Steven: Hello everyone, this is Stephen Durlauf, and welcome to the Inequality Podcast. I'm really delighted to introduce today's guest, Scott Page, who is the John Seely Brown Distinguished University Professor of Complexity, Social Science, and Management at the University of Michigan. As his title indicates, Scott is a truly integrative and deep social scientist, calling him an economist, calling him a political scientist, calling him a complex systems, research all of those don't do justice to the breadth and depth of his work. Scott is by any measure really one of the leading social scientists of his generation, and speaking for myself, he's somebody who has constantly taught me throughout his career, and so Scott, welcome to the program.

Scott: Thank you, Stephen. It’s an incredibly nice introduction.

Steven: So Scott, I actually thought we'd talk a little bit about you in the sense that you really are unique in terms of the combination of skills you have going from really rigorous microeconomic theory that derives from your time in graduate school to the integration of complex systems and other state-of-the-art mathematical methods into social science. So you can tell something about your intellectual trajectory.

Scott: So briefly, I started out getting a PhD in mathematics at the University of Wisconsin and when I was there, I was extremely lucky to take randomly drawn professors, you know, people like Buz Brock, John Rust, Ken Rogoff, who just happened to be at Wisconsin at that time. I was getting my PhD in math and taking the sort of core micro theory course just as a way to kind of expand my horizons and think about ways you might apply mathematics. And at that time, both Ken Rogoff and John Rust said to me, look, there's this new thing happening, this is like 1986, 1987, called Game Theory. We think you'd be really good at it and you'd really like it. You should go do that. And so when I transferred to Northwestern and studied Game Theory with Roger Myerson, he was your colleague at Chicago, and Stan Reiter, Mark Satterthwaite, Matt Jackson, Leo Hurwicz, just some spectacular people, Nancy Stokey, who also became your colleague at Chicago. And you know, did a dissertation on kind of mechanism design stuff. But during my fourth year there, Stan and Roger got invited to go to the Santa Fe Institute and they sent me in their place to Santa Fe Institute's complexity think tank. And so I took my first job at Caltech and ironically, even though Game Theory was kind of the shiny new object, complexity was even a newer thing and in some ways maybe even shinier. Ken Arrow was there, Murray Gell-Mann was there, Brian Arthur was there, Chris Langton. And they had this idea of complexity almost being this extension of Game Theory. So you're thinking of kind of strategic, purposeive, adaptive agents that were in networks. And whereas most of the game theory results, the mathematical results involved either like two agents or an infinite number of agents. Complexity occupied this kind of interesting in-between area. And so that was just because I did formal training in mathematics, it allowed me to kind of make this pivot to do complexity stuff. And I've ever since then, I think a lot of my work is mathematical models. A lot of my work is agent-based models, complex systems models, Markov models, that sort of stuff. And so I've just really enjoyed using a variety of models to try and make sense of the world and focusing on what I think are most of the time really important questions. Every once in a while, as you know, as an academic, you get involved in something that it's just interesting and you kind of chase that down. But it's been big fun. The career has been just fantastic.

Steven: So having been present at the creation in terms of the introduction of complex systems seeking into social science, might you reflect a bit on how things have played out? In other words, what would you say are the greatest successes and where do you see relative failings?

Scott: So I think, when it first came out, I thought like, wow, agent-based modeling is really going to take off because of this kind of in-between thing. So on the one hand, we can mathematically do two agents and we can do an infinite number of agents. We can mathematically do zero intelligent agents and we can do infinitely smart agents. We can mathematically do pairwise interactions or everybody interacting, like in a common market, like in a standard Arrow-Debreu economy. And complexity was kind of giving us that in-between space. So there's a lot of kind of early work where people had like these kind of, like adaptive learning models, early network models, those sorts of things. I think what happened David Colander, who was also in this space; he wrote a book on this and I wrote an introduction to it. And I remember saying that I think the real contributions of complexity theory, if you think of kind of the four pillars of it in a way, one is diversity, one is networks, one is adaptation, and then one is interdependence, not just non-linearity, but like both strategic and sort of externality type interdependence. And I think each of those four things has independently become an incredibly important part of economics. There's a lot of stuff I'm learning in adaptation, right? There's a lot of stuff on networks. There's a lot of stuff on externalities, right? And there's a lot of stuff on heterogeneity. So the parts have all been completely mainstream. So I think that is a massive success. I think less of a success was the complete, Gestalt understanding of what a complex system is. So in economics, we talk a lot about stability. And you can think of stability as being like, if I've got a ball in the bottom of a basin, and I push it up, does come back down, right? And so what I do is I might create a Hessian matrix and like compute eigenvalues and figure out like, okay, is this thing going to go down, or is this bigger than one or less than one, right? What happens with a complex system is if you disturb it, it's not going to come back to the same place. You think in terms of robustness and robustness means, do I maintain functionality in light of some sort of knock out of a species or change an external parameter? That sort of thing. And so there's difference between like stability of inequality, going back to the same thing versus robustness of a system. Does the system maintain functioning is a key distinction. So like my wife, Jenna Bednar wrote a book called The Robust Federation, which is particularly timely at the moment. When you think about, we've got this system of government that's much more complex than can be analyzed with kind of like a single dimensional parameter. There's a federal system, there's agencies, there's separation of power, there's courts. And you can ask if there's a disturbance like a president trying to usurp more power, is that system robust, right? Now, it's not going to be the same as it was before something like this happens, but can it maintain functioning? So I think that like the principles of networks, adaptation, externalities, those sorts of things, big successes. I think some of these concepts though, like the pursuit of robustness and stability and thinking from a system's perspective, maybe haven't been as successful as I'd have hoped. I see a big sea change happening, though, and the reason why I think one reason it didn't happen is you had to learn how to program. So when I came back from this and I said, you know, Stan Reiter and Mark Satterthwaite and Roger Myerson, yeah, this is interesting. You've got to kind of do this agent-based modeling and you've got a program back then it was like in basic or C. And you know, it's a question of like, if I'm a highly successful mathematical economist, why am I going to learn computer programming? And it isn't part of the curriculum. And so I think that we just didn't have people programming. But literally in the last two years, now with AI, it's just so much easier to program that I, and I think you're also seeing the rise of kind of like creating artificial agents instead of human agents to do experiments. I think agent-based modeling is on its way back up.

