

SM305 | 10.17.2026

Special Episode | SmarterMarkets Update: The State of AI

Dr. Tricia Wang, Co-Founder & CEO, Advanced AI Society

This week, we present *SmarterMarkets Update: The State of AI*. At SmarterMarkets, we've realized that AI, like life, moves pretty fast. If you don't stop and look around once in a while, you could miss it. So we'll be bringing you periodic updates on the state of AI interspersed with our regular podcast series.

For our first update, our guest is Dr. Tricia Wang, Co-Founder and CEO of Advanced AI Society. David Greely sits down with Tricia to discuss how we can rebuild trust in the machine-to-machine age of AI; the role of open verification and proof of control; why the AI world should stop writing off what the crypto world built, and in particular, why the crypto ethos of 'don't trust, verify' is critical now for AI safety and adoption.

Tricia Wang (00s):

We would have to work even harder now to fight for human judgment, to maintain human judgment inside enterprises. I think those are the enterprises that are going to succeed and those are the enterprises that are really going to know how to exploit AI and use that to move better than their competitors. Because they are really going to know, just like in the big data age, when to use the more abstract scale quantification, which is using your LLMs and using all the agentic work now that can be automated and executed at machine speed. So there is that whole machine speed world, and the enterprises that will really know how to leverage it are the ones who have a solid understanding of the human judgment, human insights and human behavior and how to figure out this linkage and that's my whole career has been built on working with enterprises on doing that, but I see this is going to be a new challenge. I see it in all new, it's like replay, right? Now it's new challenges in the machine-to-machine world, but most startups are just like, oh, just agenticify, do this world without, they skip this part of the human part, right? Actually understanding, like, what are the nitty gritty details of policies, of compliance and how does that tie to this human's work and this person's job and this function. If you don't map out if you don't have a solid understanding of the human model then and you just throw agents, it's all going to collapse.

Announcer (46s):

Welcome to SmarterMarkets, a weekly podcast featuring the icons and entrepreneurs of technology, commodities, and finance ranting on the inadequacies of our systems and riffing on ideas for how to solve them. Together we examine the questions: are we facing a crisis of information or a crisis of trust, and will building Smarter Markets be the antidote?

This episode is brought to you in part by Abaxx Technologies. AI agents are fundamentally changing how people and organizations interact with computers and the internet. A race is underway to establish the path that will govern the nature of these interactions. Abaxx Technologies introduces Abaxx Labs as its center for engaging with the developer community through the release of open source software to promote a path that empowers users to maintain control of their identity, data, and agents in digital spaces. Abaxx Labs' first release is the Agents++ library, a subset of Abaxx's ID++ software development kit that has been tuned for use with AI agents. Download Agents++ on GitHub today at github.com/abaxxlabs. That's github.com/abaxxlabs.

David Greely (02m 45s):

Welcome to a special SmarterMarkets Update: The State of AI. I am Dave Greely, Chief Economist at Abaxx Technologies. Here at SmarterMarkets, we have realized that AI, like life, moves pretty fast. If you don't stop and look around once in a while, you can miss it. So we will be bringing you periodic updates on the state of AI interspersed with our regular podcast series. For our first update, our guest is Tricia Wang, Co-Founder and CEO of the Advanced AI Society. We will be discussing how we can rebuild trust in the machine-to-machine age of AI, the role of open verification and proof of control, why the AI world should stop writing off what the crypto world built, and in particular, why the crypto ethos of 'don't trust, verify' is critical now for AI safety and adoption. Hello, Tricia, welcome to SmarterMarkets.

Tricia Wang (03m 36s):

Hello.

David Greely (03m 40):

Thank you so much for joining us as we start these updates on the state of artificial intelligence. You have worked at the intersection of AI security, governance and trust and have spent more than two decades advising Fortune 500 companies in areas that include data strategy and AI adoption. So we really appreciate you making the time today to bring us all up-to-date. And it's a great time because you have recently co-founded the Advanced AI Society. So I thought it would be great to start with why did you create the Advanced AI Society and why do it now?

Tricia Wang (04m 19s):

Yeah well, I have been preparing myself for this moment. I know this is like a really it's a moment of great change, but essentially over the last 20 years of digital transformation, enterprises, society, we have been really maybe unknowingly building up to this moment of where we are entering into the machine-to-machine era of our humanity, of our existence as humans and before I went into business and before I went into becoming an expert at how enterprises use data, I am a sociologist first and foremost. And as a sociologist, we study the great changes and transformations. And I think in this moment that we are alive, that I have been through two transformations. Most of us have who are listening, which is that the transition from human-to-human to human-to-digital, which is where we interface with machines as human beings. I used to say we are on the precipice, but we have now entered into the machine-to-machine era where this means everything that we built for the human-to-machine era assumed that there was a human involved. And now we have entered into this new era of machine to machine. And one of the big problems I foresaw for this era, from the machine-to-machine era, is that we need to rebuild our trust infrastructure for how do we get things done. How do we delegate stuff to machines? How do we know who did what? Who is who?

