

THE BROOKINGS INSTITUTION
The *TechTank* Podcast

“Addressing ‘insider’ AI agents with Camille Stewart Gloster”



Monday, October 5, 2026

Guests:

CAMILLE STEWART GLOSTER
CEO,
CAS Strategies

Host:

NICOL TURNER LEE
Director,
Center for Technology Innovation,
Governance Studies,
The Brookings Institution

[music]

[00:00:00] **CO-HOST NICOL TURNER LEE:** You're listening to TechTank, a biweekly podcast from the Brookings Institution exploring the most consequential technology issues of our time. From racial bias in algorithms to the future of work, TechTank takes

big ideas and makes them accessible. Welcome to the TechTank podcast. I am co-host Nicol Turner Lee, senior fellow and director of the Center for Technology Innovation at the Brookings Institution. Wow, what a couple of weeks it's been. After the explosion of generative AI and some of the news headlines as what's happening with regards to AI safety, this conversation I'm having today cannot be more timely.

[00:00:47] Companies and organizations are quickly implementing various tools around how they're going to ensure that their AI agents don't go rogue. And now, as these agentic AI systems operate somewhat autonomously, this becomes a more ubiquitous concern, and enterprises are turning to a new set of tools to introduce different and more complex solutions to these governance challenges.

[00:01:12] For those of you who may not know, agentic systems can execute decisions and actions on behalf of humans. Yes, I said it. Now listen, I'm not saying anything you don't already know. This could include scheduling meetings, making appointments, purchases, generating other outputs. The capabilities promise significant advances in efficiency.

[00:01:34] In fact, my doctor uses agentic AI for me to actually get an appointment. But this does not come without its own risk, and that's what we're gonna talk about today. I'm joined by my friend Camille Stewart Gloster, who is a researcher, writer, educator who works at the intersection of artificial intelligence, cybersecurity, and geopolitical risk in both the private and public sector.

[00:02:01] And I'm so proud to know her because she was the first deputy national cyber director for technology and ecosystem security at the White House, and the former global head of product security strategy at Google. This range of experience informs her new book, "The Insider You Built: How Organizations Stay in Control of Autonomous AI Agents," which I think, after looking at it, offers new answers to many of the critical questions that we're asking today. Camille, thank you for joining me.

[00:02:32] **GUEST CAMILLE STEWART GLOSTER:** Oh, it's my pleasure. Thank you for having me. I'm excited.

[00:02:35] **CO-HOST NICOL TURNER LEE:** I mean, listen. I know. Me too. I mean, your book was so timely, I have to say.

[00:02:39] **GUEST CAMILLE STEWART GLOSTER:** You know, I wanted it to come out six months before it did, but I couldn't have asked for a better moment where people would be more receptive to it than the one when it came out.

[00:02:49] **CO-HOST NICOL TURNER LEE:** And your book does something very unique. It sort of makes AI and cybersecurity bedfellows. So I wanna start the podcast sort of diving into why did you write this book based on these experiences you've had, and tell us a little bit more about what's in it.

[00:03:04] **GUEST CAMILLE STEWART GLOSTER:** Yeah. So when I left the White House, where we had just released the AI executive order and I was leading AI security, I started, helping companies and organizations of all sizes navigate AI deployment. And what I quickly realized was they had no idea what they were doing. They were chasing innovation And kind of sacrificing the infrastructure that we usually have when we either deploy a new capability or even more so as we move to agentic, as we allow another actor into our environment. And while agents aren't people, they can act. They, they essentially make decisions. They have permissions and an identity, at least they should, and the way they move through your environment is much like a person. And this insider risk analogy became very prominent and one that I felt needed to be articulated more broadly for organizations to understand how to actually control this entity in your organization.

[00:04:09] We often have insider risk programs in companies that, um, kind of constrain the behavior of your staff and kind of direct it. People make, uh, choices intentionally or unintentionally that can cause harm in an organization and the goal is to find that out and prevent it. And we need to do the same thing for agents. They are misconfigured or have emergent behavior or just have this by any means necessary mentality that means that what you th- the path you think they might take is not necessarily the path they're gonna take.

[00:04:42] **CO-HOST NICOL TURNER LEE:** Well, let's stop for a minute because I think you and I know agents, like what they are, what they mean.

[00:04:47] You know, can I tell you, whenever I think of an agent in AI, I think of these like minions with trench coats ... kind of walking around. You know what I mean? As, I cannot get that out of my head. I think somebody said this a couple of years ago and all I see are like these little things with trench coats, like without no human supervision. Tell folks what do you mean in the book by agents? So, because there's so many definitions, Camille, about, you know, agentic infrastructure that I find it to be quite interesting that it's gonna get bigger and bigger ... in terms of functions, you know?