Steven: I want to highlight the point you brought up about actually the Gestalt, the overall systems idea, because it seems to me that, you know, concepts such as robustness of functionality that you described or the idea of emergence; these are really philosophically essential ideas in complexity. That there's phenomena that occur at a higher level of aggregation than the description of the environment. And those clearly appear in economics. You could say that the welfare theorems have the property of emergence, et cetera. But somehow the capacity to think about new categories, new ontologies of a properties of systems that is part of the frontier.

Scott: Absolutely, let me jump on that. And so Wolfram, and his book A New Kind of Science uses a categorization that other people have used. But he talks about, if you look at a system, it could be in equilibrium. It could be periodic. It could be completely random or it could be complex, which is kind of between, Murray Gell-Mann, we're called kind of between ordered and random. So what's complex and what's ordered and what's random is, again, one of these kind of pile of sand things, right? You know, okay, it's just moved from kind of, but you can think of something going from an ordered periodic regime into a more complex regime, but then becoming so complex, it's as if it's random. And you can use ideas of information theory and measures of information theory, like entropy to say when something moves between these classes of states. And I think you nailed it. And when I think of economics, I have an equilibrium. And then I ask when I have this policy change, I move to a new equilibrium. And analogy I use all the time, if I had a dog chasing a rabbit in my backyard, I could model the rabbit R as this random thing and the dog is this D deterministic process. And I could model that as a sequence of equilibrium, treating the rabbit as random, right? If the rabbit only moved every once in a while, modeling that as a sequence of equilibrium would be completely appropriate, and it would explain most of the behavior of the system. But if the rabbit is moving all the time, as frequently as the dog, and it's responding to the dog in some way, right, then it doesn't make sense to think of that as a sequence of equilibrium. It makes sense to think of that as a complex adaptive system, right? And one thing we talk about is when you think about sometimes systems may move from order to complex or from complex to order. And also within the same system, so example I always give, if you pick almost any commodity, I want to say almost. And you look at global production of that. If you took a half of a semester of a time series class, you could predict that pretty well over a hundred-year period, total global production. If you look at the price of that commodity, it's not ordered. It's incredibly complex, right? And Lutz Kilian used to be at Michigan, has like a 175-page paper on the predictability of oil prices. It should be of the unpredictability of oil prices. So price is something that can be complex, even if quantity is not. But if you picked up a standard econ textbook, that would not be one of the first things they'd probably mention, right? I mean, we think of pricing quantities being in equilibrium, right? But yet in economic systems, typically tend to be things that where quantities have fairly smooth patterns and prices are typically complex.

Steven: Oh, yeah, I think that that's a great example. I guess I'd also say the other limitation has been the relationship between complex economic models and empirical work. I don't want to exaggerate this, but it seems to me that there is a continuing disconnect. And so, for conventional economic models, there are dimensions of structural estimation, very precise ways of mapping the theories of the data that don't yet exist from the perspective of complexity. Do you think that's fair?

