the art of boring™

EP 237 | AI Adoption: How Investors Can Tell Whether a Company is Keeping Up



[00:00] [Andrew Johnson] [AJ] Hi, everyone. Back on the podcast today is Justin Anderson, Mawer's Chief Technology Officer. Last month, Justin was on the road with a speaking tour called The AI Investor.

And today, he walks us through it. We start under the hood with what actually happens behind the scenes when you're using AI tools. And from there, we talk about the economics, some things we've learned putting this to work at Mawer, and what it all means for investors themselves.

With that, here's my conversation with Justin.

[00:26] [Disclaimer] This podcast is for informational purposes only. Information relating to investment approaches or individual investments should not be construed as advice or endorsement. Any views expressed in this podcast are based upon the information available at the time and are subject to change.

[00:44] [AJ] Justin Anderson, welcome back to the podcast.

[00:46] [Justin Anderson] [JA] Andrew, great to be back.

[00:47] [AJ] It's great to have you. As I said in my introduction to this episode, you've been out on the road doing a speaking tour called The AI Investor. I joined you on a bit of that tour, and you actually begin your talk by going back to the spring of 2025, which is actually right around the time you were last on the podcast. So I wanted to start there. Why did you choose to start back in the spring of 2025?

[01:12] [JA] Well, your initial comment there about the cross-Canada tour, I don't know if you've been on a WestJet flight recently, I'm sure you have. They have that little safety briefing where they take you through the cross-country trip, there's the Toronto scene, then the Alberta scene, then the Vancouver scene. I'm on the plane thinking, I'm doing that same tour they're doing in the safety briefing.

So that was fun. It was a great way to see the country, I really enjoyed it. I got back to your homeland in the Halifax area, Nova Scotia, actually my first time in Halifax, and it was absolutely stunning. So I had a really good time there.

But to your question: a year and a half ago, people were talking a lot about tariffs, Liberation Day had just come out, and there was a lot of concern about what that meant for markets and so forth. In hindsight, we've always got that benefit, but it turned out tariffs weren't really the thing that moved the needle in the market.

It was a number of other things, which we also discussed on the podcast, and one that really sticks with me is that build-versus-buy tension we talked about. I think we referred to it as build-in versus build-out. The concept



there was essentially what's happening with LLMs and how that's changing that dynamic, where the build option becomes a lot more attractive to companies internally as the cost to build comes down dramatically. They can still build that custom software, but they can now do it in a competitive, and in some cases even cost-efficient, way relative to some of the legacy vendors they've been using. So that theme played out over the next year and a half.

We heard about the SaaS apocalypse; I think that thesis hit home for a lot of people. It led to this thought, when we were doing the cross-Canada tour, that this thesis-driven approach to investing is going to become increasingly important: looking for these kinds of inflection points, buy versus build, and that dynamic changing, especially as the LLMs and the AIs get really smart on historical data, where they're not as smart about how the landscape is changing looking forward. So that was one of the main points of the presentation, emphasizing where humans have an edge relative to AI in that ability to look for these inflection points and, from first-principles thinking, identify an investment opportunity and try to invest around it.

So that was one thing we emphasized. The other piece, I relistened to the podcast recently, we talked about how at Mawer we were going to act as if AI is a revolution for our business. The reason we gave in that presentation, and on the podcast, was that there's an asymmetric payoff here: if we put some resources into this and it turns out AI is a revolution, we're going to be in a really good position. So we did that over the last year and a half, and we've put a lot of emphasis on setting the foundations for how we can scale the use of AI at the firm more economically.

And we're starting to see the payoff of that. I was really excited listening back to that podcast, because those are a couple of themes that tie directly into these presentations, and into some of the investments we've been making.

[04:01] [AJ] And I want to get into some of that detail. But most people listening to this have used ChatGPT or something like it, very few, and I would certainly put myself in this camp until I heard you talk about it, have thought about what's actually happening when you type something in and hit enter. So can you walk us through what's happening behind the screen when we do that, and ultimately why that matters? I think that'll be really useful context for the rest of the conversation.

[04:27] [JA] Yeah, it's a great question, Andrew, and I'll take a couple of minutes to answer it, because I think it's really important to understand a little bit of the detail on how these models function, in order to start thinking about how this is going to roll out across the industry over time, and where we're focusing. So it'll set the foundation for us. Listeners, if you can, come along with us for the journey.

So when you ask a question, say we just talked about buy versus build, so the listener might type into ChatGPT, “summarize what buy versus build means.” What happens when they hit enter on that request? The first thing that happens is the LLM is going to break up that sentence into tokens. Essentially, think of every word as roughly one to two tokens each. The job of the model is then to predict what the next token should be.

