think the type of models I've been looking at are not supposed to represent reality very faithfully. I think they're there to shed light on these universal mechanisms that you've emphasized yourself. I agree it's a very different philosophical approach, but they should be complementary in the end. The Occam's razor type of argument is: well, if we can get universality without assuming that much about the behavior of agents—just random phenomena that get amplified or mitigated by exponential growth rates or exchange taxes and so on—this can bring one quite far already. Is this good enough for thought? Is it capturing something that's real? Or do we need to impose on the model some kind of economic environment, some kind of rationality, some kind of equilibrium conditions? Or is that not needed?

Steven: One of the important ideas that you brought up is the distinction between thinking about inequality dynamics as driven by randomness, stochasticity—shocks hitting individuals, luck versus differences in something more persistent, such as skills. And the reason I emphasize that is that in a number of models—and I'll say this as a compliment to the work you're doing—rates of return for the more and less affluent can differ, with mechanisms that would include access to the market, access to information, the whole process of network formation if we're thinking about neighborhoods and the way that information is communicated across individuals. So I think a useful area to complement the work you've done would be the work on social phenomena, which has emphasized the role of segregation in creating rates of return differences across different parts of the population.

Jean-Philippe: Yes, I totally agree.

Steven: And I raise that because I emphasize this is a new research area and I think there are many exciting things to do there.

Jean-Philippe: Well, just to bounce on that, the very first model that we looked at in 2000 already had a variety of network effects that we could look into. And clearly, you can imagine these types of networks representing some kind of social segregation, or the type of social

networks that we know about with very heterogeneous degrees, that can lead to enhanced inequalities compared to a much simpler type of exchange network where everybody exchanges with everybody else. Which might not look very realistic, but if you think about taxes and redistribution, in effect, that is what it is—everybody contributes to a common pot, and this is redistributed over the whole population. This can be shown to lead to minimum inequalities within this set of models. But as soon as you have strong heterogeneity in the network, you do generate these extreme inequalities.

Steven: I thought maybe we could talk about some of the work you've done in financial markets. It doesn't directly speak to inequality, obviously, but issues of heterogeneity in rates of return and access to financial markets and the like have deep inequality implications. I was hoping you could identify some of the key findings you've developed.

Jean-Philippe: Well, this could be a long conversation because there's a lot. Let me maybe focus on the ones that are directly related to inequalities, which I think is a sub-part of things that can be very interesting from a general point of view. What we've been working on is the fact that we don't believe in efficient market theory, and so we've tried to identify various channels through which the efficient market hypothesis falls to pieces. We've worked a lot on price impact—the fact that whatever people do has an impact on prices, even if people are irrational and have no information. The fact that they trade in markets generates volatility, and this volatility is not a short-time phenomenon—you do impact markets significantly for long times, even with random trades. This is related to some very new work by Gabaix and Koijen on what they call inelastic markets. But as I said, this is a long story. We prefer calling it the order-driven view of markets. Maybe we can set that aside in detail for a moment, except for the fact that all these ideas lead to a possible way of explaining why markets are so volatile, why there is excess volatility in financial markets, and this excess volatility is, I think a driver of inequalities. Because again, referring back to these toy models that we've been discussing from the beginning of our conversation, the more volatility there is, the more uncertainty, the more noise, the more the lucky ones will be benefiting from that noise. So the idea that there is excess volatility and markets are inefficient is also a very strong contributor to wealth inequalities. I tend to think that financial markets are maybe ninety percent luck and ten percent skill. So it's not something that should be neglected if one wants to understand inequalities. And actually, this is the point of the model that we set up with Roger Farmer, although initially it was not meant to be a model to understand inequalities at all. It was a model to try to understand how people can disagree in the long run. That's why we coined it a quasi-non-ergodic model, where people—although they are trying to guess from empirical models the value of some fundamental probability, if you want—the fact that their assessment of probability

feeds back on the value of probability itself leads to a very interesting and complicated dynamics that we've been able to decipher. This was mostly the idea of trying to build a model where people could disagree forever, because of a kind of feedback of beliefs into reality—the fact that people believe something may in fact lead to different realities. But then at the end we realized that plugging this model into a market where people could exchange their views led—to our surprise, because really it was not meant to be a model of inequalities—to these fat-tailed distributions of wealth. And we were very happy because I think there is some truth to it. Again, the model is too simple to be taken at face value, but the fact that large disagreements lead to inherent inequalities is an important message, I think.

