



Licensed by the California Department of Corporations as an Investment Advisor

## Investment Newsletter – December 2025

### Executive Summary

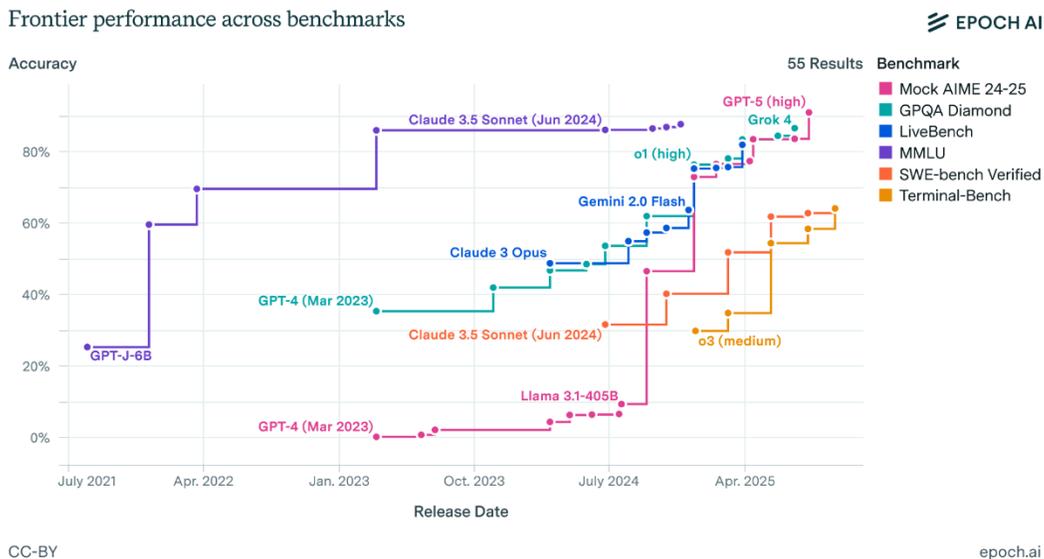
Artificial Intelligence models can help us do many tough tasks yet can also fail on simple things: they are becoming more human. We need to “interview” the various AI before we pick one for a job. The new Anitgravity agent platform from Google is the Swiss Army Knife of AI, albeit tricky and risky to use.

The AI race has transformed the technology sector into big capital spenders. This new paradigm seems likely to produce much lower returns for the sector than in the recent past. Caution is imperative.

According to Goldman Sach’s calculations our International Equity strategy earned 32% for the year.

### Artificial Intelligence 2025 Update: Capabilities, Limitations, Expectations & Implications

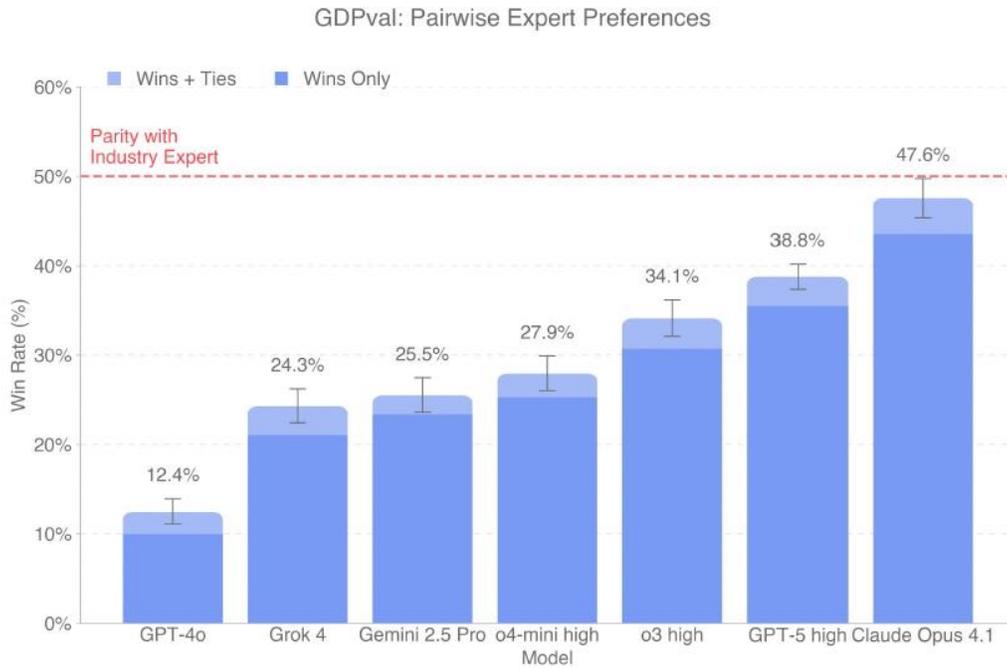
The technology industry uses a variety of testing benchmarks for measuring the smarts of the large language models that we call Artificial Intelligence (AI). These are similar to exams (or trivia games). We can track progress this way:



