

Progress Over Perfection

Episode 6: Overcoming the AI Replacement Narrative

MOVING FROM WORKFORCE FEAR TO OPERATIONAL MOMENTUM

Host: Jack Moore, Managing Partner at Evans

Guest: Jonathan Aberman, CEO and co-founder of Hupside

Series: Progress Over Perfection

Introduction

Jack: Hi, Jonathan. Welcome to the show. You've got a lot going on between Hupside and your book, I'm sure any number of things running a company. Give me a sense of what you've got going on right now. What's your focus?

Jonathan: My big focus right now is manifesting everything I'm doing, which I think comes down to this: here in the DC region, there's an enormous opportunity around AI, centered on humanity and the human value add. It's always been the case that here in DC, we've made our money through the knowledge in our heads. We're the biggest concentration of high-value knowledge workers in the world, probably between government, government contractors, and the supporting industries around them. And AI is an industrial trend driven by a model of technology replacing human knowledge and human labor, which runs completely counter to the lifeblood of this community. I'm part of this community. I've always been a homer — everybody knows that about me. But I also believe fundamentally that AI is a great efficiency engine, not a great differentiation engine, and we'll talk about that. So everything I'm doing — the teaching, the book, Hupside — may look broad to somebody who doesn't know me, but for me it's all the same mission: advocating for humans and human knowledge in a big way.

Jack: Yeah, I think that's one of the reasons that when I met you, I connected immediately. It's the lifeblood of our company too — it's all about human-centered adoption, whatever it is: organizational change, technology, that same kind of thing. I know one of the things you hit on — here's my copy of the book.

Jonathan: Yeah, I appreciate that. Thank you for quoting me in the book, and for saying nice things about it.

Jack: You're welcome.

The Replacement Story: Naming the Narrative

Jack: So we'll get into some of the details here. I want to start with the replacement story. There's an amazing story you've told me, and you tell it in the book, about leadership launching this new thing they want everybody to use. I'll talk a little later about some mistakes we made at Evans and how we handled that. But go back to where this starts — some of the messaging that's been out there, and how it makes people feel like it's threatening them, their lifeblood.

Jonathan: What I call the replacement story in the book is shorthand for what I'd say is the dominant narrative around AI in our society right now — approaching it purely as a technology, one that can only generate the margin improvement needed to justify company valuations through human replacement. There's no other place in the value chain to capture that much profit than off the back of labor. It's just math. And if you add in the fact that there are people in the Silicon Valley ecosystem who genuinely believe AI is the next life form, it's not a big jump for people outside the AI industry to think: wait a minute, you're commercializing something whose primary purpose is replacing me. That narrative colors everything. It colors how our employees experience AI. It colors how government workers have experienced DOGE. It colors how people react to data centers. It's coloring the current reaction to AI-versus-human and AI detectors. It's all the same issue: the way the industry has gone about commercializing itself couldn't be more unproductive for actually getting people to accept it. My argument is you have to come up with a way to describe the current narrative, and the simplest way to do that is to call it the replacement story. It puts it in a box. It's neither good nor bad — it's just a way to look at the world.

Jonathan: I used to do a podcast like you're doing now, and I had an author on named Richard Powers, a great writer. We talked for half an hour about the importance of telling a story — how you can change the world by telling one. What we realized is that for a lot of us as leaders — regional leaders, national leaders, business leaders — we lead by telling stories. And if we don't create a story that's compelling for our people, there's a story out there already, and we'll lose out. That's what I've been seeing for the last couple of years around AI, and it's why CEOs are saying, wait a minute, I'm spending all this money on technology and not seeing ROI. Or, I'm seeing all this money spent and getting workslop. Or, I shouldn't use this at all. On and on. That's because we're not putting people in a place where they can feel AI is actually good for them. So they resist, passively or actively. And I don't blame them.

Jack: Not at all.

Predictable or Different: Why AI Can't Be Both

Jack: It's really interesting if you juxtapose that with the lack of a replacement story around something like the smartphone. You go into this in the book, but it's such a feel-good story. I remember the excitement of getting my first iPhone — my camera, my Palm Pilot, my MP3 player, all in one device. Eventually corporate email too. It brought real productivity, but there was none of this replacement feeling around it. It's a very different tone than the AI-centric talking heads out there talking about reducing headcount.

