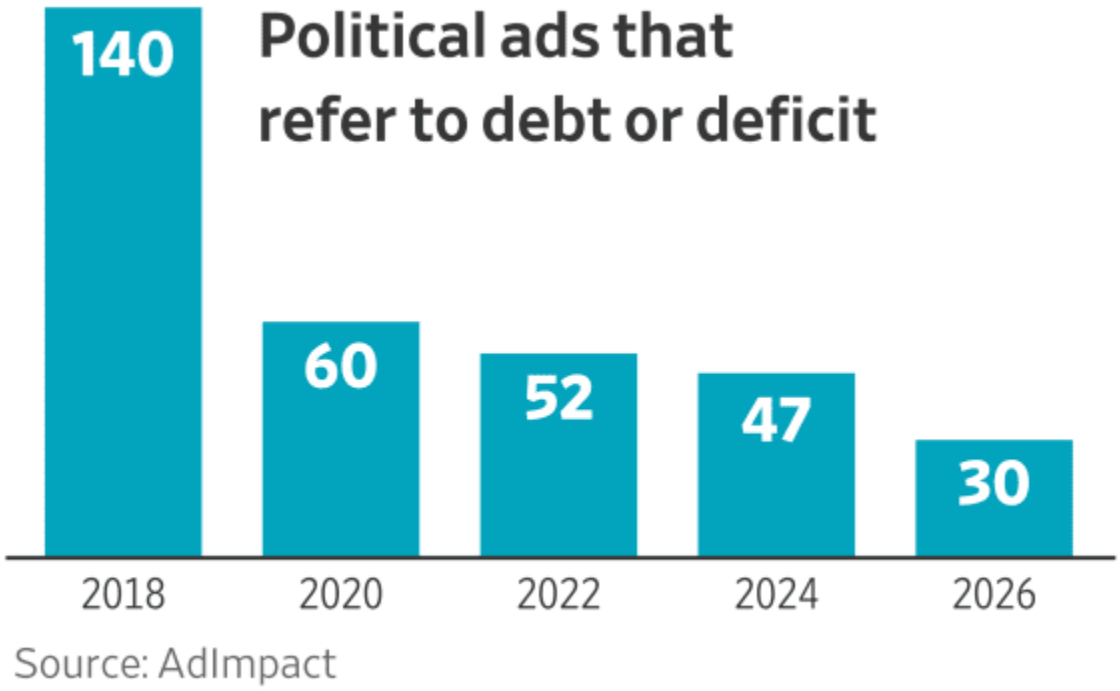
October 2026



From WSJ's The 10-Point on Saturday:



Something is missing from this year's midterm campaigns: America's soaring debt. In 2018, 140 political ads mentioned the debt or deficit. This year, just 30 have.

The silence is striking. Federal debt held by the public reached 100% of GDP for the first time in 80 years, while interest costs now exceed military spending. Yet the politicians once willing to sound the alarm—the fiscal hawks—have largely vanished.

What happened? Some retired, lost elections or simply changed the subject. Rising debt without a financial-market revolt robbed hawks of credibility and made prominent economists rethink the conventional wisdom on debt. Richard Rubin and David Luhnow examine how Washington stopped worrying about debt—and what it could cost us.

The lead story in this weekend's WSJ:

Nasdaq Sinks as Tech Rout Deepens

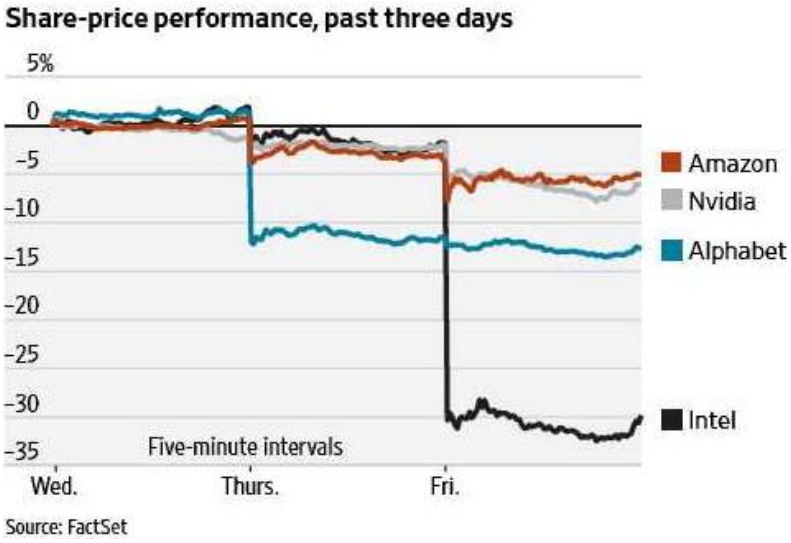
Index logs its worst month in over a year as market takes 'a major gut check'

BY DAVID UBERTI AND HANNAH ERIN LANG

A selloff in artificial-intelligence stocks snowballed throughout the week and culminated Friday with the Nasdaq composite sliding nearly 3%, marking the index's worst month in more than a year and its worst day since June.

Intel was one of the biggest victims with shares crashing 28.5%, the chip maker's worst daily session on record. The plunge vaporized more than $17 billion in combined paper holdings by two of the company's largest shareholders: the U.S. government and Nvidia.

The vulnerabilities in Wall Street's favorite trade have been building for months, periodically appearing around smaller companies drowning in debt. Each time, the largest firms behind the U.S.'s AI build--out have reassured Wall Street by printing profits, suggesting the show will go on even amid higher interest rates.



This week was different. After tech stocks had treaded water for weeks, AI giant Anthropic momentarily rejuvenated investors with a filing that showed soaring revenue ahead of its planned initial public offering. But then publicly traded tech firms reported earnings—and markets got a reality check.

First came a dissatisfying quarterly report from Alphabet, which showed growth in Google's search-advertising business slowed—in part because of the impact of "AI mode" queries. Traders punished the stock with a 13% decline Thursday that marked its worst day on record, hammering major Alphabet shareholders such as Berkshire Hathaway.

The bad vibes got worse late Thursday when Amazon reported plunging free cash flow and revenue that missed analyst estimates. Then, Intel forecast a prolonged slowdown in personal- computer sales, pressured by memory chip prices sent skyrocketing by the AI boom.

The losses bled across the tech sector Friday: every member of the Magnificent Seven closed down, while shares of memory providers, semiconductor manufacturers and neocloud companies also declined. Bitcoin prices, which often move in tandem with riskier stocks, slumped more than 5%. As the day wore on, the selloff gained steam, dragging the Nasdaq composite to a 2.9% drop on the day. …

The selloff has left some investors warning the pullback will continue and others waiting for expected IPOs by Anthropic and OpenAI. At the same time, next week's midterm elections could decide the extent to which data-center backlash nationwide translates to restrictions on development.

Since OpenAI's 2022 launch of ChatGPT kicked the AI race into gear, investors have often piled in indiscriminately, leaving indexes top-heavy with trillion-dollar giants. But the Federal Reserve's September rate hike—and Wall Street's expectation of more to come—has changed investors' calculus.

"Interest rates affect—directly, or indirectly—almost everything," Fed Chairman Kevin Warsh said this week.

The specter of higher borrowing costs recently battered one of the riskiest corners of the AI trade, upstart companies including CoreWeave that rent high-end chips to train and run AI models. Still, many investors believed that Silicon Valley's larger cash-printing machines would be immune— to the extent that some even sought shelter in megacap tech stocks as bond yields marched higher in September.

Now, traders are grappling with the prospect that AI powered profit growth at the market's leading tech firms will only cool off further. Following Amazon and Alphabet's reports, the Magnificent Seven are now expected to report slower earnings growth for the third quarter than the remaining members of the S&P 500, according to a FactSet analysis. …

Wall Street had been largely permissive of big spending because large companies printed cash to finance massive expenditures on chips, data centers and more. …

Shrinking cash stockpiles have left tech giants turning to investors for help, from equity offerings announced by Alphabet and Meta in recent months to debt offerings across the sector ballooning to hundreds of billions of dollars a year.

… President Trump has signaled the U.S. could launch military campaigns against Iran after the midterm elections. That could spur additional disruptions of energy supplies from the Middle East. Front-month Brent crude futures closed at $117.64 on Friday. …

From BCA Research's Peter Berezin on Friday:

US profit margins stand at record-high levels, and analysts expect them to only go up (**Chart 1**). Projected long-term earnings growth is the highest in history, higher even than at the peak of the dotcom bubble.

As we discussed in a Special Report earlier this year, there are sound reasons why profit margins have risen over the past few decades, including increasing corporate concentration and diminished consumer price sensitivity. However, there are also several temporary reasons for today's abnormally high margins.