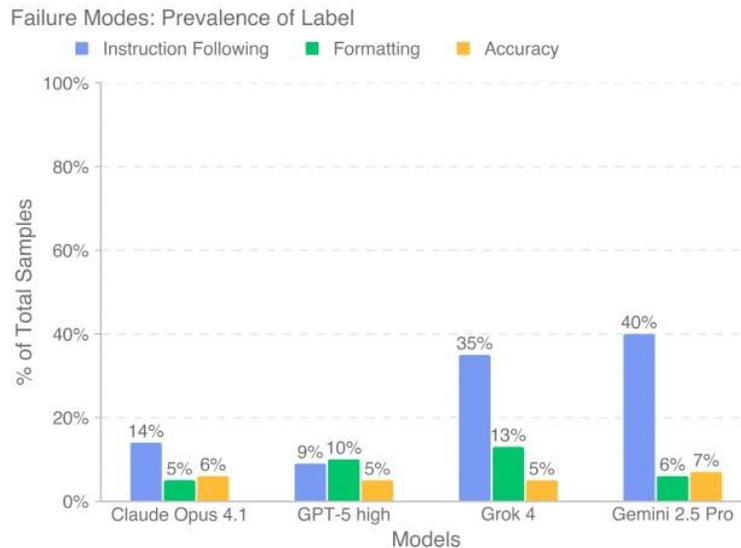
# Berkeley Investment Advisors

## Investment Newsletter – December 2025

What this means for real world usefulness is another matter. A recently published paper from OpenAI proposes a new metric: *"GDPval, a benchmark evaluating AI model capabilities on real world economically valuable tasks"*. This is more useful; it will help us understand how Artificial Intelligence compares to human intelligence. They created real tasks for real industries and recruited human experts from each industry to perform the tasks. Then AI did the same tasks and the results were judged by a different set of industry experts (without knowing what was from AI versus a human). Here are the results from the paper:



Overall, humans are better at their jobs than AI, and, the Claude Opus model was the top AI model. This study also looked at why the AI failed for each task. The results give us a good picture of where things stand when using AI for real work:

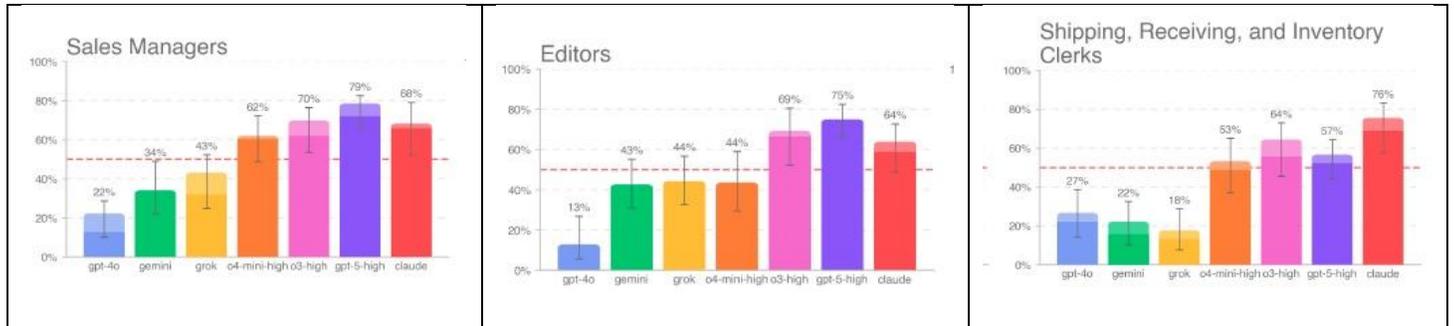


# Berkeley Investment Advisors

## Investment Newsletter – December 2025

All the models have trouble following directions. In June 2024 I explained that results varied a lot depending on the quality of the prompt you used. This has become much less important over time, but we still see that AI is challenged to understand what we want it to do. I think that this problem can and will be solved, as will the formatting problems. But notice the accuracy error rates – none is lower than 5%. It is not clear at all that the current approach can push this rate much lower. For a lot of real-world tasks (such as finance!) this is too high to be useful.

There are, however, many uses for AI where accuracy problems don't have much impact. Looking at model wins for a few vulnerable jobs, we see:



Claude and ChatGPT bested the humans more than half the time in these three occupations. Conversely, engineers, financial managers, and accountants are unlikely to lose out with the current technology. Note that google has since launched Gemini 3 which is a far better model than was 2.5.

In the June 2024 newsletter I first wrote on this topic. Artificial Intelligence (AI) models have advanced significantly since then. No technology industry test can tell you whether AI is useful for what you need to get done. You have to test it on your own tasks. Hence this update in which I'll show what it can and can't do. AI is getting more human-like in that it has strengths and weaknesses. As with any job candidate, we must interview our AI to see who will work best for the job we have for it. Therefore, I'll go through my experiences for my use cases – which function as my AI "interviews". But first, let's look at who will be interviewed.