And so I started I co-founded the Advanced AI Society with my co-founders because we fundamentally felt that there is a new type of AI that is out there that is really great for this new trust infrastructure, but no one knew about it and it's called verifiable AI. It really focuses on the core need for this era, which is around verification, verifying what happened, who did what, where did things go? And it's very different from this current AI that we have, which we call claims-based AI, which is you use the AI and it makes a claim that's going to protect your data or whatever, but then it's just a claim. There's no way to verify it. And we call that 'trust me, bro' AI.

And so Advanced AI Society, the Co-Founders, Michael Casey, Sheila Warren, and Bettina Warburg, we all saw this need for that. Everyone is scared of AI, but if they only knew that there was another way to use it, that we have something beyond just 'trust me, bro' AI, that for higher risk AI situations and contexts and deployments, that we have this whole verifiable AI world. How about we live in that world where it's trust no bros. Let's just trust the math. Let's trust the verification technology and these technologies I knew and my co-founders, we all are very intimate with because between the four of us, we have been investing in this space. We have been researching it. We have been writing about it, creating the narrative for creating policies for it, and that world is called decentralized technology. It's also referred to as cryptography or cryptocurrency because there is a lot of these kind of new policies that arise. And a lot of people associated all those policies and those products with essentially douchebags because the crypto world has really screwed itself over. You know, it has not done itself a favor of really explaining to society the value of machine speed transactions, and so most people associate crypto with being very scammy. And I would say that's pretty correct. If you look at the news, those are the headlines that dominate.

But if you actually disentangle all the terrible news of how people have been screwed over and scammed and hacked, there is actually this incredible layer of technology that's a latent market of primitives and protocols that have been built to allow the things that we want to do online with machines to actually be verified in a way where you don't have to trust anyone. And crypto has been working on that for a while, but we were like, obviously no one knows about this whole new group of all these new products and these builders because it didn't start off well with the crypto industry as a reputation. So we felt that we needed to make a society to essentially say, hey, you know, a maturing industry has a organization that represents the interests of the builders and the deployers and the ecosystem to explain what it is to advocate for and to really just parse through the noise. There is a lot of noise when you look at decentralized tech or whatever word you want to call it. So we wanted to make a society that says, here are the real players. Here is the vetted mature players that are operating this market. Abaxx being one of them, you are all a member of a society. And these are the people who are creating real tools, real solutions that enterprises, that society, that civil society, that governments can trust.

David Greely (09m 11s):

Fantastic and I would love to dig into the tools a little bit more deeply in that as you said, I think people come to this and AI equals LLMs and cryptography equals crypto, but as you said, there is a lot of tools that were developed and, you know, first put into applications in

cryptocurrencies, but they can be used for so much more. And I wanted to ask, how do you see those tools being used to take us from ‘trust me, bro,’ the world of today to a place where trust is very verifiable?

Tricia Wang (09m 47s):

Yeah. So what you’re asking is how do we take all these technologies and how do we mainstream it, right and how do we make sure that people know what to do with it? And so our big thesis is that right now we have been in, for AI adoption inside enterprises. It’s really just like we are right in the initial phase where companies are trying to figure out how to integrate this. But as you see, if you look at the trend and if you talk to people actually inside enterprises and you go and you say, are you moving past pilot? Are you actually scaling what you are doing? Are your agents actually crossing clouds? Are you allowing them to, let’s say, complete a trade or transaction? What you hear from the CISOs, the security leaders, right, the people who are actually making decisions are on the line. The answer is no. There is a lot of experimentation and of course, like I think a lot of the AI growth has come from LLM use and token use, but we haven’t even seen what agents could do yet because a lot of that is not happening inside enterprises.

As an individual, there’s a lot more happening because the risks are a lot lower, but for enterprises that’s regulated, especially if it’s about finance, one bad decision can land you in jail because the liability is unclear. But it’s right now falling on the deployers. And so what needs to happen is that enterprises need to be able to say, okay, we now know that the AI market, the AI industry to really trust the industry, to trust AI is that the most of the focus has been saying, can we align a model? Can we say this model is trustworthy? Do we know what it’s done? Can we say at the point that we tested the model function as like lower hallucination, that has been the main quest of AI safety. If I want to boil down all the research papers is pretty much like, can we make this model safe? Can we trust this model more? And guess what, everyone, I don’t think AI safety people always get excited when they, or like they are not too happy when they hear this, but, and I don’t think it’s a bad thing, but the reality that anyone working inside AI safety testing models will tell you is that testing, getting a model, aligning a model is absolutely necessary. It’s critical work, but it’s insufficient for deployment. And so what do we do? We need to focus on actual deployment, getting agents to be useful. Where, how do you lower the risk and make them valuable? That means you need to let them be non-deterministic. You don’t want to constrain them. You want them to operate with intent. You want them to be able to make decisions on the fly because that’s what you want an agent to do. But then if you lock them down too much, they’re not going to be that useful. So you are really in a bind. We call this the risk to value bind, is that if you want to lower risk, then you are almost right now forced to lower the value of them.