First, business sales are growing much more quickly than unit labor costs. The gap between the two series correlates almost perfectly with economy-wide corporate profits.

Looking ahead, the outlook for corporate America appears less rosy. Nominal GDP growth slowed from 8.5% in Q2 2026 to 5.6% in Q3. It is expected to decelerate further to 5.1% in Q4, and average 4.4% in 2027. If this were to happen, nominal GDP growth – which tracks business sales very closely – would fall below the 10-year yield. The economy has often succumbed to a recession when growth drops below the long-term bond yield (**Chart 3**).

On the flipside, our wage survey indicator points to some reacceleration in wage growth over the coming months. This could create a toxic brew of slowing sales and rising costs for businesses.

CHART 1

Analysts Expect Record-High Earnings Growth From Record-High Profit Margins

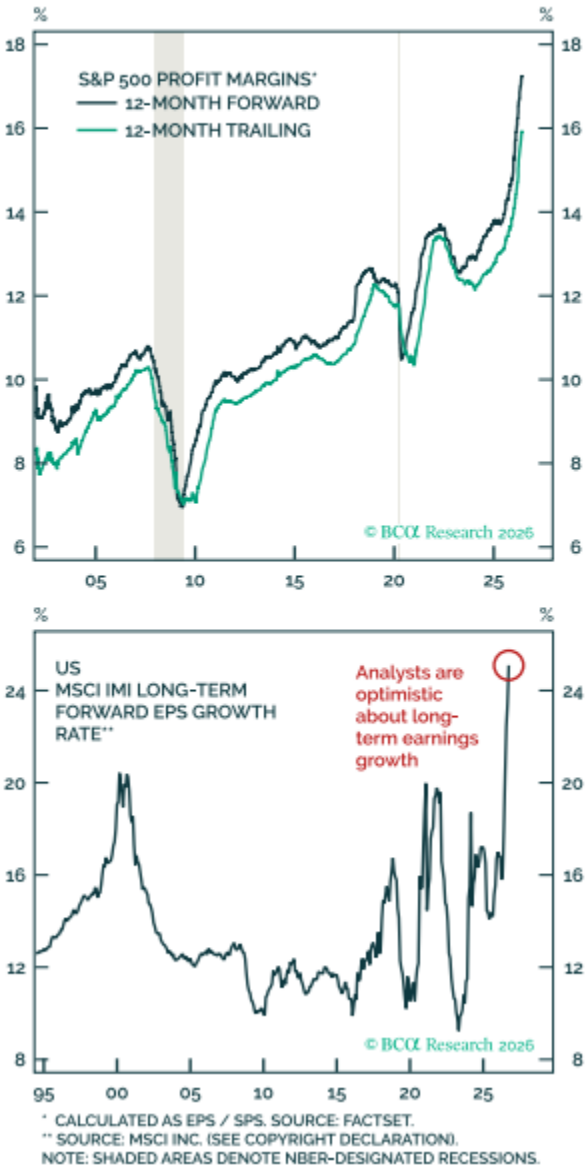
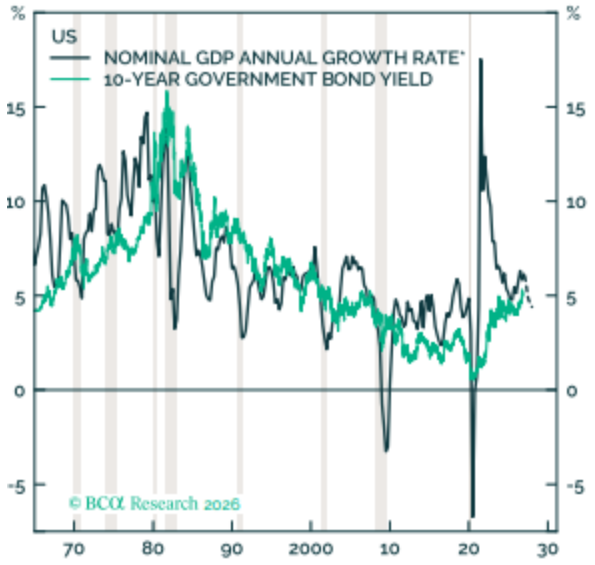


CHART 3

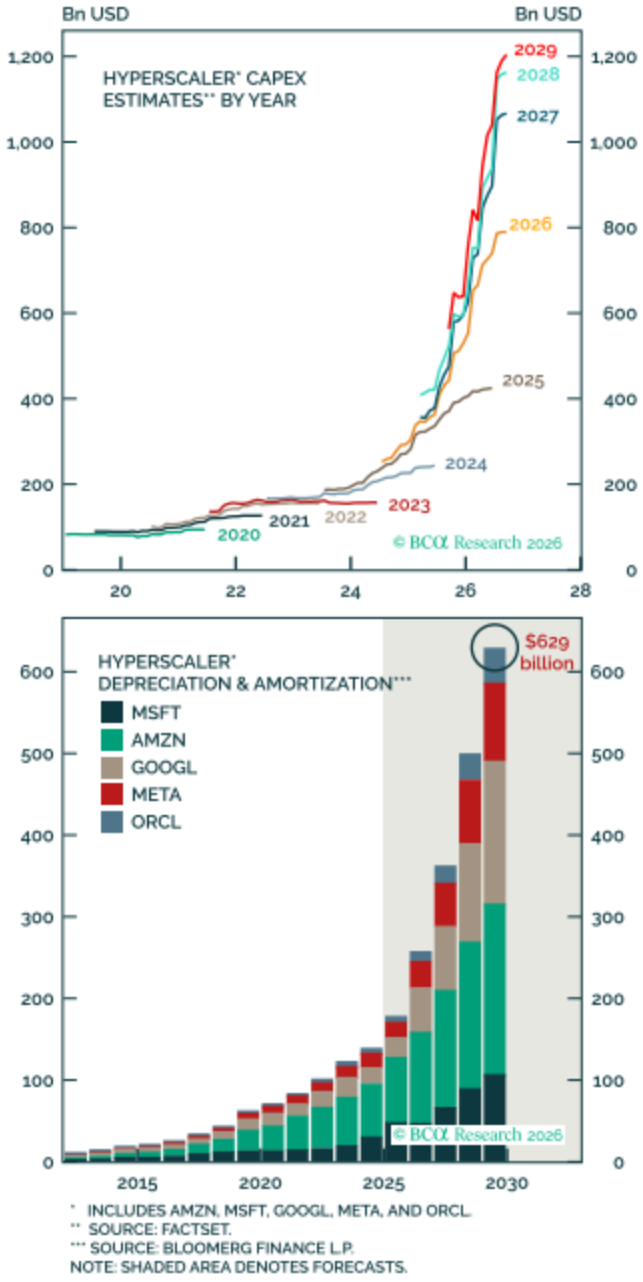
Nominal GDP Growth Is Set To Fall Below The 10-Year Treasury Yield In 2027



* INCLUDES ESTIMATE BASED ON BLOOMBERG FORECASTS OF REAL GDP GROWTH AND INFLATION.
SOURCE: BUREAU OF ECONOMIC ANALYSIS (BEA).
NOTE: SHADED AREAS DENOTE NBER-DESIGNATED RECESSIONS.

CHART 5

Hyperscalers Are Set To Spend More Than $1 Trillion On Capex In 2027



* INCLUDES AMZN, MSFT, GOOGL, META, AND ORCL.
** SOURCE: FACTSET.
*** SOURCE: BLOOMERG FINANCE L.P.
NOTE: SHADED AREA DENOTES FORECASTS.

Meanwhile, the AI capex boom is looking increasingly shaky. The five hyperscalers are on track to spend more than $1 trillion on capex in 2027 (**Chart 5**). In 2029, these five companies are expected to incur $630 billion in depreciation and amortization charges – more than the $565 billion in operating profits that they are likely to generate in 2026. Unless sales grow rapidly, it will be difficult to prevent tech earnings from sputtering.

We are not calling for an imminent capex bust. GPUs remain in short supply. In fact, rental rates for older Nvidia chips have risen this year. Memory prices also remain firm. But we are probably in the latter innings of the capex boom. This means that volatility will increase against the backdrop of less-stellar equity returns.

Six from WSJ's Markets A.M. by Spencer Jakab. On Oct. 30th:

If History Repeats, Small Caps Will Win in a Landslide

No misses since 1946

What happens when you combine three sweet ingredients?

In the stock market at least, there have been some delicious outcomes. The next six months could be the equivalent of s'mores for investors in smaller U.S. companies.

Between the November of a midterm election year and the following April, small-capitalization stocks have had an almost unbelievable average total return of 42% annualized over the past century, according to data from Leuthold Group. By contrast, they've lost 4% annualized during the preceding six months.

