

TOP *of* MIND

GENERATIVE AI: HYPE, OR TRULY TRANSFORMATIVE?



Since the release of OpenAI’s generative AI tool ChatGPT in November, investor interest in generative AI technology has surged. The disruptive potential of this technology, and whether the hype around it—and market pricing—has gone too far, is Top of Mind. We speak with Conviction’s Sarah Guo, NYU’s Gary Marcus, and GS GIR’s US software and internet analysts Kash Rangan and Eric Sheridan about what the technology can—and can’t—do at this stage. GS economists then assess the technology’s potentially large impact on productivity and growth, which our equity strategists estimate could translate into significant upside for US equities over the medium-to-longer term, though our strategists also warn that past productivity booms have resulted in equity bubbles that ultimately burst. We also discuss where the most compelling investment opportunities in the AI space may lie today, and the near-term risks investors should most watch out for.

“

We’re entering the era of what I think of as “Software 3.0”... companies don’t need to collect nearly as much training data, which suddenly makes the technology much more useful, accessible, and less expensive.

- Sarah Guo

The intelligence of AI systems is being overhyped... those who believe artificial general intelligence (AGI) is imminent are almost certainly wrong.

- Gary Marcus

AI probably isn’t in a hype cycle... this technology cycle isn’t being led by upstarts, which makes it less likely to fizzle out or take a long time to get going.

- Kash Rangan

The vast majority of the companies that have outperformed the broader market over the last several months on the AI theme are still trading at relatively reasonable multiples to GAAP EPS.

- Eric Sheridan

”

Allison Nathan | allison.nathan@gs.com

Jenny Grimberg | jenny.grimberg@gs.com

Ashley Rhodes | ashley.rhodes@gs.com

WHAT’S INSIDE

INTERVIEWS WITH:

Sarah Guo, Founder, Conviction, former General Partner, Greylock

Gary Marcus, Professor Emeritus of Psychology and Neural Science, New York University

Kash Rangan, US Software Equity Research Analyst, Goldman Sachs;
Eric Sheridan, US Internet Equity Research Analyst, Goldman Sachs

AI’S POTENTIALLY LARGE ECONOMIC IMPACTS

Joseph Briggs, GS Global Economics Research

US EQUITIES: GAUGING THE AI UPSIDE

Ryan Hammond and David Kostin, GS US Portfolio Strategy Research

MARKETS AROUND PAST PRODUCTIVITY BOOMS

Dominic Wilson and Vickie Chang, GS Markets Research

WHAT WE’RE HEARING FROM PUBLIC INVESTORS

Peter Callahan, GS Global Banking & Markets

...AND MORE

Investors should consider this report as only a single factor in making their investment decision. For Reg AC certification and other important disclosures, see the Disclosure Appendix, or go to www.gs.com/research/hedge.html.

Macro news and views

We provide a brief snapshot on the most important economies for the global markets

US

Latest GS proprietary datapoints/major changes in views

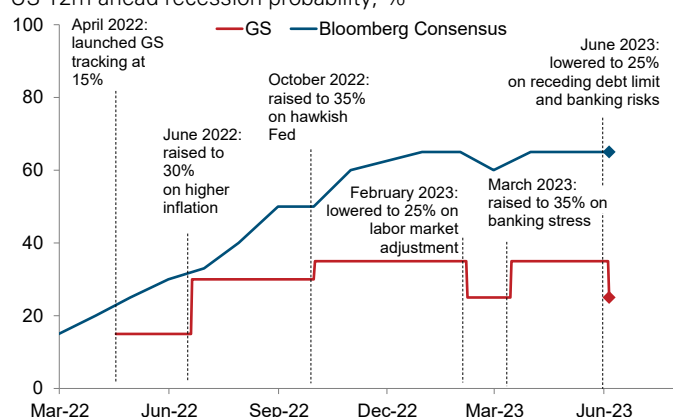
- We recently raised our Fed terminal rate forecast to 5.25-5.5% (added a 25bp hike, most likely in July) given our above-consensus growth forecast and signals from Fed officials.
- We recently lowered our 12m recession odds to 25% given the passing of debt limit tail risk and our increased confidence of only a modest GDP drag from tighter bank lending.
- We recently lowered our Dec 2023 core PCE inflation forecast to 3.5% (yoy, from 3.7%) given our expectation of renewed declines in inflation this summer and beyond.

Datapoints/trends we're focused on

- US labor market; we estimate that labor market tightness will only ease to pre-pandemic levels in early 2024.

Receding US recession risks

US 12m ahead recession probability, %



Source: Bloomberg, Goldman Sachs GIR.

Japan

Latest GS proprietary datapoints/major changes in views

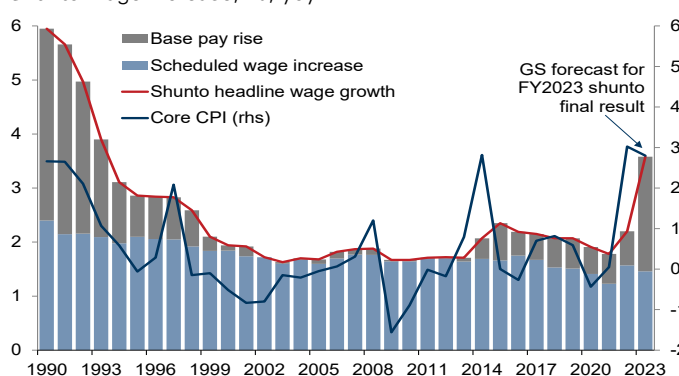
- We recently raised our FY23 New Core CPI inflation forecast to 3.8% (from 3.6%) reflecting our updated FX assumptions.

Datapoints/trends we're focused on

- BoJ policy; we continue to expect a yield curve control (YCC) adjustment in July, with a shortening of the target maturity to five years from 10 as the most likely outcome.
- Japanese wage growth; higher wage growth could yield more sustainable and stronger wage-price dynamics, but a divergent wage-price spiral is unlikely, in our view.
- Japanese consumer sentiment, which has risen sharply.

2023 shunto wage hike could translate into close to 3% macro wage growth in FY23

Shunto wage increase, %, yoy



Source: JTUC-RENGO, Goldman Sachs GIR.

Europe

Latest GS proprietary datapoints/major changes in views


- We recently raised our ECB terminal rate forecast to 4% (added a 25bp hike in Sept) given updated ECB inflation projections and no talk of a "pause" at the June meeting.
- We recently raised our BoE forecast for August and now expect a 50bp hike (vs. 25bp before) after the BoE's 50bp hike in June, which we believe signals more concern around UK inflation and a shift in the BoE's reaction function. We expect a final 25bp hike in September for a terminal rate of 5.75% (vs. 5.5% before).

Datapoints/trends we're focused on

- EA core inflation, which we expect to fall to 3.7% yoy by YE.

UK wage growth has moderated, but remains high

UK private sector regular pay growth, %



Source: Haver Analytics, Goldman Sachs GIR.

Emerging Markets (EM)

Latest GS proprietary datapoints/major changes in views


- We recently lowered our 2023 China real GDP growth forecast to 5.4% (from 6%) on persistent growth headwinds (property slowdown and lack of consumer confidence in particular) and constrained policy responses.
- We recently raised our CY23 India real GDP growth forecast to 6.4% (from 6%) on a net export boost.

Datapoints/trends we're focused on

- China property sector; with no "easy fix", ongoing property weakness will likely be a multi-year growth drag for China.
- EM cutting cycle, which may be approaching, with LatAm likely leading the way, even as most DM central banks continue hiking.

China property sector likely a multi-year growth drag

Housing contribution to yoy GDP growth, pp yoy



Source: Haver Analytics, Goldman Sachs GIR.

Generative AI: hype, or truly transformative?

OpenAI's November release of ChatGPT—a generative artificial intelligence (AI) tool that creates content using natural language prompts—followed by AI chipmaker Nvidia's substantial upward revision to revenue guidance in its Q1 earnings report, has triggered a surge in investor interest in generative AI technology. Indeed, Nvidia's share price has risen over 30% since the revision, and the handful of large tech companies building the foundational large language models (LLMs) at the heart of generative AI have substantially outperformed the broader market. But has the AI hype gone too far? The disruptive potential of generative AI technology—and whether it warrants the current investor enthusiasm—is Top of Mind.

We first explore what's differentiating about generative AI technology that's captured investors' attention. GS US software analyst Kash Rangan explains that the technology's ability to create new content in the form of text, image, video, audio, and code, and to do so via natural language rather than programming language, are its key transformative features.

Sarah Guo, Founder of AI-focused venture capital firm Conviction, further explains that whereas prior iterations of AI technology required humans to write deterministic code to perform specific tasks ("Software 1.0") or the painstaking collecting of training data to train a neural network for a specific task ("Software 2.0"), the now wide availability of foundational models (via open source or APIs), which have natural language capabilities, reasoning, and general knowledge of the world, has reduced the onus on companies to collect training data, ushering in an era of "Software 3.0" in which companies can leverage these "out of the box" capabilities much more easily and inexpensively to transform or enhance their businesses.

Generative AI's transformative potential has already begun to translate into reality. Developer productivity in some cases has increased ~15-20% by employing generative AI tools, says Rangan. And as their use becomes more pervasive, Guo sees a range of future applications, especially as traditional service markets, including legal, data analytics, illustration, and voice and video generation, are increasingly served by AI. GS GBM US TMT sector specialist Peter Callahan notes that public investors, for their part, believe this technology has all the makings of a platform shift, with the potential to transform almost all aspects of the enterprise and consumer experience.

According to GS senior global economist Joseph Briggs, that transformative potential could have far-reaching macro consequences. He estimates that its use could raise annual labor productivity growth by around 1.5pp over a 10-year period following widespread adoption in the US and other DM economies, and eventually raise annual global GDP by 7%. And GS US equity strategists Ryan Hammond and David Kostin argue that such a productivity lift could turn what has up to now been a relatively narrow AI-led US equity rally into a much broader one over the medium-to-longer term, boosting S&P 500 fair value by an eye-popping 9% from current levels.

But even if AI technology ultimately proves transformative, has the hype around what the technology can actually deliver—and what the market is pricing—gone too far at this point? When it comes to the intelligence of AI systems today, Gary Marcus, Professor Emeritus of Psychology and Neural Science at New

York University, believes the answer is "yes". He clarifies that the oft-touted neural networks of current AI tools function nothing like the neural networks of human brains; while AI machines can perform reflexive statistical analysis, they have little to no capacity for deliberate reasoning. And while these machines can learn, this learning largely revolves around the statistics of words and proper responses to prompts; they are not learning abstract concepts and, unlike humans, have no internal model that allows them to understand the world around them. Artificial general intelligence (AGI), Marcus says, will probably be achieved eventually, but we are very far from it today, and no amount of investment is likely to change that.

And when it comes to markets, GS market strategists Dominic Wilson and Vickie Chang point out that during past innovation-led productivity booms like those following the widespread adoption of electricity (1919-1929) and PCs and the internet (1996-2005), sharp increases in equity prices and valuations became bubbles that ultimately burst.

Even today, Guo sees some areas of mispricing in the private markets as a large cohort of investors anchors to the same investment heuristics while they gain a deeper understanding of the space. And she warns that misjudging the timing of such transformative shifts is a common pitfall in investing. That said, as an early-stage investor, she is a bit less focused on valuation than on choosing markets, products, and entrepreneurs that she believes have meaningful upside.

GS US internet analyst Eric Sheridan, for his part, is somewhat comforted by the fact that the vast majority of companies that have recently outperformed on the AI theme are still trading at relatively reasonable multiples to GAAP EPS. And Rangan makes the case that unlike other large technology cycles—such as the shift from distributed systems to cloud computing—where objections from established players slowed adoption, the most powerful technology companies in the world are driving this shift, so AI probably isn't in a hype cycle.

So where are the most compelling AI investing opportunities today? Rangan and Sheridan argue that the large tech companies developing the foundational AI models, as well as the "picks and shovels" businesses serving the space—semiconductor companies, cloud computing hyperscalers, and infrastructure companies—are well positioned to capture gains during the current "build" phase. Guo agrees, but also sees opportunities across the stack, and is most excited about the application layer that the broader investor base seems less sure about today.

Finally, what risks should investors most watch out for? Sheridan is closely monitoring the prospect of changed consumer computing habits, which could upend existing business models. Guo warns that discerning between AI marketing and AI reality amid the current enthusiasm could prove difficult. And Rangan worries that the more pervasive the technology becomes, the less valuable it might be.

Allison Nathan, Editor

Email: allison.nathan@gs.com
Tel: 212-357-7504
Goldman Sachs & Co. LLC



Interview with Sarah Guo

Sarah Guo is Founder of Conviction, an AI-focused venture capital firm. Previously, she was a General Partner at Greylock. Below, she argues that AI progress is ushering in a technological paradigm shift that presents a rich investment opportunity set, especially as software engineering undergoes a transformation to “Software 3.0” and the realm of traditional services is increasingly served by AI. But Guo also warns about the risks of investors misjudging the timetable of such a large technology shift, as well as the difficulty of distinguishing between AI marketing and AI reality.

The views stated herein are those of the interviewee and do not necessarily reflect those of Goldman Sachs.



Allison Nathan: You have been a long-time investor in artificial intelligence (AI). What’s attracted you to the space?

Sarah Guo: As a technologist, it’s hard not to be interested in AI. But from an investor perspective, I began paying attention to AI early in my time at Greylock, my former investing firm,

because machine learning (ML) drove so many of the businesses we know and love as consumers—Google, Meta, Uber, Instagram, LinkedIn, TikTok. Those are all algorithmic businesses that leverage ML for recommendations, pricing ads and services, and detecting spam and fraud, among many other applications. Seeing how important the previous generation of ML was for those businesses naturally led to exploring how these classical ML approaches could be applied in other domains. An obvious one was the cybersecurity space, with companies such as Awake, Abnormal, or Obsidian Security, because the goal there is often to find signal from noise. We also identified opportunities in areas like call centers, as well as completely new use cases such as self-driving cars and delivery robotics. And as we anticipated that companies beyond the internet giants would want to leverage these ML capabilities, that led to seeking opportunities in picks and shovels like next-gen developer tools and infrastructure.

