

RPM—Ep. 48 | The future of Responsible AI with Suzanne Tavill and Paul Fehlinger

Anna Marcus: Welcome to RPM, the podcast where we explore the world of private markets. I'm your host, Anna Marcus. In today's episode, you'll hear a conversation between Suzanne Tavill, Partner and Head of Responsible Investment at StepStone, and special guest Paul Fehlinger, Director of Policy Innovation and Impact at Project Liberty Institute. They'll be discussing AI, the opportunities and challenges that come with it, and why what we see right now is just the tip of the technology iceberg.

In his role at Project Liberty Institute, Paul spearheads strategic initiatives shaping the future of the data and AI economy. A highly respected voice in digital policy, he has scaled international organizations, advised governments and top VCs, and been featured in several publications.

In this discussion, Suzanne and Paul build on our July 2024 whitepaper in which we introduced the emerging concept of Responsible AI and its role in delivering trustworthy AI systems. Now they take that thinking one step further as they consider what "good" looks like in this emerging data revolution, how we can best manage the risks involved, and why LPs will be key to shaping and driving emerging best practice in responsible AI.

Suzanne Tavill: So, welcome today to our discussion together about responsible AI and how this is the tip of the responsible technology iceberg.

Paul Fehlinger: Hi, Susan, it's a pleasure, to be talking about this big topic today. As you said, we're recording this right next to the business day of the AI Action Summit, where the world's elite are from investment from startups and also from policy, from all around the world actually gather here. My name is Paul Pfenninger, I'm the director of policy, innovation and impact at Project Liberty Institute, and it's a great pleasure to talk about this today. The session we were just in together and was a milestone in a transatlantic process, probably the first one on an ongoing basis that we together at Project Liberty with Venture-ESG, convened both in Europe and also in North America. And this process brought together more than 50 LPs institutional Investors, and the data is in. Venture-ESG conducted more than 26 in-depth interviews in addition with LPs, and we launched a major survey. And it's quite revealing. One out of four LPs, or only one out of four LPs, does not see the massive, negative externality risks that are associated with this. In our survey, 74% said this is a major risk to our investment thesis, and at the same time, only 19% of LPs that represent a critical mass of actors in the ecosystem, I believe more than 1.5 trillion of dollars in assets under management, roughly, only 19% feel that today they have the right expertise to address this. And yes, I think LPs today realize that technology is not neutral. What does good actually look like?

What is a good data economy and what is the associated risk landscape and all those best practices. They are emerging as we speak. So, I think it is clear today that LPs have a very important role to play. And you at Stepstone, you are in many respects a pioneer among LPs in addressing responsible AI last year, which is very early, if you look at the developments in the industry, you already published a white paper, on Approaches to Responsible investment in AI, and you are thinking very heavily about frameworks. So, I'm curious, at what point did you actually decide that you need to do something on those topics?

ST: Yeah. So, you know, our first observation is that technology is the lifeblood of all organizations today. And what we were particularly paying attention to was the transition that occurred with the availability of ChatGPT, but the arrival of ChatGPT was something that took it out of the hands of specialists and into the hands of the layperson. And so, what we saw across all our businesses, not just venture and venture capital and growth businesses, but all our businesses, real estate, infrastructure, big buyout businesses, is that all started to become aware of the availability of this technology. And very quickly we started to see companies move to but actually create and adopt. So now what we have is businesses, you know, developing applications, applying pre-existing applications, maybe that someone else has built. And quite quickly we can see benefits coming through. Um, but clearly we also are well aware of the risks. And so, when we put on our sort of going back to first principles, we put on our hat as a fiduciary, our fiduciary duty says that we have to be focused on maximizing financial risk adjusted return for our clients. And I've just said that something like I can impact both on the cost as well as the revenue line. It is both an impact for the risk side as well as the opportunity side. And we deem it material in many, many cases. So, that's why we started early, because our job is to see things that are coming towards us and to try and be as prepared as we can, as we go out and deploy capital. You know, one of the things that we see today is that the issue that you highlight of what does good look like is the critical question, both at the LP level, the GP level and at the Asset level. You know, we sit in a fortunate position with one of the widest networks and relationships across the GP community globally. That allows us to see emerging best practice. When we take that, we combine it with the work that me and my RI team does around what's happening at the government level, what's happening in the ecosystem, the collaboration level. You know, we see that we're at a pivotal point where we can help disseminate best practice through our ecosystem.

