Memo to: Clients

From: Howard Marks

Re: A Layman Thinks About AI

---

**Should we invest in artificial intelligence, and can we do so intelligently?** That's what I'm going to address in this memo.

On one hand, I'm not a tech person, and I don't claim to have above-average knowledge of artificial intelligence (or technology investing, for that matter). On the other, AI is the biggest development of our time and could be one of the most impactful things the world has ever seen, so investors don't have the option to abstain from thinking about it.

With that disclaimer, I'm going to share my deliberations, put forth the questions I'm wrestling with (not all of which are directly relevant to investing in AI), and provide my thinking about them. While my limited knowledge prevents me from definitively answering these questions, I hope to provide a service by assembling them in one place. And while news on AI comes hot and heavy these days, with security prices for AI and AI-adjacent companies gyrating up and down, I'm not going to try to keep up, since what I have to say relates to the long-term outlook for AI, not what security prices will be tomorrow.

* * *

Investing consists of just one thing: positioning capital to benefit from future events, and most investors cope by relying on forecasts. But not everything is equally predictable, and not all forecasters are equally skilled. Thus, to invest in AI based on expectations about the future – and I think that's all we have – we must consider both what we think the future may hold and the chance we can get it right.

I've written in the past about the fact that people who engage in top-down investing have to be right about the outlook for the economy, the industry, the company, and the market. If they have a 60% probability of getting each one right – which is probably generous – what's the probability of getting all four right? The answer is 13% ($0.6^4$), assuming the outcomes are independent of each other.

**In the case of AI, there are an unusually large number of important variables, each of which complicates the decision of how to invest.** An investor in AI has to be at least somewhat right about the evolution of the technology, the demand for compute, the competitive structure, the beneficiaries of the technology, the availability of financing, and the timing of all the above. Assessing these key elements seems to me to be unusually challenging. In that case, what's the combined probability that any thesis is correct? That's just something to think about as you read.

**Lastly, bear in mind that (a) we can never know everything about an investment, (b) the biggest mistakes are made when people think they do, and (c) the incompleteness of our knowledge regarding the future can't be permitted to completely prevent us from taking action**. It's the investor's job to bear risk in the pursuit of profit. How much and in what form are decisions we each must make. But given its vast potential, we can't decline to participate in AI just because the future is unclear.

© 2026 Oaktree Capital Management, L.P.

**Question one: How does AI do what it does?**

Before getting more granular, I'm going to start with a big-picture, non-financial question that I find particularly interesting. Since I have absolutely no knowledge on this subject, I went to Anthropic's Claude for some answers. Here's some of what it said:

> We know what AI is made of and *how it was built* almost perfectly, but we know surprisingly little about *how it actually does what it does*. . . .
>
> We also know the training recipe: show it enormous amounts of text, have it guess the next word, and nudge billions of internal settings slightly each time it guesses wrong. Repeat that trillions of times. . . .
>
> Nobody wrote the rules for how it answers the question. Those rules emerged from the nudging, spread across billions of numbers that mean nothing individually. The best phrase I know for this is that these systems are grown rather than designed. Researchers in "interpretability" are making real progress. They have found, for example, that a model writing a rhyming poem picks the rhyme word before it writes the line, which means it plans ahead. But they can explain only small slices of what goes on. It's like having a complete record of every trade in the market and still not knowing why prices moved. . . .
>
> When I explain how I reached an answer, research suggests my explanation doesn't always match what actually happened inside. A model asked how it did a sum may describe the schoolbook method while internally doing something very different. So the honest summary is that we've built a mind-like thing we can inspect completely and still can't fully explain. And when it explains itself, it may be no more reliable than we are when we explain ourselves [interestingly, here Claude talks like it's a person]. . . .

As you think about answering the questions that follow below, consider these: Does anyone truly understand how AI works? How will AI evolve and what will it be able to do and when? How will users put it to work? What impact will it have on users' profitability? Investors may not have to know all these things, but it's my understanding that few people have significant insight into these matters. If that's the case, can anyone reliably answer the questions that follow?