And so the only way to lower the risk and increase the value of your agents is to adopt open verification technologies. And we define open verification technologies, we rebranded how I told you, like, no one wants to hear about crypto and Blockchain and ledgers and distributed tech and decentralized tech, runtime attestation, TEEs, W3 credentials, audit entrails, decentralized identity, secure enclaves, like no one wants to hear about it. No buyer is going to be like, tell me about identity binding, right and digital signatures. It’s not that they can’t understand it. It’s just initially they’re practical. They’re like, tell me the problem you’re solving. Don’t sell me your technology. Tell me the problem you can solve for me and then we can talk with a tech. And so what is important is that we have to focus on the problem to be solved is trust and what crypto helped us understand is ‘don’t trust, verify,’ verify the math. And so what we rebranded all this decentralized tech, amazing tech is to say what all that is, is really called open verification and what open verification is, is something that every enterprise needs in deployment, which is where you can openly verify the action. Meaning the root of trust is not with the operator.

So in a real world use case, what would that mean? It would mean that let’s say I am entrusting this brokerage to make a trade for me. You are telling me, and I know you are using agents, right to complete that work. Cause that’s how you are going to get the best transaction, right? I am choosing you because you are promising me that you’re going to operate at machine speed and you can operate faster than this other competitor and I am like, great. I want that. I want you to operate at machine speed and make trades for me and I want you to get the best market deal, right? But you tell me you made the trade. How would I know that the trade was actually done? How do I verify? How do you even know the trade was done? And so what we need to move to is a structure where you can openly verify where I, independent of you and all the operators and deployers and entities involved to say, oh, the trade was done and it actually landed in my account. It was done on this time. None of it was fabricated and this is critical because agents are non-deterministic, meaning they are like five-year-old savants in the same way that you don’t know what a five-year-old that’s going to do the next minute, whether they’re savant or not. I mean, children are always just brilliant.

Like they are always like trying to like find ways out of the rules, right? They want to like get to their snack, their clear goal, but they’re going to find any way to do it. Even if you tell them, no, they are like, can I get a snack at this time? Can I get a snack earlier? What if I did this, right? They are trying to like reach the goal and find ways to break it. Right and that’s what agents are. They are really smart

children that we just keep imbuing with more intelligence and that means we don't really know what they are going to always do. You can't, there is going to be a certain element of non-determinism, meaning you can't predict it and you want them to be that way. You want them to be useful. You want them to make these tool calls and kind of get over these barriers. However, this means that they can also lie. Agent logs are not trustworthy. It's like, why would you trust a five-year-old to tell you what did they do with the candy? It's like you wouldn't. So we cannot trust agents. Open AI Hugging Face incident, I think when they did the audit, it was like 7%. Don't quote me on it. Maybe it was 7%, but there was a solid percentage of the log that was fabricated by the agents. And the reality is like a human cannot read over an agent file, an agent log. The way in a human logs, like I did this thing at this time or at this moment and this geographic location, and I checked, that's not how an agent log looks. You ever look over an agent log. It's not that it's incomprehensible is that it is beyond human capacity to comprehend and verify each statement, each entry, right. And then you imagine, that's one agent log. Then you want to times that by however many, you know, I have spoken, I have worked with companies in my companies where they have like 20,000 agents, 30,000 and they expect to get to hundreds of thousands of agents.

So we are at the point where in the machine-to-machine era, for the first time, we were way over the crossover point, which is that machine capacity has surpassed human capacity. The output of machines has surpassed it. And so in that world, you have to ask yourself, well why would we put the human, where do we put the human in the loop? And I think that's the wrong question. Everyone's like, oh, we got to keep the human loop. It's like, no, we need to keep human judgment in the system but in fact, we need to take the human out of the loop. The human is way too slow. And I think sociologists and ethnographers and anthropologists will attack me for saying this, but I don't want a human in the loop. I want machines that I know agents are going to watch out for me be awake 24x7. I want my own agent swarm protecting me, verifying what's happened and of course I want to maintain human judgment. That's critical is to maintain the human judgment, to design the agent swarm, to design what the goals are, the outcomes are, to design the rules of the system, what needs to be verified, what needs to be openly verified, the level of security. I will consult my agents on that, but we need to take humans out of the loop. We are so screwed if we keep humans in the loop, we need to keep humans in judgment.

David Greely (17m 58s):

And getting that judgment at the right level.

Tricia Wang (18m 01s):

Yeah, at the right level.