Currently the main U.S. models are:

- Gemini 3 from Google
- ChatGPT 5 from OpenAI
- Claude Opus 4.5 from Anthropic

Perplexity, which I'll also review, is an AI application that uses a mix of these and other international models.

China also now has some prominent models. Early in 2025 DeepSeek R1 was announced. This model, which is now on version 3.2, has capabilities that are very competitive with the top U.S. models. It shocked the AI world because it was developed with a fraction of the computing resources. On top of that, it was open source meaning anyone could download it and run it. (I installed this on a virtual machine and used it to run a translation application I created using Claude). Running this model can be 10 times cheaper than the proprietary U.S. models. Therefore, it is a direct challenge to the profitability aspirations of the U.S. AI

industry. However, it is not yet as polished as the U.S. models. In this article I'll focus on the U.S. models.

These models can do some amazing tasks yet fail at tasks that would seem to be far easier: AI models have a "Jagged Frontier" of capabilities. In some areas AI can do Ph.D. level work, but then it can fail at something a 4<sup>th</sup> grader could handle. The term "Jagged Frontier" was coined by a professor at Wharton named Ethan Mollick. Professor Mollick studies entrepreneurship, innovation and artificial intelligence and writes a blog about AI called "One Useful Thing". If you are interested in following the field, reading this blog is a good start. He provides many examples of the Jagged Frontier – as will I.

In June 2024 AI agents had just been launched. Now there are many startups building agents for various specialized tasks. These agents (which I refer to as derivative agents) use the main models as their brains but use custom made system prompts (also known as agent instructions). Github.com has a page with a whole list of AI agents: <https://github.com/e2b-dev/awesome-ai-agents> . I have tested two of these derivative agents extensively for finance work: Manus.im and Shortcut.ai. There are many I have not tested and it's not that hard to create your own agents.

The most powerful (and risky) agents are coding agents that you can run in your own computer environment. Whether you use the basic AI models via a chat or through specialized agents depends on exactly what you're trying to do. As before I will analyze the current state of artificial intelligence using specific use cases that I think are common and useful for gauging where we are.

### **Writing and Editing**

Writing is a strength of AI. If you do a good job of describing the details of what you want it can write a good draft and then it is relatively easy to edit yourself or use further prompting to modify the document. I successfully used Claude to write an agreement for neighborhood drainage improvement work. One of the neighbors had a lawyer look at it and I got no request for changes. I've also used it frequently as my editor to point out ways to improve my writing. This is a great use case for AI.

### **Document Summarization and Analysis**

A new regulation from the department of labor came into effect recently which requires very specific processes for Registered Investment Advisors to accept retirement plan rollover assets. The two main Federal Register documents regarding these requirements total 177 pages! Most of this does not apply to my case but some key parts do apply. I used AI to extract and detail the relevant requirements. I confirmed what it came up with by reviewing legal articles and reading the cited parts of the regulations. Then I submitted my draft compliance template to ChatGPT for review to check whether my approach would satisfy the regulatory requirements. I got back some useful recommendations that I included in my template. This took quite a lot of time but without the AI tool, it would be far more costly to comply with these new regulations. Another huge time saver.

## Research

As I explained in June 2024 Perplexity has mostly replaced google search for me because it is much more thorough and provides its references. Google has now put its Gemini AI at the top of its search results but the reference links are not as easy to use. I also think Perplexity has better answers, but this might be more my personal taste. AI answers are still sometimes wrong so you cannot just accept them without verifying by looking at the source.

On the other hand, an attempt to use AI to read and analyze Medicare Part D drug insurance documents led me off the “Jagged Frontier” of AI capabilities. The key documents for analyzing risks are the formularies. These are 100 to 200-page pdf files listing every drug they cover and their rules for it. It sounds perfect for AI but it turns out that no AI can actually read these files completely and accurately. When I pointed out errors, they explained their limitations (see the Appendix). A suggested solution was extracting to text files, but ultimately, they could not read these either. I tried all 3 major models and they all failed. In that process they suggested many work arounds, always assuring me that this time it was going to work, but nothing they suggested enabled them to access the complete information in the files. I.e. they sent me down the AI rabbit hole. AI can be useful but not always.

## Data Gathering and Financial Analysis

I requested Claude and Shortcut.ai each to build a spreadsheet to calculate the build up of retirement savings to reach a goal. The Shortcut.ai spreadsheet had several problems that any finance student would probably avoid. Claude produced an acceptable result on the first try, though some obvious improvements were in order. When I asked for enhancements, it successfully added them. I plan to test even more complex spreadsheet building in the future using Claude.