**Question two: What will AI be able to do?**

What AI is great at is evolving. Just a few years ago, what AI could do was generally summed up as knowing everything that has occurred before, identifying historical patterns, extrapolating them, and using this ability to perform tasks. This has already made huge changes in some areas, like coding, and promises to alter many aspects of life as we presently know it. But now, AI has gone well beyond that, to something that probably can fairly be described as "reasoning."

Just this one area poses several questions that are both stimulating and daunting. AI's full potential will be subject to the answers to some of these:

- Can AI's process be made to reflect our preferences, values, and goals, or will it apply its own?
- Can AI be depended on to be honest, humble, ethical, and loyal? Will it be accurate and objective or subject to influence and manipulation? Will it admit when it's wrong or doesn't know an

answer? Can it give us skepticism rather than sycophancy? Will it become a threat to humans and seek to advance its own interests instead of ours?

- How will it do in areas such as creativity, originality, aesthetics, subjectivity, and intuition? Will it be able to perform what I call "second-level thinking" – thinking in a way that's different from that of most people and at a higher level – leading to superior solutions?
- Will it be able to think about the future and "blue sky" to the same degree as humans?
- Can it be instructed to furnish primary sources (as opposed to published citations) and provide counterarguments?
- Will one model be able to do it all, or will different models be best at different things? If so, how will we know which one to use?
- Will artificial general intelligence – the ability to do everything a human can do – be the standard? Will AGI mean the ability to do the same things as (a) the average person, (b) a person in the 90th percentile, or (c) the smartest person in the world? In other words, will there be varying levels of knowledge and things AI doesn't know?

All these things bear on what work will be entrusted to AI.

**Question three: Who will profit?**

As Warren Buffett pointed out about the internet during the tech bubble, transformative technologies don't necessarily reward investors. We know AI will increase efficiency, but will it increase profitability? And whose? If several AI models can become equally smart on a given subject, will their service become an undifferentiated commodity that sells on the basis of price? If it does, will profitability and ROI be ground down, in which case the economic benefits may go more to the users of AI than to the vendors? What does all this imply for the return on the trillions of dollars being invested in AI infrastructure?

Or, on the other hand, will some AI companies acquire capabilities that make them monopolies whose products are essential for companies' success and even survival? If so, think about the prices they'll be able to command. (The big chip producers – Nvidia and TSMC – are racking up the big profits these days, demonstrating the value of a strong competitive position.)

Lastly on this subject, I haven't heard much about whether and how AI will be taxed. Given that one of AI's most likely effects will be replacing workers, and given governments' need to replace the tax revenues that previously came from taxing the people AI puts out of work, **what's to say AI won't be taxed like people?** Many of today's top earners are taxed at combined rates of 50% or more. What tax bracket will AI be in?

**Question four: What will be the competitive landscape?**

Can we know which companies will succeed? With what degree of certainty? How strong are the moats surrounding today's leaders?

Perhaps the most startling characteristic of AI is the speed at which it makes progress. **In that connection, I'm struck by the fact that today's leading model labs, OpenAI and Anthropic, were unknown to most of us just a few years ago. If companies can forge ahead that fast in the world of AI, can they likewise be supplanted by others, or at least challenged?** If it's possible, consider the vast resources and potential of the so-called "hyperscalers": Amazon, Microsoft, Alphabet, Meta, and

Oracle. Who will win among these incumbents? Or will it be companies we've never heard of? How about companies that don't exist yet?

Think about the probability 20-25 years ago that the leaders then would still be the leaders now. Microsoft, Amazon, Apple, and Google came through that era as winners, but Napster, Yahoo, and Netscape fell by the wayside. It takes a lot to maintain a lead in a rapidly changing technology. Add in AI's novelty and speed of progress and ask, "what's the probability today's AI leaders will still be the leaders in five or ten years?" **If the competitive picture can shift as I suggest, how much is a stake in today's leaders worth?**

**Question five: How profitable will the companies be?**

The value investor estimates a company's earnings or cash flows in the future, discounts them back to the present, and compares the result with today's stock price. To predict earnings in the future, we need answers to questions like these:

- How many units of its product will it sell?
- What will be the selling price per unit?
- What will each unit cost to produce?
- How much of the profit will need to be reinvested in developing the next product?
- What will be the company's resulting return on investment?