That seems to stem from the combination of three winning patterns—the "best six months" plus midterms plus the "size effect." Their combined benefit isn't only strong but also remarkably consistent. One has to go back 80 years to find a midterm year when small caps failed to gain.

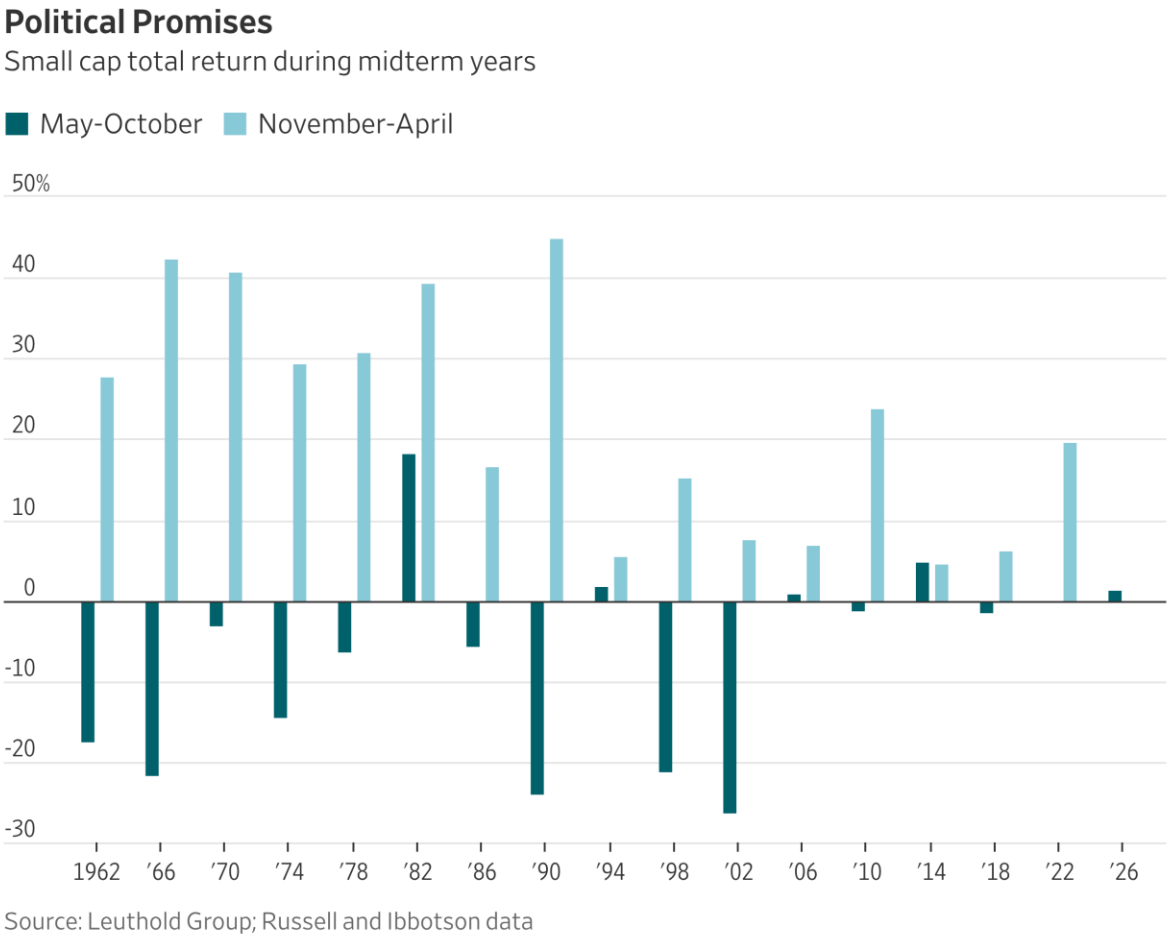
The relief that investors feel once midterm elections results are known has held steady no matter the political outcome. The main thing seems to be that the uncertainty is over.

American elections just happen to be in November, but the stock market's six-month pattern seems to reinforce it in other markets too. It's believed to have its origins in the seasonality of money demand when the world economy was dominated by agriculture—planting during spring and harvest in the autumn.

The advice to "sell in May and go away" isn't as reliable as it once was, but it seems to hold up in this particular case. Why is this effect so strong for smaller companies, though?

One is that they're more impacted by market liquidity. When it ebbs during the summer or periods of greater investor caution, small stocks tend to get hit harder. The flip side is that they more than make up for it when investors put fresh money to work. The January effect, for example, tends to be strongest for small caps.

The stars have been megacap tech stocks in recent years, but small caps have a superior historical track record overall. Depending on which index one consults, a dollar invested in U.S. small-cap stocks a century ago grew



to five or six times as much as one invested in the large-cap S&P 500.

There'll be lots of champagne on ice around the U.S. for victory parties or drowning of sorrows Tuesday night. Investors might have something to celebrate regardless of the outcome.

On Oct. 26th:

Financial Wisdom From Cartoon Apes

Monkey see, monkey buy

Five years ago today, Brandon Buchanan's investment firm paid a record $3.4 million for a drawing of a primate.

Feel free to gaze at his or any of the 9,999 other JPEGs in the Bored Ape Yacht Club. You can even make a copy for personal use, but officially owning one on the blockchain will cost you. The "best offer" for Buchanan's is about 97% lower than in 2021, which is typical.

To most people, Bored Apes and their nonfungible token brethren, such as Pudgy Penguins and CryptoPunks, are a curiosity. Their wild rides hold some useful investing lessons, though, such as what made them so valuable in the first place.

One is "social proof." Celebrities endorsing a product is an age-old tactic, but buying one with their own money creates its own kind of validation. And many of the news stories about Justin Bieber, Snoop Dogg, Paris Hilton and Jimmy Fallon's Bored Apes mentioned specific dollar amounts paid (in cryptocurrency equivalent).

Once broadcast, they seemed reasonable. That can anchor our price expectations for speculative stocks, but it's especially relevant when valuing something like art that will never produce cash flows.

Then there's identity. Some buyers consider meme stocks or crypto as a sort of initiation into a group—an emotionally appealing feature cleverly tapped with the apes' "club" moniker. "We're part of the same community," said Fallon to Hilton as they compared apes on "The Tonight Show."

They were among celebrities sued for failing to disclose financial incentives tied to their ownership (plaintiffs' case is on weak ground since Bored Apes probably don't qualify as securities).

Hype aside, collectibles are a substantial asset class and diversifier among the wealthy, NFTs included. In an email interview, Buchanan points out that tokenized art "can be admired by others, and the transferability is seamless."

You might be capable of drawing a nifty cartoon, but good luck selling it on the blockchain. The secret sauce is a combination of buzz and scarcity. There's only a loose connection between an artwork's beauty and its price.

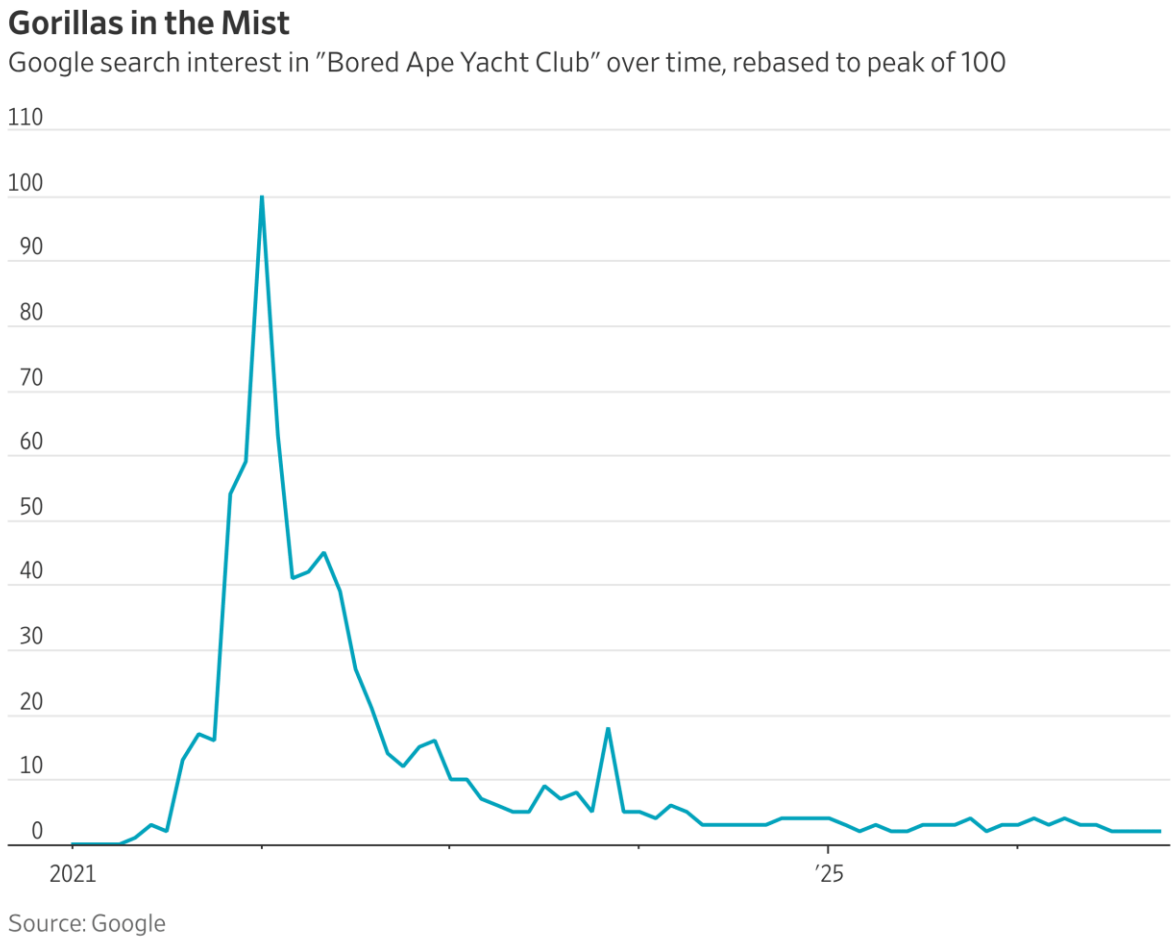
Leonardo da Vinci's "Salvator Mundi" sold for $125 in 1958 when it was thought to have been painted by a student. Then it smashed price records in 2017, fetching $450 million. Fewer than 20 da Vinci paintings exist.