In this case, “summarize what buy versus build means,” it's going to predict the next token, which is essentially the next word, and once it predicts that word, it takes all of those back and repeats the process. It keeps doing that and eventually has a sentence, and eventually a paragraph or a whole page. So that's essentially how they're working: token-prediction machines. When you ask that question, in this case it's a relatively small number of



tokens, maybe 10 tokens on ”summarize what buy versus build means.”

But if you're asking more complicated questions, like, “give me an update on the latest news in AI,” now you're asking a question the model hasn't been trained on. Maybe the model was trained three months ago. So to answer that more complicated question, you have to do something like a Google search, retrieve all the results, and pass that into the context window. Now you're going from 10 tokens to maybe 100 or 1,000. And eventually, with much more complicated operations, maybe tapping databases, you get into the thousands, or tens of thousands, of tokens.

The thing with these models is the cost, and essentially the compute required to process a request, is directly related to how many tokens are involved. If you're using 10,000 tokens versus 10 tokens, you're looking at orders of magnitude more expensive tasks. So the game in this business is really to understand token use and figure out ways to rein that in, bring it down, so you're minimizing your token use to get the maximum output.

So that's why we set the stage here: understanding the context window is really an important part of how AI is going to roll out through the industry, and the game companies are going to try to crack is how you drive down the token spend.

[07:01] [AJ] Good foundation there. If I'm an executive at a business, Justin, and I'm trying to think through the economics as you've just laid them out, can you help me understand that side of things? Where should I be looking, and is there a framework I can apply?

[07:15] [JA] Yes. So today you hear in the news all the time this comment around what's the ROI on these AI investments, and there's a thesis out there that says there's limited to no ROI, and that's going to cause a CapEx bubble. There may be some potential truth to that, or something to that story. But the bottom line is ROI is really in focus for the industry today because it's going to define the risk of a bubble. It's going to define the longevity of these technologies as they roll through the enterprise. You have to get value out of them.

So then we think about, okay, what is the ROI calculation? If you break that down into its component parts, ROI is a ratio. On the numerator, it's the value you're getting from using the technology. Any time you ask it to write a report, explain something, or do some kind of job, you're getting some kind of value for your business. Some of those values become measurable things, where you can see you're doing more with less, a productivity improvement. Maybe you're expanding your marketing reach or engagement, so you can measure some of these things.

In many cases, they're very hard to measure. What's the value of being more prepared in a management meeting when you're doing investment research? You can't really quantify that directly. It's going to express itself through better returns over time, hopefully, but it often becomes a qualitative piece. That's the numerator: what's the value proposition of using the tokens.

Then there's the denominator, the cost of using the tokens. The way we break that down is: how many tasks do you ask the AI to do, multiplied by the number of tokens per task you end up having to use to serve it. That goes back to our comment about the context window being so important, because if you're using lots of tokens, that's directly, linearly related to the cost you're paying.



And then finally, there's the dollars per token: what are you spending per token consumption? That's really a function of the model, for the most part, of whether you're using the right model for the job. We talk about open-source models as another lever you can pull to drive down that dollar-per-token spend. But that's really the equation we're talking about, and something we spent a lot of time emphasizing during the tour.

[09:26] [AJ] So don't maximize tokens, maximize the effective tasks associated with the use of those tokens.

[09:33] [JA] Exactly.

[09:34] [AJ] One of the terms you used throughout the talk, whether in the prepared remarks or the Q&A—and there was lively discussion in a lot of cases—you kept coming back to this term "orchestrator culture." What do you mean by that, and why do you put it on par with, or even ahead of, the technology?

[09:48] [JA] So coming back to our ROI calculation, value versus cost: the orchestrator is, well, that's you and me. You've been building some applications yourself, and some folks on your team. The idea of the orchestrator is that we've moved, with this technology, into a world where everybody at the company can now actually be an orchestrator. They can work directly with the technology to build little applications or have chat experiences very specific to their domain.

One of the challenges is that, as organizations, we're not really designed for this world. We've been designed to be functional specialists: if you want to build an application, you go to a business analyst, who helps you write out the requirement, we put it through a governance process to see if we should be spending firm resources on it, then it goes to a developer who tries to implement it, and it goes up and down the chain. That becomes a pretty expensive process to manage.

The idea of the orchestrator is that you can collapse a lot of those steps into a single person who can have the vision for the product they want to build, and then build it, deploy it, share it, and really leverage their knowledge base. But getting there isn't easy, because organizations are organized around functional specialists. So how do you get there? It's a number of different things: incentivizing people, and driving the culture in a way that encourages people to build an orchestration culture. Not a trivial task, something worth spending focus on.

