

# Assessing the Implications of a Productivity Miracle

What happens when cognitive tasks can be done at zero marginal cost?

NOVEMBER 30, 2023

**GREG JENSEN**  
**LAUREN SIMON**  
**JOSH MORIARTY**

**A**re we on the edge of a productivity miracle? This question is on our minds as we both observe and use cutting-edge machine-learning tools. As machine-learning solutions surpass human performance at many cognitive tasks that had previously been the exclusive domain of people, the odds are rising that the marginal cost of many economically important cognitive tasks will fall close to zero.

In our AIA (Artificial Investment Associate) Labs, **we are seeing firsthand** both the challenges and incredible opportunities that recent technological breakthroughs present. As we simultaneously use AI/ML tools and theorize about their future implications, two points are becoming increasingly clear:

1. In most areas, there are significant bottlenecks to using existing tools to perform any complex cognitive task without a fair amount of human intermediation. While machine-learning tools can score highly on complex assessments like the LSAT, medical board exams, and Bridgewater's own internal investor tests, translating this ability to doing the full scope of work these occupations entail requires overcoming many hurdles.
2. Many of these obstacles are likely tractable, and with time—in some cases, just a few years—many cognitive occupations will be performed more reliably by machine intelligence than human intelligence.

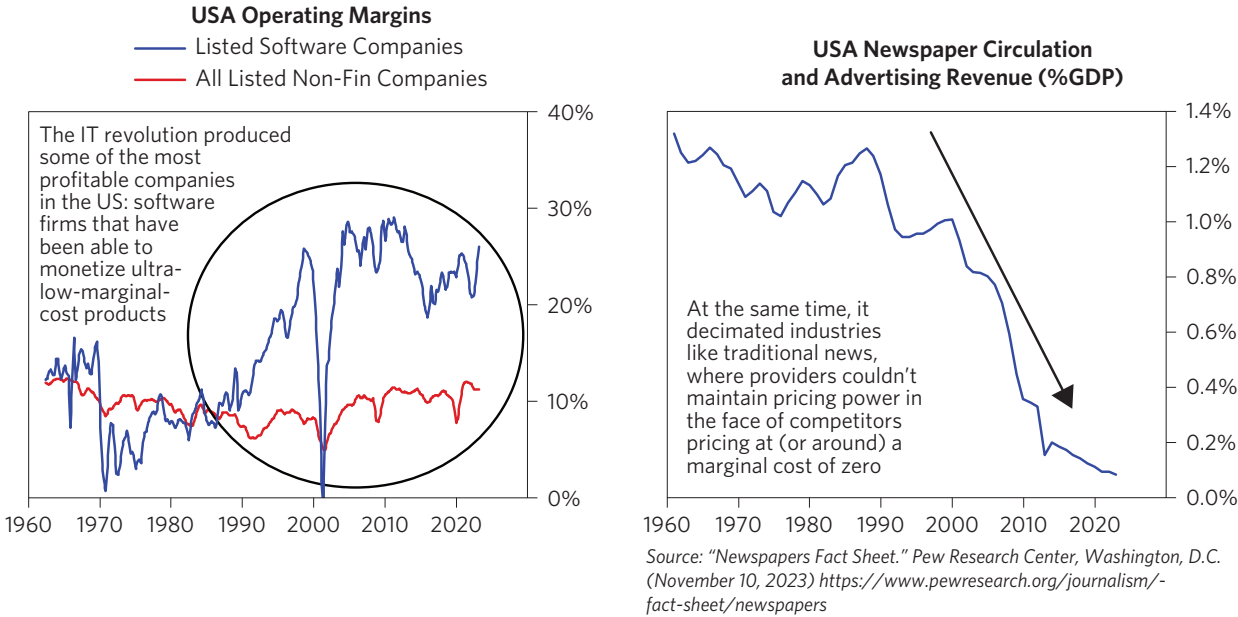
The range of potential outcomes from here is very wide, as **we've discussed in the past**. **We won't touch on most of the possible paths today, but we do want to share some thoughts on the potential world in which a large share of white-collar work can be done at low or zero marginal cost.** First, we consider the range of impacts this could have on companies and on the distribution of spending and activity across the economy. We then explore the extreme scenario in which the marginal cost of innovation itself—the cognitive work underlying long-term growth—falls close to zero, potentially producing “explosive growth.”