PF: Now what I think you raise a super important point. This is a very agile ecosystem. So, at Project Liberty Institute we really look at this question of what is a fair data economy across entrepreneurship, investment policy, innovation and digital infrastructure, very important point. And we let fascinating consultations together with the Aspen Institute, Aspen Digital across five continents. So, in Europe and North America and Africa and Latin America and in Asia. And so, we went out and convened policymakers, international organizations, business executives, investors, civil society experts, leading technical people, even people like Vint Cerf from all around the world to ask exactly the question, what is responsible? Technology Innovation? And the answer is there's not a list of the 12 or 21 things you need to keep in mind. This is an evolving sector and by virtue venture is a risk industry, this is investing in cutting edge technology. And so, we really need to look at this adopt a holistic approach so to say and look at how we build technology. Invest into technology commercially deploy them, and also, of course, regulate technology. And in this innovation cycle, LPs have a pivotal role, and I believe them a very important role across your your GPS, because they are the ones who invest in the next generation of data and AI technologies. And we at Project Liberty Institute also convened something that was called the Fair Data Economy Task Force that brought together luminaries, brilliant minds, executives across the spectrum of stakeholders from more than ten countries, including Nobel Prize economics winner and what not, and they also said in their blueprint for innovation and growth that the role of LPs can have an impact and will have an impact, like we saw, for instance, in climate and other fields on the trajectory of the data economy.

And, you know, we better get this right, because we are talking about a multi trillion-dollar industry. And those questions that you just raised about what is responsible AI and data. This is fundamental to the growth of GDP and to the innovation capacity to entrepreneurship all around the world.

ST: But like at the crux is the question of how do we risk manage cutting edge technologies?

PF: It's a mouthful of a word, but there's something that is called the Colling Ridge dilemma. So maybe not a lot of people have heard about this, but I find it one of the most important concepts to understand what is going on. It basically says, in very simple terms, that you cannot measure and understand and comprehend all the impacts of cutting-edge technology until it is mainstream adopted, because it just doesn't have the scale to have all the impacts on the markets, on, the social fabric, on democracy, on our institutions and our systems until it is deployed. So, think about it in very simple terms. If you have a speedboat, right? That is small, you can course correct very fast, no problem at all. But once you have a certain maturity in markets, it's more like a container ship. So, it's very slow just changing a bit. The course to the left or the right. It's an effort. And, oh my God, better not try to stop because this is like, you know, once it's in motion, it's in motion. And technology is a bit like this as well. Um, so there's no easy answer. But we need to be agile and adaptive as we learn. And you said it yourself, there are both the social dimensions and the lessons from web 2.0. So, there's social polarization, there's algorithmic amplifications, there are implications on democracy. But everybody who's in venture should also deeply care about the economic dimension. Who owns the data? How competitive is the entire market? What is the impact on the distribution and the generation of economic value? And what kind of digital infrastructure is built now that is probably going to stay as a market structure? This has an impact on returns, on how innovative you can be also at an early stage in the innovation cycle, and how open the market is to entrants. And this is what a lot of apps, especially in the venture segment, are in the business of, and I think there's a third point that a lot of people often forget, which is you mentioned yourself how important it is to talk about both regulation and also governance, and I believe a lot of the governance is actually embedded in the technology itself, how it is set up and, and, and but at the same time, we need policy innovation for those new technologies as well. And very often you find a situation where it's not, it's everyone in the ecosystem that struggles to understand the effects and the negative externalities.