[00:05:16] **GUEST CAMILLE STEWART GLOSTER:** I'm so glad you asked that. I liked your minions.

[00:05:19] **CO-HOST NICOL TURNER LEE:** Yes, it's, it's a minions.

[00:05:21] I know I said it first, but I, I'm not the owner of that analogy, okay? Okay. Somebody else said it.

[00:05:25] **GUEST CAMILLE STEWART GLOSTER:** Okay, fair. Fair. Um, you know, one of the things I talk about in the book is that I don't want people to quibble over the definition. To your point, there are a wide range of definitions of what is agentic, and I encourage people to focus on the point at which The tool is taking action without immediate direction. So not an if-then statement. You know, we all remember being able to set up, like, if this happens, then that happens. Not that, but where a tool is actually synthesizing information, analyzing it, and can make a bevy of choices to get to an outcome. When you have that, I don't care if it's got the harness, if it doesn't have the harness, if it...

[00:06:06] You are in an agentic territory and should be thinking about how you govern that behavior and understand that behavior and have visibility into that behavior as if it is the most sophisticated of agents.

[00:06:19] **CO-HOST NICOL TURNER LEE:** Well, and it's so interesting to me. I mean, there are high numbers of industry right now that are deploying agents, right? Like I mentioned, whether it's the appointment agent, right, or it's an agent that is helping to sort of synthesize a variety of data inputs for the, uh, company, you know, wants to make sure they have better client efficiency. I find it so interesting that companies have sort of moved quickly on the agentic side to some respects. Not all, but to some respect. So the question I have for you is, like, work- are companies, like, prepared for this? I mean, I thought we were still just trying to figure out predictive and generative AI and now- ... you know, it's sort of like they've leapfrogged to the use of agents. I mean, where are companies when it comes to the use of these technologies? Because I think your book is very relevant around, you know, managing security risk as well.

[00:07:08] **GUEST CAMILLE STEWART GLOSTER:** Yeah, I mean, one of the things the book talks about is kind of breaking down these silos between security, safety, enterprise risk management, all of those things because the largest instance of shadow AI right now is increasingly agentic.

[00:07:23] **CO-HOST NICOL TURNER LEE:** Wow.

[00:07:23] **GUEST CAMILLE STEWART GLOSTER:** You know, many of these generative AI capabilities, you give your, your team Claude and it's got Claude Cowork, ChatGPT Work, all of these capabilities that can look across systems, pull data, analyze

that data, make decisions, take actions on your behalf. Most organizations were kind of chasing or following the innovation, trying to experiment and didn't wrap their, take the time to wrap their heads around the, the full scope of the problem at first. And there are organizations that are really leaning in to continue to learn and be agile around, okay, well what does this new discovery mean? What does this new action or behavior mean for how we govern and understand agents in our environment? And some are just kind of hoping for the best, which is the scariest of territories.

[00:08:11] **CO-HOST NICOL TURNER LEE:** Well, and what I like about your book is, I mean, you're a policy woman as well, right? I mean, this is not just, you know, Camille as the cyber expert. You've had a lot of experience thinking about the policies that move any type of technology. 'Cause I'm sure when you were in the Obama administration, you weren't talking about agentic AI.

[00:08:28] Uh, okay, right, I was gonna say, I think we were just talking about cyber- We weren't even there yet. We weren't even there, right? I mean, I was, like, thinking about it the other day when I knew I was interviewing you. So let's think about it for just a second. I mean, most of the framework that's been out there when it comes to technology, um, the, well, let's just say the safety and security of technology innovation, has been within a cybersecurity framework or resilience framework. I'd love to hear from you before I go a little bit more into the details for those who are listening around, you know, what you're suggesting they should do, uh, to make sure that they're, uh, managing the risk of agentic AI. But, like, why is it so important, Camille, to have this conversation around what cyber looks like today in the age of AI? 'Cause I think the book really brings that out, too, in terms of governance is really required as well.

[00:09:14] **GUEST CAMILLE STEWART GLOSTER:** Yes. So first, I decided to focus on organizations because we are in a policy moment where policy is not moving fast enough. Um, we're not wrapping our heads around it enough. And in certain places, like here, we're paused on action at a national level in terms of a concrete AI policy. And that means that industry is on the front lines of writing the rules and the societal norms and the way this technology is going to impact our lives. And unwittingly, organizations have stepped into a role that they weren't designed to fill. But also, even if you think about this from a maximizing business value, getting the m- the best return on investment, these decisions that they need to make, these investments in governance, and I, I'm trying to expand the term governance 'cause when you hear governance, you think compliance, risk, policy only. But when I say governance, I truly mean, like, wrapping your hands around this capability. And it... That's why it's such a cross-functional discipline. You need security tooling, security frameworks, um, security

mindset, but you also need to understand enterprise risk management, all of the safety features, um, the ethics questions.

