the art of boring™

EP 239 | AI Investment Research: Uses, Limits, and Where the Edge Remains



[00:00] [Andrew Johnson] [AJ] Hi, everyone. ChatGPT launched almost four years ago, and since then AI has gone from novelty to something a lot of us use every day, including investors.

Today, Jeff Mo is back on the podcast to talk about how it's changing the way he works. Jeff is a portfolio manager here at Mawer, investing across both Canadian and U.S. investments. We get into all the ways AI is proving useful in his role, where it currently comes up short, and where the edge is when everyone has access to the same tools and models.

Here's our conversation.

[00:32] [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.

[00:49] [AJ] Jeff Mo, welcome back to the podcast.

[00:51] [Jeff Mo] [JM] Thanks, Andrew. Always great to be here.

[00:53] [AJ] Thanks for joining us. We're going to jump right in. I was thinking about this before jumping on with you—ChatGPT was released right around this time of year, actually, back in 2022. They really set off this chain of advancing AI uses for the general public.

I wanted to start by asking, when did you first start noticing that AI might be something that could really impact not only the markets, as we've experienced over the last few years, but how you work?

[01:21] [JM] I certainly used ChatGPT 3. I think that was the model that really caught everyone's attention back four years ago, but at the time it didn't seem strong enough to really be helpful in our workflows. It was something that we continued to monitor.

I would say probably early to middle of last year, I'd gotten into this habit of writing really long prompts. This would have been ChatGPT 4, maybe, or 4.1. I was writing a prompt that, if you copied and pasted it into Word, was 8-10 pages long. It was funny because ChatGPT would work for 20 minutes and then freeze because it had run out of context, or it wasn't handling memory in the back enough. I don't know if skills had come out yet, but certainly the concept had eluded me, and I was making it much too hard on AI. But when it did run and go through, I thought the output was decent. It was a little bit wordy.
I remember maybe 16 months ago, I caught it fabricating an entire factory for a company I was looking at. I looked at it, and it had made up a factory in a specific city with a certain square footage. It was very precise.



Nevertheless, I think the gains have been sufficient in terms of allowing me, in that case, to ramp up on a company I didn't know, that I thought, okay, this could really be a game changer. Curiously enough, I think that's also when we started to look for more AI stocks to invest in, thinking that, actually, maybe this is not quite the bubble that some people thought at that time.

[02:55] [AJ] Take me back before these tools were available. How did you decide which companies were worth a real look and which ones just got that 30 seconds? And maybe you can compare and contrast: what does that first look feel like now that you've got these tools available to you, and they're getting better and better every month, it seems?

[03:12] [JM] It's completely transformed some workflows, for sure. In the old days, the pre-AI days, whatever you want to call it, we would start with screening, and this would literally be probably a 30-second view. Everyone has their screening methodology.

Our idea generation screening methodology probably had 20 or 30 metrics that aligned with Mawer's investment philosophy. A few metrics would look at the potential for a competitive advantage existing in the business: things like higher returns on capital, returns on equity, or returns on assets. A high margin sometimes could also signal, although not always, a less competitive advantage.

Then valuation metrics, and growth metrics. That was pretty manual. It would literally be myself or teammates scrolling through a spreadsheet, or it could be hearing about an idea somewhere, manually talking to someone, or it might be from a teammate.

Then we would go into preliminary analysis. I used to actually go through and train some of our new analysts on this. I said, hey, one key success factor in investing, I think, is to be able to reject an idea in 20 minutes rather than 60 minutes. If you can get efficient at that, that's a competitive advantage: you can look at more ideas. I would do a screen share with that analyst and go through how I would look at and analyze a company in 20 minutes rather than 60 minutes, to try to come to a quick go/no-go decision.

If an idea passed that initial preliminary screen, we usually went to an intensive analysis. Sometimes there would be gates, but if we went through the entire intensive analysis process, where we did our full report and a discounted cash flow model using our Monte Carlo simulation method, that could take 2-4 weeks. Sometimes you might pass on an idea after one day, because you realized, okay, based on what I've learned so far, that's just not very attractive. But generally it would be very manual, and it is definitely biased toward spending time and understanding things in a very bespoke way, but still very much reliant on human judgment.

To answer the other part of your question, what does it look like today? I would say screening and that initial analysis is really outsourced, heavily reliant on AI workflows now.

We created a skill that I think has over 350 reference files. The main skill.md file itself has 2,000 lines, and altogether I think there's probably pushing five or even six digits of lines in that skill. A lot of these reference files aren't called frequently, but they analyze each industry or sub-industry that we think is relevant. We call it an initiation skill. It will also write a delta report, but it will give you that old, Mawer initiation that took 2-4 weeks.