David Greely (18m 03):

And you recently held a conference on this event, right? You held an open source and sovereign AI event, exploring who controls that AI stack from the models to the data, to the infrastructure that they all depend on. I am curious, what was the conversation like at that event? Are these the issues people were talking about?

Tricia Wang (18m 21s):

Oh, it was so exciting. I wished you were there in the room. It was during SF Tech Week and here we are in the land of these large frontier labs where it's all about, it's the opposite, you know, if you are at a frontier lab right now, which I don't think it should be, it doesn't have to be the, you know, clashing build principles. But right now the reality is AI is in a claims-based world, right? We don't have a way to openly verify when AI is doing, so that could change. And so here we are in San Francisco during SF Tech Week, where it's like the whole world has descended upon the Bay area and really you already see everyone's moving here. There's so much excitement. Everyone is trying to close deals and everyone's trying to look for like the right VC and it's like, do you just start up a one person or two people? That's the level of debate. You know, it's like how many less, like what's the smallest group you can have? What's that new ratio.

And so there is a lot of fear and excitement in SF right now and so during that whole week, the only conference that was hosted on this topic of sovereign AI, of open source was TrustedRouter's conference. It's hosted by Joseph Perla, was organized by Joseph Perla and his new company called TrustedRouter, which allows you to openly verify the data that you put in into a LLM is not going to be used by LLM. It's using encryption technology. So they are an open verification tech that allows you to maintain your agency. So he organized this conference and I co-hosted this with Advanced AI Society, co-hosted this with Apollo Commons and pretty much how we organized, we brought together incredible builders of sovereign AI tooling. These are not like some like hackers that don't know how Silicon Valley operates. One of the companies who is our member, DC3, run by Craig Thayer. He is ex-security at Roblox and he is building a new type of agent shell for decentralized compute use and this is incredible. These are people who have incredible experience in Web 2, this human-to-machine world that they have been building, right. They know that tooling and security measures in and out and now these

are the people who are at the precipice of understanding that machines have surpassed human capacity. And so this is where the builders are moving into who have this insight. Like we have to build for this machine-to-machine world. But I think the builders here had extra insight was they are saying in this new world. I am not trying to tell you can trust me, bro. I am not just trying to create another app that like going to kind of go by the same rules of Web 2 of this human-to-machine world is that they are building for this new era where they know that it's all about trusting the math and so it was incredible to see the kind of builders that were in the room. Many of our members were there. We had Vouched. They flew in and they have been working on what's called CHAOSS, a type of know your identity protocol that's now at Linux Foundation Decentralized Trust.

We had launched the announcement of a Proof-of-Control, which is the standard that is being built by all of our members, along with 100 plus security leaders from Fortune 500 companies all around the world to have signed up to say, I need a standard that will give me proof of control over my system. And these are just people who are like, I need to do this on my free time and this is the kind of primitive that's missing. Then I think that's the radical thing we have done. If I can brag about anything we have done at the Advanced AI Society is that we have really just made this simple thesis to say, hey, instead of freaking out about whatever percent you think of catastrophic risk is going to wipe out humanity, which is all you hear about in Silicon Valley, right. The people who are building the tools are like we are doomed. You know, there is like 10% risk or 20 or 40 or whatever, 50% risk, whatever you think we are so doomed to die and so we are like okay, great, fine. That could be there with AGI, but how about in the meantime, we just focus on the controls that we need to agree to, to put on systems and then we can use these systems and they will generate wealth, they will generate all these KPIs and it's going to expand our companies and make society whatever, make things easier and why don't we just agree to what those controls are? So if they are broken, one, you will find out, two, you will find out right away and three, if you implement the most robust form of Proof-of-Control at tier four, at verifiability level tier four, is the action will not even run because if the math does not work, the action will not unlock.

And so we have created this mechanism, this standard that's a type of open verification that says you can openly verify predefined controls that you put into place and we can all agree to what that evidence package looks like and it really shows you clearly when you want to move from tier one, which is 'trust me, bro,' claims-based, assertion-based AI to tier two, which is 'trust those bros,' attestation and a third party evaluator, like, oh, trust the third party evaluator to say your system was great. It's like awesome, it's like the new version of SOC 2, great, annual. How cute, you know, whatever, whatever level of like certification is like trust those bros. I think we are now moving to quarterly, right? AIUC has really pushed industry forward with offering that, but our whole premise is that for a secure, agentic society where people have rights, we have agency over our futures and we are going to live in the machine-to-machine world where we can actually have any chance to flourish, we have to have the right to open verification. And you can reach that and Proof-of-Control tier three and four, which says, hey, tier three just says how tier three and four both say trust no bros, trust the math and that tier three, you will know if your control broke, you will find out right away. So in the Open AI Hugging Face incident, they let down their guards with their sandbox. So obviously the agents escaped. It's like telling a five year old, like, don't leave the house, but your doors are all open, right? So they let down the gates and so tier three, you would have found out if the agent had left the sandbox right away, everyone would have been alerted. At verifiability tier four, you would have found out if the agents even tried to leave the door and then the mechanism, the math, the equation would not have matched to allow that door to be unlocked. So the agent would have never escaped because you are tying it to hardware, the root of trust.

