

## **The AI Drop | #2: Planning in Quicksand: Strategy Under Radical Uncertainty with Geoffrey Schaefer Transcript**

Anna Gressel (00:01):

Hey everyone. This is Anna Gressel and you're listening to The AI Drop. Your briefing on how AI is reshaping the future of business law and industry. We bring you fresh drops each week so you never miss a beat. Welcome to the podcast. Hey everyone. Today we're tackling something that keeps every C-suite executive up at night. That is how do you build an AI strategy when the ground is shifting under your feet every quarter? And to help us unpack that, I'm thrilled to have on the podcast Geoff Schaefer, who I'm lucky to call a friend, but is also the vice president of AI strategy and governance at Leidos. One of the largest companies in the world operating at the intersection of defense, national security and commercial technology. And that means Geoff doesn't really have the luxury of taking it slow on AI, not when the stakes are so critical.

(00:48):

So Geoff, welcome to The AI Drop. I'm excited to have you on the podcast to talk about AI strategy, which is this topic that I am super focused on these days. And you're one of my favorite thought partners. I wanted to invite you in and we can have this conversation that we would've otherwise had over coffee in front of this bigger audience. But I also wanted to tell you what Claude called this episode. I was like, "What should we title this episode?" So we give titles to everything. And it said Planning and Quicksand, Strategy Under Radical Uncertainty, which I kind of really love. Geoff, welcome to the podcast.

Geoff Schaefer (01:18):

Thanks so much, Anna. That's not a bad title. I'll take it. Quick Sand is definitely, I think, the opposite of the Tara firmer we wish we were all on these days.

Anna Gressel (01:27):

For sure. And Geoff, I thought it would be helpful to just, before we kind of dive into things, have you tell us a little bit about your path and particularly how did you get this title VP of AI Strategy and Governance, which I don't even think it would've existed a few years ago.

Geoff Schaefer (01:41):

Yeah, it's been a pretty eclectic career so far. I got started in the intelligence community doing financial intel and counterintel. That was fantastic. Did that for a few different agencies, but the last one I started switching from being an Intel analyst to doing more strategy and innovation work almost by accident. I made a left-hand turn into finance after that, worked for a large hedge fund, a crypto startup and the New York Fed, and really got to extend some of that work into those domains and pick up some new areas and tech risks and operations. And that really just started compounding into a body of work that I wanted to focus on more and more. And so in 2018, I pivoted into the AI space directly. I came back to initially Booz Allen Hamilton and wound up starting their responsible AI practice and leading that. I was their chief AI ethics advisor as well.

(02:42):

## **The AI Drop | #2: Planning in Quicksand: Strategy Under Radical Uncertainty with Geoffrey Schaefer Transcript**

And then I started becoming more and more interested in the broader strategic questions. I found myself doing public appearances where I was more excited to talk about the possibilities of AI than I was about the risks and the harms. And so as I was looking at where do I take all this work? I happened to get connected to Leidos and they really were great in shaping this position for my interests and skills. And we're really looking to take all of this to the next level. So you're right, this is not a title that has been sort of common if in existence at all in the last few years. I think we're starting to see it pop up in more and more places, but I feel incredibly lucky to have this role here. And it's, as you might imagine, intensely dynamic at this stage.

Anna Gressel (03:32):

Yeah. I mean, I don't doubt that. And I feel like they're so lucky to have you. There are very few people who I think have the muscle to do both strategy and governance, and yet I think they're so paired. They're really important to understand together because in some ways, I think some question we forget to ask on governance sometimes is governance for what purpose? And to really have a view on that, you need to have a thesis on strategic alignment. So we'll talk a lot, I think about that in the episode, but I want to kick us off with something fun and I feel like panels often do lightning rounds at the end. I actually thought it would just be fun to kick us off with a lightning round. And for people listening, Geoff has no idea what I'm going to ask him.

Geoff Schaefer (04:09):

I'm in for trouble here. These

Anna Gressel (04:10):

Are not questions he's had previewed in advance for him. So if you're game, Geoff, you want to jump in?

Geoff Schaefer (04:16):

Let's jump in.

Anna Gressel (04:17):

All right, cool. So I'm going to give you a question and you give me a prediction. Two years from now, what is one thing that most large enterprises will be doing with AI that very few are doing today?

Geoff Schaefer (04:29):

That's a great one. The gut response in the spirit of the lightning round is a small majority of companies will double in size with agents.