In addition, architectural advancements in AI over the past decade in academic and industrial labs —i.e. convolutional neural networks (CNNs), generative adversarial networks (GANs), reinforcement learning (RL), etc.— were incredibly compelling. In particular, scaled-up transformers-based models have shown to be shockingly capable and very general. The research acceleration over the past five years solidified my conviction that developments in AI are ushering in a paradigm shift—certainly the biggest technology shift that I’ll see in my investing career—with the vast majority of investment opportunities still ahead of us. We’re in the first inning.

Allison Nathan: What’s differentiating about generative AI technology that’s receiving so much focus today vs. previous developments in AI?

Sarah Guo: Recent progress in AI isn’t just more of the same. These new more general and more powerful capabilities expand the relevant scope for ML and enable very different product user experiences. Prior to ML, we had “Software 1.0”—deterministic code written by humans, function by function, to perform one task at a time. In 2017, Andrej

Karpathy, the technologist that led the autopilot team at Tesla, coined the term “Software 2.0” to describe ML-driven software development, where the main work was no longer actually writing the software, but collecting training data to train a neural network for a specific task. However, the traditional ML development cycle of labeled data collection and engineering to achieve individual tasks at an acceptable level of quality is very labor-intensive and expensive, which has been an impediment to its widespread adoption.

Today, we’re entering the era of what I think of as “Software 3.0,” in which many capabilities are available “out of the box” with a foundational model either available in open source or offered via an API. These “base models” have natural language capabilities, reasoning, and general knowledge of the world. In this paradigm, companies don’t need to collect nearly as much training data, which suddenly makes the technology much more useful, accessible, and less expensive. Any company that chooses to invest in AI can now invest in adapting these models to enhance or transform their businesses.

Allison Nathan: Even if generative AI is very promising, is the current hype overstating the technology’s capabilities?

Sarah Guo: Misjudging the timetable of large technology shifts is a common pitfall in investing. I am all-in on a fundamental bet that this shift will drive substantial value creation, but this is a decade+ transition. In the meantime, areas of mispricing have certainly surfaced. In the private markets, a large cohort of investors is trying to figure out how to gain exposure to this technology, or at least how to think about the risk profile around it. And while they’re developing a deeper understanding of the space, the tendency has been to anchor to investments with more obvious heuristics. For example, many investors seem to be assessing startups based on whether the people leading them are former researchers at OpenAI or DeepMind, because that’s a much easier question to answer than whether a particular product or research thesis will be successful. Similarly, because databases are a known and well-understood category of software, vector databases are receiving substantial investor attention.

That said, I am already seeing some investors becoming more skeptical because most enterprises haven’t yet adopted generative AI, but this seems short-sighted. Remember that ChatGPT only launched in November; the average enterprise planning and execution cycle tends to be longer than six months. So, investors will need to be patient. As with the internet, mobile, and cloud, some winners emerged immediately, but others only emerged a decade later;

discovering the use cases and building great software takes time and entrepreneurial ingenuity. You wouldn't have wanted to stop your internet investing with Napster.

Allison Nathan: But do the high valuations in the space today concern you/give you pause?

Sarah Guo: The recent lessons across all stages of the technology markets should not be lost on investors. All companies are eventually valued on a multiple of cash flows. But as an early-stage investor, within some bounds, we can be less focused on valuation and more focused on choosing markets and entrepreneurs to identify breakout winners.

Allison Nathan: So where are the most compelling investment opportunities in the space today?

Sarah Guo: We're investing full-stack. First, we have picks and shovels investments; infrastructure, data infrastructure, and engineering workflows are being reimaged. Demand for Nvidia GPU capacity is insatiable, but cloud management and delivery of GPU clusters is still quite immature, and lags significantly behind CPUs. We are investing in making AI infrastructure friendlier to enterprises and application development easier.

And then there are the models themselves. Some of that business will remain centered in the very large labs: OpenAI, Google, DeepMind, Anthropic, etc. But big model opportunities remain, for example in actions/agents, image, voice, video, and robotics. We are particularly excited for the democratization of software development through better code models. Generally, open-source language models are becoming increasingly capable, and that will likely continue to be the case, in part due to contributions from large companies like Meta. So, a range of model providers will exist. Leveraging these models against company- or consumer-specific data is nontrivial, so there are large opportunities around intelligently labeling data, data management for AI applications, and better understanding and orchestrating these models.

The opportunity I'm actually most excited about is the application layer. Many investors are unsure about this layer, believing a narrative that all the value is in the model training itself, but there's a huge amount of creativity and work in getting non-deterministic models to work in production use cases. Both startups and incumbent application companies will leverage these capabilities across many areas: everything from observability to security to customer relationship management (CRM) and to markets that were traditionally services, including security services, legal, data analytics work, illustration, and voice and video generation, that now can begin to be served by more software. We're excited about the democratization effect that AI brings, and expect its second-order effects to become investable, too.

Allison Nathan: Do incumbents have an advantage here?

Sarah Guo: Within the broad opportunity set, incumbents certainly have some advantages—their distribution and data—and so there will no doubt be huge incumbent winners. But incumbent advantages aren't always as valuable as they may seem. For example, we are invested in a stealth AI security company aiming to automate a labor-intensive part of the

cybersecurity workflow in enterprises today. The company needed a certain set of training data to "fine-tune" or customize its models and initially sought to partner with an incumbent that possessed the data. But no incumbent had collected the data in a form useful for training the model. So, while some extremely valuable data is sitting at incumbents today, some just doesn't exist yet, and it will be a free-for-all to figure out how to efficiently collect it.

All told, these companies are competing on the many dimensions of building a software business, and I don't think that AI fundamentally swings in favor of incumbents or startups. My personal bias is that early-stage investing is the best place to be gaining exposure to this technology right now, partly because the space is quite young, so pure-play public opportunities don't exist just yet. But the dislocation is a huge opportunity for any investor that can distinguish signal from noise, be that in public or private markets.

Allison Nathan: What are the biggest risks to the AI investment opportunity today?

Sarah Guo: Distinguishing between AI marketing and AI reality will be hard work for investors. This is a highly technical field, and the state of the art changes every week. The rapid leadership commitment from public companies to the AI trend has been extraordinary, but painting your earnings calls and company statements with AI marketing isn't going to do much good if it doesn't translate into margin improvement, better products, and new revenue. Resourcing AI efforts, as well as dealing with the innovator's dilemma that AI automation could displace a significant amount of human work or reduce the cost of offerings, is complicated territory for public companies.

It is hard for large companies to make dramatic changes fast, but that's what this shift requires. Enterprises are needing to disrupt themselves, grapple with privacy and data use concerns, rapidly staff up AI product teams, and think creatively about pricing and packaging of new offerings. Amid such rapid change, companies cannot build everything from scratch, and choosing the right partners will be a strategic advantage.

Allison Nathan: How should investors navigate this risk?

Sarah Guo: My advice to investors is to focus on choice of technical partners, concrete plans, and outcomes. When AI products account for a significant share of incremental revenue, it's hard to argue with that performance. Or, in consumer businesses, if the metrics investors typically use to assess a company's performance—engagement, transactions, ad inventory, etc.—materially improve after introducing a new AI product, that's what you want to see.

Another significant risk is public and regulatory backlash against AI technology due to concerns around abuse of these technologies in the areas of bias, disinformation, cybersecurity, etc. Just like the internet, general tools like generative AI can be used for good and for bad, so investment in risk mitigation must occur alongside investment in innovation. But given the important advances in critical areas like science, education, and healthcare that this technology could enable, it would be a shame if we ended up regulating the industry to a halt before the technology has had a chance to really deliver on its immense potential.

Interview with Gary Marcus

Gary Marcus is Professor Emeritus of Psychology and Neural Science at New York University. He has done extensive research on artificial intelligence (AI), including in his latest book, *Rebooting AI: Building Artificial Intelligence We Can Trust*. Below, he argues that the intelligence of AI systems is being overhyped and, while we could get there eventually, we are currently nowhere near achieving artificial general intelligence (AGI).

The views stated herein are those of the interviewee and do not necessarily reflect those of Goldman Sachs.



Jenny Grimberg: How do generative artificial intelligence (AI) tools actually work today?

Gary Marcus: At the core of all current generative AI tools is basically an autocomplete function that has been trained on a substantial portion of the internet. These tools possess no understanding of the world, so

they've been known to hallucinate, or make up false statements. The tools excel at largely predictable tasks like writing code, but not at, for example, providing accurate medical information or diagnoses, which autocomplete isn't sophisticated enough to do.

Jenny Grimberg: Some observers have argued that these technologies can learn/understand because they employ similar neural networks to those in the human brain. What's your view?

Gary Marcus: The neural network of a human brain works nothing like the neural networks AI tools use. And, contrary to what some may argue, these tools don't reason anything like humans. At most, AI machines do some of what the Nobel laureate Daniel Kahneman characterizes as system 1 thinking—reflexive statistical analysis—and very little system 2 thinking of deliberate reasoning. AI machines are learning, but much of what they learn is the statistics of words, and, with reinforcement learning, how to properly respond to certain prompts. They're not learning abstract concepts.

That's why much of the content they produce is garbage and/or false. Humans have an internal model of the world that allows them to understand each other and their environments. AI systems have no such model and no curiosity about the world. They learn what words tend to follow other words in certain contexts, but human beings learn much more just in the course of interacting with each other and with the world around them.

Jenny Grimberg: So, is the hype around generative AI overblown?

Gary Marcus: Yes and no. Generative AI tools are no doubt materially impacting our lives right now, both positively and negatively. They're generating some quality content, but also misinformation, which, for example, could have significant adverse consequences for the 2024 US presidential election. But the intelligence of AI systems is being overhyped. A few weeks ago, it was claimed that OpenAI's GPT-4 large language model (LLM) passed the undergraduate exams in engineering and computer science at MIT, which stirred up a lot of excitement. But it turned out that the methodology was flawed, and in fact my long-time collaborator Ernie Davis pointed that

out around a year ago, yet people still proceeded to use it. Narratives about how current AI systems will displace scores of workers, and some people are worried that robots will soon take over the world. But AI isn't nearly smart enough for that in its current iteration. Four years ago, I joked that, should you ever find yourself in a situation where the robots are coming for you, just close the door. And it's still true that robots can't open doors, nor can they drive cars reliably. We are nowhere near achieving artificial general intelligence (AGI). Those who believe AGI is imminent are almost certainly wrong.

“Those who believe artificial general intelligence (AGI) is imminent are almost certainly wrong.”

Jenny Grimberg: Has something gone wrong on a conceptual/technical level in AI that the technology is so far from general intelligence?

Gary Marcus: To a degree. Something has gone wrong from a sociological standpoint. Giant approximation machines, which are essentially what LLMs are, are relatively easy to build and monetize, so people are focusing on them rather than on other ideas that may have more merit but are harder to implement and monetize quickly. So, the dynamics of capitalism certainly aren't helping, and have probably slowed technological progress relative to what's theoretically possible. That said, the problem of intelligence is an extremely difficult one, and most of the efforts to work on it in a computational context are less than 75 years old, which isn't very long for the development of a science. People often talk about intelligence as if it's a magic number, like an IQ score. But intelligence is comprised of many aspects: being able to follow a conversation, fix a car, learn a new dance move, or perform just about any kind of interesting human activity requires intelligence of many different sorts. And expecting a machine to master all of those in just 75 years probably isn't realistic.

Jenny Grimberg: Is it possible to ever develop truly intelligent artificial systems?

Gary Marcus: I believe so. I think of the current stage of AI as akin to the age of alchemy, during which people knew they could get something to happen, but didn't yet have a theory of chemistry. People today can conceptualize what AGI might look like, but have yet to develop a sophisticated enough understanding about how to build intelligence into machines. I see no reason to think we won't get there eventually. Some people argue that intelligence simply isn't something that can be built into machines, but I don't buy into that view. Even if

machines are never able to feel pain, for example, they may be able to understand what a person is feeling when they're in pain and what they might do as a consequence of feeling that pain, such as taking medicine or going to the doctor. In that vein, machines will probably eventually develop a much clearer understanding of human beings and become much more reliable and truthful. The question is when. I'm often cast as the pessimist, but several months ago I had a debate with prominent software architect Grady Booch in which he took the pessimistic position that AGI won't happen in our lifetime, the lifetime of our children, or even the lifetime of our children's children, whereas I took the optimistic position that AGI would be achieved sometime in this century. That said, it will likely take another several decades given the current state of AI and how much work is left to do.

Jenny Grimberg: Couldn't the inflection point for this technology potentially occur much sooner given the significant amounts of money large companies are pouring into AI research/development?

Gary Marcus: Not necessarily; throwing massive amounts of money at a problem doesn't mean it will be solved. In 2016, I warned that driverless cars were being overhyped, in the sense that solving them would prove much more difficult than many thought. The key issue is outliers. Driverless car systems basically work through memorization, so when they encounter a new situation, they're often at a loss for what to do. A good example of this, courtesy of Wired's Steven Levy, is what happened at Google's automatic car factory in 2015—the cars had just learned to recognize that it was acceptable to drive over piles of leaves on the road, because that particular situation wasn't in their training set. Since then, a hundred billion dollars have been invested in driverless cars. And yet, in April 2022, a Tesla "summoned" across a parking lot at an airplane trade show ran directly into a \$3.5mn jet, because, again, its training set didn't contain jets, and it had no abstract understanding that it shouldn't drive into large, expensive objects. And so, it did, none the wiser. That should serve as a stark reminder that just because the money is there, doesn't mean the results will be.

Jenny Grimberg: What, then, is required from a technical/policy/societal standpoint for intelligent artificial systems to become a reality?