And as a regulator, if you embrace cutting edge technology, there needs to be some form of experimentation, of sandboxing, of testing new things. And often they have limited bandwidth and resources, also from their side to embrace this agile iterative approach that we actually need. So, we must remember that there will not be one single framework that will figure it out. I think on this question of what good looks like, we can agree on a number of things, because there are learnings from web 2.0 and you mentioned things like market concentration and so on and so forth. But it is a collective learning journey, I think. But LPs have, and you have, I believe, where you sit, a very attractive value proposition because you have a sort of bird eye view across your portfolio of GPS and you understand what your GPS are doing, what works and what doesn't work. And I believe in peace. And I'm curious how you see this from where you sit, have really a role to play to help GPS figure out this highly complex topic that everybody is facing in the industry right now.

ST: Yeah. It's interesting. You know, what we see consistently, actually, is that it's very often founders that recognize the power of the technology that they're creating, and they are really struggling with

their legacy. What is this business that they are creating? What does it say about them? What does it say about their legacy? And so, you know, there's some pretty profound sort of implications of some of these technologies, right?

You've said it yourself data rights, then you're impacting, you know, creatorship, you know, security, you know, you're touching people's health records, education records, that could live with them for a lifetime, right, the implications of your business model, you've talked about democracy. These are big, meaty issues. And very often the most natural place for founders to turn is to their trusted GP or the GP syndicate. And it's really, you know, they're looking to their GPs for guidance about how to navigate this complexity. And one of the first things that we really see, and we've seen this before in, in other guises, you know, the first step of getting governance right around complexity, setting up the appropriate policies that set out the envelope of the issues that need to be thought about, setting up the appropriate steering committees with the right people there, to help lead the ship and navigate through those issues. You know, very often it's about having the right people who know the right questions to ask, even if they might not know all the answers. Once you have that governance in place, then we see a natural progression to implementation. In our paper, we certainly cite some of the work of NIST, the National Institute of Standards and Technology out of the US, that has a very considered risk management approach, particularly targeted at the company level. And we've looked at adapting that to the GP and up again to the LP level. But essentially, what they've done there is talking about, and this should come as no surprise, you know, implementation with regard to development, with regard to evaluation, testing, auditing, feedback loops. None of this is surprising, right? But doing it well within this context and doing it well for specific applications, you know, that is where the learning curve is emerging now. We think that GPs are very well positioned around this. They can get appropriate, you know, real specialists in particular areas, however narrow involved in the discussion spreading that cost across their portfolios. And absolutely, this whole effort is going to take a lot of time and a lot of cost. But long term, we do believe that it's going to unlock a lot of value. It also should protect against downside risk, right. As an LP we don't want to see litigation.

PF: Yeah.

ST: Is that you know we need two stacks to move together. And that is...let's call it the entire development stack from those creating the foundational layers to those integrating to those creating applications. That entire stack has to move together because you can't have one part of the layer trying to do effective risk management. Everything builds on the next layer, and it really is the same, we sort of see with the ownership stack between the LP to the GP to the company, we can be a reinforcing ecosystem to try and drive, as I said, both stacks together.

PF: I personally believe, and I'm curious what you think about this, that for a piece leaning into what is good governance, what is good regulation? What is responsible conduct? Is actually a genuine strategic advantage in investments in data and AI, because it's fundamentally about de-risking investments over the long term and making sure that there are there's risk mitigation and better outcomes for economy societies and ultimately returns. So, from where you sit, do you think that

those topics will become even more important given that this is an industry that is entirely about, well, risk mitigation, of course, taking risks but also mitigating risks.

ST: You know, I think so. But one of my concerns is that when things get overly classified, for want of a better word in the risk management basket, it often gets conflated with overregulation or like compliance. It does not get appropriately sort of attributed towards value creation. So, it really is sort of navigating this path through to ensure that these actions are clearly linked with value creation. The other concern that I have is, you know, we're seeing different pathways across different countries, as they think about the governance of these topics and to Europe and, you know, from our side, what we're particularly concerned about is that from...if we take it from the consumers perspective, the consumer, you know, here is, you know, is using an application that could have been developed anywhere. You know, we've got people downloading Deep-Seek, right? There's no clear way for the consumer, whether that's an individual or corporate, really, to tell about the health of that system, really for to allow them to gauge the risks that they're taking on, as they're embedding it within their own company, you know, their own sort of HR systems or healthcare systems or governance systems, right? And so, you know, the fragmented nature of how regulation is emerging globally really is at odds, which are at odds with what is fundamentally a global technology. And, obviously, one concern is that companies look at regulation arbitrage, right? Trying to drift towards setting up businesses in what is perceived is in a lower regulation, environment...

