the art of boring™

EP 236 | AI Build-Out: What We Know, What We Don’t, and What’s in the Stock



[00:00] [Rob Campbell] [RC] Coming up on The Art of Boring, my colleague, Irena Petkovic, joins me to separate what we know about the AI build-out from what we don't. We cover the latest hyperscaler earnings, why supply now looks like the bottleneck, and where the midterms and financing fit into the market reactions that we've seen over the past few weeks. And channeling Charlie Munger's advice to invert, always invert. Irena makes a case, process-wise, that we should always be asking what's in the stock, in addition to what it is worth. Irena Petkovic, up next.

[00:31] [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.

[00:48] [RC] Irena, welcome back.

[00:49] [Irena Petkovic] [IP] Thanks for having me again, Rob.

[00:51] [RC] We're recording this on Monday, November 9th. It was a heck of a week last week. Maybe just share, for those who are unaware, what we saw in markets.

[01:00] [IP] There was particular volatility last week among a number of AI-related names, but I think that the volatility last week was the continuation of a lot of volatility we've seen in the AI trade over the past month, I'd say. The reason for that has been a number of things, including anxiety around financing. Then, last week, the market was very nervous around what the midterm elections in the U.S. could mean for data center construction. That caused a lot of volatility, and things have been swinging around a lot. I think when markets move this much, there's a temptation to have a confident answer for everything.

But similar to what Paul was saying when he was on the podcast a couple of weeks ago: when we approach the market with a little bit more humility, that can enable us to play good offense and good defense. I think part of that is slowing down and acknowledging what it is that we know to be true, or at least have good evidence for, and what we don't know. Then given what we don't know, how can we actually manage those risks?

[02:03] [RC] When you see markets move the way that we did last week, part of me thinks there might be genuinely new information that's causing prices to reset. Another part of me thinks there's a huge amount of uncertainty and there's some price discovery that's going along with that. Can you share what it is that we do know today, and then let's get into what we don't know.

[02:23] [IP] In terms of things that we do know, I'm going to sound like a bit of a broken record since I was last on



the podcast. We did get another set of hyperscaler earnings, and when I'm talking about hyperscalers, I'm talking about Amazon, Meta, Microsoft, and Google. We continue to see very strong results out of their AI infrastructure businesses.

Much like last quarter, we continue to see an acceleration of revenue growth. Amazon's AWS reported their sixth straight quarter of revenue acceleration. Google's cloud business doubled year over year, which is up from an already very impressive 80% growth last quarter.

I think, most importantly, that revenue acceleration was accompanied by margin expansion as well. That, taken together, is further confirmatory evidence for us that these hyperscalers are generating an attractive return on the CapEx that they've invested in building these data centers. Another thing that was different about this set of earnings, and was a message that they talked about last time but they put greater emphasis on it this time, was that now demand for AI infrastructure is not the bottleneck. It's actually supply.

We heard executives from all of the hyperscalers talking about the growth of their cloud businesses actually being limited by the rate at which they can bring new data centers online. Given what I said at the beginning around the market being nervous around what the midterms could mean for data center construction, that caused a lot of anxiety in the market as well.

Another piece of new evidence that we got since I was last on the podcast was that Anthropic filed their S-1. An S-1 is the most meaningful piece of disclosure that a company gives before they become a public company. And Anthropic, in particular, is important because they are disproportionately selling to the enterprise, whereas OpenAI is more targeted at consumers. I talked a lot on the last podcast about how it's important to see AI diffuse through the enterprise, and that we see evidence of broader adoption beyond certain departments, etc. And Anthropic's S-1 really gave us some confirmatory evidence that there's not only room to grow within existing customers, but also a very long runway of potential new customers.

They talked about how the number of customers spending more than $100,000 per year on tokens has grown from 200 in 2024 to almost 12,000 as of September 30th, 2026. That's a huge amount of growth, but on an absolute basis, 12,000 isn't that many customers. In terms of growth within existing customers, they shared that more than half of the Fortune 500 is spending more than $100,000 on tokens, but that spend represents less than 0.1% of their total expense base. Their best customers are spending more than eight times that. There's room for even those best customers to increase their token usage, and that points to continued adoption of AI in different departments and for different functions. We don't have any evidence that AI is reaching any type of saturation point.

Those are the big things we know and that we got new evidence for.

Something that I think became more true, especially in the eyes of the market, is that the continued build-out of these data centers is going to rely on external financing. I talked last time about how the hyperscalers have depleted their internally generated cash flows, but now you're starting to see these companies actually come to market with greater frequency. Meta came to market after their quarterly results for their largest equity raise ever. Google came to market over the summer. We're seeing that become less of a theoretical and more of an actual truth.



[06:03] [RC] To repeat something you said a few minutes ago, this is really the supply side, whether it's in terms of components or the financing of those things. You're saying that the market reaction that we've seen has been primarily related to that. Am I hearing you right?