While answering these questions for any company requires dealing with uncertainty, doing so in connection with AI introduces particularly novel considerations like these:

- What will its product be capable of doing, and how will customers put it to work?
- Will the product be priced as a commodity or a unique, essential good?
- How long will chips last, and how fast will data-center depreciation have to be charged off?
- For how long will today's breakneck spending race continue?
- Will enough money be made to produce a decent return on today's capital expenditures?

I could certainly be wrong, but I don't remember puzzling over uncertainties of this magnitude in my days analyzing the conglomerates in the 1970s, or in connection with the internet and e-commerce in 1999-2000. Those businesses, as novel and complex as they were, seem simple and straightforward compared to AI.

So, I wonder: can we project 2030 earnings for today's hyperscalers and leading frontier labs (let alone companies that mostly consist of conjecture)? Do we have to be able to predict earnings reasonably accurately to invest? If we come up with estimates, how much confidence should we have in them? And if we can't, how can we think about investing in AI companies?

Here's the test I propose nowadays: Tell me what you think the e.p.s. for a given AI company will be in five years. **Will you bet that the eventual answer will be within +/- 50% of your figure?** (E.g., if you guess $10 per share, how confident are you that the actual figure will be between $5 and $15?)

**Question six: What's the outlook for data centers?**

Historically, data centers were digital warehouses that stored data and hosted computer applications and websites. The term data center (or "AI factory") has evolved to refer to the workplaces of AI. Analogous to auto plants, they're where AI is produced, ranging from storing information and training models to providing time on the system (called "compute"). Data centers are a critical element in the development of the AI industry. The construction of data centers, and thus the performing of compute, is subject to constraints – physical and otherwise.

While the future of AI is exciting, some people might prefer to take a cautious approach to participating in it. This suggests investment via tangible infrastructure rather than developers and vendors of the emerging technology, a strategy akin to selling picks and shovels to gold miners rather than trying to figure out which miners will find gold. In theory, this makes determining the ultimate winners less critical since whoever ends up winning would still need to use the infrastructure. Compute is in short supply today, adding to the attractiveness of data centers for investment. However, a number of factors must be considered:

- When we think about "infrastructure," the phrase "bricks and mortar" often comes to mind, and it's that solidity that can make data centers appear the safer AI option. But a relatively small part of the cost of a data center goes for the building itself – the "shell." Chips or "GPUs" make up most of the cost, and that (and the risk of obsolescence) is generally borne by the center user.
- Large building sites will have to be located and purchased, and data-center developers will have to overcome regulatory hurdles and resistance from citizens opposed to data centers in their backyards.
- Connectivity to users of inference via available fiber will be a necessity.
- Data-center builders will have to deal with potential bottlenecks in terms of the needed electric power, cooling equipment, and construction services.
- Limitations on the availability of financing for data centers could constrain the pace at which they're constructed, and the cost of financing for future projects could surprise.

Interestingly, the potential bottlenecks mentioned above could complicate life for would-be builders of data centers and limit their growth. On the flip side, any constriction of the growth in the supply of "compute" – as implied by the success last week of a number of moratoriums on data center construction – should make the ones that presently exist or that can be built in the future more profitable and thus more valuable.

The key here is that data centers aren't a "layup." The intelligent investor has to find opportunities that are insulated from the uncertainties listed above and won't be subject to overbuilding and thus oversupply of compute. As in all forms of investing, there'll be successful data-center investors and less-successful ones. Given the dynamics of AI, however, I think the successes of the winners and failures of the losers may be even more pronounced than is usual in "infrastructure."

**Question seven: Should we be worried about our unusual reliance on this sector?**

Today, AI is responsible for a disproportionate share of what's going on in the investment environment, introducing considerable uncertainty for all investors (and all participants in the economy), not just for investors in AI:

- AI represents a small part of the U.S. economy but is responsible for an outsized part of economic growth.