Gary Marcus: The proper incentives have to be developed and funding needs to be allocated in the right directions. Attitudes and mindsets will also have to change. People in the machine learning community are overconfident. They're convinced that they've discovered the one true way to develop intelligent systems, and aren't very receptive to advice from field practitioners in cognitive science, psychology, or linguistics. History has shown that scientists and engineers can become fixated on ideas that ultimately don't work, which significantly slows down progress. In the early 20th century, scientists endeavored to identify what genes were made of. Gregor Mendel had proven that a biological basis for heredity existed, and scientists were convinced that basis was proteins, so they spent decades trying to identify which proteins. That was the wrong question; they instead should've asked, what biological *thing* are genes made of, which turned out to be DNA. Once

Oswald Avery figured that out, the field progressed very quickly. The field of AI is very similar. People are currently dogmatically pursuing the idea that LLMs are the answer to achieving AGI. I consider them to be a frustrating distraction—LLMs may be *part* of the answer, but they are almost certainly not the *whole* answer. So, the machine learning community will have to reorient at some point. I expect that such a reorientation will eventually occur and the machine learning community will find the right answer, at which point progress towards AGI will happen very quickly.

“Large language models may be *part* of the answer, but they are almost certainly not the *whole* answer [to achieving AGI].”

Jenny Grimberg: Given all that, what's your main message to investors interested in the AI space?

Gary Marcus: Be wary of the hype—AI is not yet as magical as many people think. I wouldn't go so far as to say that it's too early to invest in AI; some investments in companies with smart founding teams that have a good understanding of product market fit will likely succeed. But there will be a lot of losers. So, investors need to do their homework and perform careful due diligence on any potential investment. It's easy for a company to claim that they're an AI company, but do they have a moat around them? Do they have a technical or data advantage that makes them likely to succeed? Those are important questions for investors to be asking.

Jenny Grimberg: What concerns you most about AI today?

Gary Marcus: I'm concerned that we're giving an enormous amount of power and authority to the small number of companies that currently control AI systems, and in subtle ways that we may not even be aware of. The data on which LLMs are trained can have bias effects on the model output, which is disquieting given that these systems are starting to shape our beliefs. Another concern is around the truthfulness of AI systems—as mentioned, they've been known to hallucinate. Bad actors can use these systems for deliberate abuse, from spreading harmful medical misinformation to disrupting elections, which could gravely threaten society.

I've raised these concerns with a significant number of government officials around the world. Almost all of them agree that something must urgently be done, but nobody is entirely sure what that something should be. I believe we need to establish a global agency for AI, with buy-in from national governments, large technology companies, non-profits, academia, and society at large, to collaboratively find governance solutions and vet new technologies before they're deployed at scale, akin to what we have in medicine. Fortunately, this seems to be where the world is headed. Several government leaders and heads of large tech companies have recently argued for this. It's hard to predict how this will all play out from here, but this is an important start on the road towards safe, secure, and peaceful AI technologies.

Artificial intelligence, explained

Artificial intelligence (AI) is the science of creating intelligent machines. AI is a broad concept that encompasses several different subfields, including machine learning, natural language processing, neural networks, and deep learning.

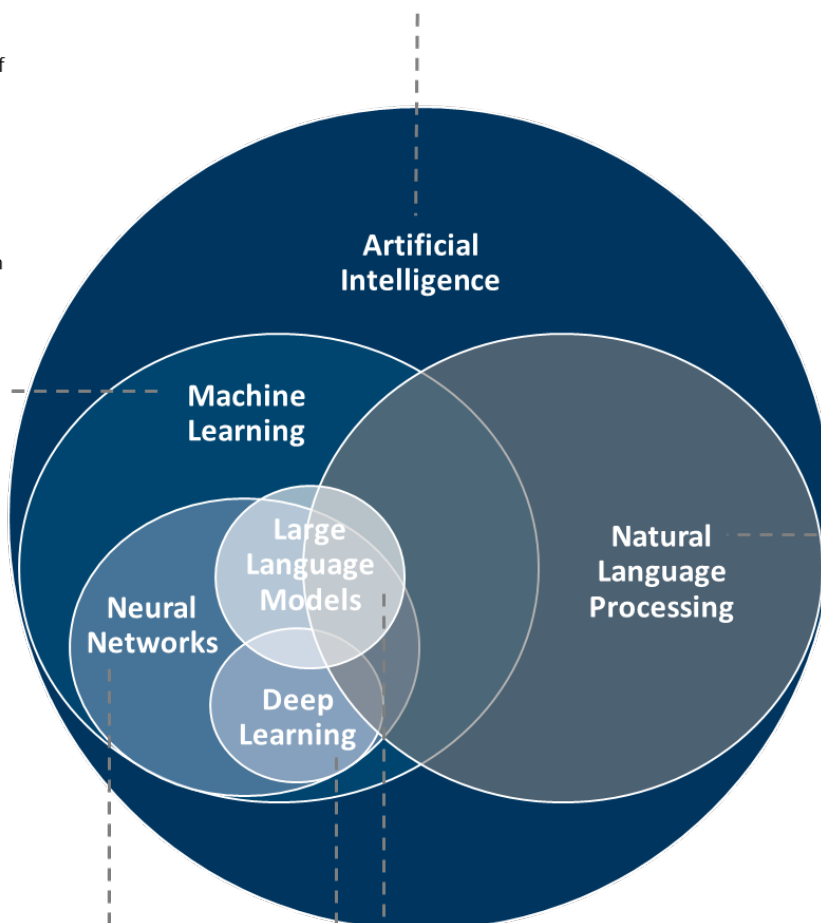
Artificial intelligence (AI) is the broadest term used to classify machines that mimic human intelligence. There are three main categories of AI: artificial narrow intelligence, artificial general intelligence, and artificial super intelligence. Artificial narrow intelligence is considered “weak AI”, which is trained to perform specific tasks, like voice or image recognition. Artificial general and super intelligence are considered “strong AI”, which has cognitive abilities equal to/greater than those of humans. No practical examples of strong AI exist today.

Generative AI is a type of AI system that generates text, images, and other content in response to natural language prompts.

Machine learning is a subfield of AI that focuses on developing models and algorithms to help computers improve their performance through experience. Large amounts of data are fed into a computer, which then discovers patterns in that data and uses it to make predictions and decisions. There are three major types of machine learning models: (1) supervised machine learning, which uses labeled datasets to train models, (2) unsupervised machine learning, which uses algorithms to analyze unlabeled datasets, and (3) semi-supervised machine learning algorithms, which sits between the first two.

Transformer models are machine learning models designed to process sequences of elements. The premise of the model is an attention mechanism, which enables the model to learn and understand the relationship between words in a sentence.

Neural networks are a subfield of machine learning. They are mathematical models inspired by the human brain structure. Each neuron, or node, of the network takes an input, performs a computation, and creates an output. If the output of any individual node is above a specified threshold value, the node is activated and sends data to the next layer of the network. One of the best-known neural networks is Google’s search algorithm.



Natural Language Processing (NLP) is a subfield of AI focused on giving computers the ability to understand text and spoken word in a similar manner to human beings. NLP uses computational linguistics combined with statistical, machine learning, and deep learning models to enable computers to understand human language. It does so using two techniques: (1) syntactic analysis, which identifies the structure and relationship between words in sentences, and (2) semantic analysis, which focuses on the meaning of the words and their context in the sentence. Google Translate is one example of NLP technology in the real world; chatbots like Siri and Alexa also rely on NLP.

Large language models (LLM) are a type of machine learning model that are trained on large amounts of unlabeled data using self-supervised learning or semi-supervised learning to perform NLP tasks. LLMs use deep neural networks to generate outputs. ChatGPT is the most well-known example of an LLM.

Deep learning is neural networks with three or more layers. Deep learning differs from “classical” machine learning by the type of data it works with and methods by which it learns. While machine learning algorithms leverage more structured, labeled data to make predictions, deep learning doesn’t necessarily require a labeled dataset, and it’s less dependent on human interaction. Deep learning is used in many applications, such as speech recognition and autonomous driving.

Source: IBM, Goldman Sachs GIR.

Special thanks to GS equity research analysts for graphic. Original version published in *Americas Technology: Generative AI – Part I: Laying out the investment framework*.

A discussion on artificial intelligence

Kash Rangan and Eric Sheridan are Senior Equity Research Analysts at Goldman Sachs covering the US software and internet sectors, respectively. Below, they discuss the recent rise of generative AI, which companies and sectors stand to benefit, and what that means for investors.



Allison Nathan: The recent emergence of generative artificial intelligence (AI) has attracted a lot of focus. Why is there so much excitement around this technology, especially when AI has been around for a while?

Kash Rangan: Generative AI differs from traditional AI in two main ways. One, it's capable of *generating* new content in the form of text, image, video, audio, and code, whereas traditional AI systems train computers to make predictions about human behavior, business outcomes, etc. And two, it allows humans to communicate with a computer in their natural language, which has never been done before; traditionally, computers were

prompted using programming languages. The implications for personal and professional productivity from the advent of generative AI are tremendous—if a computer can generate quality content, people can spend the time saved on higher value-added activities.

Eric Sheridan: People are very focused on AI right now because the consumer and enterprise imagination has caught up with the technology. Alphabet first described itself as an AI-first company at its 2017 developer conference, and AI has been embedded in most everyday products, such as search algorithms and recommendation engines, for some time. But the generative AI tool ChatGPT has captured people's imagination, just as the iPhone did when it was introduced even though smartphones already existed, allowing it to scale very quickly. Those are moments of "unlock". While it took several years for the iPhone to become a consumer-adopted device at scale, ChatGPT was the fastest application to reach 200mn monthly active users (MAUs) that we've ever tracked. So, AI's unlock moment has happened, and in a more intense way than we've ever seen in the past.

Allison Nathan: Is all this hype around generative AI warranted, or overblown? What differentiates it from the hype around previous technologies that seem to have fizzled out and/or were slower to take hold than expected?

Kash Rangan: AI probably isn't in a hype cycle. For one thing, this technology cycle isn't being led by upstarts, which makes it less likely to fizzle out or take a long time to get going. The shifts from mainframe to distributed systems in the early 1990s and from distributed to cloud computing in the early 2000s took a longer time to take off than many people expected as larger, established companies were critical voices against them. IBM argued in favor of mainframe systems against the distributed architecture of Oracle, a relatively small company at the time.

And existing on-premises systems and technology providers argued against the shift from distributed to cloud computing, warning that the cloud wasn't safe or economical, didn't scale well, etc. It took years for those objections to be overcome and for cloud to find its footing. Only when the larger, established companies had operating clouds was there a harmony of voices telling buyers that the technology was acceptable.

In contrast, driving the AI tech cycle are some of the most powerful technology companies in the world, who are building the foundational models at the heart of generative AI. When a unanimous verdict exists among the technology providers that a technological shift is actually happening, it's real. And when customers start to become interested, it's not hype. And customers are interested. We're having discussions with the CIOs of global corporations who are amazed at the productivity benefits this technology could bring if deployed internally. And all of this is occurring at a time when the market is rewarding productivity gains. So, this doesn't feel like a hype cycle.

Allison Nathan: So, this isn't a bubble?

Eric Sheridan: While you never know you're in a bubble until it pops, the vast majority of the companies that have outperformed the broader market over the last several months on the AI theme are still trading at relatively reasonable multiples to GAAP EPS. Bubbles are typically about enterprise value to eyeballs/clicks, addressable market dynamics, or sheer euphoria as a driver of valuations as opposed to what the right multiple on net income is to pay. So, this feels very different from previous tech bubbles.

Allison Nathan: Even if the hype isn't overblown, how long could it take for this technology to really have an impact on companies, workers, and consumers?

Eric Sheridan: Following the introduction of ChatGPT (OpenAI) and Bard (Alphabet), consumer internet companies are now moving into the build phase in which they're building foundational models, some of them for specific businesses/industries. Once built, some of these solutions will need to be deployed in the real world to see what works, scales, and gains adoption. A good analogy for thinking about the timeframe of impact is the shift from desktop to mobile computing. It took companies like Alphabet and Meta four years after the iPhone's introduction to begin referring to themselves as mobile-first companies—that's how long it took for the infrastructure to be built out, at which point companies could start talking about application disruption. So, over the short-term—the next 6-12m—many companies are viewing AI technology as potentially productivity-enhancing internally, and they're building, testing, and learning to see how the technology may be additive to their businesses externally on more of a three-year horizon.

Kash Rangan: Productivity gains right now are concentrated at the developer level. Anecdotally, we've seen 15-20% boosts to

developer productivity through the automation of some of the manual and rote process of writing code, with the software behind that starting at around \$10/developer/month. From a cost-benefit standpoint, that's incredible. The next population of workers that will be testing this technology will be those in sales, marketing, and customer support, which combined account for about a third of the professional working population in developed market economies. Software companies are readying products that should have a positive impact on these populations in CY2024.

Eric Sheridan: I'll add that the impact on consumers will probably take a bit longer to manifest because it typically takes multiple years to change ingrained search, shopping, consumption, etc. behavior in a big way. For as much as ChatGPT has scaled to 200mn MAUs faster than anything we've ever tracked, the volume of traditional search queries spread across Google and Bing, for example, has yet to be impacted. So, we're still quite a while away from AI technology having real effects on the consumer application front.

Allison Nathan: The build phase sounds like a lot of money in, little money out. When can we expect companies to monetize the gains from generative AI technology?

Eric Sheridan: The build phase doesn't mean that nobody is garnering revenue. The companies selling the tools necessary to build out AI technology certainly are; Nvidia recently issued revenue guidance of \$11bn in the second quarter of FY2024, which led to a dramatic rise in its stock price. The timeline for when other companies will monetize gains varies by the type of company. Monetization for consumer internet companies typically only occurs once they have scale of users and of deployment of purchasing power. A good analogy is the app store. For Apple's App Store and Google's Play Store, the two application distribution mechanisms in the duality that is the mobile OS layer in computing, the first three years were spent building to scale in the application developer community. During that time, hundreds of millions of users weren't playing games monetizing at a couple dollars a month.

Kash Rangan: Monetization could happen sooner for software companies. Monetization for software companies will come from being able to steer their customers towards premium SKUs of their products that exclusively have generative AI capabilities. At the very outset, companies may charge customers a low monthly fee to play around with the technology, thereby lowering the barriers to adoption. Once these products evolve and become more sophisticated, companies will have much more pricing power.