Scott: Well, I think it's more than fair. I mean, there's a sense early on that with complexity theory that like game theory, that it was going to give us a set of kind of like conceptual understandings or ideas. You mentioned emergence, things like long tails, things like small world graphs, things like, you know, people talking about adaptation to the edge of chaos, which was a bit of a buzzword, but yet it's true in the sense that like systems do tend to become more complex over time. But once they become random, then people simplify their rules and they pull back into the complex realm. So there's a whole bunch of kind of like frames where they actually can see the world in broad strokes that have been incredibly useful. But then the question becomes, okay, can those things be taken to data in a way that they can be proven? And so there's some things like power law distributions of firms and cities some distribution of results that yeah, people were able to take to data. But if you think about like predictive broad strokes sort of stuff, the simpler economics models that are based on equilibrium, those are so much easier to estimate. But then we have to ask ourselves what do the R squareds look like, across those, I mean, some of them are very high, some of them are less high. But if I'm going to say something's complex, you're caught in a catch-22, it almost means that, you know, let's go back to the commodity prices that it's going to be hard to predict. Now what I think we get interesting is let me give a case where I think it's where you see this kind of cutting edge on things, think of the efficient market hypothesis, right, which is a clean, classical economic equilibrium sort of thing that prices should be random. If you look at Andy Lo's work on adaptive markets, I would classify that as complex systems thesis, which would say, and this is, Rothschild-Stiglitz, which is also, famous paper economics that like things are going to stay random, but there's going to be movement towards random, but with prices are random, there's no reason to do any research anymore than there'll be patterns than it should be random, right? So that's so complex systems, it would be there should be short-term patterns. They shouldn't be very big. Occasionally, there's a big one that those are going to get eaten up over time, like the January effect. That's true, right? That's been true. So I think that there's cases where it's worked. Where it's really worked, I think as well as neoclassical economics, is domains where the behavioral dimensionality or behavioral degrees of freedom are really, really low when COVID came out, like there's, you can think of kind of like agent-based models that included all this stuff versus simple mathematical models, traffic models. With disease models, it's kind of like, do I have the disease or not have the disease? Do I go out, not go out? In traffic, it's like, when do I leave to go to work? I mean, so that there's not a lot of kind of like behavioral dimensionality. In those cases, when the behavioral dimensionality is low, but yet you can then add all the kind of like what people would call kind of model verisimilitude or, you know, kind of like micro level modeling that's very accurate that you do with complex systems models. They've actually been pretty good, but people are so heterogeneous and we learn in so many different ways. It's a high-barred, expect-to-complex systems model to work there. Like Colin Camerer's really, I think fantastic on this, like learning models and game theory models work really, really well when you've got like sophisticated agents, right, stakes are high, and they get to do it a whole bunch of times. Complexity, by definition, is somewhat perpetually novel. You're not going to expect it to happen, right?

Steven: Well, and I think you bring up actually, a very deep issue in thinking about individual behavior. And one of them simply is that in complex environments, the idea that all of the actors are going to form common beliefs is especially problematic. And so let me give an example of that. There's very good evidence that small interventions before freshman year of college have persistent effects on educational outcomes. As a hypothetical, if students, it's emphasized that college is hard. People have disparities in the high school background, et cetera. And so this is a new world you're entering. The example I want to give is what's the interpretation? The first time I take a test in a mathematics class as a freshman, I get a lousy grade. Do I interpret that as am I not smart? Did I not work hard enough for the disparities in the background, et cetera? And I guess what I'm putting on the table is I think something deep you brought up is in the context of complex systems, the capacity for the data to simply speak for themselves, as opposed to requiring narrative or interpretation, is qualitatively different. So I will at least suggest that as a dimension one might want to think about.

Scott: No, I absolutely agree. And if you think of, you'll know the name of the study better, but Matt Salganik worked with, there was a bunch of early childhood data and you looked at kind of how people's life course played out. And they had a whole bunch of social scientists trying to predict like what's the likelihood that you graduate from high school, you graduate from college, you go to jail, you do all these things. And all these people using traditional models from sociology and economics, they just couldn't predict. I mean, we can explain really well, but we really, we can't predict. And in fact, one pushback from the complex systems people against kind of the more structural models would be that we're doing a lot of explaining, but not a lot of out-of-sample predicting. One thing that complexity theory might say is that it might be very, very hard to predict individual cases, but you might be able to predict the movements, the general patterns. And this is, you know, a mergentist kind of swarm of bees argument. If you try to pick the behavior of an individual bee, you're not able to do it, but you can predict the behavior of the swarm. And this is again, one of these kind of complex systems ideas that like the swarm is kind of an emerging thing. Can individual interventions have a big effect? They can, because if you kind of can push a whole bunch of the, some percentage of the bees a little bit in a particular way. And then you get kind of like, you know, systemic feedbacks through interdependence. Small interventions can have a large effect. Want to be careful there, though. I think one of the kind of false promises from complex systems, not created by complex system scholars, but really more from translators. You know, it's just exciting to think about like the butterfly effect. And this idea that I could just do one tiny intervention and then crime goes away. Poverty goes away, right? I think that's probably overly hopeful.

Steven: Yeah, so I would unpack a couple of things you said. One of them is I think you actually gave a very constructive vision of empirics and complexity. And one of them is to identify what is salient in these systems. And really the examples you gave, they all had to do with scales and what phenomena to see. So in other words, the idea that of emergence, the counterpart to that is you can see things at different levels of aggregation. Some of the more successful empirical examples fall in that category. The other, it seems to me that warrants more consideration are simply timescales, the idea that in certain environments, depending on the timescales you're looking at, you're going to see qualitatively different low frequency than high frequency behaviors. You know, that's something I've personally been interested in in the context of thinking about economic history. I recently wrote some, I'll say, speculations or hopefully constructive speculations on how complex systems could help structure asking questions about such as understanding dynamics of industrial revolution, the like.