Anna Gressel (04:39):

Geoff didn't actually mention this in his intro, but worth noting that he is working on an actual book on the kind of agentic enterprise and the workforce around AI, the workforce transformation that is happening, that could be and maybe should be a whole additional podcast. I agree with you. I think there are going to be companies of all sizes

## **The AI Drop | #2: Planning in Quicksand: Strategy Under Radical Uncertainty with Geoffrey Schaefer Transcript**

that are going to really be able to scale that way, but which companies and how they do that is going to be a big question mark. Okay, second question, what's the AI bet you're most confident in right now? And that could be like the investment or strategic direction where you'd say, "I'm willing to be held to this."

Geoff Schaefer (05:14):

" I think the audience will understand pretty quickly in this episode that it all comes back to agents. And I think that will be the locus of so much of our conversation here because my answer is agents are the real deal. This is the form of AI that we've been promised without actually understanding that this was the form we were waiting for. But for AI to actually be maximally valuable for us, for organizations writ large, we need it to be able to act autonomously and take actions on our behalf and achieve outcomes in the real world. We can't all be sitting in front of ChatGPT or any of the other models and actively querying and going back and forth and literally chatting. All day, we need to be able to task these AI systems to go achieve goals on our behalf. That is the ROI that we've been promised and I think agents are the form factor that's going to bring that to life.

Anna Gressel (06:10):

Yeah. And I think that gets to the point that you and I talk about all the time, which is how does governance then enable that? Because you've got to create friction points in the right places, but you can't disrupt that entire thesis. And so that balancing act is, I think, a particularly tricky one. But right, we're going to continue with the lightning round. What about the flip side of that question? So what's the thing that everyone seems certain about in AI right now that you think might be wrong?

Geoff Schaefer (06:33):

I am increasingly wondering about the infrastructure piece of this, not to pick on data centers, but it just seems incomprehensible to me that we will remain on a path where we have to paper over half the world in these massive bodies of infrastructure that no one really wants and that of course suck up a ton of resources. So there's a world in which we make those incredibly efficient and we need fewer of them of smaller size, but it just seems like we're going to have to stumble in or we will stumble into a different paradigm. I have no idea what that is, so I won't make that prediction. And I'm not also making the prediction that we're in a bubble. I actually don't think we are. I think we may see some trimming of some of the investment. And there is a big question on how we actually manifest the ROI, again, that we're promised across society.

(07:34):

And I think that is agents and adoption and the very classic sort of boring, hard, sweaty change management work that never really goes away no matter what the technology is. So I think we're going to see some, again, some trimming of the investment or some refocusing of that investment, but I don't think we're in a bubble, but I also don't think this is the final sort of version of the infrastructure that is going to allow this to get built out in a way that I think we're all excited to see get built out in various ways, and certainly the ROI will demand.

Anna Gressel (08:09):

## The AI Drop | #2: Planning in Quicksand: Strategy Under Radical Uncertainty with Geoffrey Schaefer Transcript

Yeah. We just did an episode that folks can go back and listen to on like, are we in an AI bubble right now? And I think the point that you made is exactly the right one, which is what do we actually think the long-term horizon is of these technologies? We may not know how we get there, but we're kind of undoubtedly in a transformational moment and that could unlock value in ways that we can't even think through at the moment. I think that's probably one of the biggest parts of the promise here. Okay, final question in our lightning round, and then I want to talk strategy, but let's transition into it this way. If a CEO came to you tomorrow and like from a different industry, they have no AI strategy in place whatsoever and their board is starting to ask them hard questions, and that CEO said to you, Geoff, "I've got 90 days to show progress on my AI strategy." What would you tell them to do first?

Geoff Schaefer (09:00):

That's a really interesting question because I struggle more and more the more I focus on AI strategy on what good AI strategy looks like. The analogy that I think I keep returning to is a good AI strategy is analogous to rebar and concrete. So you strengthen your corporate differentiation by using AI in the right ways, in the right places for the right reasons. And so you still need a strategy to do that, to understand what are the right systems and technologies that we want to use, how are we going to use them internally for customers, et cetera. But I think that is all in service of the broader enterprise strategy, especially as these systems paradoxically become more and more capable, but more and more commodified. So if you have a commodified technology, almost by definition, you can't build a strategic moat around it. And so what does that come back to?