And so that's what we have outlined and we have developed this just quickly outside and I am not bashing on any of the standards groups like IEEE or Linux Foundation, where we actually have moved the standard, right? So we moved our Proof-of-Control standard, Linux Foundation Decentralized Trust, but there is so many standard bodies out there and we just said initially, how about it was like a renegade group of people where we all just came together and said, why don't we just predefine those controls? So it's amazing to have Michelle Dennedy be one of our distinguished reviewers. She has been involved from day one on predefining those controls from her perspective as being the first Chief Privacy Officer, but also from Abaxx's perspective of what do their clients need. And then we have Bruce Schneier, who is one of the most well-known cryptographers. We have now Charles Iheagwara, who is the Head of Cybersecurity AI at AstraZeneca. And then our co-Chair is Ken Huang, who has released countless AI security standards and he is saying this is the one standard that is missing for deployment, for making sure that you have control of your agents. And we have 100 plus contributors who have just said over the summer, which is when does anything get done? Never. Nothing ever gets done. But we did.

I think we achieved the impossible, which is we got to a version 1.0. When you release this, there's still time to comment on proof of control. It's version 1.0. We now are hosted at LF Decentralized Trust. So we don't even control, you don't even have to trust us. But it's sitting in this neutral space and what we are doing at the Open Verification Lab is really building out, I'm excited that this is why this is the beginning of working with Abaxx is that we will be building out with our members, with our founding contributors, the kind of

research and the tooling that's going to be needed to ensure that proof of control is actually implemented. Because the power of a standard dies if no one uses it. So we know making the standard alone is not enough. We need to make sure deployers adopt it, vendors adopt it, insurers adopt it and regulators make it legally enforceable. Which is why we are operating in all those areas where we have our policy institute, where we are advising on the EU AI Act, introducing open verification into the EU AI Act, working closely with ALMA-AI at University of Bologna. We have, seven of our members have gotten into, the first verifiable AI companies have gotten into the EU AI sandbox. And we're also working closely with the Chief Privacy Officer, Christopher, at the state of Utah. And at the SEDI conference that he is organizing in November, every state in the US has all eyes on Utah, looking at what are their policy, what are they going to say in their legislation for digital identity and for sovereign management of identity.

So this is why it's so exciting to be working with Abaxx, you know, as a member to be collaborating because you all have been so on top of this, so ahead of this with your ID++ protocol and Agents++. And so I have been hearing about this all from Michelle, you know, she is always, when she is advising us and looking at the requirements, she is like, well, this is what we do at Agents++. And I think that's, this is really important that we have deployers and people like Abaxx who have these tools advising on the standard, because the reality is no one wants a vendor controlled standard, no one. And so this is why Michelle and many other leaders are taking the best from their companies and saying, let's all agree to a common standard at the base of what counts as verification, open verification and then we will all build our special IP for our products on top of it. But it's really going to allow a whole industry to flourish, a whole ecosystem that's going to be critical for making this machine-to-machine era, one, a place that's trustworthy that we are all going to want to live in.

David Greely (28m 17s):

And I would love to dig into that aspect just a little bit more with you, because there's always there is the math and there is the money and as you said, nobody wants a vendor controlled ecosystem. But we have all kind of been accustomed to things like Web 2.0, where the large tech companies end up controlling the stack and controlling the standards. And I was curious, you mentioned a few ways a band of highly accomplished renegades can get around this. How important is open source for that? Does that enter into this or is that separate?

Tricia Wang (28m 52s):

For us, this is, it's good that you asked this because you know, you had originally asked me about the open source conference, Open Source Sovereign AI Conference, right, and that was the main topic last week in SF. So open source is for us, it's a moot point. We are like, open source is important, but open verification is an extension of open source ethos, purpose built for the agentic age and so open source as an ethos is incredibly important, because open source says things have to be openly inspectable, you know, there's a whole licensee structure around this. And it means that no one can control it. And so open source is critical. And so in the same way that open source has evolved into open data, open hardware, open standards, we're saying now for the agentic era, for this machine-to-machine era, open verification is that next evolution of open.

Now when we are practically talking about an artifact of the supply chain, your software supply chain, let's say open source software, then our whole and open models, open data, open weights, or are you doing local models, right? We are saying that's a fit for purpose decision and as a deployer, as a vendor, or as an individual who is building out your agentic stack, your AI stack, those are your decisions to make and it doesn't matter whether you use a open or closed source model, or open weights, closed weights, open data, closed data, local inference, whatever it is, that's great. Those are all decisions you made pre-deployment for your supply chain, right. But what happens afterwards is that your agents are going to be in the wild, and they are going to be useful. They are going to be doing things, they're going to be crossing clouds, and you are going to want to make sure that you can openly verify the high risk actions. And we are very agnostic, we are not even saying you have to always openly verify.