A really useful AI agent would be able to update my various financial data spreadsheets which are currently done by an assistant. This is not a difficult task for a human (except maybe the boredom of it). But here again, I’ve found a task beyond the frontier of AI capabilities. We’ve tried agents powered by all the main models but none can do this task with enough accuracy to be useful. The best agent for this was Manus.im (bought by Meta) but it still had an error rate in the range of 20% and this agent costs more than my human assistant. Here are some of the common problems for AI gathering financial data:

- They fill in “simulated data” – i.e. they make it up (can AI be lazy?)
- They lie – saying they are working on it when in fact they are not
- They give back a spreadsheet with only a few cells filled in.
- They can’t handle dynamic web content that relies on JavaScript - such as dropdowns, tabs, or buttons like “Select Investor Type.”

## Solving Problems via Coding Agents

Claude’s coding agent is useful, especially for folks like me who are not expert enough to build stuff on their own. In February I was able to use Claude code to build a proof-of-concept translation program that used a locally installed DeepSeek R1 AI to power it.

Berkeley Investment Advisors  
Investment Newsletter – December 2025

---

In May I decided to use ChatGPT to build an AI agent to pull closed-end fund data for me. If you are really, really persistent you can likely get things done. But, more often than not, the AI sends you down a rabbit hole and you eventually have to give up. To see what I mean you can read my session with ChatGPT where I try to get it to create an agent to gather closed end fund data for me:

<https://chatgpt.com/share/68382f47-b234-8005-a436-402f1a44e3dd>

I never got this approach to work – though I did learn a bit about the technicalities of building software.

In October I decided to build a tool using AI that could enable users to analyze the risks of Medicare Part D drug insurance plans: the potential cost of drugs not in plan formularies. My first attempt failed: I discovered AI could not read adobe acrobat files. In that process Claude suggested I use its coding agent to turn these into text files. It was quick and easy to use Claude code to convert the files. I count this as a coding success, even though I then learned that AI could not read the text files either<sup>1</sup>. Unfortunately, this very useful agent was disabled on my Linux virtual machine by an update to Claude code in November.

In late November Google launched the new Gemini 3 model and, along with it, an Integrated Development Environment (IDE) called Antigravity which lets you use Gemini, ChatGPT, or Claude as an agent that writes applications for you. This is a Beta version for testing and therefore free to use. I agree with Professor Mollick that this is going to be a big deal so I'll quote his blog post at length:

*"A fundamental perspective powering AI development is that everything you do on a computer is, ultimately, code, and if AI can work with code, it can do anything someone with a computer can: build you dashboards, work with websites, create PowerPoint, read your files, and so on. This makes agents that can code general purpose tools. Antigravity embraces this idea, with the concept of an Inbox, a place where I can send AI agents off on assignments and where they can ping me when they need permission or help."*

Gemini 3 was getting rave reviews and I was very optimistic. So, I decided to use Antigravity and all three frontier models to build a Medicare Drug Plan Risk analysis application. When setting up a coding agent you want to control the agent's access as much as possible. I asked all four AI partners, Claude, Gemini, ChatGPT and Perplexity, to help figure out exactly how to set it up securely. In the end, the level of security I wanted turned out to be impossible with Antigravity. I settled for a slightly less secure setup but not before wasting a lot of time and AI use credits. Lesson learned: AI does not check for roadblocks to the end goal; it will lead you step by step down the rabbit hole to the dead end.

My goal was for an AI agent to help users understand their risks given the plans available to them. I prompted each of the three top AI to give me a detailed plan for building my system and used a combination of these to instruct the code building agents. Using Claude I built a sharable "Artifact" that was the user facing interactive agent. Claude also mapped out the data sources needed to make this work. The suggested plan was for me to build an online database using the data

---

<sup>1</sup> Apparently, the problem has to do with their algorithm for "chunking" long documents. This destroys table formatting and renders the documents only partially readable.

files from the Medicare system and then have the Claude Artifact access this data via an application programming interface (API). I used Claude inside Antigravity to build the database and API. However, once I had the API working, I found out that a Claude Artifact cannot access any such API due to networking limitations.

I was too far down the rabbit hole to give up, so I decided to use Antigravity to build an application to be hosted on a free service. The whole architecture of this system was a result of question-and-answer sessions between the AI agents and myself. I had Claude in the browser, hand off its code to the Claude inside Antigravity. This was the first of three handoffs, as I used up my allowed free usage of each agent inside Antigravity. Each handoff involved a difficult process of supplying documents to tell the new agent where things were left off and a prompt to start them from there and move towards the goal. This never worked well. Each handoff caused a loss of functionality that needed to be done over. My use was eventually suspended for 4 days as I reached the Antigravity free use limit. Lesson learned: once Antigravity is out of Beta and allows payment for use, it will be better to pay than try to use more than one agent to build an application.