[11:15] [AJ] I'll put you on the spot there, Justin. You're the Chief Technology Officer here. How are we applying this, whether it's the technology or the orchestrator culture, at Mawer? And what are some of the things you've learned along the way?

[11:26] [JA] You may have heard of Jevons paradox. The idea of Jevons paradox is that as the cost of some technology comes down, you might think the technology company makes less money from it over time. In many cases, that's not the case; it actually makes the market much bigger and drives up the ultimate value they're able to capture.

The philosophy I see with these technologies is very similar: there's an opportunity for us in technology to drive down the cost of using it. If we can get our dollar-per-token and tokens-per-task down to cents on the dollar, that opens up all sorts of different use cases our business users can tackle, because the business users are really



more in charge of the numerator part of that ROI calculation, finding the opportunities, since they have the domain knowledge. They've been serving clients, doing investment research, marketing, working in finance, and they have that specialist knowledge of their field. They're best suited to figure out how to use AI to help drive their field forward.

Our job in technology is the denominator: making sure the cost for them is as low as possible without compromising quality, so they can maximize the number of use cases per Jevons paradox. So that's something we've been very laser-focused on. In terms of how you do both of those, we could do a whole podcast on the value side. It's things like holding hackathons and having a weekly AI moment where people showcase what they're working on. It's really important that leadership drives it, that the leadership of each department figures out how to drive the value use cases for their respective areas.

On the denominator side, there are a number of things, but some of the ways you can tackle it are open-source models, which give you a lever to drive down your dollars per token. Something we've done at Mawer is design and build our own harness. We call it Nova. The job of this harness is to serve that context window, those tokens, as efficiently as possible, so that when our users are doing tasks, they're using exactly the tokens that are necessary and no more, in order to capture that Jevons paradox cycle. So those are some of the ways we're trying to tackle this ROI calculation at Mawer.

[13:42] [AJ] The harness, this would be akin to, it's almost like an air traffic controller, or a stoplight or directional light on a train track. It takes what you're asking and routes it to the most efficient use on the other end.

[13:55] [JA] That's it. So our harness now has access to essentially all the data at our company, and it's all governed under who should have access to what. So when someone's asking a question, it's determining whether this user has access, and which data is needed to serve the request, doing that at lightning speed, and only pulling the data that's necessary. So it's handling all those data operations in the background and trying to do it as efficiently as possible.

The other thing you get into, as you're trying to drive down your token spend that might be interesting for listeners: you eventually get to a place where you might want to start fine-tuning and doing post-training on your own models. That sounds complicated, but the reason you want to go down that path, from a cost perspective, ties back to that token use we opened with. As you have more token consumption, you realize: what if I moved some of the knowledge from the databases I have to pull, and instead baked it into my model weights? Then I wouldn't have to do those database calls, or pass it into the context window. I could save an order of magnitude of tokens. That's something I think more advanced companies, the AI natives, are going to get into more and more, as they figure out ways to get a cost advantage relative to their peers, even fine-tuning models to drive down their token use.

[15:16] [AJ] Interesting, and we're going to come back to that AI native term shortly, because I want to ask you more about that. But just before we move on: if you're not an AI native firm currently, what's more important at this stage, if you aspire to be? Is it to just get in there and use the technology and not really worry about cost? Is it to apply a cost-conscious framework at the same time? Are you trying to balance those?

[15:40] [JA] The number one thing I heard when we went around and did this tour was, constantly, the feedback



was: we want to get more into this stuff, but our governance or compliance creates a lot of hesitation to bring the technology in. So if I was putting a word on it, I'd say governance is a major friction point for people getting into this technology. And for good reason: this technology has a lot of power, people recognize that, and they want to appropriately govern it, because you're dealing with your business and you don't want to put it at risk. So there's a sense of, let's take this gradually and do it right.

At Mawer, we've spent about a year really refining our own governance process. It's quite a boring topic, but we like to be "boring, make money," so it's right up our alley. It's really about how you set up your data foundations, your data governance levels, your app governance, and then eventually your node governance, which is really about controlling how people access your data, and then eventually AI governance itself. So we've built an entire governance framework end-to-end that takes into account the unique risk profile AI brings.

[16:46] [AJ] Justin, in your talk you sorted companies into three groups. Can you walk us through those?