- AI's contribution to GDP growth currently comes overwhelmingly through capital expenditures, primarily for construction of data centers. These are rising rapidly while the rest of construction is in decline.
- For 14 consecutive quarters through June 30 of this year, earnings grew faster for the "Magnificent 7" – mostly the hyperscalers – than for the rest of the S&P 500.
- Some AI companies have been generating large gains from equity investments in fellow AI-related companies, making them top contributors to S&P 500 earnings. FactSet reported that in the second quarter, Alphabet and Amazon notched a total of $151 billion of gains on equity investments, mostly on Anthropic and SpaceX (which now includes xAI).
- AI's impact on aggregate earnings doesn't stop with the sector. Sellers of associated components are reporting strong earnings as well, and there are knock-on effects as AI workers and AI-adjacent workers spend their pay on rent, food, cars, and other things far removed from AI itself.
- AI has disproportionately contributed to stock market gains, with a group of 42 AI-related companies generating around two-thirds of the S&P 500's price gains since January 2024, according to J.P. Morgan.
- The sum of these trends has left the S&P 500 highly concentrated in a few AI and tech companies.

The key observation is that the factors our economy and markets depend on today are relatively undiversified. Exposure to AI constitutes an unusually powerful "fault line" that runs through portfolios and has the potential to dominate all aspects of the investment environment. Negative developments with regard to chip life, corporate strategy, or demand for compute could have far-reaching consequences. **It's incontrovertible that disproportionate reliance on a single sector makes the economy and stock market more vulnerable than they would be if the factors underlying them were broad-based.**

When Michael Milken pioneered the issuance of sub-investment grade debt in 1977-78, he ushered in an investment-industry-wide willingness to accept risk as long as it's viewed as being accompanied by enough potential return to compensate for the risk. His thesis rested on two cornerstones: (a) incremental yield that potentially compensates for the losses on the bonds that default and (b) **enough diversification so that a few defaults won't spoil the overall portfolio outcome.** But today, our stock and debt markets are heavily reliant on a single factor: the success of AI.

**Question eight: What are the implications for the credit market?**

AI-related debt issuance has totaled nearly $600 billion globally so far this year, according to Goldman Sachs, comprising debt issued by the hyperscalers, in data-center/project financings, and by other AI-linked borrowers. Around $400 billion of this went to the U.S. investment grade debt market, representing over one-fifth of total IG supply and taking the share of AI-related bonds in the Bloomberg USD IG Corporate index to around 12%. For the U.S. high yield bond market – into which the hyperscalers don't directly issue – the weighting is more like 6%. Clearly, AI-related debt issuance is one of this year's major capital-market developments.

Going forward, AI seems likely to require an average of $1 trillion of new capital per year in the period ahead, and hundreds of billions of that are likely to continue to be raised each year through the issuance of debt. What will be the impact on interest rates of AI's demand for debt capital?

Since expanding capacity and developing new products are imperatives for the hyperscalers and frontier labs – and since they don't want to see competition crop up – these companies likely can't afford to slow

their investment in AI just because interest rates may be higher than they expected. That means they have to be "takers" of financing, even if it is in short supply and increases in cost.

In the past, economists used to talk about "crowding out," in which aggressive government borrowing makes it hard for the private sector to find the debt capital it needs. At present, it feels to me that the demand for debt from rate-insensitive hyperscalers willing to pay whatever it takes has the potential to elevate the interest rates the government must pay. Hasn't that contributed to the rate rise we've seen to date?

In the same vein, will the equity markets be capable of providing the hundreds of billions of new equity investment required each year, supplementing the companies' cash flow? The amount raised in traditional U.S. IPOs was below $50 billion in eight of the last ten years (the exceptions being almost $80 billion in 2020 and $140 billion in the boom year of 2021). I suspect there's just so much new capital available in a given year; won't AI's appetites for both debt and equity ultimately compete against each other for it?