[00:10:32] All of those things come to play in an agentic environment because it is very reminiscent of an organizational environment writ large. And you've got all of these disciplines, right? You- when you hired people, you've got HR, you've got legal, you've got recruiting that brought people in. You've got a manager to oversee them. How do we create that infrastructure for this new actor in our environment? And cybersecurity is at the center of that, yes, but it does not act alone.

[00:11:02] **CO-HOST NICOL TURNER LEE:** Yeah, I mean, it makes sense. I mean, at the governance level, is this a conversation that policymakers need to be having now, right? Because it appears with all of the things that have been going on lately in the news, uh, like, you know, OpenAI agents recently, I think, hacked into a UN, um, website and- Yes Australia was actually a victim. I mean, how do policymakers sort of reconcile what security now looks like, you know, both from a technical standpoint but also a policy standpoint?

[00:11:29] **GUEST CAMILLE STEWART GLOSTER:** We have to get better situated to have more ali- agile policy development. We love to get it right, and we are letting perfection be the enemy of good. And one of the things that I think we've started to see in how the EU does policymaking is they write something, and then they know you're gonna react to it, everyone's gonna react to it. And they make tweaks to it and adjust. And That adjustment needs to continue even after the legislation has gone into effect. And that is a model we need to start to adopt more broadly, that we take some time to understand and collaborate with industry, collaborate with academia to understand the long-term impacts, this kind of cross-sector development of policy. But we've got to develop something, and we definitely have to articulate the red lines, and that is work we have not yet done. And so our organizations, our, our industry is kind of writing its own rules and actually does want that guidance.

[00:12:34] **CO-HOST NICOL TURNER LEE:** Yes. Yes. Yes. But then it's so interesting, too, if I can, um, you are a cybersecurity person, and I don't know, the cybersecurity people have been around for a long time and have their own norms and standards, right? Yes. And then in comes AI. So from the standpoint of that community, are there any changes or adaptations that they need to continue to make? Because I, I just, I just think that AI has not come with a set of clearly defined, um, norms, uh, which makes it so hard, right, to do policy. You know, I was just talking about this with somebody. It's not like the internet, you know what I mean? It's something different. So I'm just curious from you, like, the cybersecurity community, who's so much leaning into ISO language

and other things- ... Camille, right? I mean, how are they, like, handling and managing this? And, like, on their terms, like, what would be ideal for them to be better stewards in this, uh, new, innovative period?

[00:13:28] **GUEST CAMILLE STEWART GLOSTER:** Yeah. I mean, everybody's gotta be more agile. So yes, the discipline will change. A lot of the foundations remain the same, though. One of the capabilities that organizations need to develop is platform engineering. Engineering with AI agents is going to be an even bigger part of advancing security and safety because you've gotta be able to actually make the adjustments needed. Governance requires runtime enforcement, which means enforcing your governance principles, your rules at the point of action. It cannot be that you just write a policy and it lives out in the ether. You've gotta combine that technical apparatus to say, as an agent is making decisions, we're able to flag that this one needs additional review, this one needs, can move forward, and this one we need to just shut down immediately. And, um, and so those are some of the the changes that I, I would like to see, as well as this kind of moving more quickly. You know, in security we've got these 30, 60, 90-day patch cycles to patch software and to close security bugs. Those things take too long. Continuous learning, this, this aspect of seeing an incident unfold, seeing an emergent behavior in an agent, uh, make itself available, seeing a misconfiguration and changing it across all of those different s- parts of the stack, organizational, operational, and technical, as quickly as possible is muscle memory that organizations don't come to easily, but we have to adjust to.

[00:15:00] **CO-HOST NICOL TURNER LEE:** So you talk about this in the book, um, this proposed framework, the authority-centered enforcement and attribution, or ACE framework. Explain to us what that is 'cause I think what you're talking about and why I really like the book is that you're giving people tools to be able to navigate this space of uncertainty, which I think is really helpful, right, at this point in time.

[00:15:20] **GUEST CAMILLE STEWART GLOSTER:** Yes. ACE is that governance in action that I'm talking about.

[00:15:25] **CO-HOST NICOL TURNER LEE:** Mm-hmm.