[16:52] [JA] You kind of got at this in your previous question, asking about the evolution of bringing someone in who's very new to AI and evolving them. A little demeaning, but it's not meant to be, it's just to define these groups. You start with the AI dinosaurs, someone who's maybe resisting it, or has never used it. We all start as dinosaurs.

You start using the technology, get excited about it, and evolve into what I call, again somewhat derogatory, the AI rat racer. If you go back to that ROI calculation, the rat racer is very much focused on the top-line value. They're creative, coming up with all sorts of great ways to use the technology and different use cases. Maybe they're even engaging in a little bit of AI slop. You've got to be thinking and working as partners with it, not just handing it all over, and that's part of learning that orchestrator culture, to call back your previous question.

But that's the next phase: the rat racer spinning on the value they're trying to create, and not spending a lot of time thinking about Jevons paradox or the cost side of things. Eventually the bill comes in, and you look at it and say, well, I was doing a lot of cool experimentation, but I'm not sure it was worth this much money. That's where a lot of rat racers are today, starting to question that, starting to say, experimentation time is over, how are we actually adding value here?

That leads you to the next phase, which is also thinking about that numerator value creation. That's the AI native, but also considering the cost side, being obsessive about cost, realizing that driving costs down to cents on the dollar means you open up Jevons paradox, you open up all sorts of new use cases, and you start to actually drive a very positive ROI.

I think companies are going to almost self-sort into those buckets. The vast majority today are probably in the dinosaur or rat-racer bucket, so if you're one of those, don't feel too bad, you're with most of the crowd, that's where the world is today. But to get to AI native, I think it does mean thinking a lot about these issues: fine-tuning models, the harness, using open-source models to drive down the cost structure of the technology.

[18:55] [AJ] If we shift gears a little and step back to talk about what we do at Mawer, or any investment management firm, as our core business: a big part of our job is to assess the management teams of our investments or potential investments, and seemingly every management team out there says they're all in on AI



right now, we're even talking a little bit about that right now. If you put your equity analyst hat back on for a few minutes, how would you assess, or actually tell, a dinosaur from a native from a rat racer at this point?

[19:27] [JA] I always love the opportunity to put that hat back on, so thank you for the offer. Let me take a shot. I bring this one up because I think it's a really good example of the role of the human in investing. There are many roles, but one of them is to figure out these types of questions. We talked about these buckets, going from dinosaur to native. How do you get at that? I think it really is a people question. You're not going to be able to ask an LLM, "is this company AI native?" It's not going to tell you that; it's going to give you whatever they've curated in their press releases, and you're not going to get real information that way.

The only way you're going to get it, in my opinion, is to hit the ground and go talk to these companies, interview them, and get into the weeds of how they're maximizing the value they're getting out of these technologies. What are the different use cases? Have them demo it to you. Show me the different ways you're using the technology. That's only going to get you to rat racer, "oh, this guy's using it a lot, that's great, I'm learning."

Then you want to start getting into: do you have an open-source model? How do you use it? Do you have your own harness? You may not want to ask it in a leading way, like I'm doing right now, because they might just say, "of course we've got a harness," so figure out a more neutral way to see if they're going after the cost side of the ROI calculation. So those would be some of the techniques: trying to figure out if they take the cost side seriously, if they can articulate the role, if they can show you a dashboard with all the token use across the company and the dollar-per-token spend, and how they've been systematically driving it down. Can they articulate all of that and get down into the weeds?

I think that's the mechanism you'd use to sort that out. I don't see how you could ask that of an AI, because I don't think it's going to come across unless you talk to the people.

[21:05] [AJ] Excellent point, and we're going to have our colleague Jeff Mo on the podcast in the coming weeks, and I'll put that exact same question to him, since he sits in a portfolio manager position currently. I did want to stay on the investment side of things just briefly. If we think about the economics generally, when a technology like this takes hold and an industry gets more productive, in the past a lot of that usually ends up going to customers through lower prices, coming back to some of the things you were talking about earlier. Is that what you expect to happen here as well?

[21:40] [JA] It's another one where I've firmly got that equity analyst hat on, so I appreciate it. The history of innovation and technology is such that, in most cases, if you looked at the base rates, the value ultimately accrues to the consumer. The technology goes up, people start adopting it, and eventually you get the airlines of today: they don't make a lot of profit, but consumers have never been better off, they get tons of access to travel and so on.

There are cases where that doesn't hold. Look at the internet era and internet companies: those businesses were able to carry a lot more of the value. Of course consumers also got a lot of value, but I think the producer, in that case, was able to retain more of the economics, for various reasons, around concentration of services and that sort of thing.