In the past, the highly profitable non-AI businesses of the hyperscalers produced massive free cash flow that rendered them debt-free, cash-rich, and highly rated. Things are different now, and these companies are spending their cash hoards and current cash flow, and they're borrowing more. Their strong investment-grade ratings enable them to easily finance their expansion in AI. But as this process wears on, might their creditworthiness and/or ratings change? What are the possible implications for those who lend against the high ratings today and watch as the companies increase their leverage?

Some of the data-center financing taking place is facilitated by the investment-grade ratings of the hyperscaler tenants. Other debt is collateralized by chips. But investors have to be very conscious of "out clauses" in some data-center leases that can enable tenants or guarantors to terminate their obligation under certain conditions – such as the failure of the center to operate at 99.9% reliability – and of the potential risk that collateral in the form of chips won't last as long as the debt.

What are the implications of the industry's so-called "circular transactions," in which companies finance purchases of their products? We see frequent instances of lending and related buying between chipmakers, hyperscalers, frontier labs, and data-center builders. Sometimes money goes from A to B and back to A. Sometimes it goes from A to B, and then to financing for C, which spends it with A. Are these transactions the parties would enter into on an arm's length basis, or are vendors subsidizing buyers to support demand and protect market share for their products? What are the risks associated with this interdependence? And what is the impact on the hyperscalers' creditworthiness of borrowing on the part of special-purpose vehicles that keep debt off balance sheets – is the parent behind the debt or isn't it?

Finally, what are the effects of borrowing in a time of exuberance surrounding an exciting technological advance? In the end, we have the confluence of several factors:

- a massive desire to borrow on the part of AI companies,
- investor excitement surrounding AI, and thus a willingness to lend,
- heavy concentration on AI in the capital markets, as described above, and
- so-called "circular transactions" between participants in the AI industry.

Lenders to AI face a new era, marked by debt concentrated in a single industry, rapid change in industry fundamentals, and a high degree of reliance on the level of future revenues from sales of compute. **Lenders will have to be exceptionally vigilant regarding the creditworthiness of the borrowers they back and the structure/terms through which they do it.**

**Question nine: How predictable is AI? How answerable are these questions?**

In memos and in my book *Mastering the Market Cycle: Getting the Odds on Your Side*, I've written that for someone attempting to deal with the future, a prediction isn't enough. In addition to your prediction, you need a view on the probability your prediction is correct. It's my view that with regard to AI – given its technical complexity, unprecedented nature, and high rate of change – that probability is unusually low.

I wanted to quantify AI's predictability, but I worried that doing so was well beyond my competence. Thus, I asked Claude to help by ranking AI on a scale from 1 (unknowable) to 10 (fully predictable). Its answers were highly informative. Claude summed up on the subject, reflecting some of the uncertainties discussed above (emphasis added):

> **My overall judgment: AI as a whole scores about a 3.** That number hides a split. The *direction* is fairly predictable, maybe a six or seven: the technology will keep improving and spread through the economy. The things investors get paid for are much less predictable, around a 2: how fast it happens, how big the profits are, and who keeps them.

I also asked Claude to rate subsectors of the AI industry on predictability and identify parts of the non-AI economy it considers comparable in that regard. Here are what it describes as its "judgment calls":

| AI Sector | Predictability Score | Comparable Industry | |
|---|---|---|---|
| Regulated power utilities serving AI load | 6 | Regulated utilities | 6-7 |
| Data centers leased long-term to big tech companies | 5-6 | Large industrials | 5-6 |
| Hyperscalers | 5 | N/A | |
| Chips | 4 | Airlines, shipping | 3-4 |
| Speculative or "neocloud" data centers | 2-3 | Telecom/fiber '98-'01 | 2-3 |
| Frontier labs (OpenAI, Anthropic, xAI, etc.) | 2 | Early-stage VC | 1-2 |

For me, the surprise was the frontier labs' score of 2, which I take to mean "practically unknowable." These seem to me to be impressive companies with big leads, but Claude notes as follows: "These businesses are close to binary. Capital needs are huge, prices per unit of capability keep falling fast, and it's unclear whether a lasting moat exists." (Ironic, isn't it, given that Claude belongs to Anthropic?) Intriguingly, these most uncertain propositions according to Claude seem to be priced at a premium in the private market. It'll be interesting to see how the shares of OpenAI and Anthropic, both of which have filed for IPOs, are priced publicly and how they trade afterwards. And note that much of the equity gains that goosed S&P 500 earnings in the second quarter – described on page 6 – were related to this sector, which is described as "least predictable"; how lasting will they prove to be?