Scott: Let me push on the state of time. I think it's so important what you just said there about speculations. So a lot of my teaching has been on just the use of models and what they do. And I have a set of reasons why you model, right? You read model to reason, you model to explain, you model to predict, that sort of thing. But one of them is you model to explore, right? Just to think through things, and one of the amazing things about kind of complex systems modeling, there's originally it was often called artificial life, artificial worlds were terms that we're thrown out a lot more early on than later. But it's kind of only got like one economy we can really see playing out one world system, one ecosystem. And one thing that's happened in complex systems all along is people creating things like manufacturing and evolving ecosystems on a computer, manufacturing, creating models of entire cities and histories, and using these to just kind of like explore whether it's policy changes or—so what would, what would generically be true of any system? Like you see power law distributions of species and things like this, right? Long-tailed distributions of species is that generically true, right? Or you see like a power law against some networks like on the internet, you know, you see like there's a power law distribution of degree, where then other networks, you see things that have like, much different sort of distributions. And you can just start asking questions, what if I make different micro level assumptions? So one of the, I think the key things of complex systems as you imagine, I make micro level assumptions, macro phenomena occur. Those macro phenomena feed back down to the micro level, right? That isn't a new idea. That's Coleman's boat, right? But what complex systems is kind of paddling Coleman's boat down the stream, right? It's a macro to macro, macro to micro, micro to macro, you know, it's back and forth, back and forth, back and forth. And I mean, one thing you're talking about is, yeah, that sometimes at the macro level, you're getting something different than just averages, right? You're getting cliques, you're getting social movements, you're getting collapses, you're getting, right? And that's because of the interdependence; if everything were just additive, and one of the things when you think about like, the structural models and economics, there's a lot of summation signs and a lot of integrals, right? And implicit in that is oftentimes some additivity, right? Not necessarily. You can, you can obviously do systems effects with integral signs and things like that. But there's a fundamental predisposition towards kind of additivity and separating things out as opposed to thinking of kind of like interdependence between what's going on. It's that interdependence that creates kind of the interesting stuff at the macro level. And then that feeds back to the micro level.

Steven: Yeah, I think that's essential in another concept that's worth mentioning. Simply is that a phase transition, the idea that systems are in one type of configuration, certain features change, and you get possibly discontinuous changes in the qualitative features. In my own opinion, that's probably not a bad metaphor for the emergence of revolutions. In other words, societies go through some set of fluctuations in the nature of—you could call it legitimacy, could be an application of it, but let's say mathematically, you know, parameters have described how closely interdependent agents are, and then certain shocks can then hit the system and cause a qualitative change.

Scott: Absolutely right. There's always a tension between some of these ideas in complex systems. Michael Cohen, who I actually dedicated the models book to one time referred to Santa Fe as like a festival of like amazing metaphors, right? So phase transition being one, and there was so much fun early on like people constructed these models. I mean you know these much better than I do. Like a checkerboard grid that has a bunch of what's called in trees for a second, right? And so it can either have a each square in the checkerboard can have a tree or not have a tree, right? And then as a budding young arsonist, you get to come in and like light a tree on fire randomly. And then you ask, does the whole forest burn down or not? And then you got this parameter on the bottom, like what percentage of the trees are filled in, right? And what happens is, depending on your network topology, but let's just take a grid search. It turns out like at 50%. If you're less than 50%, there's no way the fire spreads if you have a big grid. And if you're more than 50% it totally spreads. So if I then brought a graph of probability of a forest fire and the horizontal axis, I jad density. If I'm less than 50%, the probability is zero. And then all of a sudden, zoom, phase transition once they get above 50%, forest fire, right? So in the early days of Santa Fe, and again, it's hard to believe this is true, but 30 years ago is the first time he really started having simulation models on computers that we could share over a network and all see them, right? That's a relatively new thing. So we're all seeing these things, and we're just going, this is so cool, right? This is oh my gosh, right? This is a phase transition. And then one person says, that's a forest fire. And Steve Durlauf says, that's a revolution, right? And then Steve Durlauf and Buz Brock say, that's a technological change from one type to another. Or, you know, somebody else comes along and says, this is percolation of a technology. Right? Like each one of those things is an idea in the space of silicon chips. And suddenly, once you get enough filled in, everything kind of spreads, you know, so there's, there's just so many. This is that, you know, Michael Cohen metaphor. Like once you suddenly see this idea of, if I have things that are connected in some way and there's a density to them and things have to flow between things, those sorts of systems have phase transitions. And then you start asking yourself, okay, wow, I've been running a lot of regressions. Well, if I run regressions on things, if I'm still always below 48%, I'm going to say this doesn't matter. So let's suppose that like, I'm your colleague John List, I'm anybody trying to make the world better by doing experiments. If I'm experimenting by moving from a delta of density of trees is 43% to 45%, I'm going to see no effect. But if I move it from 45% to 53%, I'm going to see a huge effect, right? So let's suppose instead of a tree, that being a tree on a grid, that's a student who gets a good breakfast, who feels safe at school, right? Who wants to work hard. Other squares are students who don't have those things, but they're socially influenced by other students. Then there's the possibility of phase transitions in student performance. Again, though, that view is a kind of overly optimistic view. But again, this was what made the early days of Santa Fe super, super fun because you're just being exposed to things that ex post seem obvious. But before you walked in there, it just didn't occur. You were like things slope up. A marginal change has marginal effect as opposed to marginal change has threshold effect.

Steven: So I want to tease out two points. I think they're very important there. One of them is, complexity really does challenge in many ways the calculations that have been done about policy effects. And so an example that comes to mind is that there was for a long period of time, arguments that changes and per capita expenditures on students didn't really matter for outcomes in education. Now, I think that that better work on exogenous variation has shown that's not correct. But the point I want to make is that extrapolating the variation we saw in per capita expenditures to large-scale changes simply doesn't work when you have a complexity perspective because of all the implicit non-linearities. And then the second point that I think is essential, you brought up actually has to do with econometric tools. In other words, that part of what complex systems do is that they indicate the limits of conventional tools. And so an example that I've worked on is thinking about intergenerational mobility. Well, the conventional models that are used are based on linear regressions. Through that lens, you can never see a poverty trap. You can never see in an afluence strap. An unexploited value of complex systems thinking is really to go back to the essential econometrics of measurement of phenomena such as intergenerational mobility. And ask what can't you see from the tools we use, given what we know from the ergodic mathematical theories.