This thinking is still several steps removed from potential market impacts and considerations for investors. We focus here on what might occur in the economy from a bird's-eye view. And while destructive uses of AI are a pressing short-term risk, in this report, we focus on the productive uses that will need to be built with care over a longer time horizon.

# A Low- or Zero-Marginal-Cost World for White-Collar Work Creates Massive Opportunities and Risks for Profits

What happens to companies when the marginal cost of producing something they sell falls near zero? There’s a lot we can learn from past cases—**the economics of durable margins have not changed, even as their technological basis has shifted.**

Pricing power determines whether the marginal cost of a given product falling to zero is a best-case or worst-case scenario for its sellers. We saw this in the IT revolution of the 1990s and 2000s, which brought the marginal cost of production or distribution to zero for many businesses involved in creating and disseminating information. Besides consumers, who benefited from a profusion of free and low-cost products, the “winners” included software companies that earned eye-popping margins by charging a premium for products that cost nothing to distribute. The “losers,” such as newspaper publishers and camera companies, could not compete with zero on a marginal-cost basis; they shrunk dramatically as consumers opted for much cheaper (often free) alternatives.



What, then, are the enduring sources of pricing power in low-marginal-cost industries today? The list below is a non-exhaustive starting point, though AI could evolve to interact with these advantages in unpredictable ways.

- **Barriers to entry from high up-front costs.** The classic example of this phenomenon is utilities, which distribute their products at very low marginal costs—but only after making large up-front investments in infrastructure. This huge capital barrier creates tremendous pricing power, which is a reason that utilities’ profits are so highly regulated.

Providers of large language models benefit from this dynamic because their models are very expensive to train, creating a hurdle for aspiring competitors (which may increase as models get ever more sophisticated). As their products become increasingly important to the economy, these companies may eventually face analogous regulations. But until that occurs, they have a great deal of pricing power.

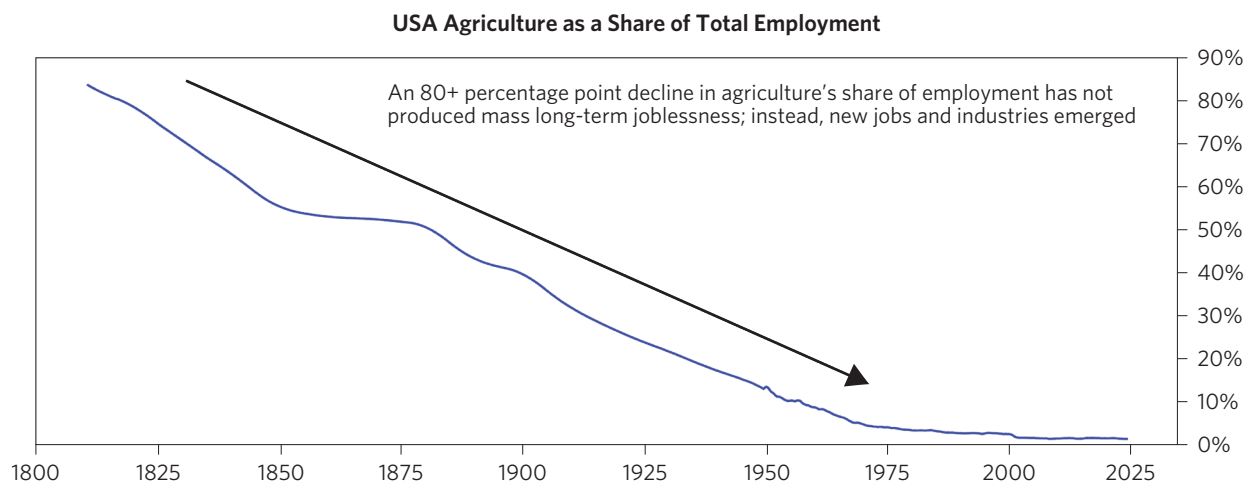
- Quality differentiation.** People tend to pay for products they find uniquely valuable, even in the face of lower-quality, cheaper alternatives. This is one source of pricing power for some software providers with ultra-low marginal costs; it's similarly helped some flagship news organizations preserve or even expand their margins in recent years, while many outlets with less ability to differentiate themselves have folded. (Of course, that is only after a decades-long, industry-wide decline.)