Finally on the subject of predictability, I'll take this opportunity for a word about geopolitics, which with regard to AI generally means "China." Nearly all the leading-edge chips come from Taiwan's TSMC. The U.S. restricts sales to China of the most advanced chips. This encourages China to offer low-cost models that can undercut sales of more powerful U.S. models. What happens to the U.S.'s competitive position if this changes? Will chips from China catch up with or pass today's leaders? Will China ever be in position to discourage or even prevent sales of Taiwanese chips to the U.S.? And, will worries about China's intentions blunt efforts to impose safety guardrails on the AI industry?

Clearly, there's a lot investors don't know and can't predict about AI. **The question is how much risk will you take and in what form? There are many parts to the AI industry, and (according to Claude) they range from the frontier labs, with regard to which we're largely in the dark, to securely leased data centers and their power providers, with predictability roughly comparable to that of large industrials and regulated utilities.**

**Question ten: Is AI a bubble?**

Worry over some of the questions discussed above has begun to surface in recent weeks. A confluence of negatives in the investment environment arose, centering around inflation, the federal deficits and debt, the resulting increase in interest rates, and the seeming intractability of the war in Iran. To these general issues were added some considerations specific to AI, such as:

- competition from new models; cheaper models with fewer bells and whistles; low-priced Chinese models; and open-source models that are free to download,
- declining prices per "token" (the unit in which model usage is priced),
- the possibility of slowing growth in the use of compute,
- reported sales growth that falls short of expectations, and
- concern about some companies' high leverage levels and heavy dependence on future financing.

Questions that are easily dismissed in good times can worry investors when their equanimity is shaken. And then we find, as I say, that whereas in real life things fluctuate between pretty good and not so hot, in the minds of investors they go from flawless to hopeless. The things AI investors are worrying about these days are nothing new, but a changed investment climate has rendered them sufficient to darken the mood.

**Even though the subject has been debated for well over a year – including by me – people keep asking: is it a bubble?** First, we need a definition of a bubble, and I want to supply mine: irrational exuberance leading to imprudent behavior, excessive investment, and too-high prices for participation. I think an AI bubble, like any possible bubble surrounding a technological advance, has to be considered in at least three ways. **Is the technology itself overhyped and likely to disappoint? Are the amounts companies are spending excessive and fated to produce inadequate returns? And are investors taking too much risk and paying too much for stakes in this field?** Reaching conclusions on these subjects relies heavily on the answers to the above questions. I don't feel much has changed in these regards since I last wrote a memo on AI in February:

- I still think AI itself will prove powerful and consequential and change corporations, the economy, and life speaking broadly in ways we can't even imagine. So no, the technology isn't an illusion.
- A great deal of discussion is taking place regarding the revenues that will have to be generated in the future to justify the capital expenditures. But until we know something about how AI will be used and who will profit, we can't assess the wisdom of the companies' spending.
- And until knowledge of these things enables us to reasonably estimate the companies' earnings and the profitability of the infrastructure, we can't judge the appropriateness of the prices being paid for stocks or the yields being offered on debt.

While there are a lot of questions swirling around AI – some of which I've described above – the key boils down to this rather complex one: **will AI provide sufficient value to users to warrant sales of**

**compute and prices for compute high enough to result in adequate returns on the trillions of dollars being invested in it (and for which companies)?**

The conclusion I draw from my conversations with people I consider knowledgeable is that, at this point, answering this question is largely a matter of faith. That reminds me of a quote attributed to American economist and writer Stuart Chase:

> For those who believe, no proof is necessary. For those who don't believe, no proof is possible.