I still do think, that the intelligence, and especially the context and understanding of AI is not quite as good as a human. The way I would characterize it is to say, hey, in the old days, let's say 0% is you don't know anything about the company other than its name, and 100% is you know enough that you can make a decision. You can always know more about a company, that's never ending, but you've got enough to make a decision. I would say 100% of that was manual back in the day. Now I would say this initiation skill gets us to maybe 50-60%, perhaps 65% of the way there. That probably saves a week, a week and a half, and it really allows you to be sharper in the questions you ask.

You can get 65% of the way there spending five minutes of human time, probably 1-2 hours of machine time, and it probably costs us $100 bucks to run this skill. You can come to the table and start immediately asking sharper, deeper questions and really get a sense of what the key debates are and how you might go about looking for evidence.

I think we continue to work on monitoring as well, especially earnings monitoring, but also everyday daily monitoring. That's a little bit harder, because the sources of information are a little bit broader than when you do an initiation report. We've found lots of uses for AI. It's giving us feedback on whether our trade decisions were good or bad. It's giving us feedback on the portfolio itself in terms of risk management and portfolio construction. I would say it's touching every workflow, but probably that initial deep dive is where it's really compressed the timeline for us and the team.

[07:47] [AJ] I want to dig a little deeper on all of that, Jeff, because you're getting into, or at least you alluded to, some of the questions that I have going forward. It sounds like you're seeing more companies, you're going deeper on each one, and you're able to do that in a shorter amount of time than previously. I'm curious to understand what that actually changes for you later on, and you touched on a little bit of this. But maybe more specifically, do you think you end up with better ideas or just more ideas on the shelf?

[08:13] [JM] I would say that we now feel more confident analyzing businesses that maybe we didn't feel as confident analyzing in the pre-AI days. One technique we've used is to start the initiation and the investigative process by first asking AI to suggest and then download 30, 50, 100 documents that might be relevant to the company at hand. For example, if you're looking at a defense company, you may want to download a lot of the communication and budget data from the Department of Defense.

In the past, we would also read that, but it was a case of where do you get to the 80/20? You can't spend a year analyzing a company. You probably have to do it in a few weeks. Is the ninth document in the Defense Department's budget that important? Often the answer was no. But AI can get 100 of these documents and analyze them all simultaneously, so why not at this point?

Industries that I think we now feel more confident analyzing would be ones that are highly technical or highly technological: industries like semiconductors, biotechnology, and any other area where there previously were a lot of technical barriers, where we felt, as generalists, we don't know if we are necessarily the last man at the table, so to speak. I won't say that you don't need any domain expertise, but I think AI has allowed us to ramp up the curve a lot more quickly.



Another aspect is this is not just for initiation skills. You can have very specific monitoring points now. On earnings, you could ask: is this particular company selling more connectors, or is it being designed into more connectors when AI data centers go to 800-volt direct current? That's something at least a little bit technical. Because we've scraped every public statement that NVIDIA CEO Jensen Huang has ever made, and we've downloaded all the technical documents and got AI to look at it, that has really allowed us to answer some of these questions in a much richer manner. Whereas in the past, we might have said it's too difficult to become an expert in high-voltage interconnects. It's not really something I went to school for.

[10:28] [AJ] It's fascinating, because this is changing a number of parts of the process. We just spent a little bit of time there on the initiation, the delta, and the monitoring of individual holdings. If we step up a little bit, Jeff, and think about the portfolio as a collection, we often talk about wanting contradictions built into the portfolio. I'm curious, does AI change the way that you can look at that as a portfolio manager? Has it uncovered more detail or allowed you to see things from a different perspective?

[10:57] [JM] One of the areas where AI has been a game changer is on screening. There are two ways to screen. One is your very basic quantitative screen, and that doesn't really change. When you throw a ton of numbers at AI, it costs a lot of tokens and it's not very effective. For that type of stuff, stick to your databases and what you've traditionally done. But you can pair a quantitative screen with a qualitative screen.

For example, we were recently looking at the life insurance industry, because there's a theory that with interest rates rising, and with longevity maybe going up as AI might improve health outcomes for some individuals, life insurance companies should be beneficiaries. So you can have that thesis.

We also thought life insurance was a good balance or contradiction in the portfolio, because life insurance would do relatively better if rates were higher. Life insurance companies generally have lower P/Es and more stable growth. That's a very different profile than a company benefiting from AI capital expenditures.