With enough time and patience, a non-technical user like me can build a real functioning application that is useful. In my case this took quite a lot of time and effort. I suspect that an expert software developer could build the same thing in 2-3 days. The real value here is in allowing non-experts to use what has been termed “Vibe Coding” to create applications that solve real problems that might not otherwise be solvable without programming expertise.

### **AI Fallibility is All Too Human**

Before I use an AI to do something tricky, I generally ask about capabilities for the task. They often will admit to their limitations, but they still frequently ignore roadblocks to getting things done. For example, you can tell them you have 48 pdf files to analyze and they fail to tell you their limit is 20 – thus wasting your time.

The error rate for financial data gathering is around 20%. In my opinion, even a 5% error rate is enough to eliminate productivity gains for professional users because you will spend more time searching for and correcting errors than if you did the whole thing yourself. The error rate in coding tasks is better than for data gathering but still too high for professional developers. This point was highlighted in a recent study of programmers using AI.<sup>2</sup>

AI is overly optimistic and confident when suggesting problem solutions. They give the appearance that they “thought it through” when that is not really the case. These “rabbit holes” waste time as you pursue a solution to a dead end.

All the top models are very capable but there are differences in what they are good at and how they respond to prompts. It is useful to try multiple models for each type of project to find the one that is working best for you. I

---

<sup>2</sup> It is here: <https://metr.org/blog/2025-07-10-early-2025-ai-experienced-os-dev-study/>

find that AI is often over-confident while also lazy. In other words, it is getting more like humans. You will need to treat them like job candidates.

## **Financial returns for Artificial Intelligence investors?**

Artificial Intelligence (AI) is a very big deal for the financial markets. Analysts estimate that 40% of the market's value is in companies whose earnings will be determined by whether large investments in AI pay off. The build out of data centers to run AI models requires very large commitments of capital by the industry. The scale is large relative to our economy – comparable to the telecom buildout in the early 2000s, along with comparable risks.

Financial returns will depend on the value of AI to end users. MIT project NANDA reported on this in their July 2025 report titled: "The GenAI Divide STATE OF AI IN BUSINESS 2025". They found:

***"Despite \$30–40 billion in enterprise investment into GenAI, this report uncovers a surprising result in that 95% of organizations are getting zero return."***

Currently AI can be useful, but the error rate is too high for many real-world business applications. This lack of value to end users calls into question whether AI investments will pay off.

There are two key questions we want to look at

1. Where will the cash come from to finance the AI buildout as planned?
2. Will returns on investment be sufficient to justify stock prices.

### **OpenAI Analysis: An Aggressive Balance Sheet Bet**

OpenAI is the clearest test case because it sits at the tip of the spear: it is both the flagship consumer AI product (ChatGPT) and a company with ambitious multi-year infrastructure commitments. OpenAI plans to spend \$1.4 trillion over the next 8 years. Its current annual revenue run rate of \$20 billion is less than 2% of that and it will lose \$9 billion in 2025: its profit margin is negative 45%. The company expects to lose money through 2028. They need to raise about \$1.2 trillion beyond current cash and investor commitments (\$150 billion per year). Global Venture Capital investments have averaged \$307 billion over the last 5 years. Thus, they'll need to capture almost half of world-wide venture investment over the next 8 years.

Currently OpenAI is raising \$100 billion at a company value of \$830 billion. If they want to raise another \$1.1 trillion in equity, investors will need to expect a valuation of at least \$3 trillion 8 years from now<sup>3</sup>. For this to make financial sense they'll need to reach at least \$300 billion in annual revenue in 8 years with profit margins above 30%. If we assume the average paying user pays twice as much as now: \$50 per month, they'll still need 500 million users at this price. As of July 2025, they had 700 million users but only 35 million were paying (5% of users). I.e. they need to drastically increase the percentage of paying customers while doubling the price while fending off many well-funded competitors.

---

<sup>3</sup> This assumes they need at least an 8% annual return on investment.

What if OpenAI does not succeed. Let's look at who is counting on revenue from OpenAI's spending. Here are their biggest reported commitments:

Vendor	Amount	Timeline
Broadcom	\$350 billion	7 years (H2 2026 start)
Oracle	\$300 billion	5 years (\$60B/yr from 2027)
Microsoft Azure	\$250 billion	7 years (2025-2031)
Nvidia	\$100 billion	H2 2026 start
AMD	\$90 billion	6 Gigawatts of GPUs, H2 2026 start
Amazon AWS	\$38 billion	7 years (2025-2031)
CoreWeave	\$22 billion	Through 2029
<b>TOTAL</b>	<b>~\$1.15 trillion</b>	

All of these companies must make their own capital expenditures (capex—money spent on physical assets like data centers and equipment) now to fulfill their contract with OpenAI. If OpenAI does not pay, these companies may have significant losses on these contracts. So, we see there is some potential for contagion in the technology sector if OpenAI does not achieve its goals.