Try as I may, I can't come up with a more appropriate resolution to the questions surrounding AI than this.

**Question eleven: How should one think about investing in AI?**

I consider artificial intelligence the most consequential and least predictable development for investors in at least a generation. **But investors needn't dismiss something just because it's uncertain.** The right response isn't paralysis but (a) calibrated humility regarding strategy selection, position sizing, and risk-taking and (b) a willingness to act in the presence of uncertainty.

Different aspects of investing in AI require answering different questions. An equity investor in a frontier lab may need to be right about technological progress, competitive positioning, pricing power, future profitability, and valuation. A lender financing a data center under a long-term contract may need to get fewer things right, but the things that matter – the tenant's creditworthiness, contract terms, power availability, and residual asset value – are quite different. Thus, a key question is, "which of the questions am I or my organization well suited to answer, and which uncertainties am I willing to bear in the absence of that expertise?"

Everything else being equal, a high level of uncertainty can be acceptable to investors if the price of the investment is low enough relative to its intrinsic value. At least in theory, the sectors rated lower on Claude's predictability scale can be the better investment prospects if they're sufficiently cheaper than the more predictable ones. At the end of the day, intelligent investing is largely about price. **Note, however, that in the case of AI, thinking about intrinsic value and thus cheapness is greatly complicated by the questions raised above.**

In an effort to tie this all together, I'll leave you with a list of questions that might help you reach your decisions. Note, however, that these are questions that apply to all investments, not just AI:

- How many things have to go right for me to earn an acceptable return?
- If my central assumptions prove wrong, how much could I lose?
- Which risks am I being compensated for, and which am I simply assuming?
- How much optimism is already incorporated in the price?
- How does this investment change the overall risk profile of my portfolio? Does it amplify the risk I'm already carrying, or does it diversify it?

There's an old saying, and it seems highly applicable to AI: "Two kinds of investors lose a lot of money: those who know nothing and those who know everything." I think in the case of AI, it's unusually important to accept that you don't know everything. AI is sure to produce a lot of surprises, but they

needn't be unbearable. The thing to ponder is, what's the probability an AI investment will disappoint . . . and what's the probability it won't?

I'll conclude with a great quote I happened upon the other day:

> **[The] argument is made that there are just too many question marks about the near future: wouldn't it be better to wait until things clear up a bit?** You know the prose: "Maintain buying reserves until current uncertainties are resolved, etc." Before reaching for that crutch, face up to two unpleasant facts: The future is never clear; you pay a very high price in the stock market for a cheery consensus. Uncertainty actually is the friend of the buyer of long-term values. (Emphasis added)

The perfect ending for a memo on AI, isn't it?

Note, however, that it came from one of my oldest files: timeless wisdom from Warren Buffett that appeared in *Forbes* magazine on August 6, 1979. **I'm always happy to give Warren the last word.**

November 11, 2026


***Legal Information and Disclosures***

---

*This memorandum expresses the views of the author as of the date indicated and such views are subject to change without notice. Oaktree has no duty or obligation to update the information contained herein. Further, Oaktree makes no representation, and it should not be assumed, that past investment performance is an indication of future results. Moreover, wherever there is the potential for profit there is also the possibility of loss.*

*This memorandum is being made available for educational purposes only and should not be used for any other purpose. The information contained herein does not constitute and should not be construed as an offering of advisory services or an offer to sell or solicitation to buy any securities or related financial instruments in any jurisdiction. Certain information contained herein concerning economic trends and performance is based on or derived from information provided by independent third-party sources. Oaktree Capital Management, L.P. ("Oaktree") believes that the sources from which such information has been obtained are reliable; however, it cannot guarantee the accuracy of such information and has not independently verified the accuracy or completeness of such information or the assumptions on which such information is based.*

*This memorandum, including the information contained herein, may not be copied, reproduced, republished, or posted in whole or in part, in any form without the prior written consent of Oaktree.*

---