For AI-powered products, the clearest source of quality differentiation will be access to high-quality proprietary data, which makes it possible to train tools that are fine-tuned for specific tasks. Access to top scientific talent could also be a major driver of differentiation, enabling companies to create standout products by developing superior algorithms instead of (or in addition to) utilizing better data.

- Switching costs** are another crucial source of pricing power. Tech companies like Google have even faced litigation over efforts to get their applications preinstalled on devices, a way of embedding switching costs at the outset. Switching costs may not be as big a factor for AI, as AI tools themselves may become able to do much of the grunt work of converting over to a new platform. But other drivers of switching costs, like the costs of searching for a new product or teaching human employees to use a new tool, will still apply. In many cases, existing software companies that roll out AI-powered tools to their customer bases will benefit from these hurdles, piggybacking off the switching costs already inherent to their products.

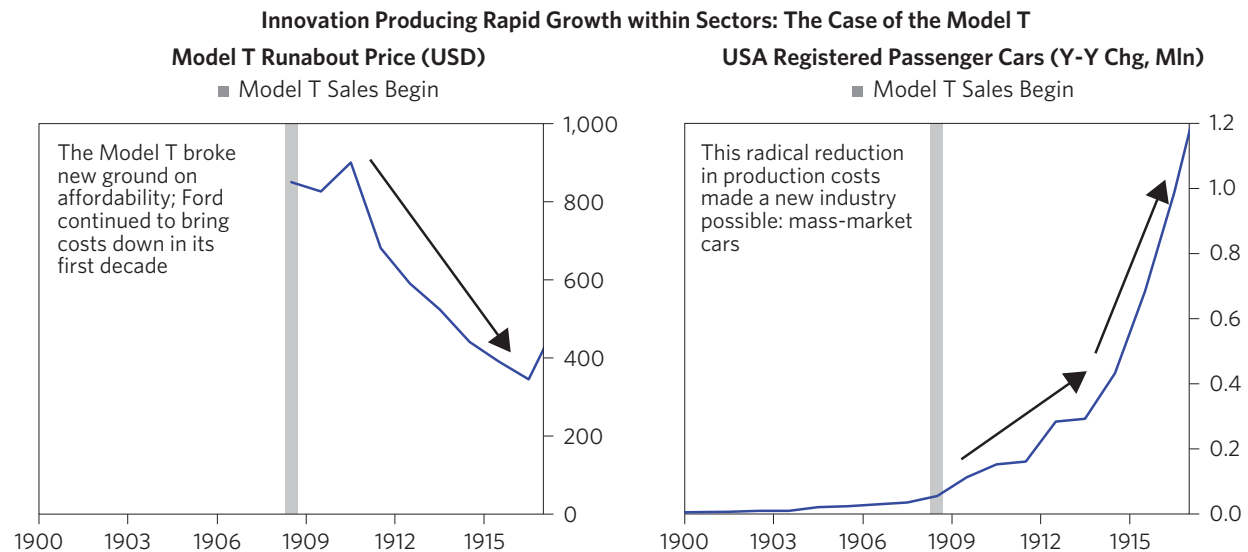
## Expect a Reshuffling of Spending and Activity as New Industries Become Economically Viable and Disposable Income Is Freed Up

Automation can produce painful short-term dislocations when it happens quickly. But over the long run, productivity-enhancing technologies free up time and money to spend in new areas. Most of the jobs that exist today did not exist 50 years ago. Looking farther back, agriculture made up close to 85% of total US employment around the start of the 19th century; today, agricultural workers make up only about 1.4% of the workforce. Our workforce gradually expanded in other sectors as fewer people were needed to farm and consumer demand shifted.