(10:05):

That comes back to simply what are you as a business? What are you trying to do? How can you differentiate in those ways? And then finally, how can AI help you do that? It won't be the way or the thing, the end goal unto itself that you're trying to achieve, but it will enable you to achieve sort of your broader set of strategic goals and outcomes as a company. So I think thinking of AI strategy as the rebar in the concrete to build sort of the larger edifice of your company is the way to think about that. And so for the first 90 days, I would refocus, I would redouble my efforts on actually updating the corporate strategy and then backing in to what the AI strategy itself should look like to help reinforce all of that. I

Anna Gressel (10:55):

Think about that a lot personally, because I've been asking myself this question like, "How do you build this durable AI strategy?" The challenges that industries themselves are fundamentally changing and that you're building a strategy at the same time that companies are often thinking about whether to make truly strategic pivots where they may be headed in the next five, 10, 15 years and AI is a piece of that, but certainly not the full story at the same time that it is potentially changing for some industries the entire background landscape. So I think in some ways it does beg the question, is it the strategy first or is it like, can you afford to just move on the fly and hypothesize and change and adapt?

Geoff Schaefer (11:38):

## **The AI Drop | #2: Planning in Quicksand: Strategy Under Radical Uncertainty with Geoffrey Schaefer Transcript**

Yeah, I think that's a really interesting question. It makes me think that really companies are in this almost quantum state right now. So on the one hand, to your point, they are experiencing disruption to their specific company, to their industry, to the broader ecosystem that they exist in, in large part or an increasing part rather because of AI. At the same time, they need to be proactive, as you said as well. They need to be focusing on what does the future of their industry look like? What are markets, new markets, adjacent markets that they can and should be moving into? What does, again, their corporate strategy need to be in the next year, two, five years, et cetera. It's really kind of this, again, this quantum state, this opportunity to be both worried and to take that sort of disruptive moment and really start asking bigger, more existential questions about what you want to be when you grow up.

(12:39):

Ledos started as SAIC. We broke off in 2013, but our total history is 50-ish years. And so if we want to be around for another 50, and I promise you we do, we also have to ask ourselves some big questions about where this is all going. What does our customer markets look like in the future? How do we move into those spaces? And all at the same time that AI is disrupting those markets, those customers themselves, our own company, while at the same time, again, we also have to be using this technology to make those moves into those new markets or those adjacent markets, et cetera. So it's this weird sort of dual disruption enabling technology that our strategies have to contend with themselves.

Anna Gressel (13:29):

Yeah. I think that's an important point. And for companies that are kind of tech native or more recent, frankly, this is for sure existential to get the future of AI right, but it's a different set of questions entirely than companies that have existed or industries that have existed for 50 years, a hundred years, 300 years, they've gone through other, frankly, industrial transformation moments. I mean, Freshfields is like hundreds of years old, actually, it's amazing. And to think about the kinds of transformational moments that industries have had to weather in the past and go through, is AI just another one of those, or is there something that makes this different? I feel like I have my own views on that, but I'm curious for yours.

Geoff Schaefer (14:10):

Think this is different. I've talked about this before. The reason why I think it is, is because we often point to the industrial revolution as sort of the canonical example of we've done this before, we've lived this before and we were successful. I think that's a bad comparison for a couple reasons. The first is that was in a relatively concentrated set of industries. It took place over 75-ish years and during that period, we still lost a generation of workers. So yes, we came out on the other side of that with net new and net more jobs and everything by and large worked out, but the disruption was real. It was painful. In some ways, we still have some lingering effects of that. But in addition, AI is different for a few reasons as well. This is disrupting virtually every sector in every industry all at once.

(15:09):

## **The AI Drop | #2: Planning in Quicksand: Strategy Under Radical Uncertainty with Geoffrey Schaefer Transcript**

The speed at which this disruption is happening is almost unbelievable. ChatGPT that kicked all of this off was just a handful of years ago. That came out in 2022. We're four years later and we're already seeing this technology disrupt jobs across companies, across all sorts of industries. And so when we have a technology that is competition to human labor, and we have that human labor increasingly feeling the pressure to sort of reinvent their own jobs and careers at a minimum to figure out how to use this technology themselves to become more valuable and productive, et cetera. And this is happening across, again, every sector and almost every country in the world, all at once on this timeframe that we're on, this seems to me sort of self-evidently unprecedented. That is before we even get into the sort of bigger questions around what happens when we have artificial general or artificial super intelligence, where this AI in theory can do everything that humans can do at minimum at the same level of quality and efficiency, if not way more.

(16:21):

So I think when I was at Bridgewater, we had a principle, if you're worried, you don't have to worry, but if you're not worried, you have to worry. This may all work out just fine, but I think that's a bad assumption that it will inevitably just manifest in that way and we'll all be living happy, flourishing lives and everything will be just fine. I think we're going to have to work at it at a minimum. And so the more proactive we are in sort of understanding how these waves of disruption are impacting individual employees, companies, sectors, et cetera, the more prepared we'll be to actually manage that disruption well. And finally, to your point earlier, this requires, again, being incredibly proactive, waiting around to see how this technology's going to develop, how other companies are employing it, using it for different use cases, et cetera, where it's succeeding, where it's going wrong, that just doesn't work for a couple reasons.