[06:17] [IP] Yes, it is related to that.

[06:19] [RC] Maybe for my own understanding, how much of what we saw in the last week was related to the midterms specifically, as opposed to this build-up of concerns around the financing, which are not new and had been bubbling for some time?

[06:30] [IP] I would say that the volatility we saw over the past week was related specifically to the results of the midterms, but it was compounding on existing fears about financing. There's some reflexivity in this too, because, like I said, the markets need to remain open and amenable to these hyperscalers raising at attractive rates.

Then you also have the labs themselves, Anthropic and OpenAI, looking to go public. Having market conditions that are amenable to that, it's a self-fulfilling prophecy with this infrastructure build-out. If people are worried about it, then the market is weaker, and it creates weaker funding conditions.

[07:07] [RC] What are the key things that, if you were to have a crystal ball, you'd love to have clarity on that would really make sense as to whether this market reaction has been appropriate or not?

[07:11] [IP] I think there are three things, two that we have already touched on, that would be really good to have a view on into the future. The first is, of course, what the actual impact of these elections will be, or honestly, any political party on data center construction. Several Democrats campaigned on slowing data center development. That was specifically concerning in the states of Georgia, Ohio, and Texas, because those states represent the majority of gigawatts that are supposed to come online in 2027-2028. We really don't know what policies they're going to put in place. What we do know thus far is that the governor-elect in Texas has said that data centers that have already been permitted are allowed to finish construction.

That largely de-risks the pipeline, at least from that perspective, for 2027 and 2028. But there's any number of ways that these politicians can try and change the process for constructing a data center. There can be different concessions they have to make, longer timelines, or an outright moratorium. We just don't know.

Again, the second thing is how the build-out will be financed. That is very much dependent on external financing and we just cannot know whether market conditions are going to remain supported.

The other big thing that we haven't talked about yet is how internet traffic will change. Increasingly, we've seen tokens being used by agents, as opposed to people using LLMs or the chat interface. That was accelerated by the launch of products like Muse from Meta.

The reason why that is so important for us is that agents don't use the internet in the same way that humans do. Agents don't browse the internet the way humans do, they don't look at or click on ads, and they also don't respond to traditional brand marketing. As agents make up a higher percentage of internet traffic, there are a lot



of questions that are relevant to many names in our portfolio.

The most obvious company would be Google, with their search business. It would also be relevant to other businesses in our portfolio, such as payment businesses and travel businesses. That's another thing that could change drastically, and it has increasingly come up this earnings season.

[09:30] [RC] I would think either for better or worse, in the sense that some companies in the portfolio might actually benefit from the way that internet traffic changes.

[09:37] [IP] Absolutely, it's an open question. Would more value accrue to whoever facilitates the transaction, like a payment network, or would more value accrue to whoever owns the agent, like a Meta or something like that? There are any number of combinations of how the economics could be split. But, again, it's an unknown in our portfolio that we have to manage.

[09:57] [RC] You mentioned humility earlier. One of the things that I think is a real asset in the way that we conduct our investment process is the way that we look at discounted cash flow models: the stochastic nature, running the Monte Carlo, thinking in terms of ranges and scenarios as opposed to point estimates. Of course, I imagine that work continues. But I know there's a supplemental piece of work, or almost a reverse engineering of the DCF, that takes place as well. I'm wondering if you can share a little bit about that.

[10:23] [IP] When you talk about reverse engineering the DCF, the question we're answering is, "What's in the stock?" This is something that we're all asking ourselves at increasing frequency, given the increased stock-specific volatility that we're seeing. The stock that you see yesterday is not necessarily the stock that you see today, because the price has changed such that the risk-reward is changing.

And what we're trying to do, as the price changes, is trying to understand what assumptions are implied by that price today. We're looking for what kind of revenue growth underpins that price today, what kind of earnings growth, and things like margins and even duration of growth. Those are things that we're increasingly testing.

And then we're taking a step back and asking ourselves, "Do we think this is a reasonable assumption for this input? Do we think this is a reasonable margin for this business?" That's particularly important when we're looking at companies in the AI complex, because it's all very much connected together.

We can look at a business and see what kind of revenue assumptions support the price today. And then we can ask ourselves, "What level of hyperscaler CapEx is required for that? Is that reasonable for the companies we cover? What does that imply for data center construction?" Those are the types of things that we're looking at and the questions we're asking ourselves.

The end result that we're trying to get to is we're trying to see where the company sits in terms of its distribution of outcomes. Is it priced for total perfection, and there's no allowance for any type of deceleration or any type of increased competition or what have you? That would create negative skew. Or do we think that the assumptions are more conservative, and that we think that it's at the lower end of its distribution range and we could have positive skew? That's the kind of work that we're increasingly doing in response to a market that's generally more volatile.