It's important to remember that we're only around five months into generative AI; the first five months of cloud computing looked like nothing. So, the ability to monetize will grow. Software companies that not only have big distributions and large customer bases, but also troves of data to train large language models (LLMs) on, are in a particularly strong position to be effective in monetizing the gains from AI through differentiated SKUs and higher revenue per user (RPU) for existing products, especially since the cost structure is already paid for in terms of distribution and product development.

Hyperscalers—which are large cloud computing companies that provide computing and storage services at scale—could also garner revenue relatively soon. Cloud computing provides the computing resources and infrastructure needed to deploy AI at scale. So, AI-driven initiatives could lead enterprise clients to increase cloud computing spend following a year in which most of them optimized their spend due to fears about the economy, and that could create positive revenue trajectories for the hyperscalers exiting 2023 and into 2024. Cloud hyperscalers are also well-positioned in the sense that they've spent 10 years and hundreds of billions of dollars building out cloud infrastructure, allowing generative AI to build on top of a solid base of cloud revenue. So, they're not starting from ground zero; they'll be leveraging a lot of what's already been put in place to augment generative AI capabilities on top of the biggest investment cycle the technology industry has ever witnessed.

All that said, gross margins from generative AI will likely be negative for the foreseeable future as capex growth exceeds revenue growth. Right now, Microsoft, Alphabet, and Amazon are spending more than \$100bn in capex, of which a large portion is on cloud computing and AI, with generative AI likely composing the fastest growing category. But generative AI's revenue contribution to the cloud industry is currently de minimis, which is typical at the front end of a large capex cycle. It took a 10-year investment cycle for gross margins in Microsoft's cloud business to go from negative to where they are today, which is well north of 50-60%. That's probably the baseline for AI.

Allison Nathan: What types of companies are best positioned to capture the gains from generative AI?

Eric Sheridan: The handful of large tech companies developing the foundational models for generative AI are clearly well positioned. Semiconductor companies and the hyperscalers in cloud computing also look well-positioned to capture gains during the build phase.

Kash Rangan: Legacy software companies with subscription business models are also poised to benefit incrementally from what we estimate will be a generative AI software total addressable market (TAM) of ~\$150bn. And infrastructure software companies stand to benefit. Running generative AI at scale will require significant compute power and data storage. When cloud computing emerged on the scene, some argued that data centers and IT workers would no longer be necessary because the cloud would do everything. But as thousands of cloud applications cropped up, the infrastructure required expanded dramatically, not to mention that data centers needed to be able to talk to the cloud, which added further complexity to the process. That's why IT spend as a percentage of capex and total revenue continues to rise. Generative AI is being thrown on top of the existing cloud architecture, and it needs to be able to talk to cloud applications, which exponentially increases the complexity of interactions and opportunity for something to break. So, infrastructure software companies that provide diagnostics, tooling, measurement, feedback, and stabilization will probably have a field day. Some of the companies that will "win" in this

space are public companies, and some of them don't even exist yet.

Allison Nathan: Will there be space for new/small companies?

Kash Rangan: When ChatGPT emerged, some venture capitalists believed that it would disrupt every company. Now, they generally agree that companies like Microsoft, Adobe, Salesforce, etc. won't be disrupted, because they have scale in engineering talent and capital, and troves of data to dominate the foundational model layer. So, the consensus seems to be that new entrants won't disrupt the foundational layer. That's not surprising; the history of technological shifts shows that typically only a handful of scaled winners garner the vast majority of the profit pool at the operating system layer, and that will likely also be the case for AI.

But the AI space more broadly won't be the land of just the giants. The application layer will be wide open for innovation. Use cases will be invented for AI technology that nobody has thought of yet. Nobody predicted that thousands of Software-as-a-Service (SaaS) companies would grow out of cloud computing or that Uber would grow out of a mobile-first world.

Eric Sheridan: Applications built on generative AI capabilities that disrupt the healthcare, education, legal, etc. industries haven't been created yet, but are being widely discussed as potential possibilities. In every computing cycle, interesting application developers sprout up; I see no reason to believe that this time will be different. And if previous venture capital cycles have taught us anything, it's that multiple new firms will attempt to create disruptive applications, and a few of them will likely succeed.

Allison Nathan: What's the current state of AI regulation, and how might it evolve?

Eric Sheridan: Typically, the regulatory curve is, at a minimum, half a decade behind the innovation curve. The AI regulation curve, by contrast, is operating almost in parallel to the innovation curve. The sheer scale of potential job displacement associated with AI technology and the tail risk of some apocalyptic outcome is sounding a lot of alarm bells for politicians and regulators, who, for most of the last decade, were also behind the curve on elements of internet data collection, privacy, and information dissemination. That has brought the regulatory forces to bear quickly this time around.

The tech companies in our coverage universe and those in the private domain want more regulation partly because they don't want to be blamed if AI technology produces a bad outcome that could be attributed to them, so they want to work in partnership with regulators to create guardrails around this technology almost simultaneously with it being innovated. A rather cynical view also exists that big tech companies are advocating for guardrails around the tech because those

guardrails will create moats around these large companies, making it difficult for smaller/new companies to disrupt them and win in this space. Regardless of the reason, regulation is at the forefront much earlier in this tech cycle than any other we've ever witnessed.

Allison Nathan: Does the regulation running almost in parallel to the innovation increase the risks to investing in the AI space, in the sense that the regulation could stifle the innovation?

Eric Sheridan: Regulation isn't a risk. Regulation typically changes the way capital is allocated and the unit economics of an industry. Regulated industries usually have lower profit margins, but the barriers to entry are higher given the cost of complying with regulation. So, while investors may have to accept lower profits, they also don't have to worry as much that the companies they're investing in will be disrupted by new entrants. Regulation can definitely slow innovation down. But it's also necessary to keep bad actors out. And ultimately, a space rife with bad actors, especially AI where the consequences of bad behavior could be severe, isn't beneficial to any investor.

Allison Nathan: What risks, then, should investors be concerned about?

Eric Sheridan: The potential for changed computing habits is a risk both if it doesn't, and does, happen. We've lived through multiple cycles of people arguing that a certain thing is going to disrupt the search engine. Mobile was supposed to do it, social media was supposed to do it, and even Amazon vs. Google was a debate among investors several years ago in terms of whether the Amazon search box would be the end of the search engine, yet it lives on. And if AI is the thing that finally does disrupt the search engine, that could have significant consequences for existing business models. Whole industries are built on certain elements of aggregated supply and generated demand, and if consumer behavior shifts away from the search engine because of AI, or the search engine needs to change because of it, that could lead to wildly different economic outcomes than what investors are used to. So those are both significant risks to watch for.

Kash Rangan: The investor risk I worry the most about is that generative AI technology becomes so ubiquitous that it becomes commoditized. And if it's not special, how can companies charge a premium for it or monetize it? Today, the technology isn't ubiquitous because the expertise needed to train LLMs is scarce. LLMs require some supervision, yet very few computer scientists are currently specialized in generative AI, which limits how quickly LLMs can learn. And they do indeed learn, because LLMs are neural networks, which are modeled after the human brain. But if LLMs end up learning very quickly, the technology could become widely diffused. And at that point, the technology may no longer look valuable.

A snapshot of generative AI tools

Category	Tool	Description/Features
Chatbots	Bard	A generative AI chatbot developed by Google, initially based on its LaMDA model
	Bing Chat	A chatbot powered by Microsoft Bing
	Character.AI	Simulates conversations with real and fictional characters
	ChatGPT	A generative AI chatbot developed by OpenAI
Text Generators	Copy.ai	Generates blog posts, social media posts, and emails
	Frase.io	Produces slogans, summaries, introductions, articles, titles, and product descriptions
	Jasper	Provides users with content templates and enables collaboration
	Peppertype.ai	Offers ready-made templates for creating meta descriptions, articles, and emails, enabling commercial use of the produced content
	Rytr	Creates titles for SEO optimization, produces blog posts, articles, emails, and social media ads
Code Generators	K-Explorer	Makes code completion and custom model suggestions
	PyCharm	Provides users with code completion, highlights errors, and enables automated refactoring in Python
	Tabnine	Provides users with whole-line code completion and learns coding patterns
Image Generators	Artbreeder	Creates collages and generates images with the option of manipulating a subject's age, gender, etc.
	Craiyon	Converts text-to-image (not suitable for creating larger images)
	DALL-E	Creates, edits, and varies images, offering the commercial rights to created content
	NightCafe	Art generation with different styles and resolution options
	starryai	Enables the creation of artwork with different style options, aspect ratios, etc., giving full ownership of produced content
Video Generators	Elai.io	Allows for the conversion of text to video, offering 25+ different avatars
	Flexclip	Supports video creation and offers editing tools (adding transitions and filters, removing backgrounds)
	Lumen5	Offers templates to create original videos based on text, articles, and blog posts
	Synthesia	Enables text-to-video conversion, providing 70+ different avatars
	Veed.io	Video generation and editing, adding subtitles, removing background noise, and resizing videos
Design Generators	Colormind	Creates color palettes based on movie scenes, artwork, and other images
	Designs.ai	Generates logos and banners, provides design templates, and enables the export of produced content to different formats
	Khroma	Creates custom color palettes
	Uizard	Creates designs for mobile applications, websites, and landing pages based on sketches
Voice Generators	Lovo.ai	Enables text-to-speech conversion and generates realistic AI voiceovers
	Murf	Creates voiceovers for different contexts, enables adding punctuation, and provides the commercial rights to the content
	Play.ht	Provides AI-generated voices for various commercial purposes in 140+ languages
	Replica	Enables text-to-speech conversion, offering AI-generated voices
	Speechify	Allows text-to-speech conversion while enabling the adjustment of reading speed
Music Generators	AIVA	Creates authentic music based on preferred style, granting copyright of the produced content
	Amper AI	Produces royalty-free music based on preferred genre, length, instruments, providing perpetual license
	Evoke	Generates AI-generated and royalty-free music collection
	Jukebox	Creates authentic music with AI-generated lyrics, providing users with different genre options
	Soundraw	Enables original music creation and commercial use of the produced content

Note: Table does not constitute an exhaustive list of all existing generative AI tools.

Source: AIMultiple, tool websites, Goldman Sachs GIR.

Special thanks to GS equity research analysts for this table. Original version published in *Americas Technology: Generative AI – Part I: Laying out the investment framework*.

AI's potentially large economic impacts

Joseph Briggs finds that widespread adoption of generative AI could potentially significantly boost global productivity and GDP

The recent emergence of generative artificial intelligence (AI) raises the question of whether we are on the brink of a rapid acceleration in task automation that will significantly save time and labor costs, lead to a productivity burst, and increase the pace of economic growth. Although significant uncertainty still exists around the capabilities and adoption timeline of current generative AI models, we find that generative AI could potentially raise annual labor productivity growth by around 1.5pp over a 10-year period following widespread adoption in the US and other DM economies, and eventually raise annual global GDP by 7%.

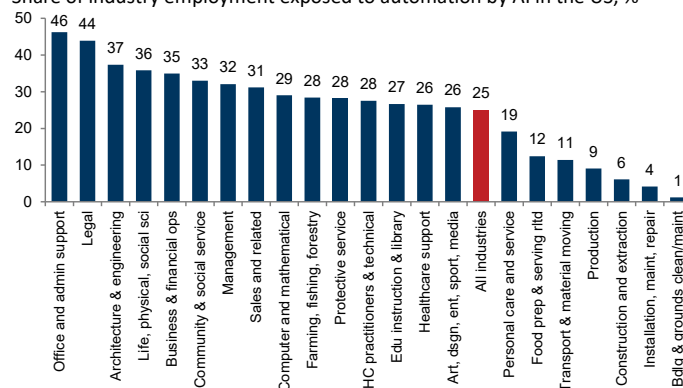
The labor market impact of generative AI

The largest effects of generative AI on the economy will likely come from its impact on the labor market. To assess the potential size of this impact, we use data on the specific work tasks that are undertaken in a typical work week for over 900 occupations in the US and 2000 occupations in the Euro area. These data contain measures on the importance and difficulty of various tasks associated with each occupation, which we combine to estimate the share of total work exposed to labor-saving automation by AI. Specifically, we select work activities that are most exposed to AI automation based on our review of probable use cases of generative AI and assume that AI will ultimately be capable of completing moderately difficult tasks (up to a difficulty level of 4 on a 7-point task complexity scale). We then take an importance- and complexity-weighted average of essential work tasks for each occupation to estimate the share of its total workload that AI could potentially replace.

Our key finding is that a lot of workers spend a lot of time performing tasks that AI models are well-suited to automate. In particular, we estimate that roughly two-thirds of US occupations are exposed to at least some degree of automation by AI, and that of those occupations which are exposed, most have a significant—albeit partial—share of their workload (25-50%) that can potentially be replaced. After weighting our occupation-level estimates by the employment share of each occupation in the US, we estimate that a quarter of current work tasks could be automated by AI, with particularly high exposures in administrative (46%) and legal (44%) professions and low exposures in physically-intensive professions such as construction (6%) and maintenance (4%).

Using European data, we estimate that a similar 24% of work tasks in the Euro area could potentially be automated by AI. Although detailed work task data are not available for other countries/regions, reweighting our industry-level AI exposure estimates by country-specific industry-employment shares suggests that generative AI could eventually automate around 18% of global work, with larger shares in DMs than EMs.

A quarter of work tasks in the US could be automated by AI
Share of industry employment exposed to automation by AI in the US, %



Source: Goldman Sachs GIR.

The productivity boost from generative AI

The large share of employment exposed to automation from generative AI raises the potential for a boom in labor productivity that significantly increases global growth. We see two channels through which AI-driven automation could raise global productivity and GDP.

First, most workers are employed in occupations that are partially exposed to AI automation and, following AI adoption, will likely apply at least some of their freed-up capacity towards productive activities. This dynamic is observable at firms that have already adopted AI, with studies¹ generally finding that AI adoption led to a 2-3pp annual boost to labor productivity growth for several years afterwards.