(17:22):

The first is that means you're in, like you said, a reactive position, and so you're already at a structural disadvantage if you're waiting to see. But the second is, no one knows what's going to happen here. We're all going through this together. And so the more you engage with this technology, the more you experiment with it and try to use it for different business ends and outcomes, the more familiar you'll be with how it works, where it works well, where it doesn't, and where it might actually transform your company in highly beneficial ways.

Anna Gressel (17:56):

Yeah. I think I'm really coming around to the view that a big part of the strategy question is around having a hypothesis because if you don't have a view and you don't have a theory or a thesis about what AI is going to do internally to your company, to your industry or sector, then you really are losing time on testing that. And that's the problem is you kind of end up with this time drag where if you don't have an early theory, you just don't know how off you are. And I think that's what's happening for some companies that just say, "Well, we're going to play around with it." We had the pilot stage of AI that happened early after ChatGPT and everyone's like, "Oh, we'll do a pilot and we'll see." And everyone looks grassroots bottom up some pilots. That I think is fine if you're just getting going and many companies are just getting going.

## The AI Drop | #2: Planning in Quicksand: Strategy Under Radical Uncertainty with Geoffrey Schaefer Transcript

(18:46):

But I think if you have any view that your company might be significantly affected, the problem with that kind of way of testing or way of seeing things is you just don't know if you're right or wrong and you don't know how that's going to then affect those kind of key dimensions. And so really it shortens your timeframe to adapt. And right now, if we're thinking about volatility, we might really want to price in adaptability as one of the key dimensions or key attributes of companies is to just be able to change if change is required. And we're seeing this even from ChatGPT to agents, right? Like the thesis, the theory, the adoption hypothesis, even just what the technology can do every day. Some people may have said earlier, we think we're going to get to agents, for sure. Some people were saying that, but do you know if you're right or wrong?

(19:34):

How do you know? The thesis kind of gets you to the testing that gets you to the proving. But I'm curious about your view on what are the dimensions of this? And if you're charting this out, is there a useful framework? Is everyone just inventing a framework? Where do we start with this? I

Geoff Schaefer (19:52):

Think a couple thoughts in sort of anchoring to what you were just saying first, what I think is really interesting about the pilots and that approach is pilots really connote a certain mindset. You're dipping your toe in the water, you're not committing, but if you're the average organization and you're trying to figure out, how can I manage this technology to productive ends? How can I manage this disruption well? That requires a certain degree of commitment and risk tolerance and forward momentum that if you are structuring your AI work around pilots, I think almost definitionally, philosophically, you are not able to actually capitalize on. Pilot is inherently more risk averse. It's more limited. It's easy to say, "Well, that didn't work out. Let's try something else or let's just not really focus on this technology altogether. Let's go back to our core competencies." Whatever it is, if you are taking a pilot approach, you are inherently being more small-sea conservative in ways that I think just aren't going to meet the moment.

(21:00):

And so yeah, if you are moving forward with major investments, commitments, plans, transformations, et cetera, particularly at enterprise scale, you are taking on a bit more risk. But I think when you really look at the holistic picture and what committing firmly to using AI in transformative ways looks like across your organization, I think on net, you're actually taking a posture that is much less risky for your company because you are moving forward, you're committing to using this technology in ways that may require you to figure it out and adapt and be iterative along the way, but you're committing to the transformation that I think is necessary to sort of meet this moment and to ensure that you have another 50 years or 100 years as a company. But coming back to your broader question, I think the framework gets back to the rebar analogy. So really understanding sort of from a first principles way, what is your company?

## The AI Drop | #2: Planning in Quicksand: Strategy Under Radical Uncertainty with Geoffrey Schaefer Transcript

(22:05):

What are you trying to do? Who are you serving? How are you serving them best today? Where are you falling short? How do you want to evolve and adapt your business models? Are you thinking about new products and services, new markets, adjacent markets like we talked about earlier? Again, put simply, what are you trying to do and what sort of assets do you have at your disposal to do those things? And then I think you can literally start marrying up in a fairly one for one way, different AI capabilities, particularly agentic capabilities to those products, services, markets, customers, et cetera. And not all of those sort of matchups, if you will, are going to be promising, obvious, value added, et cetera, but many will. And I think at the end of the day, taking an approach that is practical, tangible, and useful will really serve you well.