I think what's critical is that we have to take the ideology out of all of this. I am not against a centralized player. I think frontier labs, there are going to be moments when it doesn't matter. You don't need to openly verify if you are like, I don't know, like, what recipe should I make tonight? You are not like, oh my god, I need like the highest encryption. I mean, maybe you do, but I am just, for some people, it doesn't matter and for some people, it may. But what's important is that we have the tools to make that decision for each individual, each community, each enterprise, each deployment context to say, when do I want to use 'trust me, bro' verification, right? And that's okay, no big deal. When do we want to have 'trust no bro' AI, or when are we okay with 'trust those bros,' just like an independent evaluator coming in checking every quarter, right? And it's really up to the entity that's doing the deployment to make that decision, because there are costs, there are decisions, right, you have to make in implementing that kind of 'trust no bro' architecture. However, it's going to get easier. And this is what I see being in San Francisco, what I see in terms of all of our members, is

that until you actually gather everyone, it's really hard to see, because it seems like it's just scattered people, and they are really not talking together. But the whole point of early on of creating these kinds of consortiums and alliances is being able to come together and say, there is a real signal here and we as a consortium, you are really able to advocate and show that if we agree to these principles and these standards and these tools, it actually expands the market, because it makes it actually grows the market. It's not just there for the incumbents. And so this is why open verification is so key and this is also why you need to be able to make the choice of whether what kind of open you do and at what part of the stack.

David Greely (32m 32s):

And when you look on the other side with businesses that are looking to adopt AI, how well do you think they understand these issues and how much do they need to learn because it's all been happening so quickly.

Tricia Wang (32m 46s):

Yeah, the reality is every single business is different, every single business. I say the people who know the most, actually, if you look at people who are the channel integrators, they actually are highly like they're much more aware of this, like the SaaS providers, right, because they are really trying to figure out how to help their clients through this. But the reality is that this is moving faster than any of us can keep up. So it is overwhelming. It's really hard to see what's the signal and what's the noise. And so this is when we are talking to enterprises, what you see is it's really siloed, the knowledge and this really mirrored what happened initially with digital transformation with big data and really AI is an extension of data. It's not like it's anything new. The mechanisms, the algorithms have always been there, but it's just like we have more powerful products and better products now and way better models, better mechanisms.

But it's not like it's new to say, hey, we are going to use these abstract representations to help us make decisions, right and so enterprises went through the same struggle initially with the digital transformation when they moved from said, we are going to be big data and you saw this happening inside enterprises where initially it was like a couple of centers owning that expertise. It was not federated, right. And so it was like bringing in like the data scientists and build out the data center, data lakes and the data warehouses, but that data was not living in the business units. You see that same thing being modeled right now with AI is that a lot of the realizations of what you can do with it are still being figured out by this is why you need a Chief AI Officer, right. Eventually, we shouldn't need one because eventually, just like with the Chief Data Officer, that knowledge and that skill and those processes, it was success because they made their job. They didn't need the job anymore because that knowledge and skills went into the business units, right and so right now we need Chief Data Officers, those kind of rules, because you really need someone to help get all of this stuff started, create the culture, help people understand the governance and the policies and implementation and tooling.

And right now, it's not necessarily in the business units yet. Of course, everyone's using AI, but I am talking about like really using it in a way where it's driving business growth. And you see that you talk to any price leader like, yeah, it's helping us optimize things and we are definitely getting to certain data faster. But it's not clear yet how it's actually making a difference for the customer experience. I am sure there's places where it is, but you are not hearing this across the board. And that's because it's not moving into the business units yet, but it will. My hypothesis is that what's blocking it is that right now we are really seduced by AI and some people think just by using an LLM, they are using like that's the most of the capacity of what AI is. And they are also using in ways that are they're still trying to learn how to use it. I give a lot of talks. My speaking agent always says like their most popular talk is helping companies understand that their LLM answers aren't real and it's controversial because like I will come in and it's really like who is using LLMs and everyone will use it and everyone will say they built their own internal and they like uploaded all their documents in and whatever company history, all their marketing studies and then pretty much what I am trying to get them to understand in my talk is to say, like if you are just using an LLM that you are really prone to AI psychosis in a business context, which is you start to believe that the answer it's giving you is absolutely real representation of reality when it's really just a dashboard.