Second, while AI technology will inevitably displace some workers, we anticipate that most displaced workers will eventually become reemployed in new occupations that emerge either directly from AI adoption or in response to the higher levels of aggregate and labor demand generated by the productivity boost from non-displaced workers.

The reemployment of displaced workers due to the direct and indirect effects of technological change has plenty of historical precedent. Information technology, for example, displaced some workers in the early 2000s, but also directly led to the creation of new occupations like webpage designers, software developers, and digital marketing professionals, and indirectly increased labor demand in service industries such as healthcare, education, and food services.

The positive employment effects of technological change are especially clear over longer time horizons. 60% of workers today are employed in occupations that did not exist in 1940, implying that over 85% of employment growth in the last 80 years can be explained by the technology-driven creation of new positions.

To estimate how these channels might together raise US productivity growth, we combine estimates of the productivity boost for non-displaced workers, the labor cost savings of displaced workers, and a composition effect from the reemployment of displaced workers in new positions. In particular, we assume that around 7% of workers are fully

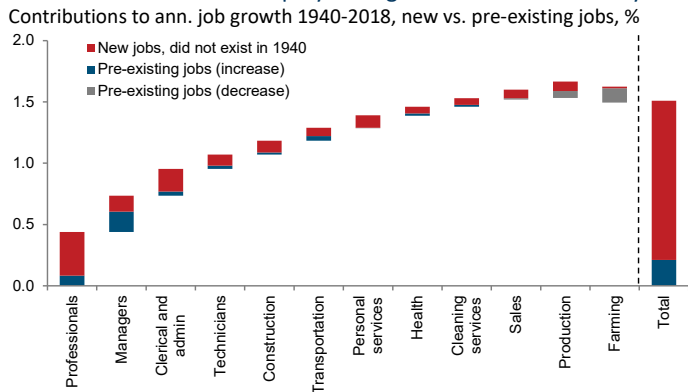
¹ Alerducci et al. (2022), Czamitzki, Fernandez, and Rammer (2022), Behrens and Trunschke (2020), Acemoglu et al. (2022), Bessen and Righi (2019).

displaced (given our estimate that automatable tasks account for the majority of the work of 7% of US workers), but that most are able to secure new employment in only slightly less productive positions, and that partially exposed workers experience a boost in productivity consistent with estimates from existing studies (those mentioned in the footnote on the previous page). For illustrative purposes, we assume that the full productivity boost of generative AI is realized over a 10-year period (but not necessarily the *next* 10 years) that starts when a large share of businesses has adopted generative AI.

Under these assumptions, we estimate that widespread adoption of generative AI could raise overall labor productivity growth in the US by around 1.5pp annually. A boost of this size would roughly double the recent pace of US productivity growth, and would be about the same size as the boost that followed the emergence of prior transformative technologies like the electric motor and personal computer.

Generative AI could also raise productivity growth outside of the US. Assuming that differences in the industry-composition of labor can account for most of the differences in the impact of AI on labor productivity growth, we estimate similarly sized boosts to productivity in other DM economies, and that that globally widespread AI adoption could boost global annual productivity growth for countries in our coverage by over 1pp annually (FX-weighted average), although the impact would likely be delayed in EM economies.


Tech innovation has led to the creation of new occupations that account for the bulk of employment growth over the last 80 years



Source: Autor et al. (2022), Goldman Sachs GIR.

Widespread AI adoption could boost global annual productivity growth by over 1pp over a 10-year period

Effect of AI adoption on ann. productivity growth, 10yr adoption horizon, pp




Source: Goldman Sachs GIR.

Large, but highly uncertain, impacts

While our estimated boost to productivity growth from widespread AI adoption is quite large, it is also highly uncertain, and will ultimately hinge on the difficulty level of tasks generative AI can perform, how many jobs are automated, and the speed of adoption. Varying our assumptions around each of these factors suggests that the boost to annual US productivity growth could range from 0.3-3.0pp, although in most scenarios the boost would remain economically significant.

The size of the productivity boost will ultimately depend on AI’s capabilities and adoption timeline

Effect of AI adoption on ann. labor productivity growth, 10yr adoption period, pp



Note: A much less powerful AI scenario is where, for example, generative AI can only “skim a short article to gather the main point” (difficulty score 2) rather than “determine the interest cost to finance a new building” (difficulty score 4). A much more powerful AI scenario is where, for example, generative AI can “analyze the cost of medical care services for all US hospitals” (difficulty score 6). Source: Goldman Sachs GIR.

We also see the timing of any macroeconomic impact from generative AI as particularly hard to predict based on the evidence from past technological breakthroughs. The burst in productivity due to the electric motor and personal computer, for example, occurred around 20 years after the key technological breakthrough, at a point when roughly half of US businesses had adopted the technology.

It is possible that the surge in interest in generative AI could speed up its adoption and lead macroeconomic impacts to materialize sooner. However, AI adoption rates by US firms were only 3.2% in 2019, and though many major companies are currently exploring how to use AI, only ~20% of CEOs expect that generative AI will lower labor needs in the next 1-3 years. For example, companies still need to navigate several barriers to adoption like data privacy before most start incorporating generative AI in their everyday workflows. We therefore suspect that the effect of generative AI will probably not be visible in aggregate productivity data for at least several more years.

Nevertheless, the significant work task exposure to AI automation, combined with our sizable estimates of potential productivity increases, highlight the enormous economic potential of generative AI if it does deliver on its promise. Indeed, applying the estimated productivity boost to countries in our coverage, we find that widespread AI adoption could eventually drive a 7%, or almost \$7tn, increase in annual global GDP over a 10-year period, and therefore view generative AI as a significant upside risk to our medium- and longer-run global economic growth projections.

Joseph Briggs, Senior Global Economist

Email: joseph.briggs@gs.com
Tel: 212-902-2163

Goldman Sachs & Co. LLC

US equities: gauging the AI upside

Ryan Hammond and David Kostin argue that potential AI-related productivity boosts could lead to more upside for US equities

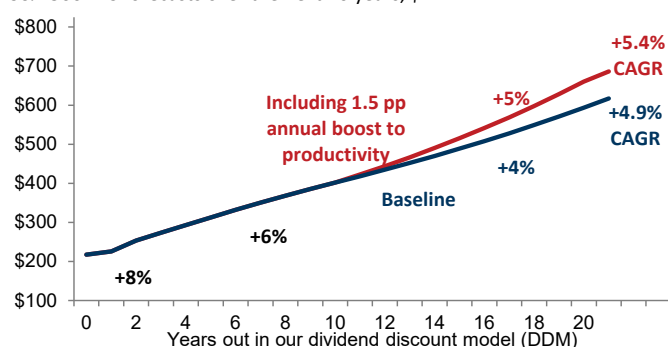
With a surge in focus on generative artificial intelligence (AI) driving recent outperformance of US tech stocks and sending some of them to record-highs, how much more AI-driven upside remains for US equities? Despite the recent gains, we estimate that potential AI-related productivity boosts could lead to significantly more upside for S&P 500 earnings and stock prices over the medium-to-longer term, although substantial uncertainty and risks remain.

A potential boost to US productivity, earnings, and equities...

Our economists estimate that widespread generative AI adoption (which we assume occurs in 10 years) could boost US productivity growth by 1.5pp annually over a 10-year period and lift trend real GDP growth by 1.1pp for 10 years (see pgs. 14-15). Under these assumptions in our dividend discount model (DDM), we estimate that S&P 500 EPS CAGR over the next 20 years would be 5.4%, 50bp greater than our current assumption of 4.9%, and S&P 500 fair value would be 9% higher than current levels, holding all else equal.

Widespread AI adoption could lead to S&P 500 EPS in 20 years 11% greater than our current assumption

S&P 500 EPS forecasts over the next 20 years, \$



Source: Goldman Sachs GIR.

...but an uncertain one

That said, the range of potential AI impacts on the S&P 500 is wide—and therefore unlikely to be fully priced by investors in the near term—for four key reasons:

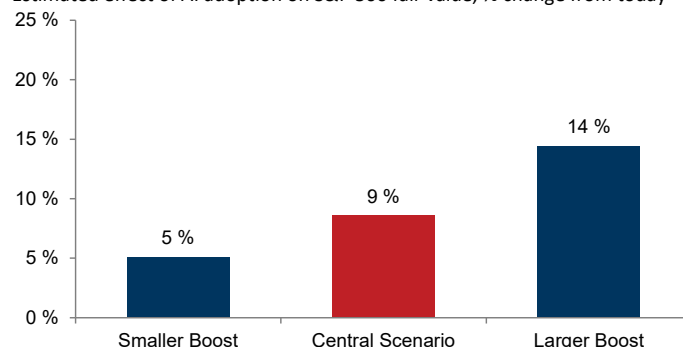
- Our economists' estimates of the impact of AI adoption on productivity growth vary** from 0.3pp to 3.0pp annually, depending on the speed of adoption, the power of AI, and the breadth of labor displacement. Based on this range of productivity scenarios, we estimate that the upside to S&P 500 fair value could be as small as 5% and as large as 14%. And the potential upside could be even larger if the uplift in GDP and revenue growth is also accompanied by an increase in corporate profit margins.
- Policy responses could constrain the ability of companies to retain the additional profits generated from AI.** Corporate profits as a share of GDP stand at elevated levels relative to history, while wages as a share of GDP remain near historic lows. If AI adoption leads to increased corporate profits at the expense of labor,

policymakers could respond by raising corporate tax rates. The effective corporate tax rate would need to rise by 8pp to fully offset the 11% increase in the stream of future S&P 500 earnings that may otherwise occur as a consequence of corporations embracing AI.

- A higher interest rate environment could negate much of the potential increase in S&P 500 fair value.** While a productivity boom that leads to lower prices could be disinflationary and put downward pressure on rates, our economists note that AI could increase investment demand and in turn lift estimates of the neutral rate, a key input in monetary policymakers' decisions. We estimate that interest rates would only need to rise by 30bp from current levels to fully offset the upside to fair value from AI adoption, all else equal.
- S&P 500 prices are more clearly tied to near-term cyclical dynamics,** even if AI adoption could provide a boost to the S&P 500 index in the long term. If economic data weaken and a recession becomes more likely (with the consensus of forecasters already assigning 65% odds to a recession in the next 12 months versus our estimate of 25% odds), S&P 500 prices would likely decline, regardless of the long-term impact of AI.

The potential productivity boost from AI adoption could lead to significantly more upside for the S&P 500 index

Estimated effect of AI adoption on S&P 500 fair value, % change from today



Source: Goldman Sachs GIR.

The perils of euphoric expectations

At the index level, the current equity risk premium and long-term EPS growth expectations are roughly in line with historical averages, suggesting investor optimism on AI adoption is not at extreme levels. However, at the stock level, the current valuation of the largest AI beneficiaries, like NVDA, is similar to the valuation accorded in the 2000s to some of the largest Dot Com Boom beneficiaries (MSFT, INTC), though not as high as the most extreme example (CSCO). Historical precedent from the Dot Com Boom shows the perils of high expectations. Even though most TMT companies were still able to generate strong sales growth between 2000 and 2002, the failure to meet lofty investor forecasts led to a sharp 50%+ contraction in P/E multiples and a plunge in share prices. Euphoric growth expectations, therefore, are another risk worth watching.

Ryan Hammond, Senior US Equity Strategist

Email: ryan.hammond@gs.com
Tel: 212-902-5625

Goldman Sachs & Co. LLC

David Kostin, Chief US Equity Strategist


Email: david.kostin@gs.com
Tel: 212-902-6781

Goldman Sachs & Co. LLC

The AI craze in pics

After its launch in November 2022, OpenAI's ChatGPT became the fastest application to surpass 100mn users


Number of months taken to surpass 100mn users



Source: Goldman Sachs GIR.

...and many company management teams are increasingly focused on opportunities from AI on earnings calls


Mentions of "AI" in selected companies' earnings calls



Note: Includes mentions of "AI" in analyst/journalist questions. Source: Company data, Statista, Goldman Sachs GIR.

...and some AI-related stocks are trading with high price-to-earnings multiples

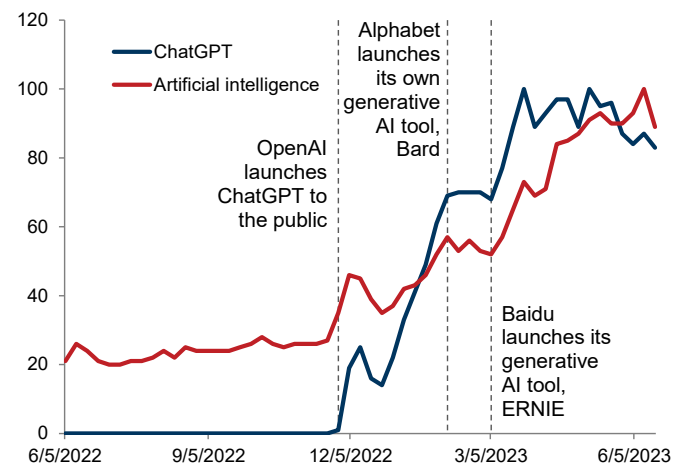
NTM P/E multiple



Source: Bloomberg, Goldman Sachs GIR.

Since then, several other generative AI tools have emerged and interest in artificial intelligence has increased significantly...

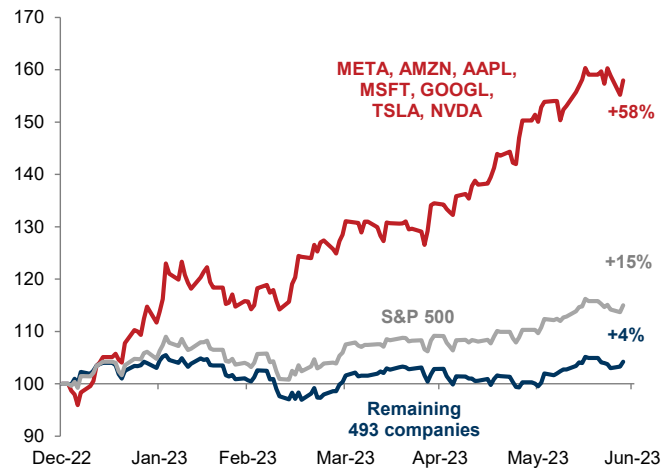
Google search trends



Source: Google Trends (<https://www.google.com/trends>), Goldman Sachs GIR.