People who are passionate about collectibles aren't as swayed by soaring prices, making buying opportunities rare. Pitching them as an easily-transferable investment short-circuits that, though.

Buchanan says he'd sell his rare Bored Ape on behalf of his fund if its price rose tenfold, but probably not one he owned personally. He expects them to be shown in museums.

"I'd hate to see NFTs…just be talked about as monkey photos that people speculated on that went bust."

Collectibles can still be a good investment without being investments, and many actual investments can still be bad. The more their value relies on psychological anchors, the greater the chance of them making a monkey out of you.



On Oct. 23rd:

Elite Colleges Can Give You an Investing Lesson

It's vanity

You don't have to be a rocket scientist to make average investing returns. Sometimes it helps not being one.

The Massachusetts Institute of Technology says that its massive endowment, managed by executives who earn more than the university's president, earned 10.3% during its fiscal year ended in June. That's less than half of what an S&P 500 index fund earned.

MIT's results are more surprising during a tech boom that it might seem uniquely equipped to harness. Harvard and Stanford's endowments, which held big stakes in SpaceX acquired before its IPO, jumped 27.1% and 31.7%, respectively. The University of North Carolina, which had more than a tenth of its endowment in that stock, did even better.

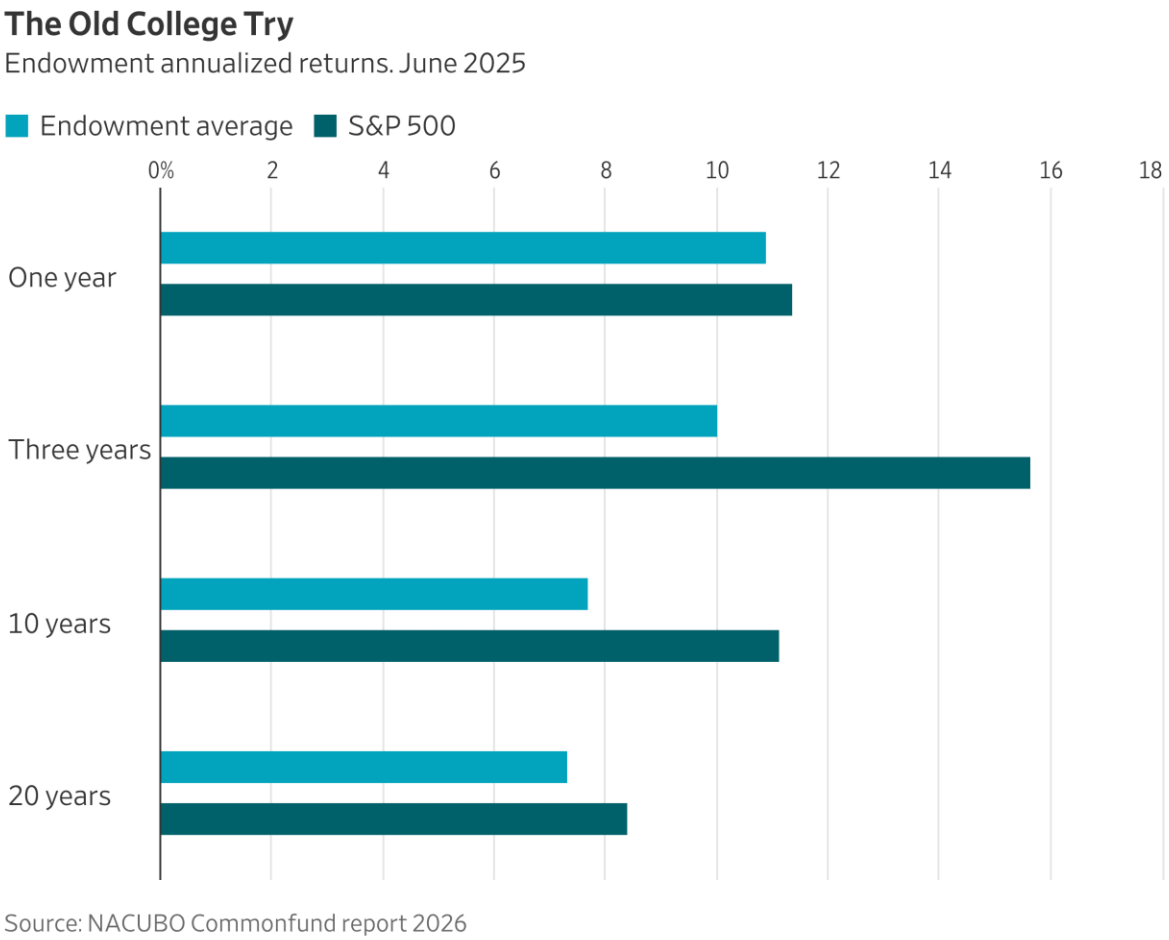
As uplifting as one-year moonshots are, big college endowments offer tuition-free lessons these days about how *not* to invest. Special access to mostly so-so but costly private equity, hedge and venture-capital funds has seen their returns lag what mom and pop investors could have made over the past decade. Meanwhile, infrequent fund value updates give endowments a false appearance of stability. Only 14% of Harvard's is directly in publicly traded stocks and just 3% in bonds.

Once upon a time, the "Yale Model," pioneered by that university's chief investment officer David Swensen, got great results this way. He could take a very long view, and illiquid investments were cheaper back then.

Now, though, those funds are seeking money from a new audience—you. Lobbyists are pushing to allow alternative investments into 401(k) plans and already offer them outside retirement accounts. Broker Robinhood earlier this year listed its first "Venture Fund," which holds stakes in unlisted companies such as OpenAI, Oura and in now-public SpaceX.

"Opening up private markets will resolve one of the greatest longstanding inequities in capital markets today," said Robinhood CEO Vlad Tenev.

That fund sports a 3.1% expense ratio, 100 times the cost of an S&P 500 index fund.



An annual report of university business administrators from 2025 points out an interesting divide between the haves and have-nots of higher education. Institutions with assets of less than $50 million allocated 13% of their portfolios to “alternatives.” At the other end of the spectrum, those with more than $5 billion in assets allocated 63%.

Why? Partly it’s vanity, according to investing veteran Richard Ennis. Harvard University gets better access than Harding University. Surely their money should be managed by the best and brightest.

But the only sure thing is expenses. Ennis writes that endowment managers “mostly operate as if nothing has changed” since Swensen’s heyday. Allocating 60% to alternatives costs them about 3% annually, or about what Robinhood charges, according to his math.

Even those of us who couldn’t get into MIT can spot the problem.

On Oct. 20th:

Stocks Face Headwind That Ended Other Bull Markets

Supply and demand

Investment manias don’t end the way they used to.

The common historical thread is that some amazing new technology attracts lots of capital. Eventually so much of the new thing gets built that supply outstrips demand.

Then the price of the wonder tech crashes, along with the shares of companies building it. Railroads, which were so revolutionary that multiple routes were built between cities in the 19th century, are the classic example.