But it looks so misleading. It's like as not real as a dashboard of numbers are right. But a dashboard of numbers, you can at least be like, oh, that's a number. That's not the person right saying it's right and of course, people believe the numbers and my whole career before this moment is really built on helping companies understand the quantification bias and how you want to avoid destroying your business because you just made decisions based off numbers, but that was easier for me to help leaders understand. I helped the leaders understand, like avoid being like Nokia, avoid being like Kodak. And I can tell you because I was in these companies leading teams where I worked with executives and now you don't hold a Nokia phone because they literally believe their numbers and could not deal with the data that I was bringing them, you know, because it was not in the form of a number. But now it's actually harder for me to talk to executives because they are like, but what do you mean it's not real? This is what the AI told me and I will be like, show me

the source of that study. And then there have been moments where they cannot find the study because their LLM is not pointing to the original study because it's a prediction. It's a generalization of the thing. So they can't. They will be like, well, the customer said this like two years ago. I am like, what customer, show me the customer. Show me the report. Who did the report? Who talked to those customers? Oh, those people are gone. They are not working the company more. So then I will be like in the room of executives when I finally get to them to the safe space. I'll be like, well, when is the last time anyone spoke to a customer?

David Greely (38m 03s):
Pin drop?

Tricia Wang (38m 05s):

Yes, and because everyone is busy speaking to synthetic data, busy speaking to their LLMs. And I am not even critiquing that. I'm like, use all these tools, you know, use the LLMs, use the sensitive data. If you want to make a focus group of using synthetic profiles, awesome, but you cannot believe that is reality, that you have to complement that with what I call the thick data, with your human judgment, where you get to the source. Because anything that's already been captured has already happened. It's done and even if it's giving you a prediction, it's based on a prediction. It's a probability of what the token says is going to happen, right. That's not reality. You have to go and verify. So I can only imagine like if you are really as a and you are a trader and you are in financial services and you are trying to really understand like the movement of markets, movement of markets that's really about movement of human intention, movement of human emotions. You are trying to figure out at a very macro level what is happening, but the best macro economists and traders I have ever worked with, people who really understand, they are able to quickly understand macro trends and they're able to tie it to real life stuff, emotions. Things are happening on the ground in this region that's going to impact the shipping lane. This trend is going to impact this thing because the oil because of this politics stuff, like they actually understand the intricacies of human emotions and human players, institutions and how that all ladders up and to institutions and how that then ladders up to numbers.

And it's all abstractions that we are dealing with, but when you get to the ground, you still have to figure out how to get to the ground level. I am not even saying that's truth because I'm like all of this truth, all of it is a representation. But you have got to get as close as possible to the source of where that data came from, which is the thick data, the thickness that defies quantification and that's the role of human judgment. We would have to work even harder now to fight for human judgment, to maintain human judgment inside enterprises. I think those are the enterprises that are going to succeed and those are the enterprises that are really going to know how to exploit AI and use that to move better than their competitors. Because they are really going to know, just like in the big data age, when to use that more abstract scale quantification, which is using your LLMs and using all the agentic work now that can be automated and executed at machine speed. So there is that whole machine speed world and the enterprises that will really know how to leverage it are the ones who have a solid understanding of the human judgment, human insights and human behavior and how to figure out this linkage.

And that's my whole career has been built on working with enterprises on doing that but I see this is going to be a new challenge. I see it in all new. It's like replay right now. It's new challenges in the machine-to-machine world but most startups are just like, oh, just agentify do this world without they skip this part of the human part, right. Actually understanding like what are the nitty gritty of policies of compliance and how does that tie to this human's work and this person's job and this function? If you don't map out, if you don't have a solid understanding of the human model, then and you just throw agents, it's all going to collapse and you are so at risk of AI psychosis. AI psychosis is real. In a civilian setting, we see suicides. We see people really believing they all of a sudden have a new girlfriend or a new lover. You know, you see all kinds of AI psychosis and I see people getting angry.

So I have experienced individuals going through AI psychosis in civilian settings, but I have started to see AI psychosis inside business settings. And it's scary because it has a different kind of impact. If you are a publicly traded company and if you are completing transactions that actually have material impact, consequential impact on people's lives, this is high risk. And we have yet to even see the kind of liabilities and errors. And this is on top of what we haven't talked about yet, which is in a machine-to-machine era, we are going to be so prone to zero days. So we have to also deal with new kind of cybersecurity risks. But this is why our belief is that the primitives, the tools that were used to build out crypto that really was building for this machine-to-machine world, crypto is really about machine payments, programmable money and it's not just about money, right? It's making anything programmable and readable to a machine. It's that all the tools that we use for this world now have to be used to fight the tools that are now readily available for any hacker.

David Greely (42m 57s):

Well, thank you so much, Tricia. We could talk for hours on these topics and barely scratch the surface, I fear. So I really hope that you will come back again and be our guest on the podcast in the future. But before you go, I just wanted to ask you one last question, and that's what are the next steps that you are going to be focused on at the Advanced AI Society?