Scott: Yeah, and on that point, this is speaking about your own work. If you think about a community flourishing in some way, that means you're doing well on a whole bunch of those dimensions. You can't separate that out as the additive sum of those effects. Each one of those effects has positive feedbacks in and of itself, and they also have interdependent effects. So I could come in and say, what if I improve schooling a tiny bit? No effect. What if I reduce crime a little bit? No effect. What if I make the water a little bit cleaner? No effect. What if I add a couple of businesses downtown? No effect. Minor effects. But if I were to do all those things, that I were to make crime go down a measurable amount, make the schools measureably better, make economic opportunities measureably better. Do moderate effects on five things all at the same time, I might get, what you're talking about before, this kind of like phase transition, where you get the help the community escape the poverty trap. And have the whole thing leap up. In the models book, I joke about this as kind of like big coefficient thinking. In the sense that what I do is I gather all this data, maybe I do some actual experiments, maybe I get lucky and there's something just happens where I get a regression discontinuity. And then in a pure, wonderfully, intended Gates Foundation way, I say, here's the biggest coefficient, making class sizes smaller. Or here's the biggest coefficient, improving teacher quality. I want to figure out which one of those is bigger. And I say, oh, it's teacher quality. So I put all my money there. Well, if there's five forces interacting that are mutually reinforcing in positive and negative ways, improving one of those isn't going to lift the system out of the poverty trap.

Steven: So another area where your work really has important inequality implications is the way that you've written, in much of it with Jenna Bednar, on the relationship between political and economic institutions and their joint interdependence with culture. Could you describe some of that work?

Scott: Yes, you know, really quick: one of the things that I found most fascinating, when I first started studying economics was, you know, Debreu’s theory of value, when you just suddenly get this difference between general equilibrium and partial equilibrium. And how important that is taking into account general equilibrium effects. And you know, I was trained by Leo Hurwicz, who visited for a year. So, you know, Leo and Stan and Roger and Mark were foundation, Mark Satterthwaite, were foundational people in mechanism design. My wife, Jenna Bednar, we've written this book called The Robust Federation, thinks of a federal system as like all these different institutions. And when we were talking about it, and this was actually, you know, some conversations with Lin Ostrom, what if you model all the institutions, right? So if you think of like, when I design an auction for the federal airwaves, that's a partial equilibrium analysis and institution space, right? So I'm not thinking of all the other institutions. And so then you can say, okay, well, that probably does most economists would say, well, that doesn't matter because I'm just like picking off each allocation problem and designing an optimal institution. But let's go with, think of what complex systems do, and say, my beliefs, my network, my norms, all these things, for lack of a better term, dark matter that we might lump up as culture, that affect levels of trust, other-regardingness, all those things are a function of the institution. So the things we take when we think about the environment and the standard mechanism design, like this is the environment, these are beliefs, this is the network, these are trust levels, these are preferences, all this stuff, those are functions of the set of institutions that exist. And so then we started saying, okay, well, so what if we just wrote models where when I choose an ensemble of institutions, that then gives me a culture, and then that culture determines which institutions will work well? So imagine, and this relates to the stuff you're talking about before about kind of economic history or economic growth generally, over time we add and subtract institutions. Well, as we do that, when you start adding new institutions, those institutions are leveraging the existing culture. And so maybe inside its simplest form, I think the main idea from this work is the institutions we choose are the ones that best leverage the existing culture, like this institution is going to be better than this one because it's leveraging our current culture. But that institution is also going to produce the culture that we have to leverage tomorrow. And so if in each period, this is in some sense, a very simple idea, but if in each period, I choose the “optimal institution,” what I'm doing is choosing the “optimal institution” to leverage the existing culture and not taking into account its effect on future culture. And so then you can get all sorts of interesting things like path dependence, and you can also get the result that like you would do better by not necessarily choosing the institution that works best in the moment, but you should also take into account what its effect are going to be. And this leads to, I mean, one of the things I really love, I think it's been one of the great gifts about being trained as a kind of formal social scientist, is it really kind of prevents you from becoming overly ideological because when you start writing these things down, you realize you come to some conclusions that say, well, if I have institutions that actually like help people build levels of trust, maybe have higher effort levels or something like that, right? Then that can be good for society. That almost seems like a very right-wing thing, but then you also get results that say, well, I need institutions that actually create cross-cutting cleavages in the network so that I have trust and generalized trust when I wanted to do these other institutions, that comes up what seems like a very left-leaning sort of thing. But if you end up sort of saying the ideas I carry around in my head and the things I advocate in my life are driven by some ideology as opposed to sort of first principles thinking, I think you're probably not; I think that's not going to be successful. So I think that like this particular research agenda has been fun for me because it's constantly challenging me, right? When you get results that like maybe you're predisposed to like or not like, then you become more aware of your predispositions.