[12:11] [RC] I know we look at things business by business, but, in general, as you do that work today, in terms of what's in the stock, does it make you excited? Does it make you fearful? We're in that continuum overall, have you and the team gotten to with respect to what's in the stock?

[12:26] [IP] One thing that we have observed when we've been doing this type of analysis is that we take greater comfort in businesses where we feel that there are multiple ways to win. We're trying to stay closer to option-rich businesses. As I talked about last time, the hyperscalers have multiple ways to monetize the capital that they've invested in these data centers. We're increasingly leaning away from businesses that may only have a certain path with which they can win. It's a path dependency or a reliance on a certain set of customers. Those are things where we have a little bit more apprehension, and we continue to look for things that are more option-rich, as Paul would say. Having multiple ways to win, really.

[13:10] [RC] Have you found that the current environment has provided a need for much adjustment in the portfolio?

[13:19] [IP] In terms of the last few weeks, honestly, it doesn't require that much adjustment to our process. I think that when you say "the current environment," that actually extends over a much longer period, and there have been structural changes to the market and, frankly, in the world that we've just had to adjust to. As you've heard many other people on the podcast talk, there's greater volatility geopolitically, commercially, even AI is disrupting businesses and creating new industries at a faster rate than ever before. That has created more volatility in the market.

You also have changes in market structure. You increasingly have multi-manager hedge funds that are significant participants in the market, and they're more focused on short-term catalysts and earnings revisions. We have just had to adapt and do things like continually reassess what's in the stock and get tighter on that existing process. I don't think that that has changed specifically over the past few weeks. It's been an evolution of our process that we've had to make over the past few years.

[14:22] [RC] Some of the points that you outlined earlier, I'm curious: what have been the spiciest points of debate within the team? Understanding humility, but sometimes you have a view. What have some of those points been on over the past couple of weeks?

[14:35] [IP] I think a big point of debate across multiple asset classes has been the way in which internet traffic will change as a result of AI. In many of the large-cap portfolios, Alphabet is a significant holding, and there was a lot of debate going into their latest earnings report, whether the price was reflective of any type of even deceleration, not even a decline in search revenues. And we had anecdotally been hearing about declines in click-through rates. We had all been experiencing when you perform a Google search and you get the AI summary on top and you stop there instead of scrolling through. It started to percolate in our minds that this may be a risk, and then we had to look as to whether this is in the price or not.

And there was a lot of debate as to what this actually means for Google, but looking at our models, it suggested that any type of even slowdown in search revenues was not in the price. Even very simply, the last time that search was a concern for Google was in 2023, and the share price was $150. Even though earnings have grown



since then, obviously, when the stock is over $300, you can generally say that it's probably not something that is in the stock.

That's something where there was a lot of debate on what's going to happen and what we know and what we don't. Ultimately, it came down to what's in the stock and what our balance of probabilities was. We felt that it was most prudent to actually trim our position back, and that ended up being a good decision. Again, this is all about thinking in bets and thinking probabilistically, and it could have turned out to not be a good decision in hindsight, but this time, it was.

[16:15] [RC] A lot of this conversation about what's in the stock. There's a long history in investing. And, in fact, we've talked about this on this podcast going back several years: the notion of variant perception. Knowing what the market assumes, and then if you can figure out that you've got a view that's different and well justified, that's your cue to act. You talked about what's changed in markets and how we've adapted. How has the concept of variant perception changed? Or, put differently, what are the sources of genuine variant perception today, whether it's as it pertains to the hyperscalers, the AI trades, or anything else?

[16:49] [IP] I think that the concept of having a variant perception hasn't necessarily changed. The path to getting there can look different. When long-onlys were the marginal price setter in the market, it was more likely that you had a smoother path to realizing or testing your hypothesis around a longer-term variant perception. "I think this business's ultimate profitability will be X in five years." That would be an example of variant perception.

Now, since there are more participants in the market who are reacting to shorter-term catalysts, and we're more short-term focused, the path towards testing that hypothesis can be more volatile. It doesn't mean that your variant perception is going to be wrong or right a different percentage of the time. I don't think it has actually changed whether your variant perception is true more often than not, but I think the path to getting there can be more bumpy.

The way that we've responded to that in our portfolios is by running more names and trading generally in smaller increments. That gives us room to react when we feel that there's a dislocation, and we feel that the fundamentals are still intact, and we have confirmatory evidence of our variant perception. It also enables us to change our mind more quickly and more easily if we feel that we've gotten some evidence that our variant perception is perhaps not true. I think it's caused us to be more active in our trading and to stay more agile.

[18:21] [RC] And as you said earlier, to make sure that we're being humble. Irena, thanks so much.

[18:26] [IP] Thanks, Rob. It was a pleasure.

[18:27] [RC] Hi, everyone, Rob 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:

Amazon
Meta
Microsoft
Google
Anthropic
OpenAI