Mega-cap tech stocks have rallied sharply year-to-date, outperforming the broader S&P 500 index, driven by optimism about the potential benefits to companies from AI...

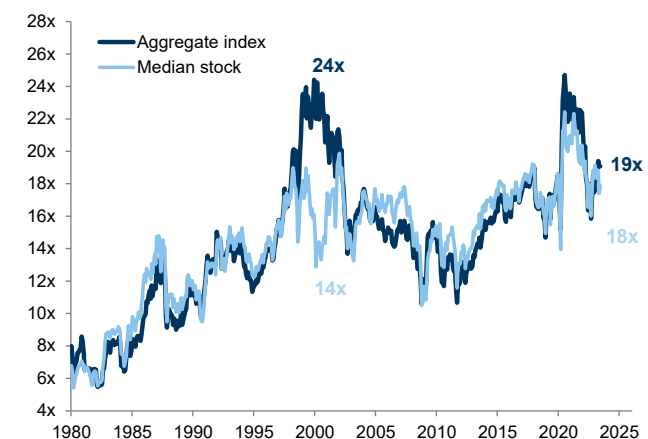
Indexed returns, 12/31/2022=100



Source: FactSet, Goldman Sachs GIR.

However, the key beneficiaries of AI adoption have not pushed aggregate index valuation to the extreme level of the Dot Com Boom

Aggregate vs. median S&P 500 NTM P/E dislocation



Source: Compustat, Goldman Sachs GIR.

Markets around past productivity booms

Dominic Wilson and Vickie Chang assess the impact of past innovation-driven productivity booms on markets and what that could mean for the potential AI productivity boom ahead

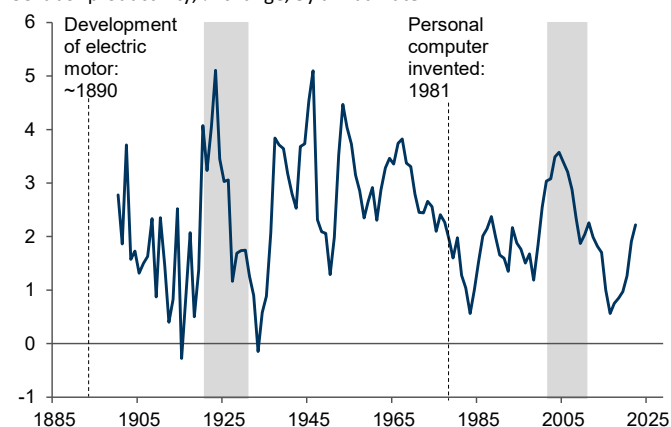
With generative artificial intelligence (AI) potentially ushering in a period of sustained substantial productivity growth (see pgs. 14-15), a key question is how that might impact major macro markets. We turn to history as a guide on the macro market impacts of innovation-driven productivity booms, focusing on two major episodes: the widespread adoption of electricity after World War I (1919-1929) and the broad adoption of PCs and the internet in the late 1990s and early 2000s (1996-2005).

Of course, factors beyond the productivity boom also drove markets during these episodes. The EM crises of 1997-1998 significantly impacted the global economy and asset prices, and after 2003, China's accession to the WTO prompted big shifts in manufacturing. Similarly, the start of the 1920s productivity boom overlapped with the transition from a wartime to a peacetime economy. Limited data availability and changes in financial markets prevent a full comparison across the major markets in the two episodes.

But markets around these prior productivity booms nonetheless shared some commonalities: Both booms had the biggest impact on equities and equity valuations—which rose substantially—and both ultimately ended in bubbles and subsequent busts. We find that the potential AI productivity boom ahead shares some of the key features of these prior periods, so could this boom/bust cycle happen again?

Major innovation-driven productivity booms occurred around the adoption of electricity and of PCs/the internet

US labor productivity, % change, 5y annual rate



Note: Grey shaded areas represent resulting productivity boom. Source: US Bureau of Labor Statistics, Woolf (1987), Goldman Sachs GIR.

The nineties boom and bust

During the period of PC/internet adoption (1996-2005), US equities posted healthy, if unspectacular, gains. Profits and earnings outpaced GDP somewhat, but the S&P 500 gains were broadly in line with nominal GDP gains. US Dollar appreciation was relatively modest and, excluding EM economies, FX was little changed on net. Similarly, both the Fed funds rate and 10-year yields declined over the period and

tracked the domestic demand cycle. Oil prices fell sharply during the EM crises in 1997-98 but moved higher by 2005.

Asset market performance during 1996-2005 productivity boom

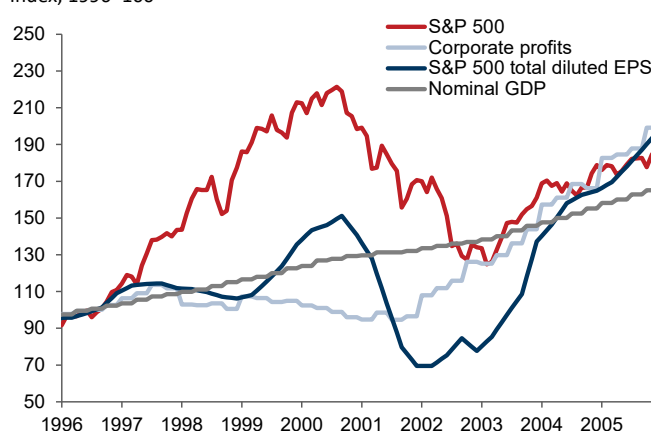
	Change from:		
	1996 to 2005	1996 to 2000 peak	2000 peak to 2005
Effective Fed Funds Rate	-140bp	98bp	-238bp
10-year Treasury Yield	-118bp	101bp	-219bp
Real Broad Dollar	9%	21%	-10%
Nasdaq 100	199%	687%	-62%
S&P 500	105%	142%	-15%
Oil (WTI)	218%	83%	74%

Source: Haver Analytics, Goldman Sachs GIR.

But these relatively modest shifts over the period mask a much larger economic and market boom and bust within the period. During the initial productivity boom, the pattern of market shifts, though not the magnitudes, match what would be expected from an (over-) anticipated productivity boom—equities rose sharply and valuations climbed to extreme levels.

A significant domestic economic boom accompanied these moves. The investment share of GDP climbed, the savings rates fell, and the current account deteriorated. Both the Fed funds rate and longer-dated yields fell over 1997-1998 as the Asian financial crisis and Russian default hit, but with domestic demand booming, the funds rate rose to a fresh cycle peak in 2000. Longer-dated yields rose too but remained below their 1996 levels as low and stable inflation held down the term premium. Significant US Dollar appreciation in the late 1990s (peaking in early 2002) largely owed to the EM devaluations of 1997 and 1998, but the Dollar—as the preferred recipient of capital flows—also rose against other advanced economies. However, as boom turned to bust, equities saw large declines, interest rates fell, and the bulk of the Dollar strength reversed.

Over 1996-2005, profits and earnings outpaced GDP, but S&P 500 gains broadly tracked nominal GDP gains



Source: Haver Analytics, Goldman Sachs GIR.

The roaring twenties

The evidence around the productivity boosts in the 1920s, as electricity adoption spread, is sparser but provides some parallel lessons. Once again, equities saw sustained gains and a sharp climb in valuations alongside the productivity boom, but the 1929 crash ultimately ensued. The story for rates and FX is

harder to map to the current context given the differences in monetary policy and exchange rate management. Inflation was extremely low over the period. But the Fed's discount rate again followed the economic and equity cycle, falling in 1924 as the economy weakened before rising steadily and hitting new peaks as the equity bubble accelerated and then burst.

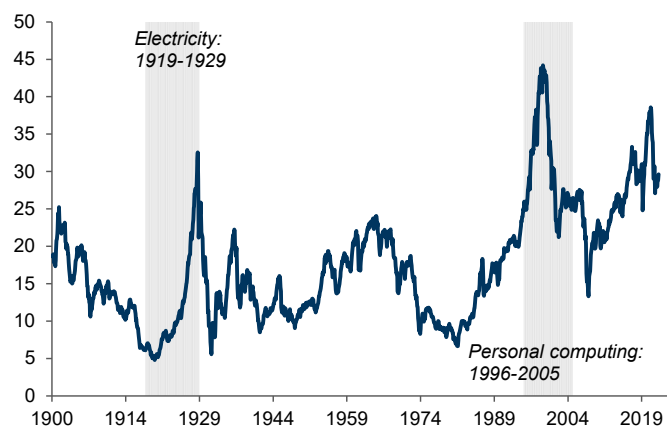
Overall, these two prior experiences suggest that the biggest impact on asset markets was felt in equities and equity valuations, which ended in bubbles both times. The behavior of rates and FX appears to have been driven more by domestic demand than by persistent structural shifts from the change in trend productivity growth, though the 1990s provide some support for the idea that economies experiencing outsized productivity gains could see FX appreciation pressures.

Bubble trouble

Consistent with history, US equities have already been the focus of expectations of AI-related productivity gains in the recent period. Our equity strategists have laid out benchmarks for the equity index upside that an AI-induced productivity boom might fundamentally justify (see pg. 16). So, will the market overshoot those valuations—creating a bubble that ultimately ends in a bust—as the historical experience suggests?

Valuations climbed rapidly during prior innovation periods before retracing

Shiller Cyclically-Adjusted P/E



Source: Robert Shiller, Goldman Sachs GIR.

Bubbles are complicated phenomena, often driven by momentum and self-fulfilling price dynamics. But several reasons explain why productivity booms can lead markets to overpay.

First, investors may fall prey to a fallacy of extrapolation. With genuine innovation, productivity gains will be real. In the short term, accelerating productivity growth can increase profit shares even at the economy-wide level. But, on average, competition or investment largely eroded those initial gains over subsequent years. This implies that a faster phase of profit growth at the start of periods of innovation tends to be “paid back” over time. To the extent that markets price initial increases in profit growth as persistent, the long-term potential shift in the earnings trajectory may be overestimated.

Second, investors can fall prey to a fallacy of aggregation. During periods of innovation, some individual companies may be capable of stretches of stunning earnings growth driven by a

new technology. But it is a mistake to assume that what can be true for an individual company can be true on aggregate. Even at the individual level, competition and market entry can ultimately limit the potential for sustained high profits. With potential “winners” sometimes more obvious than losers, investors may price a chance of increased profitability across a broad range of potential winners. The result may imply a rate of economy-wide profit growth that is unlikely to be feasible.

Third, activity fueled by the bubble itself can appear to justify the optimism. As asset prices rise, they may encourage a boom in investment and consumer spending. This in itself may provide a boost to the profitability of companies supplying those areas. But if increased revenues and profits are ultimately based on unsustainable demand that is generating economic imbalances, then those gains too will eventually unwind. In other words, a domestic boom created by overvalued asset prices can fuel the perception that higher profit growth can be maintained. For example, in the late 1990s, the domestic boom generated a major savings-investment imbalance that ultimately unwound in the bust but that generated more rapid demand growth for a period.

Fourth, to the extent that an acceleration in productivity growth leads to monetary policy that is easier than it “should” be, it can help fuel asset price overvaluation. This could happen for several reasons: the acceleration in productivity growth could lead inflation to undershoot; central banks could be slow to appreciate that the neutral rate has risen; or unsustainable current account deterioration could postpone the inflationary consequences of a boom. This is particularly a risk when a boom overlaps with other disinflationary forces, as it did for the US in the late 1990s.

The challenge of keeping it real

All that said, bubbles can form without these conditions, and not all high-productivity periods lead to bubbles. But the challenge with periods of sustained productivity improvement is that the underlying economic shifts are both powerful and real. They provide fundamental support for higher asset prices—and create the basis for dramatic gains for some companies—even if that fundamental improvement is then too widely or too dramatically priced. The coming potential AI productivity boom shares some of the key features that led to these issues in the past: a breakthrough innovation that might lead to sizable increases in productivity and profitability, which then creates the basis for substantial new investments and fuels belief in a broader cycle of innovation.

If the market does overpay for the AI productivity boom, that has the capacity to impact a broad set of asset price shifts. The 1990s history suggests that this dynamic could be associated not just with a period of unsustainably high equity prices, but also larger demand booms, greater FX appreciation, and higher interest rates in the leading countries than would have otherwise been the case.

Dominic Wilson, Senior Markets Advisor

Email: dominic.wilson@gs.com
Tel: 212-902-5924

Goldman Sachs & Co. LLC

Vickie Chang, Global Markets Strategist

Email: vickie.chang@gs.com
Tel: 212-902-6915

Goldman Sachs & Co. LLC

What we're hearing from public investors



Peter Callahan is US Technology, Media, and Telecommunications Sector Specialist in Goldman Sachs Global Banking & Markets. Below, he discusses what he's hearing from public investors on the artificial intelligence (AI) theme, how they are positioned around AI, and what they're watching ahead.

The interviewee is an employee of Goldman Sachs Global Banking & Markets and the views stated herein reflect those of the interviewee, not Goldman Sachs Research.

Q: How focused are public investors on the artificial intelligence (AI) theme?

A: AI is no doubt the dominant theme in the public marketplace, with AI mania in full flight right now. I don't remember a theme being this pervasive and important to investors—six months after the emergence of the generative AI tool ChatGPT, AI has captured the imagination of almost every client that we speak to. Unlike other recent product cycle stories that ultimately lost steam, including blockchain, the metaverse, and autonomous driving that appealed largely to a relatively narrow, specialist investor base, AI is attracting interest from investors of all shapes and sizes—specialists, generalists, macro, thematic and everything in between. That broad interest has translated into price action, with a small group of large tech stocks massively outperforming the broader market year-to-date and a handful of smaller cap stocks rising to the tune of 200%.