Jonathan: Well, listen — don't forget, we all gave up the keyboard. My wife took forever to give up her BlackBerry. But you're right, that does demonstrate that when a product's overall benefits outweigh its burdens, it explodes in popularity. Everybody loved the BlackBerry, but the UI lost. The problem is that in all this excitement, we've taken a feature and tried to explain it away, and that's where we need a narrative: the fundamental similarity of AI output. The great thing about AI, what makes it efficient, is its predictability. As it gets more competent, it can do more tasks predictably — that's great. But you can't have predictability and differentiation at the same time. You can't do both in model construction. I read yesterday that a new scholarly paper shows that when you ask AI models to be more creative and original, over the last three years they've actually gotten worse at it, not better, because of the underlying architecture. So the part of the narrative that needs to be spoken more clearly — for leaders like you, companies like Evans, and people in the federal government — is that we need to first understand where AI is good and where it's not, and have an honest conversation about that. Then understand that there's a layer of human originality, because only humans can diverge from the underlying AI baseline that exists in every function. It's a question of degree.

Jack: Right, and that's what you're seeing in implementation work already?

Jonathan: Absolutely.

Evans' Own Story: The Fear Nobody Planned For

Jack: I'll share the story about Evans. Especially after you brought up DOGE — like a lot of government contractors, our work shrank last year. We realized we needed to pivot and increase our technical fluency as a company to stay relevant, to better understand our clients' missions and the technology they're applying, because they're being asked to do the same. So we said we wanted to increase our corporate digital fluency, and one piece of that was leaning into AI. Several of us on the leadership team had started using it. We created an AI core team, started sharing some of the benefits, gave people access to the tools. But what we didn't think about was their concerns about safety — everything people were reading in the market about replacement, substitution, that kind of thing, layered on top of some worry because we'd just been through the DOGE situation. All of a sudden it was, are you trying to replace us with AI? Is the plan for Evans to become a



company that helps clients cut their own workforce? And it's like — whoa, we've never been a company like that. It felt like it came out of left field for us. From the very beginning, we'd talked about ethical use, appropriate use within our clients' constraints, that kind of thing. And that was a reaction — I know exactly what you were just talking about was one of the key elements of that initial reaction. But then we said, okay, hold on, that's not what we're doing, and we started talking about it as an enabler: your unique things are still unique. It can't replace your creativity. It can't replace your experience. It can learn about you and help you along the way. Having gone through that ourselves, it was interesting — around the time I met you and you'd shared the replacement story, that helped us understand what felt like a complete mismatch between what we were saying and what was being heard.

Jonathan: Well, it's interesting when you become friends with somebody — I've watched for the last nine months or so as you've worked through this at your company, and it's been interesting to see how you changed the narrative. You're a lot further along than you were before. And that's the point: when you're in a situation where there's a dominant story everybody hears, voluntarily or involuntarily, and the replacement narrative is so much a part of how we think about AI — if you don't have a strong understanding as a leader, and tools to actually walk the walk, you're not going to get what you want out of AI, unless, of course, you just want to get rid of headcount. But here's the funny thing: what we're finding is that the companies that leaned hardest into reducing headcount through efficiency are actually starting to hire back hybrid workers. We saw that at GM, for example — they laid off their entire quality assurance group, figuring AI would just do it. What they found was that AI was good at process but not good at catching the outlier issues. So they brought their people back, now coexisting with AI. You see it in the PricewaterhouseCoopers jobs study too: jobs where people are hybrid — using AI to be more creative, more original, highly employable — their wages are rising far faster than everyone else's, and there are more openings. The market's already adjusting. And when a market is adjusting, the opportunity for good leaders is enormous, because you can capture value before everybody else catches up. That's where great businesses get made — being at the edge. For businesses like yours, the government contracting community, the consulting industry here in DC and around the country, this is the moment: if they lean into “we have people who generate original content and output when it needs to be original, and we're wildly efficient with these tools when it doesn't,” the value-add opportunity for differentiated, value-based pricing is enormous. I really believe that's where we are as a community and, frankly, as an industry.

Jack: No, I completely agree.

Piloting Hupchecker: What Evans Learned About Itself

Jack: I'll stump for one of your tools, Hupchecker. We ran a pilot in that AI core team I mentioned, and I'll freely admit I got the lowest score in the first phase.

Jonathan: Yeah, that was concerning to me. I almost stopped talking to you.