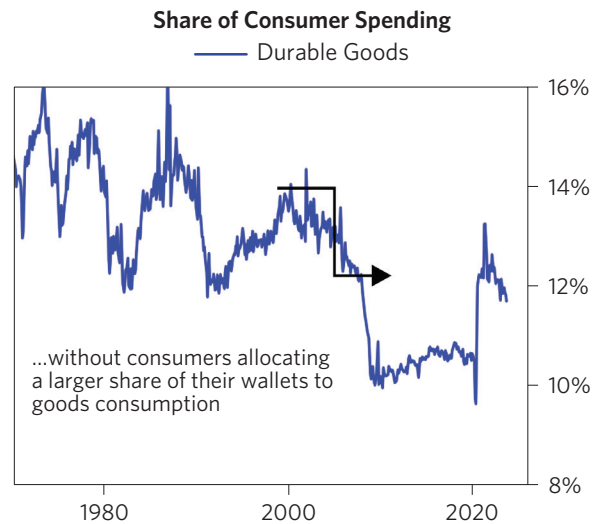
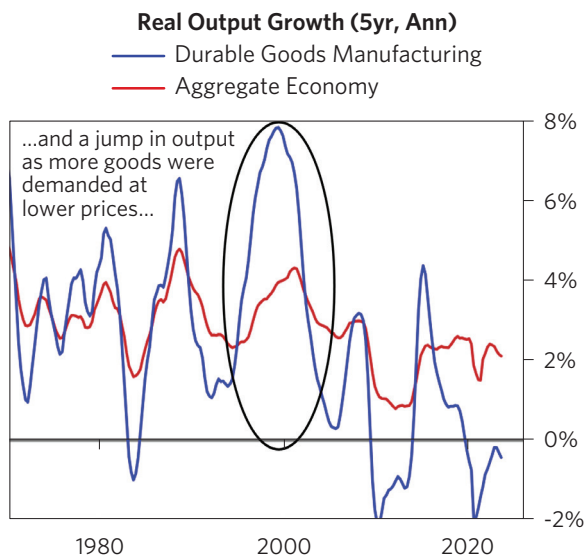
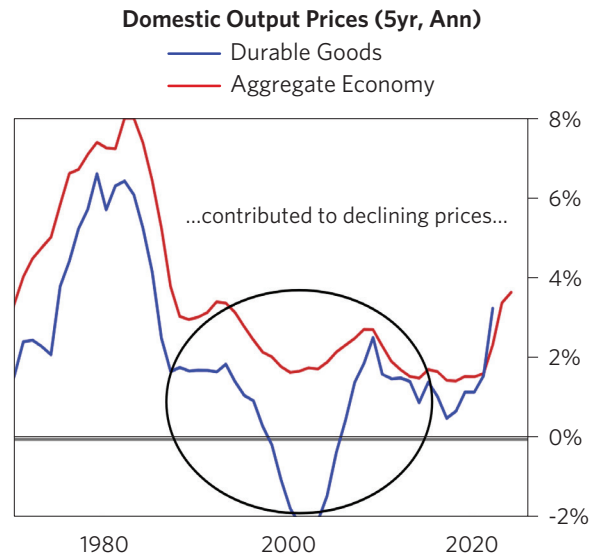
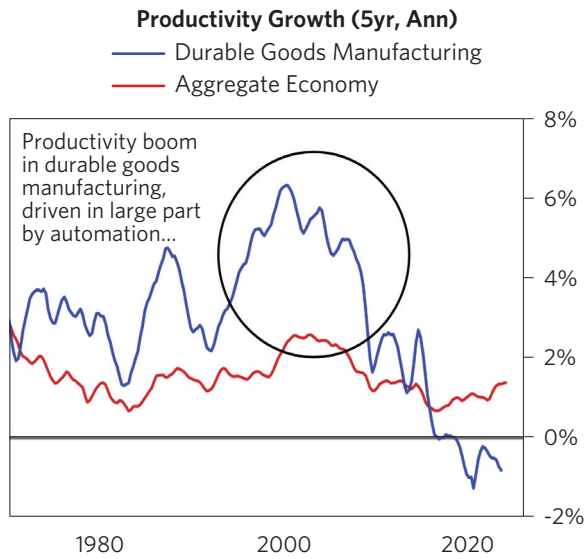
If AI can perform economically valuable cognitive tasks at very little marginal cost, it could similarly open up new horizons for spending and economic activity. The selling prices of many products will likely fall as AI brings their marginal costs down to zero, even where firms have strong pricing power (as these firms will often maximize profits by increasing sales volumes at lower prices).