Q: Given the strong performance, are concerns that the hype—and performance—are overdone emerging at all?

A: Whether too much has been priced in is always easier to answer in hindsight. For example, during the Software-as-a-Service (SaaS) cycle, several dozen companies were trading above 10x EV to sales at one point—and some companies at 50-70x EV to sales—which at the time didn't seem that striking but in retrospect was a historical anomaly. Much like then, AI excitement has fueled multiple expansion in some tech names, though generally not to historical extremes yet. AI bellwethers Microsoft and Nvidia, for example, only trade at modest premiums to their five-year relative P/E multiple averages, despite the strong YTD performance and thematic excitement.

The current market consensus seems to be that AI-led performance has legs because the technology has the makings of a platform shift akin to past transformative shifts that have tended to come along about once a decade—the personal computer in the 1980s, the internet in the 1990s, mobile in the 2000s, cloud computing and the public cloud in the 2010s, and today generative AI. Investors fully expect this technology to create and disrupt profit pools and are lining up to have exposure to the AI theme for the next decade+, even as the stories and the companies will likely change over that time.

Q: But, as you mentioned, investor excitement about many recent tech innovations—blockchain, the metaverse, autonomous vehicles, etc.—has fizzled out. Why is this time different?

A: It's true that some technologies that captured investors' imaginations ultimately disappointed. Blockchain, for all its technological merits, never found the product that investors could engage with in the public markets; the idea that the metaverse would transform our daily lives and people would soon be regularly wearing Google glasses came and went; and autonomous driving adoption has been slower than many investors expected five years ago. But it's also worth noting that even in instances when a technology has ultimately become mainstream, seeing real proliferation has taken time. The iPhone, for example, debuted in 2007, yet Instagram only emerged three years later and TikTok nine years later. So, while investors tend to get overexcited about tech innovations in the short term, they also tend to underappreciate the impact of large technology cycles over the longer term.

That said, what differentiates this technology cycle is the velocity of change, which is unlike anything we've ever witnessed. ChatGPT was the fastest application in history to surpass 100mn users—it took just two months, compared to nine months for TikTok and 30 months for Instagram. Nvidia is a prime example of the astounding velocity of earnings revisions. The company recently issued guidance that it expects \$11bn of revenue in the July quarter—roughly two to three years ahead of market expectations for that amount of revenue. So, the velocity of change is exceeding prior expectations, and certainly historical comps, forcing investors to care about the AI theme now. But at the heart of investor excitement is the potential of this technology, like other platform shifts, to touch and transform almost every aspect of the enterprise and consumer experience. The technology is not niche to a certain vertical, but rather is a horizontal platform application that is likely to pervade every industry.

Q: So, how are investors gaining exposure to the AI theme in public markets?

A: The obvious tip of the spear that public-market investors initially grabbed onto was the idea that the graphics processing unit (GPU) would be a lightning rod for AI development and the training of large language models (LLMs). But over the past few months investors have been willing to broaden out that aperture and think about secondary sources for GPUs, the tools semiconductor fabs need to manufacture them, the software required to develop them, etc. It feels like almost every week investors are taking another step forward and going up another layer in the stack to try to find new ways to interact with the AI theme. That said, right now we're still very much in the "picks and shovels" stage of AI investing—as demonstrated by the strong start to the year for Semiconductors, which are up ~44% YTD, the best start to a year for the group since 2000. Investors seem to believe that the further up the stack they go, the more uncertainty they face because it's not yet clear how consumers and enterprises will adopt AI technology.

Q: As in other instances of technological innovation, there will no doubt be companies that benefit and others whose business models may be negatively impacted. What questions are investors focused on as they attempt to discern between the two?

A: Investors are focused on many questions that sit at the heart of the current debates in the space: will spending on AI be cannibalistic or incremental, meaning, will the extra dollar spent on next-gen technology replace existing spending or add to it? Will adoption occur faster at the enterprise or consumer level; the enterprise has much to gain in terms of scale, scope, margins, and effectiveness, but the consumer has a valuable networking effect and free-flowing adoption dynamic, so who will adopt faster? What's actually a new AI product vs. a product that's been sitting on the shelf for some time and is now being rebranded as an AI product? Are public market incumbents or private market startups better positioned to reap gains in the space; incumbents certainly have the advantage of scale, distribution, data, and balance sheets, but can they rearchitect for an AI world, or will new companies "born" in AI dominate the space? Will open-source LLMs become commoditized, and what are the implications of that? So, there's no shortage of questions and no clear answers, but that's how investing opportunities are born.

Q: So, what are investors watching to gauge whether this theme really has legs?

A: The market is watching whether excitement eventually translates into concrete outcomes. As I mentioned, Nvidia saw a huge positive earnings revision following earnings in May (consensus estimates were revised up by ~75%), which helped trigger the recent AI-led excitement. But, since then, earnings revisions from other Semiconductor companies leaning into the AI theme have been much more modest. As such, as we head into the second half of this year, investors will be looking for clues about early signs of AI progress from companies—not necessarily in the form of revenue beats at this point, but more in the form of product initiatives (SKUs), pipeline builds, order trends, customer engagement and/or other KPIs, helping to set the table for 2024 and beyond.

Glossary of GS proprietary indices

Current Activity Indicator (CAI)

GS CAIs measure the growth signal in a broad range of weekly and monthly indicators, offering an alternative to Gross Domestic Product (GDP). GDP is an imperfect guide to current activity: In most countries, it is only available quarterly and is released with a substantial delay, and its initial estimates are often heavily revised. GDP also ignores important measures of real activity, such as employment and the purchasing managers' indexes (PMIs). All of these problems reduce the effectiveness of GDP for investment and policy decisions. Our CAIs aim to address GDP's shortcomings and provide a timelier read on the pace of growth.

For more, see our CAI page and Global Economics Analyst: Trackin' All Over the World – Our New Global CAI, 25 February 2017.

Dynamic Equilibrium Exchange Rates (DEER)

The GSDEER framework establishes an equilibrium (or "fair") value of the real exchange rate based on relative productivity and terms-of-trade differentials.

For more, see our GSDEER page, Global Economics Paper No. 227: Finding Fair Value in EM FX, 26 January 2016, and Global Markets Analyst: A Look at Valuation Across G10 FX, 29 June 2017.

Financial Conditions Index (FCI)

GS FCIs gauge the "looseness" or "tightness" of financial conditions across the world's major economies, incorporating variables that directly affect spending on domestically produced goods and services. FCIs can provide valuable information about the economic growth outlook and the direct and indirect effects of monetary policy on real economic activity.

FCIs for the G10 economies are calculated as a weighted average of a policy rate, a long-term risk-free bond yield, a corporate credit spread, an equity price variable, and a trade-weighted exchange rate; the Euro area FCI also includes a sovereign credit spread. The weights mirror the effects of the financial variables on real GDP growth in our models over a one-year horizon. FCIs for emerging markets are calculated as a weighted average of a short-term interest rate, a long-term swap rate, a CDS spread, an equity price variable, a trade-weighted exchange rate, and—in economies with large foreign-currency-denominated debt stocks—a debt-weighted exchange rate index.

For more, see our FCI page, Global Economics Analyst: Our New G10 Financial Conditions Indices, 20 April 2017, and Global Economics Analyst: Tracking EM Financial Conditions – Our New FCIs, 6 October 2017.

Goldman Sachs Analyst Index (GSAI)

The US GSAI is based on a monthly survey of GS equity analysts to obtain their assessments of business conditions in the industries they follow. The results provide timely "bottom-up" information about US economic activity to supplement and cross-check our analysis of "top-down" data. Based on analysts' responses, we create a diffusion index for economic activity comparable to the ISM's indexes for activity in the manufacturing and nonmanufacturing sectors.

Macro-Data Assessment Platform (MAP)

GS MAP scores facilitate rapid interpretation of new data releases for economic indicators worldwide. MAP summarizes the importance of a specific data release (i.e., its historical correlation with GDP) and the degree of surprise relative to the consensus forecast. The sign on the degree of surprise characterizes underperformance with a negative number and outperformance with a positive number. Each of these two components is ranked on a scale from 0 to 5, with the MAP score being the product of the two, i.e., from -25 to +25. For example, a MAP score of +20 (5;+4) would indicate that the data has a very high correlation to GDP (5) and that it came out well above consensus expectations (+4), for a total MAP value of +20.

Top of Mind archive: click to access



Issue 119
Daunting debt limit dynamics
May 22, 2023



Issue 104
Investing in Climate Change 2.0
December 13, 2021



Issue 118
US-China: more decoupling ahead?
May 1, 2023



Issue 103
Inflation: here today, gone tomorrow?
November 17, 2021



Issue 117
All about bank(panic)s
April 3, 2023



Issue 102
Europe at a Crossroads
October 18, 2021



Issue 116
(Japanese) Bonds, Bonds, Bonds
February 23, 2023



Issue 101
Is China Investable?
September 13, 2021



Issue 115
The Bigger Worry: Growth or Inflation?
January 27, 2023



Issue 100
The Post-Pandemic Future of Work
July 29, 2021



Special Issue
2022: 3 themes in charts
December 15, 2022




Issue 99
Bidenomics: evolution or revolution?
June 29, 2021



Issue 114
The Winter of Crypto's Discontents
December 9, 2022



Issue 98
Crypto: A New Asset Class?
May 21, 2021



Issue 113
Central Bank Tightening: what could break?
November 11, 2022



Issue 97
Reflation Risk
April 1, 2021



Issue 112
China's Congress: an inflection point?
October 11, 2022



Issue 96
The Short and Long of Recent Volatility
February 25, 2021



Issue 111
Will slaying inflation require recession?
September 13, 2022



Issue 95
The IPO SPAC-tacle
January 28, 2021



Issue 110
Food, Fuel, and the Cost-of-Living Crisis
July 28, 2022



Special Issue
2020 Update, and a Peek at 2021
December 17, 2020



Issue 109
Equity bear market: a paradigm shift?
June 14, 2022



Issue 94
What's In Store For the Dollar
October 29, 2020



Issue 108
(De)Globalization Ahead?
April 28, 2022



Issue 93
Beyond 2020: Post-Election Policies
October 1, 2020



Issue 107
Stagflation Risk
March 14, 2022




Issue 92
COVID-19: Where We Go From Here
August 13, 2020




Issue 106
Russia Risk
February 24, 2022



Issue 91
Investing in Racial Economic Equality
July 16, 2020




Issue 105
2022: The endemic year?
January 24, 2022



Issue 90
Daunting Debt Dynamics
May 28, 2020



Special Issue
2021: 4 themes in charts
December 17, 2021



Issue 89
Reopening the Economy
April 28, 2020

Source of photos: www.gettyimages.com, www.istockphoto.com, www.shutterstock.com, US Department of State/Wikimedia Commons/Public Domain.

Disclosure Appendix

Reg AC

We, Allison Nathan, Jenny Grimberg, Ashley Rhodes, Joseph Briggs, Vickie Chang, Ryan Hammond, David Kostin, and Dominic Wilson hereby certify that all of the views expressed in this report accurately reflect our personal views, which have not been influenced by considerations of the firm's business or client relationships.

We Kash Rangan and Eric Sheridan hereby certify that all of the views expressed in this report accurately reflect our personal views about the subject company or companies and its or their securities. We also certify that no part of our compensation was, is or will be, directly or indirectly, related to the specific recommendations or views expressed in this report.

Unless otherwise stated, the individuals listed on the cover page of this report are analysts in Goldman Sachs' Global Investment Research division.

Logo Disclosure

Please note: Third party brands used in this report are the property of their respective owners, and are used here for informational purposes only. The use of such brands should not be viewed as an endorsement, affiliation, or sponsorship by or for Goldman Sachs or any of its products/services.

Disclosures

Regulatory disclosures

Disclosures required by United States laws and regulations

See company-specific regulatory disclosures above for any of the following disclosures required as to companies referred to in this report: manager or co-manager in a pending transaction; 1% or other ownership; compensation for certain services; types of client relationships; managed/co-managed public offerings in prior periods; directorships; for equity securities, market making and/or specialist role. Goldman Sachs trades or may trade as a principal in debt securities (or in related derivatives) of issuers discussed in this report.

The following are additional required disclosures: **Ownership and material conflicts of interest:** Goldman Sachs policy prohibits its analysts, professionals reporting to analysts and members of their households from owning securities of any company in the analyst's area of coverage. **Analyst compensation:** Analysts are paid in part based on the profitability of Goldman Sachs, which includes investment banking revenues. **Analyst as officer or director:** Goldman Sachs policy generally prohibits its analysts, persons reporting to analysts or members of their households from serving as an officer, director, or advisor of any company in the analyst's area of coverage. **Non-U.S. Analysts:** Non-U.S. analysts may not be associated persons of Goldman Sachs & Co. LLC and therefore may not be subject to FINRA Rule 2241 or FINRA Rule 2242 restrictions on communications with subject company, public appearances and trading securities held by the analysts.