Jack: No, you didn't. So, for others out there — it was like a challenge: you get a prompt, come up with your own response, be creative, then use AI to be creative. I was probably one of the people, maybe because I'd scored

so low the first time, but in the second phase my score actually went up — seven out of eight people saw the same. And eventually, as we expanded the pilot to 80% of the company, most people's scores went down as they used AI, and it was exactly your point about originality and creativity. I explicitly asked it to help me become more creative, and — I'll give AI credit — it basically said, I can't help you with that. It told me my own experience, the small stories I weave into an answer, are what make it more original and creative. As we started talking about that as a company — unlocking what you, as an employee, bring to the table, and figuring out how to augment that to make you more productive while relying on you as a knowledge worker who AI can't replace a lot of — that really unlocked a lot of our workforce. Interestingly, the part of our workforce that didn't have that concern and leaned in fastest was our internal operations team — finance, HR, IT — they were looking for ways to improve operations, working with the consulting side of the business to design faster processes and build prototypes more quickly. It's one of those nice turnarounds. It was a challenge for a while, and now I feel like we're in a pretty good place with it.

Tool, Not Oracle: The Cognitive Decline Research

Jonathan: As we start to think about AI change management, there are a few baseline things you have to appreciate. The first is that AI is an incredible efficiency tool. There are places in your business, and people in your business, who are culturally comfortable using a tool to enhance a process without expecting a feedback loop — they expect the tool to do X, Y, and Z, and AI is really good at that. The problem is AI isn't designed just to be good at that. Chatbots are designed to be highly interactive, to create a dependency loop. Some of that is because chatbots were built off the analogy of social media, so people would already be comfortable working with them. It's not surprising most people use chatbots as search or as companions, because that's how people are used to using a browser. You can understand the design decision to have these models interact with humans that way. But unlike computers, or the iPhone, or anything else we've had, it creates a dependency loop you can fall into, where AI becomes an oracle instead of a tool. We see that in what researchers call cognitive decline: people who use AI as an oracle and accept the answers they're given show measurable atrophy in the part of the brain responsible for creativity and seeking novelty. That's demonstrable. For people who don't use process-heavy tools all the time, giving them a tool that behaves like a process tool but looks like a friend is dangerous — they start treating it like an oracle. For the people you're depending on for differentiation, you have to teach them to use AI the way they'd use a computer, a book, any tool for expanding knowledge. That means teaching them to think first: write down what they're thinking, and only then ask AI to poke holes in it, and push back against what it gives you. That's how you use AI. And what you did with Hupchecker — what you saw was you came in and said to people, what's your natural problem-solving state? People play a simple, fun game that reveals their problem-solving state, then use AI. The people who used AI naturally, as an accelerant, their scores went up. The people who used it like an oracle, their scores went down. That let you separate AI as a process improvement for the people who needed that from AI as an adjunct to your original thinking and value-add. You broke the conversation — you disabled the idea that AI was there to replace people, and then you were able to teach people to use AI properly to serve the client. That's ultimately what the originality story is: a narrative that isn't about replacement, it's about enhancement. And it's



something anyone can do using a simple think-then-apply process, which I call the originality loop — though a lot of people are already doing it cognitively without a name for it. Use AI as a tool, not as an oracle, and you're eighty percent of the way there.

What's Really Happening in Federal AI Adoption

Jack: I want to ask you — because you've got relationships with people in both the private and public sector, some pretty senior people in the federal government, civilian and defense — without naming names, what are you hearing from them about what's working with their teams on AI adoption, and just your general sense of the federal market for AI right now?

Jonathan: My sense of the federal market is that the first wave of AI as a technology — governance, data management, cybersecurity — has largely been where the focus has been. That's understandable, because we're dealing with a buyer that tends to buy technology against specs. What you're now starting to see is the beginning of “we have to train people to use the technology,” and a growing understanding that if you're going to train them, you need metrics to evaluate whether the training works. What I'm hearing is there's real demand for anyone who can say: the tools are great, but here's where you actually have to add value, and we can demonstrably teach your people to add value with these tools and prove it. There's a growing demand signal for that — in academia, in education, and here in DC. You can see it in how the conversation has moved: from “this is a cool tool,” to “what's human and what's AI,” to now, “I don't care if it's AI or human — is it valuable?” Whenever you see an industrial trend — cars, for instance — there's enormous opportunity if you get ahead of what comes next. If you have cars, you need roads. If you have roads, you need gas stations. If you have gas stations, you need hotels. If you have cars, you need tires. There's enormous money to be made right now turning all this technology spending into something useful, and that's what government is ultimately going to start consuming. The challenge for government is that government always manages, and should manage, against metrics. That's what the FAR is about, that's how government works, that's how we want it to work. So for the contractor and consulting community, it's going to matter a lot to do this work but also to come up with standards the government can point to and say, you've performed. Because it's becoming apparent that labor savings is a metric — but it may not be the only metric that matters for good citizen service, and it certainly isn't the only metric that matters for national security applications or anything where inference is required.