In some cases, we will see a “**Model T effect**”: consumers buying a lot more when the price of a previously unaffordable item is slashed, causing entire new industries to spring up. Around the turn of the 20th century, the revolutionary assembly line process used to produce the Ford Model T lowered the cost of manufacturing a car so much (and so quickly) that a whole new market for mass-market cars came into being. AI could similarly make a host of previously high-cost items very cheap to provide, such as tailored legal advice, customized software builds, individualized financial planning, and one-on-one tutoring. Consumption of these kinds of cognitive services could balloon, though pricing power considerations will ultimately determine whether the gains end up primarily with providers as profits or consumers as surplus.



Lower prices in one sector also often lead to higher spending in unrelated discretionary sectors because people effectively have more to spend elsewhere. A similar phenomenon (though not productivity-driven) happens over shorter time horizons when energy prices fall and consumers shift wallet share toward more discretionary categories of spending. Over longer periods, this dynamic generally leads to rising activity and employment in sectors facing increased demand. When supply in those sectors is inherently limited—such as in real estate—it can also produce inflation, as cost savings in one area are used to bid up the price of truly scarce resources in another.

We’ll likely see a mix of the two dynamics as AI lowers production costs: customers will buy more at the new lower price but will also spend some of their savings in other categories. Similar (though smaller scale) versions of these dynamics have occurred more recently. The automation of goods production in the 1990s and 2000s led to rising output as people bought more or higher-quality goods at lower prices, but higher goods consumption didn’t crowd out other forms of spending. In the later part of that period, durable goods even shrank as a share of consumer spending as people spent their savings elsewhere (especially in more discretionary services sectors).

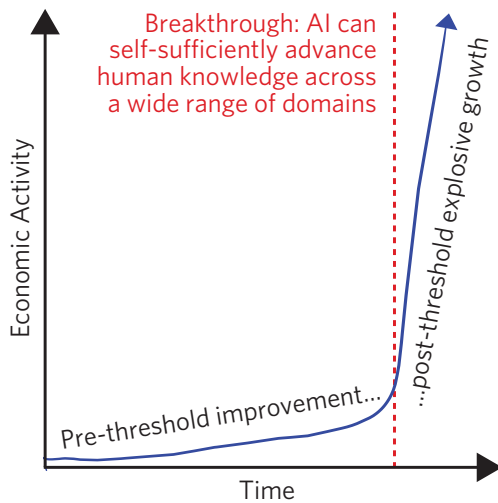


The digitization of photos may also be an applicable precedent for some AI capabilities. As producing photos became virtually free, consumption became essentially infinite, and much of the immediate surplus went to consumers at the expense of traditional photo companies. But shortly thereafter, very profitable new companies (such as social media) reliant on that infinite number of free photos sprang up.