[00:15:25] **GUEST CAMILLE STEWART GLOSTER:** There are seven steps, but they basically break down to articulating your authority. So how much authority do you wanna delegate to this tool? And that doesn't just mean I want it to do X. It also means I don't want it to do these things. Here's where it should stop. Here's how, um, much money it should be able to spend or, you know, how long it should have access to a system to be able to achieve the goal that I've given it, all of those things. It is thinking

through the detection and governance of it. So what are the broader system-wide aspects that we need? How do we understand its identity? What permissions should it have? Um, who is managing it? We need to know who is accountable within the organization for that agent, for thinking through its threat model and communicating that to the security and the IT team so that continues to evolve as you understand its ability to have impact on your customers and on your infrastructure in turn. And then it's that runtime enforcement. How do we ensure that at the point of action we can pull back a behavior, that we can identify where, um, an agent's actions are going to deviate with, from the objective that was set? And then the last part is that continuous learning. How do we make sure we adapt as we gain new information?

[00:16:48] Whether it's because you saw the OpenAI Hugging Face incident in the news and you learned some tools and tactics from that, or because within your own organization you discovered a misconfiguration or an emergent behavior or a mechanism to achieve the objective that does not align with how your organization wants to see it happen. And so the ACE framework really makes those things move with the agent so that as it takes action, as it has authority within your organization, you're seeing all of those things in real time.

[00:17:22] **CO-HOST NICOL TURNER LEE:** Well, and I think that's such a useful tool. I mean, that this is for business people, right? In many respects, you're giving them agency. You're pulling them into the process.

[00:17:30] **GUEST CAMILLE STEWART GLOSTER:** Yes. And so this book is not just for cybersecurity professionals. Right. Right. It's for everyone in the organization. That marketing person who now has been charged to create a s- customer success bot and, um, is, has on the front lines of seeing how it interacts with the customers, what kind of data it has access to, but also recognizes just how important it is to be managing this tool.

[00:17:51] This book will help you know your place in that broader ecosystem and even advocate to your leadership for more governance to build out the infrastructure that I'm talking about. There's a chapter in the book about the business case 'cause what I'm finding is whether it's security or a business unit frontline owner who owns a agent- They often are convincing their leadership that they need the resources, the restructuring, the access to information within their organization to actually be able to do these governance actions.

[00:18:23] **CO-HOST NICOL TURNER LEE:** And then I would just think, um, you know, 'cause of your background working in this space at, in the public interest level, um, what

about, um, or at the public interest level, what about people like my mother? Yeah. Are there things that they need to know? I mean, I was just in a conversation, and I think people are throwing their hands up that this AI stuff is gonna be hard to do traditional, you know, consumer education. But I do think, you know, with the amount of scams that are happening and all of the technology moving so rapidly, consumers also need to have some agency, much like the businesses. So just curious from you, like, what can consumers do as well to manage this growing agentic infrastructure? Or maybe even, like, let the companies that they, um, you know, are interacting with know that something happened with their agent.

[00:19:08] Yeah. I don't know. I'm just, I'm kind of, like, brainstorming here, Camille, right?

[00:19:10] **GUEST CAMILLE STEWART GLOSTER:** Yes. Well, so there's a bonus chapter focused on individuals and a distillation of the framework for your everyday life. Yes. So whether you're using an agent to book travel or using it to manage your life, you need to be thinking about four things.

[00:19:24] What permission are you giving it? Like, similar to what I talked about with authority. What do you want it to do? What don't you want it to do? What are the boundaries? Um, how is, how are you gonna be accountable for the actions it takes, and are you comfortable with that? So if that travel agent, um, buys 15 tickets rather than one 'cause you wanted it to check for the next 15 days for the best price, and it took that as buy a ticket every day for the next 15 days rather than just buy the best price over the course of 15 days, um, are you f- comfortable being accountable to that? And if not, then how can you make sure that it only buys one ticket or it only spends X amount of dollars? What is the broader context in which it operates? Make sure you're giving... The prompt has enough information, that you've given it enough context, that you understand the context enough to, to see what could be harmful. And then what is the evolution of the technology? And one of the things people say is, "Oh my gosh, it's moving so fast. How do I keep up and then be able to adjust my behaviors?" I mean, use the tool to help you understand the tool. Set up a prompt that tells you once a week, once a month, what are the new capabilities that I should be worried about or thinking about, or adjust my behavior relative to these tasks that I have running And then use that to guide how you're thinking about it in addition to you consuming the news and understanding the capability.

[00:20:46] So yes, I thought very much about the individual. That's why even in the organizational context, it's written so that each person could use it to navigate. And then this bonus chapter with the PACE framework, the personal ACE framework, um, helps

people think about this in context. But you're right, we have to change our mindset. We have to ... stop this.