Jack: A hundred percent. The story you just told mimics one of our own projects — helping a group at the FAA enable their workforce with AI. Their messaging up front was good, but the issues our team ran into early on were mostly governance challenges. You've got a central IT group providing tools to a workforce spread across different organizations at the FAA — they all work to the same mission, but they have different roles, so what works for one group doesn't always work as well for another. It came down to governance and tool access.

Inside the FAA Project, and the ROI Everyone's Chasing

Jack: Our consultant on that project, Jesse Lambert, worked through those obstacles — some of it still pretty early-stage at the FAA. He had to figure out what tools would actually work for that group, get them configured and access sorted, and then go team by team within the organization to understand their specific use cases and help them apply AI in a way that actually increased productivity. It took a little longer, and we don't have the full outcome of that project yet — it's still ongoing — but the metrics they're tracking are things like cycle time, cutting down their backlog on engineering updates. It wasn't some holy-grail “AI maturity” goal, it was metrics people could actually understand. It seems to be on a good track, and it lines up with what you're generally hearing about the success stories.

Jonathan: This new approach — the originality story — is at the edge of the curve right now, not the mainstream. But it's going to become the mainstream, because you're hearing more and more stories of disappointment about ROI. You've got the MIT study — 95% of AI installations don't produce ROI, that number people throw around. McKinsey has a different number. Let's suppose, for argument's sake, it's 50%, or 40%, or 30%. That's a business problem right there. There's nothing else we do in corporate America where we spend a dollar and only expect a return half the time — we'd all get fired. So the market's going to catch up to what you and I are talking about. And funnily enough, people talk about government as the follower, not the leader, of technology. In an interesting way, because of DOGE, the biggest AI transformation project in the United States is actually happening right here in DC. So it's a question of whether government wants to keep pushing the edge. If it does, it's going to land on this model we've been describing. The last thing I'd point out, as you were describing Jesse's experience — he's an example of why there's always a human value layer on top of the technology. People say this new industrial trend will create jobs like every other one did, but they can't describe what those jobs are. That's a failure of the replacement narrative. If you think about it from the standpoint of the originality narrative, there are a lot of jobs that could be imagined that don't exist yet — jobs about people doing the things only people can do in different parts of the value chain. That's why the originality narrative is the more optimistic story.

Time Capture: Who Benefits When AI Makes You Faster

Jack: It definitely is. Let's talk about another downside. Part of the ROI I still believe in is using AI to handle more of the manual, less-inviting parts of a role so you can focus on the real value-add work. But all of a sudden you're in a model where you're doing value-add work all the time — you're “on” constantly. It's cool because you're using your unique strengths and working at a higher level, but there's a bit of brain fry involved, because you don't get the breaks that come from doing some of the more routine, lower-value things. I'm curious about your thoughts on that, and how it connects to the conversations you're having with others.



Jonathan: There's a really interesting idea going around right now — one of my friends, Dave Edwards, who writes a lot about originality, calls it time capture. If you're able to do your job fifty percent faster, who gets the benefit of that time? The way we tend to think about it right now is: if you can do your job fifty percent faster, that just means you do fifty percent more work. It's a really important question — if you could capture that much value, maybe you don't ask everybody to work that extra fifty percent. Maybe you ask them to work thirty percent more and give them some white space. As a leader, you have to decide who benefits from the time and how you're going to use it. You use it so people can think more and come up with new ways to make money. Or you use it to cut labor cost. That's frankly no different from any conversation you'd ever have as a CEO with your board or your stockholders: am I optimizing for profit maximization through margin, or profit maximization through value-add? If it's value-add, I might forgo some efficiency today for more value later — I might invest in a new product.

Jack: Right. It's really interesting — I like the concept of white space, because as we're seeing higher adoption, our employees and managers are learning to help their teams figure out the right place to use AI within their work — maintaining your originality, your unique skills, and using AI to augment those, but also giving people more free time to think about new solutions. Though I do think there's something to be said for the solace of routine work, too.