The world moves faster now. When amazing breakthroughs lead to high valuations, companies rush to sell lots of shares so they can cement a dominant position. Meanwhile, freshly minted millionaires and billionaires who got in early rush to cash out. That supply of stock can pop the bubble before a glut of the new thing itself necessarily does.

Asset manager GMO’s Ben Inker and John Pease write in “A Catalyst for the AI Bubble Break” that we might take comfort in high prices of Nvidia chips or gas turbines used to power AI data centers. But a boom’s building blocks can stay expensive even after stock prices don’t. Fiber optic cable that enabled the dot-com build-out was still a hot commodity for a while after stock prices cratered.

Because of the way that investors tend to treat a big supply of new stock, GMO’s paper calculates that each 1% increase has tended to cut the overall market’s value by 4% during the next year, all else being equal. As lockup periods expire following its June offering, SpaceX alone could add around 1% to U.S. equity supply by June 2027. Anthropic and OpenAI are expected to tack on a lot more.

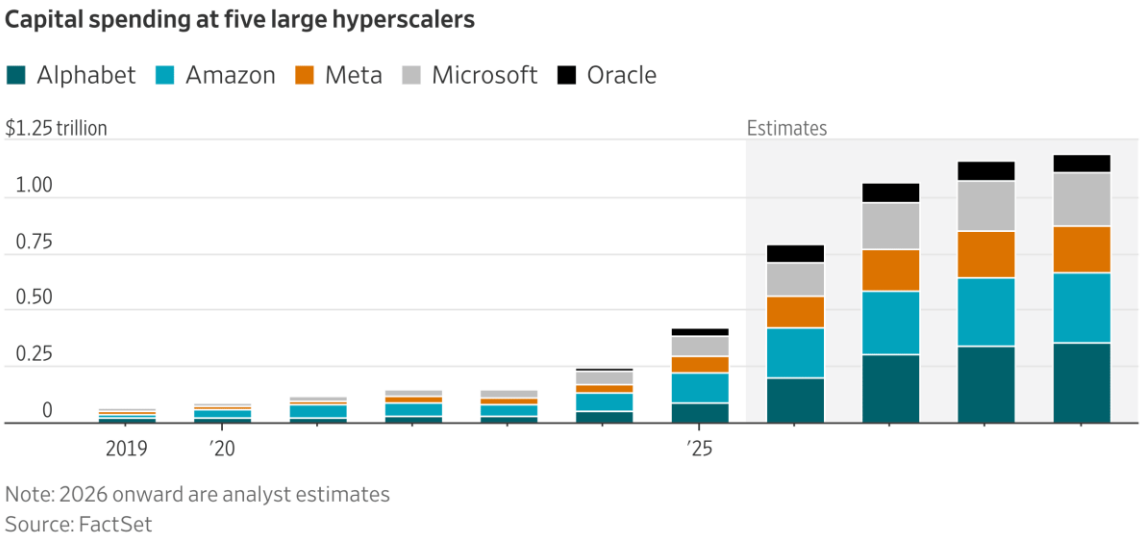
And all else isn’t equal when valuations are elevated. The S&P 500’s cyclically-adjusted price/earnings ratio isn’t far from its tech bubble record, so what might be a 20% selloff in stocks could be more like that episode, when stocks dropped by half.

“Nothing tells us a similar event is off the table for today’s market,” write Inker and Pease.

The supply equation has been a tailwind for stocks in the past 20 years. Companies bought back more shares than issued, or they bought one another. Many large ones stayed private. Asset manager Pictet estimates that U.S. public equity supply has dropped by about 1% a year since 2005, except for a brief spurt around the financial crisis.

Adding to the pressure from big IPOs, hyperscalers plowing money into data centers have hit the brakes on buybacks. Meta repurchased nearly $150 billion of stock between 2021 and the third quarter of last year, but nothing since then. Alphabet and Oracle each sold tens of billions in new stock.

Supply and demand will always matter, but investors need to focus on the asset that's easiest to manufacture these days: new shares.



On Oct. 13th:

This Investing Strategy Is So Good You Can't Handle It

Like eating a habanero

Mae West said that "too much of a good thing can be wonderful."

Well, maybe. A case in point is taking a good stock investment strategy and combining it with another good one and then another. Instead of canceling each other out, Wes Crill, a senior investment executive at Dimensional Fund Advisors, found that three well-known drivers of extra returns more or less stacked up, producing a truly wonderful outcome.

The factors are small stocks, low relative price and high profitability. Combining them resulted in an annualized return of 17.1% since July 1963.

That compares with 10.9% for U.S. stocks generally. Through this July, Crill's hypothetical combo would have grown a $10,000 investment into $208 million before taxes and fees, or more than 30 times as much as U.S. stocks overall.

But Crill invokes a different Hollywood icon—Jack Nicholson in "A Few Good Men"—to explain why nobody would have stuck it out.

"You can't handle the tracking error!"

Since there aren't too many stocks that tick all three boxes, the small number turned out to be frighteningly volatile relative to the market—no picnic itself sometimes. To earn an extra 6.3 percentage points annually, the portfolio trailed by as much as 38.4 points over a 12-month period and by a gut-wrenching 72.6 points over 36 months.

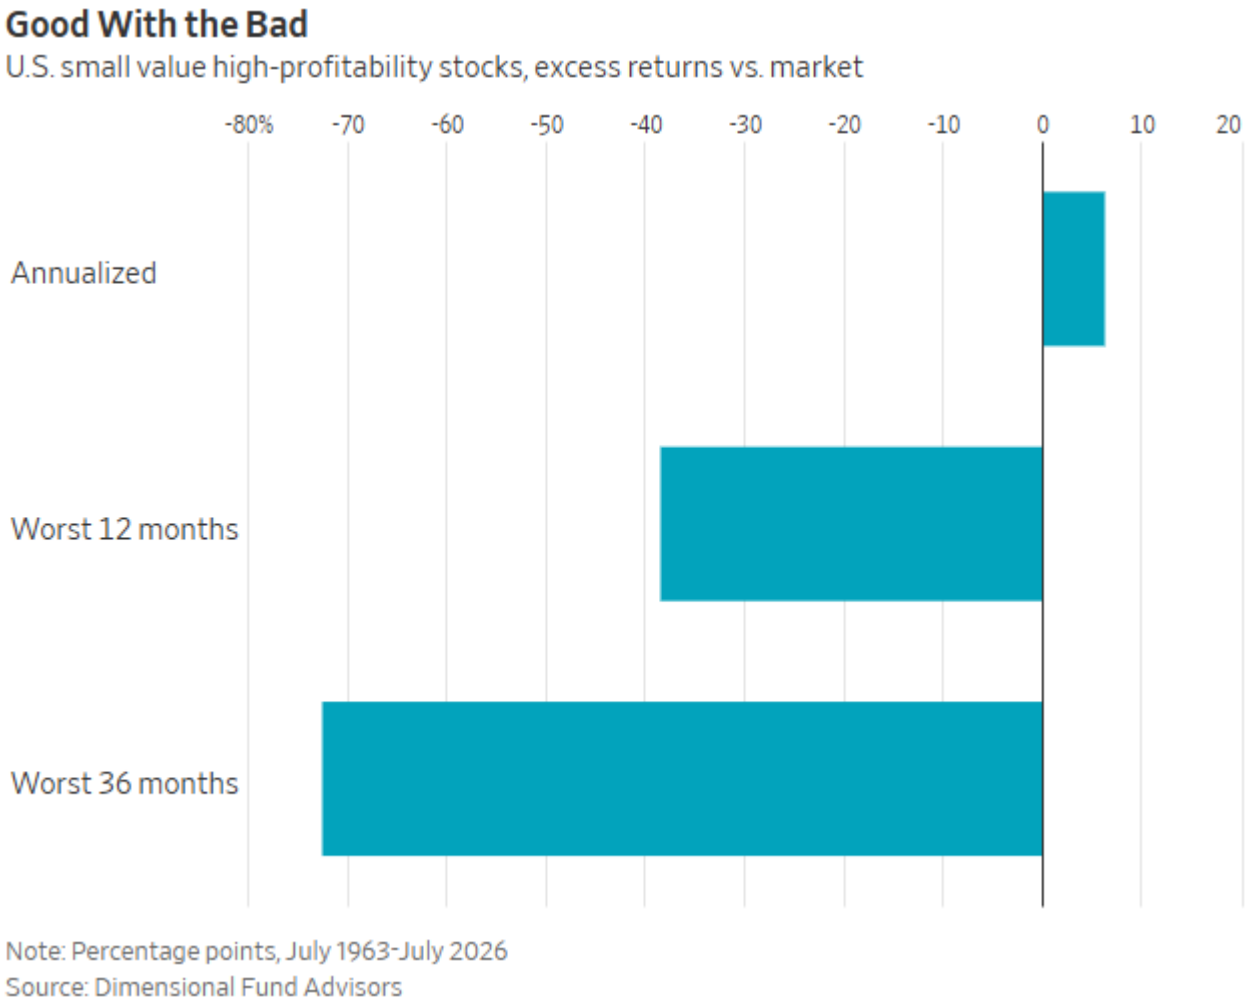
In theory, a young investor could lock away their paychecks in such a strategy until retirement day and never open an account statement. They'd have to be confident, though, that they'd never need to tap their savings in an emergency.