(23:01):

This stuff is actually not as complicated as we make it out to be. At the end of the day, this technology is really enabling and allows you to do certain things better, faster, cheaper. It allows you to do new things that you weren't able to do before. It is an amazing way to turbocharge your capacity, but all of these things are in service to other goals, to broader goals, to broader business objectives. And so I think starting with that, starting with what you are as a company and who you're serving and how you're serving them, and then backing in the technology sounds incredibly simple, but actually starting there and marrying up the capabilities with the outcomes that you're trying to drive, I think is the best way to be as practical as possible because the hard stuff is really in the transformation itself.

Anna Gressel (23:53):

I think you're right to call out the transformation exercise as one that every company in their own way is struggling with because this is a technology that has to be adopted at the individual level. You can't skip that step right now or maybe we'll get there, but that's a critical piece of the equation. But I wanted to ask you, before we get to the actual transformation, we're still on strategy and maybe we can close out on this. What are some of the things that people get wrong or the blockers that folks who are thinking about starting on a journey of putting together their AI strategy should be aware of and thinking about how to unblock and advance?

Geoff Schaefer (24:29):

Yeah. I think one of the things that we see too often is a division between the parts of a company that are seen as being more technical and sophisticated with the technology, or literally in some cases responsible for the technology. So think about CIO, IO organizations, CTO organizations, et cetera. We see probably too often a division between those organizations and the rest of the business, the more traditional lines of business. And as you might imagine, so much of the technology, the AI technology work, development, acquisition, use, experimentation, et cetera, happens in the more technological focused organizations. And so already you're starting from a position of deficit, sort of. So you have to now figure out how to have the conversation with the non-technical areas of the business about the work that you're doing, let alone how that work can and should be used to advance their goals.

## The AI Drop | #2: Planning in Quicksand: Strategy Under Radical Uncertainty with Geoffrey Schaefer Transcript

(25:37):

And so I think step number one, the clearest thing you can do is at a minimum, have a fusion of the folks in those technical organizations with the non-technical aspects of the company. And the more that this is a sort of fluid, seamless partnership between the experts or the supposed experts on this technology, because in many ways we're all starting from the same position and priors, but the more you can have a good, healthy partnership between those parts of the organization and the business, I think the more successful you'll be in identifying very quickly what are the right use cases, what are the right ways to use this technology and what outcomes are we trying to drive and how can we drive those really successfully? Because what too often happens is so much of this work is either done, developed, procured, et cetera, and very classic silos.

(26:31):

And the business looks to the more technical organizations and almost deference and ask, "Okay, what sort of AI should we be using? How do we use it? Can you upskill us? Can you enable us? Et cetera." When in fact, that conversation should almost be reversed and at a minimum sort of co-developed, co-figured out and really co-designed because the technology, again, is an enabler, is the rebar, but if it's built separately to abuse the metaphor, if that rebar doesn't get actually put into the concrete, then the rebar doesn't do any good.

Anna Gressel (27:06):

No, totally. It's interesting. Someone asked us, a client asked us the other day, how did you guys get to such high adoption rates? We have super high AI adoption rates across our organization. And one of the things that one of our teammates said, and I think is quite right to your point too, is that we have a partner-led lab within the organization that also is data scientists and it's a joint development enterprise internally. And so we do a lot of our own custom design and tooling, but it's not just in a corner of the organization. It's high visibility, it's led by partners, it involves a lot of lawyers. So it's really to your point about fusing. I think we have this kind of fusion. It's like a cool thing to see internally, how that works in practice and the creativity that that can unlock. Geoff, this has been an amazing conversation and I know we're running low on time, but what I want to do is have you back on the podcast.

(27:58):

We're going to pick this back up, continue our discussion, and it is such a pleasure to just be able to bring our coffee chats into the more public sphere. And I hope everyone had a great time. I know I did today.

Geoff Schaefer (28:11):

I agree. I can't wait to talk about some more of the things that we've been discussing over the last year and a half for your audience. So I look forward to part two.

Anna Gressel (28:19):

Awesome. Can't wait. That's today's drop. Thanks for listening. If you'd like to hear more of the AI drop, hit subscribe and join us each week for fresh insights on what's

## **The AI Drop | #2: Planning in Quicksand: Strategy Under Radical Uncertainty with Geoffrey Schaefer Transcript**

shaping the future of AI. This podcast is for informational purposes only, and it does not constitute legal advice or establish an attorney-client relationship. It may be considered attorney advertising, so prior results do not guarantee a similar outcome. The views expressed here are those of the speakers and do not necessarily reflect the views of their firms, organizations, or clients. For more information, please visit [freshfields.com](https://www.freshfields.com).