The One Thing Every CEO Should Say First

Jonathan: It's funny — I know, as we were prepping for this, you were going to ask me what the one thing is I'd recommend a CEO do first. This is a brave thing, but I think the first move, if you're serious about AI, is to get in front of everybody and say: I'm not firing anybody. I'm not firing anybody for the next six to twelve months based on AI. I may not hire, but other than the normal “are you performing” conversations, we're not making job changes because of AI. We're going to figure out ways to make AI make this business more profitable, together. I'd start there — and I wouldn't say forever, I'd say we're going to figure this out together. That's the hardest thing to do if you're a CEO, particularly with a board focused on managing costs. But every other time we've asked people to change how they do business, we've reassured them they have a place in it. That would be the most dramatic way to change the narrative I can think of. Then you can talk about tactics — but just think about how much that changes the conversation.

Jack: That's effectively what you did — not standing up and dramatically announcing “nobody's getting fired,” just, guys, I'm not looking to fire you, let's snap out of it and figure out how to work together. And you were a small enough company to make people believe it.

Jonathan: But I think that's the one thing that matters most — changing the narrative.

The Originality Loop, and Writing the Book with Claude

Jack: The other thing I keep thinking about — a lesson from reviewing one of your early book drafts — is it the originality loop? You mentioned it earlier, but expand on it. You've said you used AI to write the book itself, so it's a case study in its own right — how you used that loop to create it.

Jonathan: The originality loop is simple — all I did in the book was put into words what a lot of people are already talking about. You think first, then you write down what you're thinking — not just a couple of prompts, but a short memo, these are the things I'm thinking about. Then you give it to AI and ask what it thinks. AI comes back and tells you. You say, that's interesting, what are the holes, what am I missing, and you push back on its answer. In other words, you treat AI like an intelligent assistant: challenge my arguments, improve my arguments, what am I missing, iterate back and forth, use it as a research tool. Then, when it's done, you look at what came back and ask: am I going to own this? Is this my work product? Ownership doesn't mean did AI write it or not — that's a fool's errand to try to untangle. What matters is whether you're willing to stand behind every sentence, with the same logical weight as if you'd asked an associate or a junior person for a draft. It's the same thing. That's the originality loop: think, write, challenge, diverge, own. It's really simple. But if you read the literature on cognition and how AI creates dependency, this little cycle actually breaks that dependency. It's demonstrable.

Jonathan: Now, the book was a real example of that cognitive loop, and of breaking it. I wrote my outline, gave it to the model, and said, write the narrative for me from my content and outlines, write like me. Because I've written a lot over the years in my columns, AI does a pretty good job of writing like me. So it gives me a draft of the book chapters — really robust. I read it, give it to a couple of friends, you included, and they say, oh, this sounds just like you. I'm feeling pretty excited. So then I show it to my wife, Veronica, who used to be an antitrust litigator and has filed briefs with the Supreme Court — she knows how to write. And she looks at it and says, do you realize this has every verbal tick you've ever had, amplified on steroids? I said, what do you mean? She said, you like saying “things quietly compound” — do you have any idea, when I do search and replace, “quietly” is in this document eighty times across a hundred and twenty pages? Do you have any idea you like to use “delve”? It was eighty times — holy smokes. And I didn't see it, because it writes like me, the ultimate. So how do you break it? She and I took the draft and turned it back into English. Once I liked where it was going, I used AI again — to cite-check, to argument-challenge, to improve. Then for things like formatting and prepping for Amazon. Where I ended and the AI began got pretty blurry. But every word in that document, by the time I was done, was mine or Veronica's. That's how the book got made. My point is that even somebody who uses this thing in a sophisticated way fell prey to the pleasing nature of the model — you want me to write like you? I will write like you like nobody's ever written like you. And I was so caught up in my own wonderfulness I didn't even see it. I just thought, wow, this reads great. Of course it reads great — it's me, writing badly, at scale.

Jack: One of the questions I've been asked several times — and the answer, from that last story, was that it took longer than the small task I was trying to do — is whether it actually improved things overall. Was there a real productivity lift, beyond what you already described, and did it ultimately make the end product better?