Additional disclosures required under the laws and regulations of jurisdictions other than the United States

The following disclosures are those required by the jurisdiction indicated, except to the extent already made above pursuant to United States laws and regulations. **Australia:** Goldman Sachs Australia Pty Ltd and its affiliates are not authorised deposit-taking institutions (as that term is defined in the Banking Act 1959 (Cth)) in Australia and do not provide banking services, nor carry on a banking business, in Australia. This research, and any access to it, is intended only for "wholesale clients" within the meaning of the Australian Corporations Act, unless otherwise agreed by Goldman Sachs. In producing research reports, members of Global Investment Research of Goldman Sachs Australia may attend site visits and other meetings hosted by the companies and other entities which are the subject of its research reports. In some instances the costs of such site visits or meetings may be met in part or in whole by the issuers concerned if Goldman Sachs Australia considers it is appropriate and reasonable in the specific circumstances relating to the site visit or meeting. To the extent that the contents of this document contains any financial product advice, it is general advice only and has been prepared by Goldman Sachs without taking into account a client's objectives, financial situation or needs. A client should, before acting on any such advice, consider the appropriateness of the advice having regard to the client's own objectives, financial situation and needs. A copy of certain Goldman Sachs Australia and New Zealand disclosure of interests and a copy of Goldman Sachs' Australian Sell-Side Research Independence Policy Statement are available at: <https://www.goldmansachs.com/disclosures/australia-new-zealand/index.html>. **Brazil:** Disclosure information in relation to CVM Resolution n. 20 is available at <https://www.gs.com/worldwide/brazil/area/gir/index.html>. Where applicable, the Brazil-registered analyst primarily responsible for the content of this research report, as defined in Article 20 of CVM Resolution n. 20, is the first author named at the beginning of this report, unless indicated otherwise at the end of the text. **Canada:** This information is being provided to you for information purposes only and is not, and under no circumstances should be construed as, an advertisement, offering or solicitation by Goldman Sachs & Co. LLC for purchasers of securities in Canada to trade in any Canadian security. Goldman Sachs & Co. LLC is not registered as a dealer in any jurisdiction in Canada under applicable Canadian securities laws and generally is not permitted to trade in Canadian securities and may be prohibited from selling certain securities and products in certain jurisdictions in Canada. If you wish to trade in any Canadian securities or other products in Canada please contact Goldman Sachs Canada Inc., an affiliate of The Goldman Sachs Group Inc., or another registered Canadian dealer. **Hong Kong:** Further information on the securities of covered companies referred to in this research may be obtained on request from Goldman Sachs (Asia) L.L.C. **India:** Further information on the subject company or companies referred to in this research may be obtained from

Goldman Sachs (India) Securities Private Limited, Research Analyst - SEBI Registration Number INH000001493, 951-A, Rational House, Appasaheb Marathe Marg, Prabhadevi, Mumbai 400 025, India, Corporate Identity Number U74140MH2006FTC160634, Phone +91 22 6616 9000, Fax +91 22 6616 9001. Goldman Sachs may beneficially own 1% or more of the securities (as such term is defined in clause 2 (h) the Indian Securities Contracts (Regulation) Act, 1956) of the subject company or companies referred to in this research report. **Japan:** See below. **Korea:** This research, and any access to it, is intended only for "professional investors" within the meaning of the Financial Services and Capital Markets Act, unless otherwise agreed by Goldman Sachs. Further information on the subject company or companies referred to in this research may be obtained from Goldman Sachs (Asia) L.L.C., Seoul Branch. **New Zealand:** Goldman Sachs New Zealand Limited and its affiliates are neither "registered banks" nor "deposit takers" (as defined in the Reserve Bank of New Zealand Act 1989) in New Zealand. This research, and any access to it, is intended for "wholesale clients" (as defined in the Financial Advisers Act 2008) unless otherwise agreed by Goldman Sachs. A copy of certain Goldman Sachs Australia and New Zealand disclosure of interests is available at: <https://www.goldmansachs.com/disclosures/australia-new-zealand/index.html>. **Russia:** Research reports distributed in the Russian Federation are not advertising as defined in the Russian legislation, but are information and analysis not having product promotion as their main purpose and do not provide appraisal within the meaning of the Russian legislation on appraisal activity. Research reports do not constitute a personalized investment recommendation as defined in Russian laws and regulations, are not addressed to a specific client, and are prepared without analyzing the financial circumstances, investment profiles or risk profiles of clients. Goldman Sachs assumes no responsibility for any investment decisions that may be taken by a client or any other person based on this research report. **Singapore:** Goldman Sachs (Singapore) Pte. (Company Number: 198602165W), which is regulated by the Monetary Authority of Singapore, accepts legal responsibility for this research, and should be contacted with respect to any matters arising from, or in connection with, this research. **Taiwan:** This material is for reference only and must not be reprinted without permission. Investors should carefully consider their own investment risk. Investment results are the responsibility of the individual investor. **United Kingdom:** Persons who would be categorized as retail clients in the United Kingdom, as such term is defined in the rules of the Financial Conduct Authority, should read this research in conjunction with prior Goldman Sachs research on the covered companies referred to herein and should refer to the risk warnings that have been sent to them by Goldman Sachs International. A copy of these risks warnings, and a glossary of certain financial terms used in this report, are available from Goldman Sachs International on request.

European Union and United Kingdom: Disclosure information in relation to Article 6 (2) of the European Commission Delegated Regulation (EU) (2016/958) supplementing Regulation (EU) No 596/2014 of the European Parliament and of the Council (including as that Delegated Regulation is implemented into United Kingdom domestic law and regulation following the United Kingdom's departure from the European Union and the European Economic Area) with regard to regulatory technical standards for the technical arrangements for objective presentation of investment recommendations or other information recommending or suggesting an investment strategy and for disclosure of particular interests or indications of conflicts of interest is available at <https://www.gs.com/disclosures/europeanpolicy.html> which states the European Policy for Managing Conflicts of Interest in Connection with Investment Research.

Japan: Goldman Sachs Japan Co., Ltd. is a Financial Instrument Dealer registered with the Kanto Financial Bureau under registration number Kinsho 69, and a member of Japan Securities Dealers Association, Financial Futures Association of Japan Type II Financial Instruments Firms Association, The Investment Trusts Association, Japan, and Japan Investment Advisers Association. Sales and purchase of equities are subject to commission pre-determined with clients plus consumption tax. See company-specific disclosures as to any applicable disclosures required by Japanese stock exchanges, the Japanese Securities Dealers Association or the Japanese Securities Finance Company.

Global product; distributing entities

Goldman Sachs Global Investment Research produces and distributes research products for clients of Goldman Sachs on a global basis. Analysts based in Goldman Sachs offices around the world produce research on industries and companies, and research on macroeconomics, currencies, commodities, and portfolio strategy. This research is disseminated in Australia by Goldman Sachs Australia Pty Ltd (ABN 21 006 797 897); in Brazil by Goldman Sachs do Brasil Corretora de Títulos e Valores Mobiliários S.A.; Public Communication Channel Goldman Sachs Brazil: 0800 727 5764 and / or contatogoldmanbrasil@gs.com. Available Weekdays (except holidays), from 9am to 6pm. Canal de Comunicação com o Público Goldman Sachs Brasil: 0800 727 5764 e/ou contatogoldmanbrasil@gs.com. Horário de funcionamento: segunda-feira à sexta-feira (exceto feriados), das 9h às 18h; in Canada by Goldman Sachs & Co. LLC; in Hong Kong by Goldman Sachs (Asia) L.L.C.; in India by Goldman Sachs (India) Securities Private Ltd.; in Japan by Goldman Sachs Japan Co., Ltd.; in the Republic of Korea by Goldman Sachs (Asia) L.L.C., Seoul Branch; in New Zealand by Goldman Sachs New Zealand Limited; in Russia by OOO Goldman Sachs; in Singapore by Goldman Sachs (Singapore) Pte. (Company Number: 198602165W); and in the United States of America by Goldman Sachs & Co. LLC. Goldman Sachs International has approved this research in connection with its distribution in the United Kingdom.

Goldman Sachs International ("GSI"), authorised by the Prudential Regulation Authority ("PRA") and regulated by the Financial Conduct Authority ("FCA") and the PRA, has approved this research in connection with its distribution in the United Kingdom.

European Economic Area: GSI, authorised by the PRA and regulated by the FCA and the PRA, disseminates research in the following jurisdictions within the European Economic Area: the Grand Duchy of Luxembourg, Italy, the Kingdom of Belgium, the Kingdom of Denmark, the Kingdom of Norway, the Republic of Finland and the Republic of Ireland; GSI - Succursale de Paris (Paris branch) which is authorised by the French Autorité de contrôle prudentiel et de résolution ("ACPR") and regulated by the Autorité de contrôle prudentiel et de résolution and the Autorité des marchés financiers ("AMF") disseminates research in France; GSI -

Sucursal en España (Madrid branch) authorized in Spain by the Comisión Nacional del Mercado de Valores disseminates research in the Kingdom of Spain; GSI - Sweden Bankfilial (Stockholm branch) is authorized by the SFSA as a “third country branch” in accordance with Chapter 4, Section 4 of the Swedish Securities and Market Act (Sw. lag (2007:528) om värdepappersmarknaden) disseminates research in the Kingdom of Sweden; Goldman Sachs Bank Europe SE (“GSBE”) is a credit institution incorporated in Germany and, within the Single Supervisory Mechanism, subject to direct prudential supervision by the European Central Bank and in other respects supervised by German Federal Financial Supervisory Authority (Bundesanstalt für Finanzdienstleistungsaufsicht, BaFin) and Deutsche Bundesbank and disseminates research in the Federal Republic of Germany and those jurisdictions within the European Economic Area where GSI is not authorised to disseminate research and additionally, GSBE, Copenhagen Branch filial af GSBE, Tyskland, supervised by the Danish Financial Authority disseminates research in the Kingdom of Denmark; GSBE - Sucursal en España (Madrid branch) subject (to a limited extent) to local supervision by the Bank of Spain disseminates research in the Kingdom of Spain; GSBE - Succursale Italia (Milan branch) to the relevant applicable extent, subject to local supervision by the Bank of Italy (Banca d'Italia) and the Italian Companies and Exchange Commission (Commissione Nazionale per le Società e la Borsa “Consob”) disseminates research in Italy; GSBE - Succursale de Paris (Paris branch), supervised by the AMF and by the ACPR disseminates research in France; and GSBE - Sweden Bankfilial (Stockholm branch), to a limited extent, subject to local supervision by the Swedish Financial Supervisory Authority (Finansinspektionen) disseminates research in the Kingdom of Sweden.

General disclosures

This research is for our clients only. Other than disclosures relating to Goldman Sachs, this research is based on current public information that we consider reliable, but we do not represent it is accurate or complete, and it should not be relied on as such. The information, opinions, estimates, and forecasts contained herein are as of the date hereof and are subject to change without prior notification. We seek to update our research as appropriate, but various regulations may prevent us from doing so. Other than certain industry reports published on a periodic basis, the large majority of reports are published at irregular intervals as appropriate in the analyst's judgment.

Goldman Sachs conducts a global full-service, integrated investment banking, investment management, and brokerage business. We have investment banking and other business relationships with a substantial percentage of the companies covered by Global Investment Research. Goldman Sachs & Co. LLC, the United States broker dealer, is a member of SIPC (<https://www.sipc.org>).

Our salespeople, traders, and other professionals may provide oral or written market commentary or trading strategies to our clients and principal trading desks that reflect opinions that are contrary to the opinions expressed in this research. Our asset management area, principal trading desks and investing businesses may make investment decisions that are inconsistent with the recommendations or views expressed in this research.

We and our affiliates, officers, directors, and employees will from time to time have long or short positions in, act as principal in, and buy or sell, the securities or derivatives, if any, referred to in this research, unless otherwise prohibited by regulation or Goldman Sachs policy.

The views attributed to third party presenters at Goldman Sachs arranged conferences, including individuals from other parts of Goldman Sachs, do not necessarily reflect those of Global Investment Research and are not an official view of Goldman Sachs.

Any third party referenced herein, including any salespeople, traders and other professionals or members of their household, may have positions in the products mentioned that are inconsistent with the views expressed by analysts named in this report.

This research is focused on investment themes across markets, industries, and sectors. It does not attempt to distinguish between the prospects or performance of, or provide analysis of, individual companies within any industry or sector we describe.

Any trading recommendation in this research relating to an equity or credit security or securities within an industry or sector is reflective of the investment theme being discussed and is not a recommendation of any such security in isolation.

This research is not an offer to sell or the solicitation of an offer to buy any security in any jurisdiction where such an offer or solicitation would be illegal. It does not constitute a personal recommendation or take into account the particular investment objectives, financial situations, or needs of individual clients. Clients should consider whether any advice or recommendation in this research is suitable for their particular circumstances and, if appropriate, seek professional advice, including tax advice. The price and value of investments referred to in this research and the income from them may fluctuate. Past performance is not a guide to future performance, future returns are not guaranteed, and a loss of original capital may occur. Fluctuations in exchange rates could have adverse effects on the value or price of, or income derived from, certain investments.

Certain transactions, including those involving futures, options, and other derivatives, give rise to substantial risk and are not suitable for all investors. Investors should review current options and futures disclosure documents which are available from Goldman Sachs sales representatives or at <https://www.theocc.com/about/publications/character-risks.jsp> and https://www.fiadocumentation.org/fia/regulatory-disclosures_1/fia-uniform-futures-and-options-on-futures-risk-disclosures-booklet-pdf-version-2018. Transaction costs may be significant in option strategies calling for multiple purchase and sales of options such as spreads. Supporting documentation will be supplied upon request.

Differing Levels of Service provided by Global Investment Research: The level and types of services provided to you by Goldman Sachs Global Investment Research may vary as compared to that provided to internal and other external clients of GS, depending on various factors including your individual preferences as to the frequency and manner of receiving communication, your risk profile and investment focus and perspective (e.g., marketwide, sector specific, long term, short term), the size and scope of your overall client relationship with GS, and legal and regulatory constraints. As an example, certain clients may request to receive notifications when research on specific securities is published, and certain clients may request that specific data underlying analysts' fundamental analysis available on our internal client websites be delivered to them electronically through data feeds or otherwise. No change to an analyst's fundamental research views (e.g., ratings, price targets, or material changes to earnings estimates for equity securities), will be communicated to any client prior to inclusion of such information in a research report broadly disseminated through electronic publication to our internal client websites or through other means, as necessary, to all clients who are entitled to receive such reports..

All research reports are disseminated and available to all clients simultaneously through electronic publication to our internal client websites. Not all research content is redistributed to our clients or available to third-party aggregators, nor is Goldman Sachs responsible for the redistribution of our research by third party aggregators. For research, models or other data related to one or more securities, markets or asset classes (including related services) that may be available to you, please contact your GS representative or go to <https://research.gs.com>.

Disclosure information is also available at <https://www.gs.com/research/hedge.html> or from Research Compliance, 200 West Street, New York, NY 10282.

© 2023 Goldman Sachs.

No part of this material may be (i) copied, photocopied, or duplicated in any form by any means or (ii) redistributed without the prior written consent of The Goldman Sachs Group, Inc.