In the case of AI, it's a bit of a jump ball. There's a small number of companies really leading this from the value stack, the frontier labs, as well as the value chain of producers for the chips and so forth. You could imagine a world where they're able to retain more of the economics as they build more expertise and concentrate it, such that competition is unable to bring its power to bear. You can also imagine a world where it becomes very competitive, increasingly, as you see a lot of these open-source capabilities keeping the frontier very competitive.

I'm hopeful that's the way the world plays out, that you have leading frontier models and lots of healthy open-source options, because I think that keeps it very competitive and very good for the consumer, and for prices. If you ask me to wager, I'd bet on the latter, that the competitive environment, and therefore declining prices, is likely the way this plays out. But there's lots of lobbying going on, and lots of efforts to shut down the open-source options, so it remains to be determined how this market plays out.

[23:29] [AJ] If you had your equity analyst hat on there for a bit and were looking outward, I want to look inward now, as an investor: there was a slide tucked into the appendix of your presentation called the investor identity crisis. Can you talk about what you mean by that?

[23:43] [JA] Yeah, we faced this on the technology side of the house about a year ago, roughly when we were doing the podcast with you. Things were changing so quickly. In the case of software development, the feedback loops in that function are really tight, and the technology, the AI, evolved so quickly that it's now writing essentially all of the code for AI-native type businesses, and that happened almost within a year.

It's such a rapid change, and it's been very difficult. In many respects it's exciting to see the productivity potential, but at the same time you've got folks who maybe spent their career developing software and writing code, and you're telling them, the fact that you know JavaScript, or you're one of the best users of JavaScript, we don't necessarily value that as much anymore.

I remember priding myself on being one of the strongest Excel users out there, I almost went to a New York competition for it, and nowadays I think it's probably not that valuable, because the AI is going to know its way around that spreadsheet better than I could ever dream of. That's difficult for humans, because we build these skills and then it changes.

The privilege of the present is to judge the past. We always look at the past and say, oh, those guys weren't so smart. One group in particular I've often looked down on is the Luddites. They're anti-technology, they're breaking stuff, going into the factories, resisting the Industrial Revolution, resisting progress, what's the problem, it's terrible. But you start to realize how this goes to the heart of what it is to be human. It's identity, and these people had their identities, and their identities were being shaken, and they were being told the future was going to be different.

I think we're not in a dissimilar situation today, where, like the Luddites, we see the progress, and it's exciting what we can do and how productive we can be. I don't think anyone wants to go back to 90% of people working on farms, so that probably isn't a good idea. But at the same time, it's hard, the change is difficult, and it shakes you to your core.

You're seeing this in mathematics, and you're starting to see it more and more in investing, where it's really



shaking up particularly that analyst role of how you aggregate information, organize it, and deploy it to a human decision maker. So it's definitely something we're thinking a lot about: what is our investor identity, and how do we evolve that with this technology, while being very compassionate and careful, kind in our opinions, and realizing that not everyone's going to have the same views on these things. We've got to move towards this together, carefully, and not too judgmentally.

[26:08] [AJ] That's a really excellent point, and I was laughing a little bit at your Excel competition comment, just because, I'm probably not as good as you at Excel, but I'm certainly among the best in my peer group, outside of Mawer, in my friend group, and I was always asked to help out on Excel. Now I can't show off those skills as much, Justin, maybe for a period of time, until I know how to use AI within Excel, before they catch on too.

Maybe a last question for you before we let you go: what have you changed your mind about since you were last on the podcast a year and a half ago?

[26:41] [JA] I think the technology is moving so quickly. Maybe if you'd asked me this question naively, I might have said I'd think humans would keep pace a little more. Something that's surprised me is just how quickly the tech has outstripped the ability of humans to stay up to pace. I think that's probably just my own inexperience, and realizing it takes more time for us to adjust than some of the base rates I was assuming.

[27:06] [AJ] All right. Well, Justin, over the last few weeks you jetted across the country and back talking about this. You've now sat down here and walked us through all of it, from the technical and financial to the cultural, and what it all means, both for businesses and investors alike. So I appreciate you taking the time, and I know our listeners do too. Thank you, Justin.

[27:25] [JA] Thanks, Andrew. Appreciate the curiosity, and the chance to be with you again.

[27:29] [AJ] Hey, everyone, Andrew here again. To subscribe to The Art of Boring podcast, go to Mawer.com. That's M A W E R dot com forward slash podcast, or wherever you download your podcasts. If you enjoyed this episode, please leave a review on iTunes, which will help more people discover the Be Boring, Make Money philosophy. Thanks for listening.

Companies Mentioned:

WestJet
OpenAI (ChatGPT)
Google