Jonathan: What's interesting is, at the end of the month Veronica and I spent writing the book, I asked her that exact question. Her answer was no — it wouldn't have taken any more time if we'd started from scratch. My answer is different: it might never have happened at all. If you'd told me, Jonathan, take two weeks off from being a CEO, lock yourself in your office ten hours a day, and write an original manuscript, I'd have said, I don't have time. But I had enough time to get the data down, to do the outlines. I didn't have time to do the slides, or all the superstructure work. The book wouldn't have happened without Claude. It broke down the barriers to starting. Once I had the manuscript, using AI to fact-check, format, and prep for Amazon — all the typesetting — was genuinely helpful. So did it increase my overall productivity? Arguably. Did it increase the likelihood I'd actually produce the output? A hundred percent. That's how I'd describe it.

Jack: I'm not doing death by PowerPoint here, but overall, if we just measure this in time saved, I think we miss the point.

Jonathan: A hundred percent. I agree.

The 90-Day Plan for Leading AI Change

Jack: So, if I'm a CEO — and I think this works for a federal program leader or organization leader too — and I've said we're deploying AI and no one's losing their job because of it, that builds the safety piece. We just talked about the originality loop and how to use AI as part of it, in a way that augments your original thinking, challenges you, and handles some of the routine work at the end. So if I say all that — what should I be hoping to get out of it by Friday, instead of Monday, if I actually do those things?

Jonathan: This is why, in the book, I try to lay it out so somebody unfamiliar with change management can understand how it all fits together. It's well proven that if a leader wants to change behavior in an organization, there are steps you have to take. First, you have to build trust — that you're delegating authority downward and getting it back. Dictated change never works. People need to feel valued, particularly in the United States, where we're raised to believe we all have a right to an opinion. If you go in and tell people “do this,” it's not going to work if you're trying to drive change. So you start with day one: listen. We're going to do something together. And what we're doing together — if I'm in government — is improving our service to the citizen, because in government it's about the mission. So, number one: we're doing this together. Number two: I'm counting on you. Number three: the reason we're doing this is so we can win. That's day one. Day two is giving real thought internally to where in the organization your differentiation lives — what delights your citizens or creates profit. There are places in your operation where all that matters is efficiency, and places where differentiation matters. Where does it exist? Once you know that, figure out where and how you want to use AI in each of those different functions, because it's going to look very different — efficiency is a different conversation than differentiation. Once you've got a picture of that, which you should be able to do relatively quickly if you know your organization well, by the end of the week you should identify a team of people who are respected within the organization — people who'd want to work with you on figuring out how to bring AI in. That's your tiger team, your change team. Don't pick it purely on seniority. Seniority matters, but pick people



who are respected and enthusiastic — it's fine if they're the cool kids. If you can pull that together in the first week, you're already well on your way to classic change management: you've identified your team, you've identified the outcome, now you talk tactics. That's where the AI consultant comes in — to help you frame how to use tools to understand your people, and how to use the technology to understand where you stand. Then you develop the plan with your team and stress-test it, either by trying the ideas out on a small part of the business, or by using the technology, or by getting people to buy into the plan. That's what I call the ninety-day challenge. Any good change effort is about building a team of people who are infected with the outcome you want, who can execute it within a small experimental design, and then evangelize it for you, so you're not the only one explaining it as the boss. If you do that over ninety days, you've got a beachhead you can execute change against, and scale.

Jack: Similar to what you did — applying the same best practices. You got some people using AI, like in Oracle, and you needed to change that. You got a group of people to stop being afraid of AI and start using it properly, and now you're scaling it out across the organization. That's what change does — you always start that way.

Jonathan: What's funny about all this is there's nothing here a good leader doesn't already know how to do. It's just understanding that AI is a human resource issue as much as it's a technology issue. Once you accept that, there are tools that can help you do the rest. And then you can make money using it.

Jack: Awesome. I appreciate it. Thank you for bringing what you've learned talking to a widespread group of people across the DC area, and the core ideas from *The Originality Dividend*, your book, to the show. We'll make sure to put a link to your book in the show notes. What I'm taking away is a way to do some of this a little better than we did internally: that how you tell the story about AI matters as much as the tools you pick. If your story is about replacement, you'll get resistance. If your story is about amplification, like you were talking about, and you back it up with how you actually roll things out — the change management — you get momentum. It took us a while to get there, but eventually it helped.

Jack: To our listeners: protect your people's safety, unlock their curiosity, and build real capability. That's Progress Over Perfection. See you next time.

Progress Over Perfection is Evans' podcast for federal leaders making consequential decisions under uncertainty. New episodes explore how the best program managers, agency leaders, and modernization teams move from analysis to action.