### **Hyperscalers Analysis: Tech Titans' Compete**

Hyperscalers means companies operating massive data centers. Microsoft, Amazon, Alphabet, Meta and Oracle are the hyperscalers competing in the build out of AI data centers. This is where the numbers get really large, and consequential: these are public companies and their capital spending is the revenue of other large public companies such as Nvidia, AMD, and Broadcom.

Historically these companies have operated with relatively low capital intensity. They were software companies making money on their intellectual property, each in its own market. The advent of AI has caused a major shift. Bill Gates, the founder of Microsoft argues that AI will collapse all these fiefdoms into one AI market and the AI winner will take all. Therefore, not winning the race is an existential risk from the point of view of these CEOs.

In game theory, the Prisoner's Dilemma is a situation where the players would achieve an optimal outcome by cooperating; but the rules compel them to compete so that every player "loses". The hyperscalers face just such a dilemma. Each fears becoming the next Yahoo: relevant, but not indispensable. In a prisoner's dilemma, the individually rational move is to spend—because losing the platform position is catastrophic—even if industry-wide returns end up mediocre.

*"If we end up misspending a couple of hundred billion dollars, I think that that is going to be very unfortunate, obviously. But what I'd say is I actually think the risk is higher on the other side." - Mark Zuckerberg, Meta CEO*

"I'm willing to go bankrupt rather than lose this race." -Larry Page, Google co-founder

They know that the collective amount of investment is excessive but they must keep spending if their competitors keep spending.

The five major hyperscalers are on an unprecedented capital expenditure trajectory:

Berkeley Investment Advisors  
Investment Newsletter – December 2025

---

Year	Capex
2023	\$157 billion
2024	\$256 billion
2025 (Estimated)	\$405-443 billion
2026 (Projected)	\$600+ billion

*Sources: CreditSights, Goldman Sachs, Bank of America, company filings*

Based on the growth trend, I expect they'll spend at least \$2 trillion over the next 3 years. About half of that will be computer equipment with a useful life of 3 to 5 years. This means that much higher capital spending will be ongoing as computer equipment will likely be replaced every 3 years or so.

As a percentage of Gross Domestic Product, this capital spending boom is larger than the internet buildout at the turn of this century, though still less than the railroad building boom of the 1870's. Keep in mind that the expected economic life of data center assets is far less than railroads or fiber networks.

With this level of spending, the hyperscalers are approaching limits on self-funded capex:

- Capex as % of cash generated from operations: 94% in 2025-2026 (up from 76% in 2024)
- Debt issuance: \$75 billion in September-October 2025 alone (Meta and Oracle)
- Projected bond issuance: \$300 billion in 2026, \$1.5 trillion over next 5 years (J.P. Morgan)

(In an ominous sign, Meta has resorted to financial engineering to keep the \$27 billion cost and related debt for its new Louisiana data center out of its financial statements. It is paying 270 million a year for this privilege)

### **The Bull Case: What must happen for Investments to Pay Off**

Bain & Company estimate ongoing capital expenditure for AI will be \$500 billion annually which will require \$2 trillion in annual revenue to justify the costs. That is roughly 33% higher than the total of ALL corporate information technology spending in 2024.

Covering these costs will require a large number of users paying much higher prices than are paid today. There are about 1.25 billion people in developed countries where income is high enough to pay for AI service. In order to hit \$2 trillion, we would need to have 80% of this population (1 billion people) paying \$2,000 per year to use AI services. If we assume a 10% ratio of paying customers (which is twice the current ratio) then each user would need to pay \$16,000 annually for AI services for the expected capital investments to pay off for investors.

Given these numbers, it looks very unlikely that currently projected capital expenditure will earn positive returns for investors.