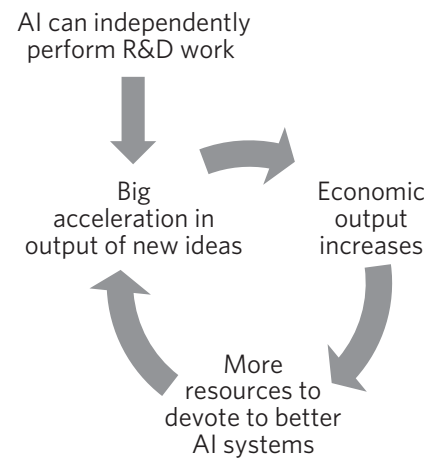
# The Most Extreme Outcomes Occur if the Marginal Cost of Innovation Itself Falls to Zero

While we have so far explored AI in the context of prior innovations, it's possible that AI fundamentally diverges from historical precedent. **Whereas prior innovations automated specific tasks or assisted human efforts to distribute and advance knowledge, a very sophisticated AI could automate the advancement of human knowledge itself.** An AI tool that could drive research on its own could add the equivalent of billions of human researchers to the R&D ecosystem, producing an innovation boom. The resulting growth would compound, as AI-powered researchers could further improve AI technology, and the rapid economic growth their inventions catalyze would make it possible to direct more resources toward creating even better AI tools. This world would be very different from those where AI capabilities do not progress as much; it could even entail “explosive growth,” i.e., an order-of-magnitude increase in the growth rate.

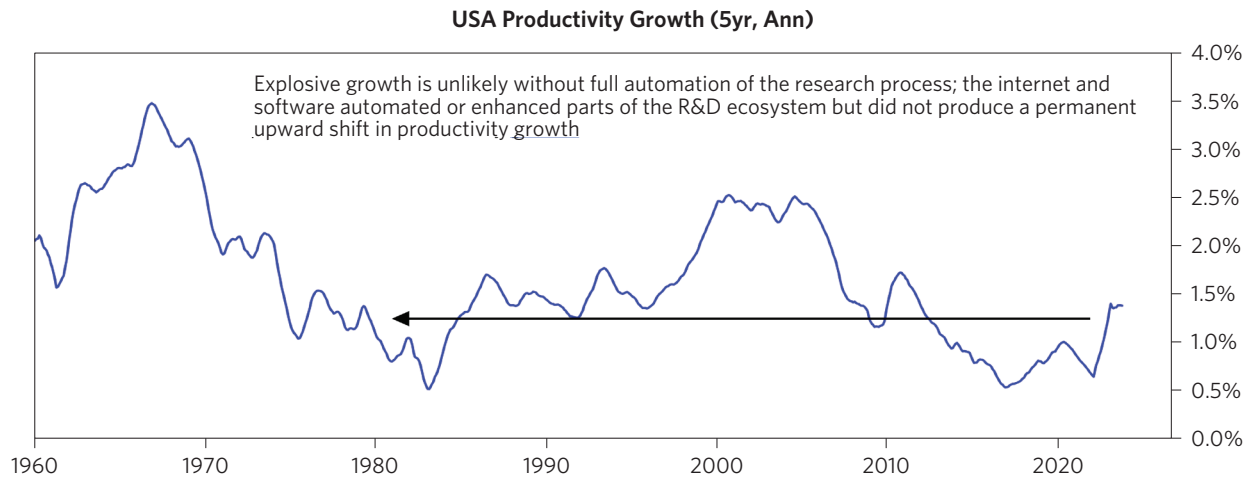
Conceptual Map: Explosive Growth



Explosive Growth Cycle



Only technology that fully automates the research process would likely produce the kind of step change in innovation necessary to see explosive growth. This would probably require human-equivalent (or better) intelligence across most intellectual tasks—what’s known as artificial general intelligence, or AGI. An AI that fell short of this threshold could still be incredibly fruitful for research but probably wouldn’t overcome some crucial hurdles to explosive growth. Past innovations with extensive uses in the R&D ecosystem—like software and the internet—have not produced a permanent shift in the pace of productivity growth. Research still faces bottlenecks, like the limited supply of human researchers and declining productivity per researcher (productive ideas getting “harder to find”),<sup>1</sup> offsetting the gains from those technologies.



Whether developing AGI is even possible is the subject of much academic debate. But even if we assume AGI is possible, the question of how far away it is—both technologically and chronologically—remains. There are no perfect metrics in an imaginative exercise like this, but we highlight a few commonly cited ones below:

**Resources**

- *Distance to traverse:* How much compute would it take to train an AGI with today’s algorithms?
- *Resources to get there:* How much money will be invested in training models?