Steven: If I've understood correctly the tenor of your argument, there's another mechanism which I find fascinating, which is that a dimension of efficient market theory is the idea that information is aggregated by prices and therefore the idiosyncratic information I have gets, in some sense, lost from the process of trading. What you find in the work with Farmer is the capacity for people not to converge based on their information sets—which itself means that my capacity to benefit from better information is, of course, only finitely true in the long run. Is that a way to think about the finding?

Jean-Philippe: Well, you see what's funny about the model is that there's a strong relation between the appearance of these inequalities and the fact that people can't converge on the true value, because in the model the true value is determined by a kind of wide average of beliefs—that's what we assume, but we could take something more complicated. The market price reveals a wealth-weighted belief, and the fact that wealth is so strongly heterogeneous leads to a permanent discrepancy between these two averages. So in a self-consistent way, people disagree—some people are closer to reality than others, some people have more money than others temporarily—and this leads to wealth inequalities, which in turn prevent market prices from giving the right answer, because the market is dominated by the wealthy ones. So the strong inequalities of wealth prevent a correct sampling of beliefs in the model.

Steven: Could you say a bit more about efficient market theory and the dimensions where you think the theory is not adequate?

Jean-Philippe: Well, I think the dogma I tend to believe in—if efficient market is one dogma I believe in, there's another one, which is what Fischer Black used to say. He wrote this beautiful paper in 1985 called "Noise," where he says that an efficient market is a market where price and value differ by a factor of two. Of course, prices are not a factor of a hundred away from value. But a factor of two is pretty large—there's a lot of space between 0.5 and 2 around

what would be the fundamental value of one. And the timescales associated with these discrepancies between value and price are very long. The idea is that if you look at markets on the very long run, on century timescales, then of course prices and value tend to track each other, but in a very loose way, and the coupling between the two is only realized on timescales that are maybe five to ten years. But in the meantime, there's a lot of space to fluctuate. And on the day-by-day basis, prices are impacted by many more things than fundamentals-in particular by social madness, by interactions, by feedback of the price on itself, by trend-following-all these things that make prices a very bad indicator of value. And thereby, again for the same reason, because there's so much heterogeneity and uncertainty about the true value, the role of noise and the excess volatility that I was talking about can express itself and lead in turn to these inequalities.

Steven: One of the things you've written about is inefficient inequality and unjust inequality. I was hoping we could turn to those and talk a bit more about your thinking on those two facets. In terms of inefficient inequality, what is being revealed from the work you've done?

Jean-Philippe: One aspect is the role of taxes. The idea of the simple model that I try to play with, although it's very simple, it's pretty rich. The idea is: let's assume that there's a private sector that is leading to inequalities, but with a higher growth rate on average. The idea would be that private investors know better, and they detect opportunities that lead to economic growth that benefits everybody. And then there's a state where, for some reason, you can argue about this, and I guess politics is a lot about this particular issue, is less efficient, because, I don't know, the decision process is more collective, so there's less risk-taking, or there's less skin in the game, whatever. So in this case, you have a public sector that leads to lower growth but more equality in a way. And the question you ask is: what is the optimal taxation rate that leads to a globally optimal growth rate for everybody? The surprising thing-and you would not think this-is that if the growth rate is stronger for the private sector, you'd think: let the private sector do its business, don't interfere, have as small a state as possible. But the result of this calculation is non-intuitive, in the sense that there is an optimal taxation rate that is somehow related to the diversification; if there's a taxation of the private sector that gives money to the public sector to invest in less risky investments, then because of this exponential mechanism and maybe portfolio analysis 101, the global growth rate of the economy will actually be better for everybody with a non-zero wealth tax. Again, I think these messages are interesting to think about. The model is too simple to reflect reality, but there are hints that sometimes the mathematics tells you a story different from the one you intuitively think about. So, as a function that you think about, it is a real thing. The difference between the growth rates of the private sector and the public sector