Tricia Wang (43m 18s):

Yeah, so our biggest goal at the Advanced AI Society is one, to get all the comments in, everyone, the industry's feedback on version 1.0 for Proof-of-Control and get those feedback in. And then at what we are going to be doing on October 30th is for the next 30 days, we will be compiling all that and figure out what do we need to do in order to get to version 2.0. So that's our first task is to launch this as soon as possible. And so I am excited for your community to be able to take a look at this and get the smartest cybersecurity AI deployers on top of this.

And then second, what we need to do is really identify who else are these verifiable AI builders and deployers who want to start these kind of pilots and integrate Proof-of-Control. So we are looking, we have an insurance working group, we have a health working group, we have agentic finance, and all of these people, we need to come together to figure out what are the next pilots. So I am here looking at Abaxx, so hopefully we will have something coming out there. But we need to now test this out, right, in real life, in a safe environment. And then really the third thing that you can be looking out for is we have about 60 days, I would say six days to six months to really make an impact on some of the massive policies in the EU AI Act and in Utah, because all eyes are on Utah and the US policy happens on a state level, right, is that we have an opportunity, we have regulators and policymakers who are very receptive to updating their laws for the machine-to-machine world and that is a really exciting opportunity for any kind of builder who's in this space or anyone who is trying to understand the macro trends is that for the first time, you know, what we are doing at the Advanced AI Society is seeing that we can align the technical architecture, which is where open verification and proof of control operates, right, the trust infrastructure, the builders, we can align that with the regulatory play, all these things are happening in the EU AI Act and in Utah, and then we can align that with market forces around insurance companies who have been increasingly dropping AI, right, they are being excluded.

But always you look at what's happening with insurance is that if we can align the needs, Proof-of-Control and say, hey, insurers, if you look at what you can get from Proof-of-Control, you can actually start pricing AI and pricing AI risk. And if all this aligns and aligns with the zeitgeist, with right now, everybody hates AI in the public, right? If you talk outside, rightfully so, because they are like, I don't want 'trust me, bro' AI, I don't trust you. So if we can then tell the public to be like, hey, actually, another world is possible, another way is possible and look at how these institutions are aligning with the builders, with these incredible primitives, this infrastructure that was built over the last 10 years, of course, by some of these terrible, scammy people but underlying all of it, they are gone. But guess what, look at this technology and look at the really mature enterprise, like adult builders who are here, who are building out honest businesses, some of which we have mentioned, and you can use these tools and they are, guess what, they are interoperable. You don't have to abandon your cloud from your centralized provider as a business. You can be like, oh, let's just tack this on to our current cloud provider because we are ensuring interoperability, it's open standards, you can openly verify these things. And then one, so it really, I'm excited for this moment over the next six months. So it will be exciting to come back maybe, you know, in six months or a year to be like, did this happen? Did this star alignment happen? This constellation of the builders with the tech, of this latent verifiable AI, with the regulatory play, with the insurers, with the deployers who have stood up already being like, we need this. We absolutely have to have, we are not going to deploy AI, agentic AI, unless I can openly verify. If we see these four forces align, we are going to see a very different environment. We are going to see much more adoption of AI inside enterprises, and we are going to see, I think, the public zeitgeist kind of shift and be like, okay there is another option and I have more choice and we have more agency in this world.

David Greely (47m 28s):

Thanks again to Tricia Wang, Co-Founder and CEO of the Advanced AI Society. We hope you enjoyed the episode. We will be back with more updates on the state of AI in the future. But next Saturday, we will be kicking off our new podcast series, Global Upheaval in Energy Markets. We hope you will join us.

Announcer (42m 04s):

This episode is brought to you in part by Abaxx Technologies. AI agents are fundamentally changing how people and organizations interact with computers and the internet. A race is underway to establish the path that will govern the nature of these interactions. Abaxx Technologies introduces Abaxx Labs as its center for engaging with the developer community through the release of open source software to promote a path that empowers users to maintain control of their identity, data and agents in digital spaces. Abaxx Labs'

first release is the Agents++ library, a subset of Abaxx's ID++ software development kit that has been tuned for use with AI agents. Download Agents++ on GitHub today at github.com/abaxx-labs. That's github.com/abaxx-labs.

That concludes this week's episode of SmarterMarkets by Abaxx. For episode transcripts and additional episode information, including research, editorial and video content, please visit smartermarkets.media. Please help more people discover the podcast by leaving a review on Apple Podcast, Spotify, YouTube, or your favorite podcast platform. SmarterMarkets is presented for informational and entertainment purposes only. The information presented on SmarterMarkets should not be construed as investment advice. Always consult a licensed investment professional before making investment decisions. The views and opinions expressed on SmarterMarkets are those of the participants and do not necessarily reflect those of the show's hosts or producer. SmarterMarkets, its hosts, guests, employees, and producer, Abaxx Technologies, shall not be held liable for losses resulting from investment decisions based on informational viewpoints presented on SmarterMarkets. Thank you for listening and please join us again next week.