We could actually do a screen and say: based on your best understanding of what makes a life insurance company high quality, and, by the way, here's a primer that has been written on life insurance to give AI some context, why don't you look at all the publicly traded life insurance companies and tell us which five or 10 you think look most attractive? Here are some quantitative metrics you should think about. Here are some qualitative factors you should think about as well, in terms of how life insurance companies typically have competitive advantages, do you want companies that sell more annuities versus more insurance policies, and so forth? That's a micro example of ways that we've screened better.

Another way to build in contradictions is that, because AI can dig into the true drivers of each business en masse, we can better understand which companies are benefiting from the electrical supply chain build-out. With AI demanding more power, and power generation growth being higher than it was in the past, that industry is having a renaissance, and we wanted some exposure there. It was much easier for AI to pick out, in some cases, a couple of companies that we maybe didn't know initially had a lot of power exposure. At the same time, you might find something that is very different and maybe contradictory. AI could give you that underlying demand driver view that isn't fully captured by classification systems like a GICS system.



[13:25] [AJ] When I think about Mawer's decision-making model, we refer to it as leader decides with input. Is AI one of those inputs now, or is this more similar to another tool in the toolbox that helps you do the work?

[13:35] [JM] I think in some cases its input is another source, compared to internal analysts, external analysts, other investors that we may have an opportunity to talk with, and management as well. Management usually is biased, wanting you to own their stock. Our model is still leader decides with input, as we talked about.

It is, I think, always going to be a human leading and making that decision, and I would say still framing the decision landscape: what are the key questions? What are the key debates on those key questions? How is the evidence balancing between positive and negative? Those things I still feel are very much a human role.

AI is both a great tool as well as an okay synthesizer. As it gets better and better, I think we would probably trust it more. But at least the way we use it today, I don't quite feel comfortable saying, hey, just go ahead and build a portfolio and trade for us, or something like that.

[14:34] [AJ] I think I know the answer to this question, based on what you've alluded to so far, Jeff. But if every manager and investment team out there has the same AI models reading all the same data, the filings, the financials, digging into the management commentary, where do you think the edge will come from going forward?

[14:53] [JM] I think it's the same edge as before AI, because you could say the same thing: if every manager out there read the same filings and financials and the same management commentary, where is the edge from? Our belief, or philosophy, is that it's based on investment philosophy and process, and also our accumulated experience laying out that process over 10,000 hours, using that Malcolm Gladwell rule of thumb.

There is a lot that goes into managing your AI. The first question you have is, what context will you give this prompt in this moment? Context would be, I think, a company filing, an annual report or a quarterly report. It could be research. It could be expert network transcripts. It could be reports from our sell-side partners. As you can imagine, very quickly that gets quite complex. So you have your data, or your context, and when you give the agent that, it's very important.

Your harness, and on each turn how it makes a decision, that's quite important. Then you have your orchestration. Earlier I was talking about how we wrote all these skills. We have a skill that will tell you how to analyze a semiconductor manufacturing company, and the specific information and context that we've given it. Because we've written that, we would ask it to think longer-term, because we're long-term investors, and so that is very much tied to how we think. Then maybe we'll say, for semiconductor companies, one source of competitive advantage, which again is something that's very core to Mawer's investment philosophy, is XYZ, so look for that in the filing, in the sell-side reports, in the context that you've given the agent.

As you can see, if we do a good job—and I'd say we're not there yet, but we're moving in that direction today—of managing our context in a hyper-focused way, managing our harness in a very bespoke way that fits our investment philosophy, and then building out these orchestrations, these skills, these workflows, in a more appropriate way, that is an added layer.



I actually think it's even more complex, or bespoke, now, because not only do you have to train your people, or, first of all, you have to define your investment philosophy, and from that philosophy create a custom investment process, and then train your people on that process. You also now have to essentially train your AI, in its workflows and its context and all the things I talked about, on your investment philosophy and process. Then, finally, retrain your people to work with your AI and understand how to manage the context and the harness and the orchestrations.

At least for us, internally, we've made a decision to build our own harness as well, so we can better control these turns. I would say today it's still, at best, at parity with Claude Code or some of these tools, but we do think that as we continue to get better and better, having our own internal harness is definitely a potential competitive advantage. And if it never becomes a competitive advantage, because people are innovating on their harnesses so much, that's only one third of the piece.