Steven: So I wonder if I might turn to your work in diversity. You know, it's certainly one of the most important sets of contribution that you've made to social science. Your book, The Difference, is now two decades old and has been re-released by Princeton as a classic. I think that that's quite justified. And so I guess I'd like you to talk a bit about the main theoretical findings you've developed on the effects of diversity in organizations and other groupings, and how your theorizing evolved over the last 20 years.

Scott: Yeah, so the point of The Difference was I found in economics and in political science, we had so many just fascinating and useful models where we talk about differences in preferences and differences in information. And so like differences in preferences, like all this stuff on, you know, utilitarianism and things like what is social welfare? All that hinges on beginning with an assumption that people want different stuff. If everybody was the same, a lot of philosophy becomes easy, a lot of economics becomes easy, but it's the heterogeneity that makes things kind of fun to think about. But then when I started doing complex systems models, I realized like, wait a minute, there's also differences in like how I categorize the world. There's differences in what problem-solving tools I have. There's differences in how I represent things, right? And so then I thought what if I were just to kind of like, for some reason, we don't think of differences in perspectives, representations, differences in heuristics, like tools, differences in models that I have or frameworks as fundamentals of social science. I mean, you could argue maybe you do some now in cognitive psychology, but at the time I thought, what if I just tried to like treat those, kind of write models the same way as like mid-level economics models from a formalism perspective? That looked at what are the implications of people having diverse representations, like I use polar coordinates, you use Cartesian coordinates, or what are the implications of you partitioning the world this way and me partitioning the world that way? What are the implications of each of us having sort of a capacity to learn, you know, 15, 20 tools, and then us going out and applying those tools. And when you do that, you find out that like the ability of a whole, like the intelligence or the capacities of a collective, really depends on two things that are very interdependent. One is how talented or how good are the parts, right? And do have models that are useful? Am I partitioning the world in a useful way, do I have heuristics that are useful. And then the other thing is the diversity of those parts, right? And what you find across whether you're looking at problem solving, predicting information aggregation, logical deduction, creating new ideas, in any one of those domains, what you find is that diversity helps, ability helps, and then ability and diversity have kind of a positive interdependence. And so it leverages a lot of very mainstream ideas from economics, things like, you know, Shapley value, Marginal value, things like that, and does a lot of kind of comparisons in terms of, you know, just what's the relative advantage, what's the benefit of ability, what's the benefit of diversity, that sort of thing. What I think I found most interesting in this space is that a lot of the, and I've spent a number of papers since that book came out on this, is that like some of the assumptions we make in economics, are because they're kind of convenient, like independence of signals, or normal distributions of signals. We're just doing that because it's convenient, and there's the micro foundations for those assumptions are very, very weak, and therefore some of the implications we’re drawn from them, probably don't hold. So for example, it's really inconceivable that a hundred people could get independent signals statistically on some binary event. Now it could be if we really were getting independent signals, like in the sense that we're looking at the luminosity of the stars, so they're independent disturbances. But if you had a 101 analyst trying to decide do I make a buy or sell on Amazon as the current price, there just aren't 101 statistically independent ways of looking at that, right? And so then suddenly, you get this deep understanding that there's a relationship between diversity in terms of how we see things, how we think about things, what information we have, and assumptions we're making about the statistics of signals that occur in our economic models, the ability of democracies to get to the truth, and all that sort of stuff. So it's really just been kind of a, this really fun exercise, a lot of it with my colleague, Lu Hong, who's at Loyola in Chicago, some with PJ Lamberson at UCLA, just kind of playing in this space theoretically. It's been really fun.

Steven: So I think one of the very important arguments that comes out of your work is that diversity is a source of efficiency. How do you see the evidence, positive and negative, in terms of that view?