The future won't look like the past, of course, and Crill didn't highlight these results to sell a new product anyway.

Instead, it's a useful way to respond to clients who ask him why lots of smart investors haven't snatched away all of the excess-return hacks that finance professors have identified. It's uncomfortable.

"That's my best explanation for why this persists," says Crill.

Just buying and holding a plain vanilla index fund is hard too. And good returns aren't guaranteed, either. An investor with truly unlucky timing could have sat tight for decades and earned less than ultrasafe bonds.



That possibility underscores why sprinkling in Crill's winning factors could be useful. Every little bit of extra return makes stock investing easier to stomach.

The analogy Crill uses is having a habanero pepper in your kitchen. Eating one is too much, but adding a small bit to a dish can spice up a portfolio without causing intestinal distress.

On Oct. 2nd:

The Amazing Stat That's a Red Flag for Investors

Aim high

This newsletter's readership is growing enough to keep your faithful writer employed—thank you—but why am I thinking so small?

The TAM, or "total addressable market," of news publications globally is $125.7 billion, according to the World Association of News Publishers. If it were possible to buy a share of Market A.M.'s future revenue then that's the sort of statistic a smart investment banker might mention.

A smart investor, on the other hand, would hang on to their wallet. There's nothing inherently wrong with guesstimating economic opportunities, but the more prominent, inflated, or specific they are, the worse of a sign it can be.

Take mattress-in-a-box company Casper, whose 2020 initial public offering prospectus cited a "global sleep economy" of $432 billion. The company initially wanted to sell stock at between $17 and $19 a share but ended up pricing the IPO at $12. Despite a pandemic-stimulus-fueled golden age for unprofitable companies, the stock had dropped to $3.55 the following year when an offer to take Casper private was accepted.

The same go-go era saw a company make even Wall Street blush. WeWork will forever be remembered for inventing "community-adjusted Ebitda," a profit measure that excluded many expenses.

Less-remembered is its estimate of a $3 trillion TAM based on almost every office worker on Earth being a potential customer. The landlord portrayed itself as a tech company and snagged a private valuation of $47 billion. But its planned 2019 IPO had to be pulled. It filed for bankruptcy in 2023.

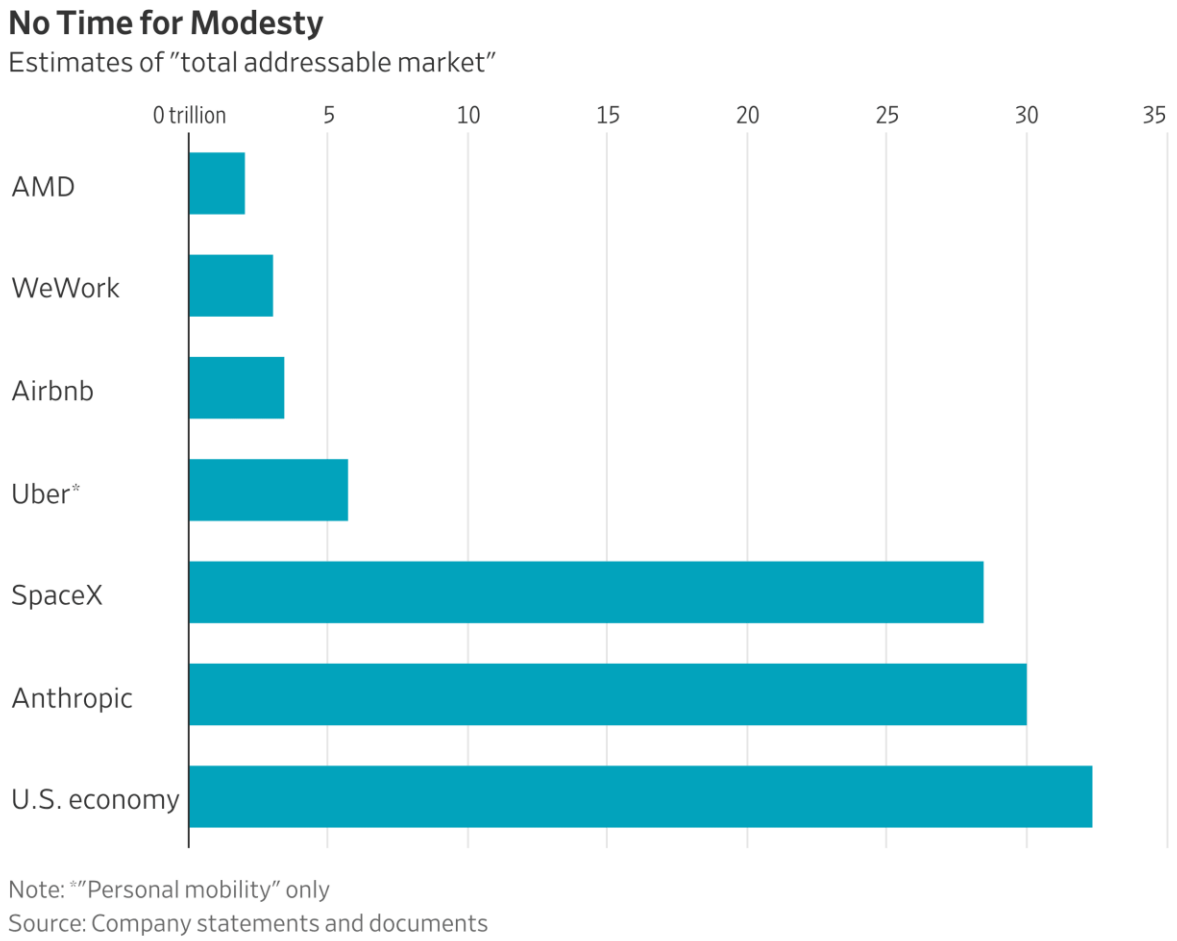
For companies that eventually turned profitable, hyperbole can be forgiven as a youthful indiscretion. Airbnb had a TAM of $3.4 trillion and Uber $5.7 trillion for personal mobility alone. Leave it to Elon Musk, though, to say "hold my beer." SpaceX, in documents for its June IPO, cited the "largest actionable" TAM "in human history" at a cool $28.5 trillion.

In his post-IPO estimate of SpaceX's value, New York University's Aswath Damodaran wrote that "the estimation of TAM has been gamified by Silicon Valley, with bloated and patently unreachable numbers."

Another thing many TAM estimates have in common is that they're oddly specific, often going out a decimal point. It's probably no accident. In his book "Proofiness," science writer Charles Seife calls it "disestimation," getting us to ignore statistical uncertainty.

Loss-making companies that go public have an especially patchy record, according to University of Florida professor Jay Ritter. They tend to do much better on the days of their IPOs than profitable ones. Individual investors rarely get to reap that gain, though. After that, the average return of the nearly 4,000 money-losing companies Ritter studied was negative 20.5% (actually -30.7%) over three years relative to the market.

Bankers often sell the sizzle, not the steak.



From WSJ Economics by Greg Ip on Oct. 20th:

Yes, the Fed Can Kill the AI Boom

Back in 2000, as the Federal Reserve raised interest rates, entrepreneurs and investors surfing the internet wave insisted it didn't matter. "Everybody knows the Fed's going to raise rates in February, but as this comes closer, VCs are throwing ever larger amounts of money at us," one told my then-colleague Jacob Schlesinger.

We know what happened next. The Fed kept raising rates and the internet-stock bubble imploded.

Déjà vu: The same arguments are being recycled today for artificial intelligence. What if "the demand for intelligence and energy is unaffected by higher rates because winning the race for super intelligence has a near infinite ROI," hedge-fund manager Bill Ackman asked on X.

John Arnold, co-chair of Arnold Ventures, argued higher rates could be "counterproductive" for inflation: They won't reduce demand for AI inputs, but do "raise business operating expenses, some of which is passed along to consumers."

Let me set the record straight: The Fed is entirely capable of deflating the AI boom and inflation.