We can still have our own orchestrations. We can still have our own context management. We're building what we're calling a knowledge base, or a knowledge brain, which is essentially a fancy way of saying database. It will allow us to ingest information into it, like a 10-K annual report, annual filing, and then it will chunk that data into pieces. Let's say you're analyzing the capital allocation of a company. You would only pull out the management piece of it, so you don't have the whole 10-K polluting that turn, and it will probably be a lot more efficient token-wise, and therefore cheaper for us to use as well.

Sorry, that's a lot, but I really do think that AI has actually made it more differentiated, not less differentiated.

[18:38] [AJ] It completely resonates with me, Jeff. What I took away from that is it still all comes back to people, humans, and how we use these tools in order to get the most out of them and apply them to the specific context of our investment philosophy and what exactly we are looking for. You've also touched on some really important points. Our colleague Justin Anderson was just recently on the podcast, and he talked about orchestration and that context window and the degradation of that context as you make those turns in one use case. The importance of acknowledging that makes you a better user of that tool through time as well.

One last question, Jeff, because you touched on management teams earlier. Every company seems to have an AI story these days, including us, as we've just talked about. But when you're talking with management teams and they start telling their story, what are you actually listening for to ensure that they understand not only how to use AI, generally speaking, but how to use it in a way that could make their business more efficient, or what have you?

[19:44] [JM] I would say go back to first principles again. Earlier, when you asked me how AI changes the investment landscape, I said the first principle is: how do you pick stocks? When you say, how do you differentiate between essentially good AI stories at a company versus maybe just a story, I would say it goes back to how do you treat management teams?

This is not the first time management teams have claimed they have some secret sauce in some new domain. In the past, they were really good at other areas because of how they modelled their acquisitions, or they were really good at lean manufacturing and lean operations. Are they really good at coding, or anything like R&D?



You need to understand the domain enough to be conversant, and then you go through your mental models. One of the first mental models that I frequently teach new hires here is the concept of inputs versus outputs. People love to talk about outputs because they're outcomes. They're tangible things that you can see. In our business, we love to talk about outputs too. We like to talk about our performance, especially when it's better than the market average, and so on.

But of course, it's actually the inputs that create those things. It's all the stuff that we spent most of this podcast talking about. We're doing all this because we're hopeful that it will lead to superior outperformance, which is an output. But the inputs are how you manage your context, what data you have, what your orchestration is, your skills.

That's the mental model I would go into when talking with a management team: to understand, do they actually understand how to use AI? Are they talking about what I know to be the basic building blocks of AI usage? Are they thinking about how to feed their models the right context at the right moments? Are they thinking about the pros and cons of using different harnesses, and also how to manage their turns during each run? Are they orchestrating intelligently?

Then also, when they sell AI tools, which I think is the first piece of it, a lot of software companies, especially these days, are now selling an AI module or add-on. I think you differentiate between one company we own in our Canadian small cap strategy that is, I would say, primarily using an off-the-shelf AI tool and reselling it, presumably with a markup, and through their own channel, essentially as a transcription service in a regulated industry. Whereas another company is completely redefining what insights you can get out of the data that the customer has, and is helping their customers find more ways, using their data, to save money or bill for revenue that they may have missed. I would say company B in this case is doing a much better job than company A of truly having an AI offering, or using AI to improve their business.

There was one company in our portfolio that reduced their staff count by 40% and said, simultaneously, that they expect to have an accelerating pace of R&D and new product releases. They're a software company. You just cut almost half your workforce, and you're saying you're going to increase your speed. But when you listen to how they talk about how they manage people—unleashing their top performers to be more decentralized and empowered to make more decisions more quickly, and giving them very large token budgets—essentially now AI can be the coding team for the decision maker on the small team, rather than in the past, when they may have had three or four maybe mediocre coders working for that top coder. That really helped us understand that this company does make sense, because they're actually moving off the dead weight. Now the AI is good enough to replace an average coder, and they're really accelerating their strategy.

Those are some examples of situations where AI is transformative for a company versus just repackaging something and calling it AI and making a buck off of it.

[23:33] [AJ] We've covered a lot of ground today, Jeff. We started by talking about how you're seeing more companies and going deeper on each one, how it's changing the way that you look at the portfolio as a whole, where the edge remains for investors using these tools, and of course what you just talked about, which is what you're listening for when management teams tell their AI story. I suspect, Jeff, that if we check back together in a year, we're going to have plenty more to talk about, so maybe we'll mark our calendars for that. This has been



really helpful, Jeff, and I just wanted to say thanks for walking us through it.

[24:04] [JM] Thank you for having me on. Always a pleasure to chat with you, Andrew. Take care.

[24:09] [AJ] Hi, 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:

OpenAI
NVIDIA
Anthropic