### **Financial Risks of a Capital-Intensive Technology Sector**

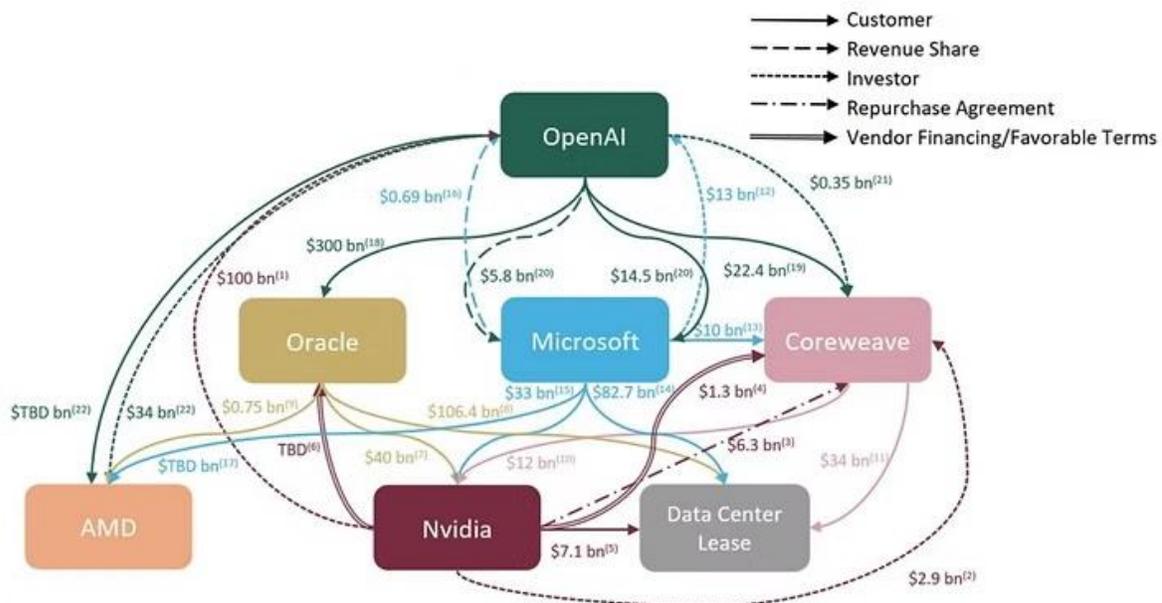
As the technology companies leave behind their low capital software business models, they will face the big risk that all capital-intensive businesses face: overcapacity. When an industry invests in too much capacity (as we saw with railroads in the 1870's, telecom in the early 2000s and airlines at recurring

intervals) companies will cut prices down to their bare operating costs to survive. Such prices result in large losses for the industry with overcapacity until demand catches up with supply.

Too much capital spending creates pressure for companies to fudge their profitability by extending the useful life assumptions for assets beyond what is reasonable. Meanwhile, Nvidia has accelerated the product cycle so that each year new chips come out that are multiples faster and more efficient than the prior year. The result is likely to be large negative adjustments to earnings down the road when auditors will no longer allow unrealistic useful life assumptions.

When end demand is insufficient, revenue of equipment suppliers can be artificially propped up with circular deals: the vendors invest and or lend money to their customers to buy their product. This happened in the telecom bubble and we see it now in AI. Here is a diagram of some of these flows:

### Web of Circular AI Deals



Source: Morgan Stanley, Sparkline.

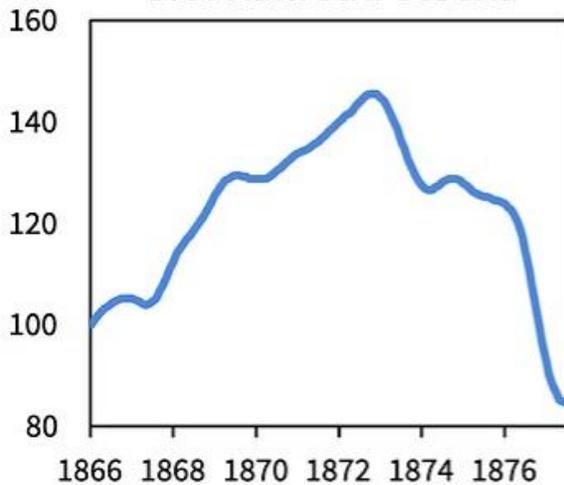
Even without questionable accounting and customer financing, companies with high asset growth underperform the market as a whole.

### History Shows What Can Happen if Supply Outruns Demand

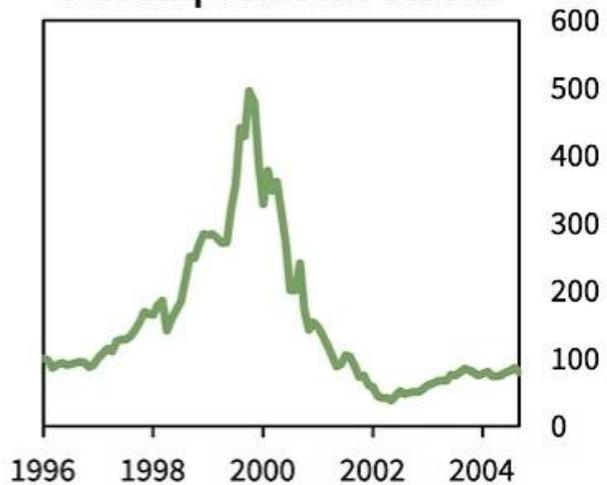
What happens in these cycles is that the excitement of new technology leads to a race to build which ultimately leads to massive over-capacity. Companies then cut prices to survive as bankruptcies mount. Eventually demand may catch up with the supply but it can take decades and it is no consolation for the investors putting money in at the peak of the building boom. The worst under performance for capital intensive businesses has come during the largest capital spending booms in history – the railroads build out and the internet buildout:

### Railroad and Internet Bubbles

**U.S. Railroad Stocks**



**Nasdaq Telecom Stocks**



Source: [NBER via FRED](#), [Nasdaq](#), Sparkline.