determines an optimal wealth tax rate, which can go from zero percent to one hundred percent depending on the difference of these growth rates. Of course, if the public sector is much better than the private sector, then it's better to tax everybody—that's a communist state—but vice versa if the public sector is much worse than the private sector, then it's better to have zero tax rate. It's the whole spectrum in the middle that I think is interesting to think about. Another way—and again this touches on the idea of fair inequalities, how much inequality is fair and how much is unfair—the idea of a simple model I looked into is that there are conditions, essentially in terms of productivity, where the whole society is better off by creating a two-tier wage grid. You would think that's not obvious, but in some cases at least the model tells you that you should create inequalities—even randomly, without reference to skill, which would be a little absurd—but by creating some inequalities you make everybody better off. It's a little bit like Rawls's idea, that what you want is not necessarily to focus on inequalities but also on the well-being of the poorest. Again, it's a very simple model. You have two types of goods: one good that is a kind of everyday consumption good, easy to produce, and so has a low price. And then there's another good which could be thought of as a luxury good, or a piece of art, or whatever: more difficult to produce, and has a high price. If the productivity of the easy good is too high, you don't need the whole population to produce that good, and so only a fraction of the population is needed, and therefore, there's unemployment. And the only way to not have that is to create a rich caste, if you wish, that is able to buy the luxury good. By making this luxury good available, by launching the production of this luxury good, you re-employ the people who are out of jobs. The idea that you must have a class that is able to buy the luxury good is something that, depending on parameters, but there's a region of parameters where it is optimal to actually create inequalities out of nothing, even if people have the same skills, because this will allow the whole population to remain employed and be better off. These toy models are arguably very far from reality, but I like models where there's a surprise at the end. When you start writing the model, you think there's not going to be anything interesting, and at the end, there's something you hadn't thought about. I think models should be there to generate ideas more than necessarily fit data, or jumping too far ahead to ask models to be realistic from the start. And again, this is something that physicists like a lot, doing these toy models that shed light on the mechanisms we want to understand.

Steven: I wonder if you might take a moment to speculate on some dimensions of inequality you haven't focused on, where you think the tools of physics might provide new insights.

Jean-Philippe: You know, as you yourself have worked on these types of models where tipping points exist, I think the idea that there might be tipping points where inequalities become completely unstable—I guess this is a little bit of a Piketty-type argument—is that in some

cases inequalities are there and they are more or less stationary, but the idea that you can have tipping points beyond which you can have runaway solutions and that inequalities become really unsustainable, both from a social and economic point of view, is very interesting. The idea of phase transitions and tipping points is a little bit the same kind of idea with different words. And as we also discussed earlier, I think the idea of non-equilibrium and the need to think about non-equilibrium and non-ergodicity is something that I hope will be key to understanding some of the stylized facts in growing inequality.

Steven: I want to thank you so much.

Jean-Philippe: Thank you for a very interesting conversation.

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 thank our producer and engineer Shane McKeon. Second, I'd like to thank our assistant director Gerardo Espinal Franco for the production oversight and doing everything that is required to bring the podcast to fruition. And finally, I'd like to thank our executive director Grace Kolavo for her support, not just for the podcast, but for every activity at the Stone Center. You may get in touch with us at stonecenter@uchicago.edu. Thank you so much for listening.

This transcript has been prepared with the assistance of AI.