Scott: Well, so there's a result in that other people have it in other places that's kind of a simplification of the bias variance decomposition, which essentially says, let me give just one example, which essentially says if I am trying to predict something in numerical value, and I'm just to average everybody's prediction and call that the crowds prediction, and I want to ask what's the squared error of the crowd? It turns out that's always exactly equal to the average squared error of the people in the crowd. So you can think of that as kind of like average ability, minus their average diversity, like how far off they are from one another, in the logic of this is straightforward, if we're guessing how many eggs are in a dozen and everybody is off by six and we all guess six, the crowd will be off by six. But if some guess 18 and some guess six, then only the average will get closer to 12. And so you could think of that as saying, oh, that's just noise, and like noise kind of cancels. But that's again taking the thing I was talking about before, this view that people are somehow just there's just a signal-generating process. Sometimes the point of the book is saying, well, let's ask what that signal-generating process is, and that signal generating process is Steven Durlauf getting a PhD, studying at Stanford, having a model, Andy Lo having a model. Each person, based on how they were trained, Elinor Ostrom thinking about collective action in a different way than Leo Hurwicz's does or something. So it's each one of us through our training having a different way of seeing things. So that term is always true. Collective error equals average error minus diversity. What was kind of interesting to me when I wrote the book the first time, the first version of the book, there were like a handful of studies that had done this. And most of them were kind of like amazing anecdotes, like people guessed the weight of a steer within a pound. Since that time and not just due to the book, but there's getting to like general movement and collective intelligence, there's been hundreds and hundreds of studies where people have done that calculation, right? People send them to me, like my friend Michael Mauboussin and sent me one or our friend Michael Mauboussin and sent me one yesterday where his class from Columbia on average was off by 50% for Jolly Beans and a jar and the crowd was off by .01%. And when you look across and again, you can you can cherry pick examples like that, but if you look at predictions by economists like hundreds of economists and 50,000 predictions, if you look at the largest prediction contest ever run by IARPA, if you look at lots of medical research, what you tend to see is that the bonus from the diversity is about 20%. You know, in some cases in Google hiring, it was like 35%, but most of the time it runs between like 12 and 20%. So it's one thing to say mathematically, there's a benefit to diversity. There has to be, but it could have turned out that in real-world situations, that's 1%. Right? If I had a place to bet, if you said to me, Scott, you're not going to believe this, but there's going to be hundreds of studies on this because there's going to be now massive data flowing over the internet in all these amazing ways. What do you think that delta is going to be from diversity? I'd have said 8% to 10%, and that would have been an optimistic view. And it's been more than double that, I think, which is incredible. Right? So if you think that five economists tend to be like 16% better than the best economists. What's, I think, also really interesting to me in this space. And I don't know how or why I haven't been as effective as I would like to have been, is, Lu and I wrote this paper a long time ago, which Lu Hong and I, which essentially said, suppose I take a whole bunch of people who are pretty good at solving a problem. And what I do is I represent them in some sort of space, thinking like a rugged landscape where there's peaks in this space. And I basically say, I can look at the likelihood you'd get stuck on any one of these peaks in a way. If you assume that there's types of problem solvers, and I have a big set I'm drawing from, if I choose the very best ones, it's going to turn out that, like I'm drawing from a space of problem solvers, they're going to end up being somewhat similar. They get stuck in the same spots. And so I'm actually, if I'm picking from above a threshold, I'm probably better off randomly picking people above a threshold and picking “the best people” who are the people who sort of do best on their own. So this was a somewhat controversial result. People accused us of like rigging the simulations and all this sort of stuff, but it's been replicated now dozens and dozens of times. And the mathematics holds. But what's fascinating is that then John Kleinberg and Maithra Raghu, one of his students, then said, a natural extension of that would be to say, when would it ever be the case that there would be a test that you could apply to individuals that the best individuals would be the best team. And they get this result that, generally that that seems like it's hard to do, but they found a couple cases, if you're putting together. a relay team, that's true. Like if I'm putting a 4x100 relay team together, I find the four fastest runners. And they found that generally that's not true. And then they said we found this other case, though, where there is a test, but the test isn't how do you do, the test was how would a team of only Steve Durlaufs do? Lu and I thought about that for a long time, he thought, well, that's really interesting. And so then we came back to that in over like a period of six or seven years and a lot of dead ends, we eventually proved a theorem that just came out like last year that proves the following: that if there is a test to pick the best team, right, that you're going to apply to individuals that gives you the best team, then that test has to be how would a team of only your type be? So the only possible test is not an individual achievement test. It's a how would a team of you do test? Now, in the case of the 4x100 relay, those are the same tests. But so this is kind of almost like a revelation principle sort of result. Like, if there's a test, the test is not how do you do alone? It's that how would a collection of you do? And then the second thing is, so then you can ask, when would that exist? We basically show that that would only exist. There's a necessary and sufficient condition. If there's a replaceability for the worst property, which means that if I rank everybody by that criteria, if it's always the case, if I took the worst person off a team by that criteria and replaced them by someone better, the team would have to do better. Now, it could be you, if you place the third-worst person by someone better, you don't have to do better. But if you replace the worst, you have to do better. So essentially what that's saying is you almost have to have it be a 4x100 relay or there to be a test of individuals. And I think this is so important because we're sitting right now as a society talking about meritocracy, right? Well, what is meritocracy? It doesn't even make sense. I mean, meritocracy, if we were hiring people to go out to everyone to run relay races, then meritocracy makes sense. But if I'm hiring people to work in teams on complex problems, we don't even know how to solve for the optimal team, but we do know that picking “the best” can't be optimal. And so, if you were to say I'm restricting you, if you were to say to the University of Chicago, the University of Michigan, McKinsey, BlackRock, Google, or Gilead, 'no, you have to have a set of criteria and only hire the best people,' You're basically saying to them, we're forcing you to get a B-plus, right? We're not allowing you to be efficient, which is astounding to me.

Steven: In some shameless self-promotion, this is actually very close to what I'm currently writing on meritocracy, which is that the defensible form of meritocracy has to be prospective, or you're assembling groups of people to do something. And therefore, the definition of who deserves the position is the person who most augments the overall objective. And so the work that you've done with Lu Hung is challenging all of the conventional definitions of meritocracy, because as you said, they're additive.

Scott: Right.

Steven: There's a criteria to put people together. And once you think about it functionally, which is the way that I'm trying to approach it, you end up with a very different conception of what it means to say somebody has merit in terms of the organization or in terms of deserts. So I really, I think that those are very important insights.