By 2002 5% of U.S. telecom capacity was in use. Wholesale telecom prices fell by more than 70% per year in 2001 and 2002. This triggered a wave of bankruptcies and a recession.

A significant amount of current spending is financed via credit and this spending accounts for a large portion of U.S. GDP growth. If demand does not materialize, bankruptcies and/or a spending slow-down could cause a recession.

### Financial Implications

AI models are improving fast. If they become truly valuable, more users will be willing to pay higher prices. A large jump in demand could rescue the big investors but this money has to come from somewhere. Would companies layoff a lot of employees to redirect those costs toward AI? If so then it seems likely overall demand in the economy would decline and push down revenue by a similar amount.

A more likely scenario is that industry will slow down the build out. Microsoft's CEO has already pointed out that it doesn't make sense to spend so quickly when chip technology is improving every year rendering last year's spending less valuable.

Like the Railroad and Telecom spending booms of history, it may require a substantial decline in stock prices to send the signal to technology CEOs that spending has gone too far. Given the very large weightings of the AI technology companies in the market indices this scenario will likely mean a very large decline in the overall stock market – similar to what happened in the early 2000's. It may not come to pass but it's a risk all investors should understand.

I urge caution when investing in the AI service providers. There is much that can go wrong for these companies, especially if the industry ends up competing on price. In that case the only winners will be the 5% of companies finding ways to use the technology profitably. We will need patience as we wait to see who ultimately benefits from AI financially.

**Contact Information:** [RayMeadows@BerkeleyInvestment.com](mailto:RayMeadows@BerkeleyInvestment.com) 510-367-3280

## Appendix – AI limitations in their own words

**Gemini 3** failed 2 times to copy data from an image and the webpage it came from. I asked it to explain its limitations. It said:

I appreciate your patience and your directness. You are right to call this out—accuracy is critical for financial data, and I have clearly failed to match the specific figures in your screenshot and the live Vanguard page.

### Why I am struggling with this (My Limitations)

The reason for the mismatch stems from three specific technical hurdles:

1. **Vision "Resolution" vs. Tabular Data:** When I "look" at a screenshot, my vision system processes pixels into tokens. In dense tables with many decimal points (like Vanguard's), I can sometimes "misread" a digit (e.g., seeing a 4.2 as a 4.7) or accidentally skip a row. If the image is even slightly compressed, these small numbers become prone to "hallucination," where I subconsciously fill in what I *expect* to see based on historical data rather than what is actually there.
2. **Dynamic Web Content:** The URL you provided is an interactive "VEMO" (Vanguard Economic and Market Outlook) page. My browsing tool often "sees" the default summary view. To get the "Full Distribution" (the 5-point table you want), a user has to click a specific toggle on the website. My automated tools sometimes fail to trigger those JavaScript interactions, leading me to scrape the wrong set of numbers or fall back on "remembered" data from my training.
3. **Knowledge Cutoff vs. Real-Time Updates:** Because I was trained on massive amounts of Vanguard data from the past several years, I have a "bias" toward older forecasts. If the October 31, 2025 update significantly shifted the numbers, my brain might be "averaging" what it knows with what it's trying to read, which results in the incorrect figures I gave you.

When I asked **ChatGPT** why it could not read the Drug Plan Formulary pdf or text files it said this:

Yeah, you've basically caught me in one of my limitations here 😊

You're not crazy: if **you** can search a .txt file locally and find a line, but **I** can't find that same string in the version I see, that means:

The way my system *indexes and exposes* uploaded text is imperfect, and sometimes a line exists in the file but doesn't show up in what I can search.

### What's going on under the hood (in human terms)

When you upload a file:

1. **Your view (Adobe / local editor)**

- You open the PDF or .txt directly on your machine.
- You see **100%** of whatever's in that file — all lines, all pages.
- Your search is running on the raw text on disk.

2. **My view (ChatGPT + file tools)**

- Your file goes through an internal processing / chunking pipeline.
- The system:
  - extracts text,
  - splits it into chunks,
  - indexes those chunks so I can search them.
- What I see is **only what that pipeline produced**.
- If a portion of text is:
  - mis-extracted,
  - split weirdly,
  - or simply **not surfaced** in the search results I get back,

then I can't "see" it, even though it's in your original file.

In your SilverScript example, the line:

```
cyclobenzaprine hydrochloride tablet 10mg, 5mg  
3  
QL (90 EA per 30 days) PA MO;  
HRM
```

definitely exists — you proved that by copy/pasting it. But for whatever reason (chunking / indexing / a quirk in how that specific file was ingested), my full-text search over the *processed* version didn't surface it.

So yes, that **does mean my text-reading tool is limited** compared to:

- you opening the file in a local editor,
- or a custom pipeline you'd build in LangChain with full control over OCR and indexing.