Skeptics assume interest rates could only slow AI companies by raising their cost of borrowing. That is a fallacy. Short-term rates are the benchmark cost of liquidity, affecting overall financial conditions. With each successive increase, tighter credit ripples out along the risk spectrum, from Treasurys to corporate and household loans, equity and so on. The leveraged investors who buy risky assets see their funding squeezed. Appetites for initial public offerings, high-yield bonds, private credit and securitizations dry up.

This is why asset meltdowns through history bear the Fed's fingerprints: the 1929 and 1987 stock-market crashes, the dot-com and housing busts, emerging-market crises in 1982 and 1997-98. As assets deflate, investment follows and sometimes recession comes next. And with recession, inflation goes down.

To be clear, the Fed doesn't want to end the AI boom. It only wants overall inflation to fall to 2%, which doesn't require that AI-related inflation do so as well.

But make no mistake: If the only way to 2% inflation is to drain the lake of liquidity sustaining AI, the Fed has the means and, very likely, the will to do it.

Follow-ups

From WSJ:

Will America Spend 9% of Its GDP on AI? The Industry Is Counting on It

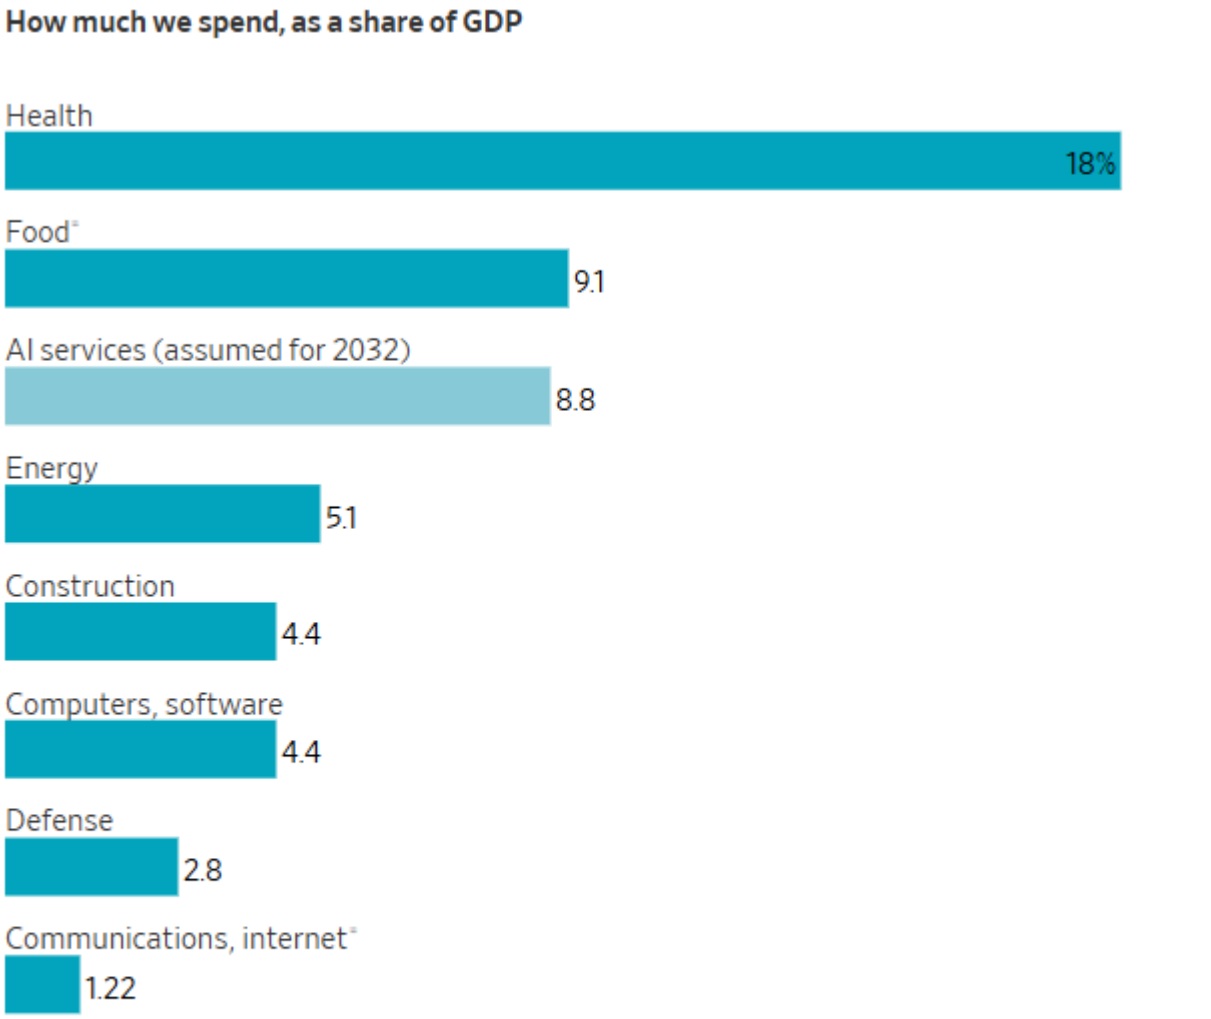
To justify staggering investment sums, Americans will have to spend as much of their income on this one technology as they do on food

By Greg Ip
Oct. 2, 2026

Believers in the artificial-intelligence boom need to take a very close look at this number: 9% of GDP.

That is how much American businesses and consumers eventually have to spend per year on the services of companies like Anthropic and OpenAI to justify the staggering sums being committed to the technology right now.

Is it plausible that Americans will spend as much of their income on this one technology as they do on food? Roughly twice what the nation pays for all forms of energy or all computers and software? Seven times what consumers spend on phone, streaming, and internet services combined?



You should be skeptical. Even the most transformative inventions eventually run into the law of diminishing returns: each additional dollar a user spends yields less additional productivity (or enjoyment) than the last. That imposes a natural ceiling.

The question, of course, is where that ceiling is. Whether or not you think 9% of GDP is right, you have to care, because this figure isn't some fever dream: it is implicit in the dollars that investors and companies are committing right now.

The figure is courtesy of Columbia University finance professor Stijn Van Nieuwerburgh. His paper, first presented at the Brookings Institution, showed that the AI build-out is now bigger than any investment boom in American history. His more intriguing, and sobering, statistic is how much revenue AI would have to garner to justify that boom.

Van Nieuwerburgh started with the actual plans of today's builders of data centers and the average cost of a data center, then assumed some data centers would be canceled, that each dollar of revenue produces 50 cents of cash flow, and that builders need an unleveraged return on invested capital of 10%.

He concludes from all this that AI revenue in 2032 would need to be $3.5 trillion, or 8.8% of GDP (which he assumes grows 4% per year, unadjusted for inflation).