Scott: Yeah, and I think we're, I mean, let's go back to our earlier discussion about complexity. So I can't remember who said this. It might have been von Neumann or something (Stanislaw Ulam), which is that like the set of nonlinear functions is like the set of non-elephants. And what's interesting about, like the point we were just talking about in terms of separability and this meritocracy notion, is so we know in the swimming example, it's separable meritocracy that exists. And we also know intuitively that if things were like super, super complex and needed a whole bunch of people to do something, then yeah, you probably can't just like have a single score. But what the paper with Lu shows and what, generalizing Kleinberg makes an assumption that like people's ideas, are distributions, and we have a purely axiomatic treatment. What we basically show is that the set of–so you think that kind of around those, like a glow. So if I'm close to separable, maybe meritocracy works, and if I'm close to really complex then I got to do something more complicated. But there's this giant area of indifference, right? And I think that what people like to think is, first of all, you don't even understand how big this set of non-elephant functions are. But what these results prove is that there's only a tiny, tiny little area around the separable functions where that logic holds. Across that giant, vast space of possibility of team performance functions, the set where the separability logic holds is minuscule. As the problems we face become more complex, the distribution across problems moves almost entirely to that region where a test doesn't work. Ironically, someone did a paper, I love this. Does an IQ test work for a team of people taking an IQ test? And it turns out the answers is “No.” In the paper, Lu and I explain why that's the case. Who scores well, we want an IQ test? Well, tends to be people who don't get things wrong, right? And so imagine an IQ test says 40 questions, and each one gets harder. So the people, some of the people who score best are going to be people who like get the first 35 questions right and miss the last five. But there's got to be other people who can get the last five because they're out-of-the-box thinkers. But they miss some of the early ones because they're way out-of-the-box thinkers, right? Like one of my friend's sons applied to a gifted school, he was five or six and they said, what's the difference between a fish and a submarine? He said, well one has only yellow cheese, and one has yellow or white cheese. And they counted it as wrong. Like that is a weird answer. Like fish don't even have cheese. And then, like they're asking the third question, and the person went back and said, wait a minute, wait a minute, wait a minute. Go back to your fish thing, and the kids are like, yes, like if I go to McDonald's and get a fish sandwich, it's only got yellow cheese on it. But when I go to Subway, I like it's got yellow and white cheese on it. Because he was thinking submarine sandwich. And I'm like, okay, that one's right now. So yeah, maybe skip the interview at this point. And so the thing is, you get people who aren't necessarily that like—really, really super creative and incredibly smart, smart in some particular way. If people who can solve really, really hard problems may not be people who get easy problems correct because they overthink them. Therefore, when you're putting together people with an IQ test, you want some people who have high scores who get all the easy questions right. But then you also want some people who are just out-of-the-box kind of super creative thinkers who maybe miss some easy ones, who get the hard ones right.

Scott: And certainly, these ideas also have implications for how to think about affirmative action in the sense that the heterogeneity of lived experiences of students matters for the collective learning of everybody. Your discussion reminded me of evidence that the composition of juries and kind of mock trials affects the outcomes. But it's because people come with different experiences, different understandings of how the world works. And so this is clearly extremely important.

Scott: No, it's super fun. Like today in class, I did this thing where I gave my students, so I got a set of international students, and I gave them like instead of eight cities, I said, put them in like pairs where each pair of cities is you know, most closely together. So they're in groups of four pairs. And I had done this because I was teaching them principal component analysis. So I put in all these data about these cities and then used PCA to plot them in two-dimensional space, and then created pairs that were closest. And then the question was like, why do you see these two as close, and why is the AI, or why does principal component analysis see it differently? And it was fascinating because there's a lot of my students saw London and New York as similar, and Tokyo and Singapore as similar. But the AI saw Singapore and London as almost exactly the same because of the heights of buildings and the number of subways, and that sort of stuff. But within my class, there was massive variation in how people saw these cities. And then if you think about it, if we were discussing what are these cities like, the richness of that conversation is definitely improved by meaningful but idiosyncratic ways in which people interpret that. I mean, I think one of the real challenges of where diversity and ability intersect is you've got to come at this with a lot of humility. You've got to recognize that the world is really high-dimensional, really complex, there's tons of stuff going on. So if you ask me about Tokyo, I'm giving a thin slice of what Tokyo is. If you ask me about Chicago or London, Chicago I'd give you a thicker slice. And so it's because of the fact that we can't see the full thing, that we see a slice, that we need to have many people. But if we're picking many people who've all experienced the same slice, we're not getting much. Again, that's very simple and metaphorical, but I think that the point I want to get is like, when you think about coming up with creative solutions—so yesterday I was doing a thing with people about, what are the ways you could sneak a dog across a border. And no one thought of more than eight, but the crowd of 12 people thought of like 20. The reason I did this, I don't know if you've seen this, this is like unrelated, but it's just so great. They opened a new border between Hungary and Bosnia. And everybody's there, they open the gate, and a stray dog walks across. And everybody does it all, they're doing somebody just starts applauding. They think it must be like the president's dog or something, it was so great. And so just for fun I would show at this group what I was teaching, do that, just to show them the diversity of things we can think of. And then to pull this back to complex systems, in complex systems, there's this notion of the adjacent possible. In the sense of you think of the things I think of as being in a network, if somebody gives me an idea I have not thought of before, there's a set of things that are adjacent to that that I can now think of. Or somebody moves a problem to a solution that I haven't thought of before. Suddenly, there's a set of things that are adjacent possible. And that's where this inner play between, where diversity and ability kind of bootstrap. If you add somebody that group who's smart and thinks differently than me, I not only get their ideas, if we're searching on a problem solving exercise, I get their ideas, but then I get the things I think of from their ideas. All right, which is super cool.

Steven: Scott, I can't thank you enough for such a fantastic conversation.

Scott: Thanks as always, it's always fun to talk with you.

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 really 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.