PF: Accepted in the minds of actors. But, in the foreseeable futures, I don't believe that there will be consensus, so, we need to have those agile iterative approaches...

ST: Is upfront. You talked about one of the big risks being social polarization, you know, the algorithmic echo chambers that have emerged. And, you know, a couple of interesting, observations on this. Yes, today this sits, according to the World Economic Forum, as one of the top global risks, which frankly, I mean, it is astounding, you know, it is a key risk that is undermining free and fair elections. But more than that, it seems to be undermining society's ability to communicate, right, and to engage in debate about really difficult issues. And, you know, when we think about everything that we've spoken about how to develop a data economy, rights and practices, and, you know, I've heard one of the most eminent researchers on AI said that, you know, "AI will become the operating system for the world", s. o, we're talking about a really important issue here, but it's complex. And, you know, one of my concerns is at just the very moment where we need our ability to communicate, to engage, to debate. You know, the risk of that getting undermined by the very technology, you know...As I've said, I have to reiterate the point that I agree with the concept of being agile, we are going to have to be iterative, but for that to work, we have to have feedback loops, we have to have debate and discussion, and we have to have the whole stacks working together, whether they the development stack or whether they're the ownership stack.

PF: Yeah. I cannot agree more. And this is like at the core of the mission of Project Liberty, and this is also something that is coming out of the Fair Data Economy Task Force that I mentioned, there is a need for, probably also some form of business model innovation, because the kind of societal risk you were just referring to, they are negative externalities of certain business models by the way, or like venture backed in majority, that optimize for engagement, and very high engagement, because well, how revenues are generated. And this does not lead to, as we now understand, and there's the



evidence in, from mental health issues to societal polarization. But there are two ways to approach this topic: one is to criticize and state the problem, and the other one is, and you used the word yourself, is to see this as an opportunity to innovate, this can be fixed by technology, so, this in itself, and I believe it's the same mindset that one should adopt for artificial intelligence and the data economy at large, but I think the social networking part that you reference, is sort of the core example really, and the foundation. If you give people more agency over their data, if you have less concentrated, more distributed models, if you innovate also at the infrastructure stack, on top of which all the applications in which the VCs invest and operate on, and you encode a certain number of, more dynamic factors that give back users and companies some form of control over the data they give to intermediaries, to models, if there's a way to distribute value in a more agile way across the stack, this can foster healthier business models that are not optimized for rage, but optimized for better outcomes that are great for investors, for entrepreneurs, and for the users alike. And I think this is the sweet spot that needs to be found also for the AI industry because this is about a market that will evolve over probably decades to come. But now the direction is of course set, but ideally it is a market that scales towards the right direction, and all actors in the ecosystem have, of course, a role to play. But I believe LPs have a very special role in this entire constellation of stakeholders, because you just have such a bird eye view on what is happening. And I believe that also through frameworks such as the one that you started developing with your white paper, you can be really at the forefront of helping GPS figuring this out. Who will make those decisions? Who will be the next technology kings? The ones that will build the next large dominant mainstream companies? And hopefully we optimize for the great outcomes, both financially but also for societies and democracies at the end.

ST: Thank you. Paul, that sounds like a good positive note to end this discussion on. Very much appreciate your time.

PF: Thank you so much for having me.

AM: That wraps up today's episode of RPM. Thanks to Suzanne and Paul for the insightful discussion, and to our audience for tuning in. For more on this topic, visit StepStoneGroup.com to download our Responsible AI whitepaper and explore our full thought leadership library. Be sure to subscribe to RPM wherever you get your podcasts.