[00:21:05] **CO-HOST NICOL TURNER LEE:** We do. We do. And look, I need you to come to my house and teach me how to do some of this optimization stuff for the bot. I haven't figured it out, Camille. I mean, I'm, I, I do think though once I do, your framework is gonna be helpful because I think we do not want to see people, you know, have the agents go out of control and then the next thing you know, you know, we're sitting here trying to figure out what to do next. I mean, which, you know, it also brings me... I just read an article with Bill Gates, essentially said, you know, a kill switch is not enough. Yeah. Because I think of all these things that you're talking about. And I was thinking for a while, kill switch should be good. Somebody should be able to have, you know, the phone back in the days where you are able to push the button. You know, somebody should push the button to say, "The agents are out, go- out of control," right? But what do you think about that? Like, actually having a kill switch is not the answer. It really goes back to what you've been talking about, right?

[00:21:56] **GUEST CAMILLE STEWART GLOSTER:** Yes. We need so much more than just a kill switch. I mean, uh- Realistically, you might be able to have a kill switch within your organization, but what would a kill switch look like across society? I mean, it would- Yeah ... essentially mean shutting down the internet, right? Yeah. You'd have to get it, no access to compute, no access to the internet and other data. Like, that's a really tough proposition. So yes, you should have kill switches within your organization and be able to shut down an agent. You as an individual should know how to shut down your agents. But we need far more than that. You need to be consistently training it to understand your objectives and your thinking, that kind of compounding learning of what you want it to do and where you want its boundaries to be. You need to be continuing to learn how to prompt better and how to, you know, provide the necessary context and understand how it's evolving, and really keep pressure testing what are the boundaries of the permission I gave it.

[00:22:50] **CO-HOST NICOL TURNER LEE:** Love it. You know, I have to ask this question as we wrap up. What was your experience writing your book? I always ask that question. Uh, you know I'm an author, too. Yeah. I felt quite lonely at some points, and so now I ask this question when we have authors that come on.

[00:23:05] **GUEST CAMILLE STEWART GLOSTER:** Yeah.

[00:23:05] **CO-HOST NICOL TURNER LEE:** What was it like to write your book?

[00:23:08] **GUEST CAMILLE STEWART GLOSTER:** Lonely and urgent. I mean, it, I was balancing everyday work, parenthood, family, all these things, with writing this book, and a recognition that the technology is moving so fast that this book could be obsolete by the time it came out.

[00:23:24] **CO-HOST NICOL TURNER LEE:** Right, right.

[00:23:25] **GUEST CAMILLE STEWART GLOSTER:** And the publishing process, as you know, is so long. It can be two years, depending on the publisher. So, um, I felt like I w- it was a race against time. I was dedicating a lot of time, and I'm just glad it hit at a moment where people were ready to receive it.

[00:23:41] **CO-HOST NICOL TURNER LEE:** The timing is just impeccable with regards to the knowledge that you're bringing to both enterprises, individuals, as well as, uh, society writ large. I mean, we have a lot of work to do in this area ... Camille. And so I really appreciated the book and I appreciate you for being that voice that is able to bring these communities together, because it's not gonna get, um, what do they say? It's gonna get worse before it gets better. Yeah. So we do need to be paying attention to this. Um, thank you so much for joining me today and discussing your recent book. Um, I ex- anticipate so much more from you. But where can people get "The Insider You Build"?

[00:24:15] **GUEST CAMILLE STEWART GLOSTER:** It's available everywhere, but if you go to insideryoubuilt.com, you can find out everything about the book.

[00:24:21] **CO-HOST NICOL TURNER LEE:** Perfect. And I know you've been out there. Uh, they can find out where you're gonna be doing book talks and all that other stuff. Yes. Uh, where do they find that out?

[00:24:28] **GUEST CAMILLE STEWART GLOSTER:** All the events are on that same site. I will be at Vital Voices on October 8th. I'll be all over the country, so please come engage in this conversation with me. And thank you for making space for this conversation, Nicol.

[00:24:39] **CO-HOST NICOL TURNER LEE:** Oh, always. Always. So thank you again. And for those of you that are listening, please continue to follow our work on the Brookings TechTank blog, uh, which is accessible at brookings.edu. Um, you can also look at the Center for Technology Innovation, the work that we're doing at the AI Equity Lab. This concludes another insightful episode of the TechTank podcast, where we make bits into palatable bites. Until next time.

[00:25:09] Thank you for listening to TechTank, a series of roundtable discussions and interviews with technology experts and policymakers. For more conversations like this, subscribe to the podcast and sign up to receive the TechTank newsletter for more research and analysis from the Center for Technology Innovation at Brookings.