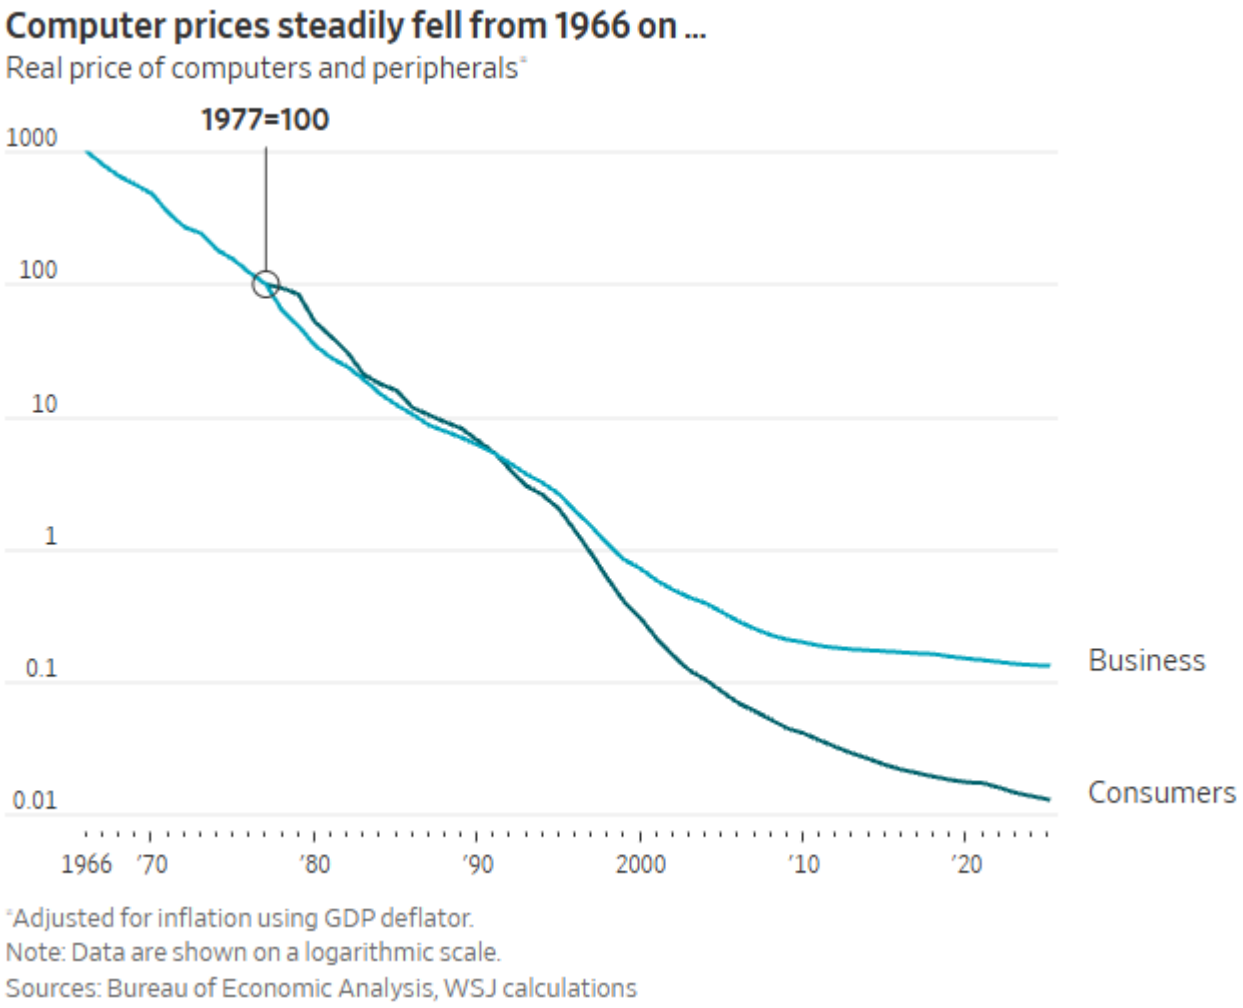
A key assumption in this revenue forecast is that the AI companies will be able to keep charging today's prices for computing power. But Van Nieuwerburgh said the scarcity of computing capacity is holding up prices and profit margins now. The companies are "basically saying we're going to keep charging scarcity pricing in 2032, when presumably there will be tons of competition and we've quadrupled our capacity."

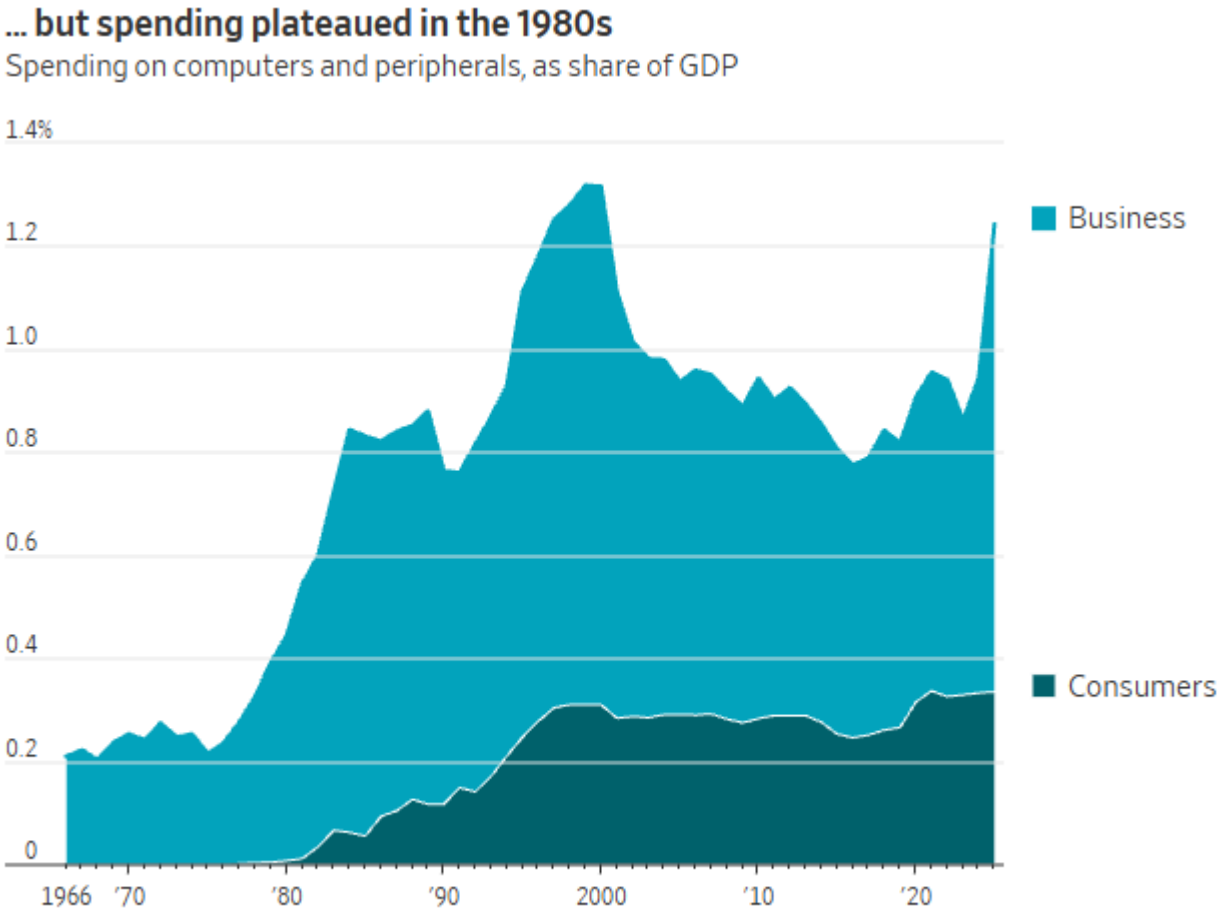
Plunging prices helped undo previous investment booms. Between 1997 and 2001, the price of bandwidth on fiber between London and New York plunged 96% as more fiber strands were laid and new technology expanded the capacity of each fiber. That price collapse accelerated the bankruptcy of numerous long-haul fiber companies.

The optimist's case

In tech, falling prices are the flip side of rising efficiency. AI models double in capability roughly every four to five months. The effective price for a given level of capability has fallen 47% per quarter since 2023, according to Epoch AI, six times faster than for computing power.

The optimistic case for AI is that a falling price and rising capability spurs more than enough demand to offset the falling price, a relationship dubbed the "Jevons Paradox" after British economist William Stanley Jevons who noted the phenomenon with coal power in 19th-century England.





That certainly seems to be the story so far. A few months ago, an expert panel surveyed by a team led by Ezra Karger at the Federal Reserve Bank of Chicago, projected OpenAI and Anthropic would have combined annual revenue of $300 billion by 2030.

That looks way too conservative: the two companies' combined run rate is already around $180 billion, and growing at double-digit rates each quarter.

AI evangelists think artificial intelligence will, in effect, become a third factor of production, alongside capital and labor. In that case, 9% of GDP doesn't seem so much. (Labor now gets about 51% of GDP.) Of course, in such a scenario, employment would plummet and governments would have to tax AI to make up for lost income and payroll taxes.

The more prosaic risk is that current growth rates won't be sustained, once adoption is widespread and the most productive applications deployed.

The computer template

From the 1960s through the early 2000s, the density of microchips doubled roughly every two years, a phenomenon dubbed "Moore's Law," and the real price of computers fell roughly 15% per year. The personal computer and killer apps such as spreadsheets fueled an explosion of demand that far offset that falling price: between 1974 and 1984, business investment in computers and peripherals tripled to 0.8% of GDP.

Thereafter, computers continued to get better and cheaper, but didn't spur the same new demand. Computers were now ubiquitous. The first versions of VisiCalc and Lotus 1-2-3 revolutionized business in ways that ever more powerful versions of Microsoft Excel couldn't. So business spending on computers plateaued, except for secondary booms fueled by the internet in the late 1990s and AI now.

AI is improving and adoption is growing much more quickly than with computers, and productivity improvements have been impressive.

Harvard economics doctoral students Fiona Chen and James Stratton reviewed software projects through the platform Jellyfish, which firms use to analyze their engineering teams. AI assistants, they found, boosted lines of code by 12%, and "pull requests," when new code is merged into an existing code base, by 5%. AI agents, which can operate autonomously from humans, boosted lines of code by 30% and pull requests by 23%.

And yet, the authors found, AI's boost to completed projects was small and "statistically insignificant." The reason was that users were spending much more time reviewing, commenting on, and changing pull requests.

The point is that AI will almost certainly raise productivity and economic growth, but not necessarily in ways that translate directly into revenue, or the fat returns investors are counting on.