**Efficiency of Deploying Those Resources**

- *Hardware improvement:* How much more powerful will chips become?
- *Algorithmic improvement:* How much better will model training get per unit of compute?

<sup>1</sup> Declining research productivity has been much discussed and shown to be empirically true across a broad range of metrics representing ideas. What this suggests about the nature of human ingenuity and innovation itself is more of an open question, and one we’d imagine the advent of super-powerful AI would throw into sharp relief.



Some models based on metrics like these suggest we could reach AGI in the coming decades. Yet these attempts assume, but cannot prove, that hardware and algorithmic improvements will continue to happen rapidly. This is not a given. Additional improvements might be a lot harder, even with AI tools enhancing the productivity of researchers, as R&D has become incrementally harder through time. For instance, the number of researchers required to achieve Moore’s law, the “rule” that the density of computer chips doubles around every two years, is over 18 times what it was in the early 1970s—and we don’t know how much harder these advancements will get in the future, or whether they will be possible at all.<sup>2</sup> More importantly, forecasts about the development of AGI rest on highly speculative and untestable assumptions about how much more powerful models would need to become to develop humanlike intelligence.

Finally, even if we were to develop AGI, would it lead to explosive growth? Critics here point out that the testing and implementation of an AGI’s ideas would likely rely on processes that naturally take time or face resource constraints, like conducting experiments or sourcing and processing raw materials. Alongside potential obstacles from regulation or a lack of suitable training data for many tasks, these delays could prevent the rapid compounding that is necessary for explosive growth. Some say that super-advanced AI tools could just innovate their way out of these problems, but this claim is very challenging to assess from today’s vantage point. **Given these considerations, full-blown explosive growth looks unlikely. But we wouldn’t rule it out at this early stage—and assigning it even a tiny probability would dramatically alter our average expectations for what the world will look like decades from now.**

---

<sup>2</sup> For this finding, see: Nicholas Bloom et al., “Are Ideas Getting Harder to Find?,” *American Economic Review* 110, no. 4 (April 2020): 1104-44.

# Important Disclosures and Other Information

This research paper is prepared by and is the property of Bridgewater Associates, LP and is circulated for informational and educational purposes only. There is no consideration given to the specific investment needs, objectives, or tolerances of any of the recipients. Additionally, Bridgewater's actual investment positions may, and often will, vary from its conclusions discussed herein based on any number of factors, such as client investment restrictions, portfolio rebalancing and transactions costs, among others. Recipients should consult their own advisors, including tax advisors, before making any investment decision. This material is for informational and educational purposes only and is not an offer to sell or the solicitation of an offer to buy the securities or other instruments mentioned. Any such offering will be made pursuant to a definitive offering memorandum. This material does not constitute a personal recommendation or take into account the particular investment objectives, financial situations, or needs of individual investors which are necessary considerations before making any investment decision. Investors should consider whether any advice or recommendation in this research is suitable for their particular circumstances and, where appropriate, seek professional advice, including legal, tax, accounting, investment, or other advice.

The information provided herein is not intended to provide a sufficient basis on which to make an investment decision and investment decisions should not be based on simulated, hypothetical, or illustrative information that have inherent limitations. Unlike an actual performance record simulated or hypothetical results do not represent actual trading or the actual costs of management and may have under or overcompensated for the impact of certain market risk factors. Bridgewater makes no representation that any account will or is likely to achieve returns similar to those shown. The price and value of the investments referred to in this research and the income therefrom may fluctuate. Every investment involves risk and in volatile or uncertain market conditions, significant variations in the value or return on that investment may occur. Investments in hedge funds are complex, speculative and carry a high degree of risk, including the risk of a complete loss of an investor's entire investment. Past performance is not a guide to future performance, future returns are not guaranteed, and a complete loss of original capital may occur. Certain transactions, including those involving leverage, futures, options, and other derivatives, give rise to substantial risk and are not suitable for all investors. Fluctuations in exchange rates could have material adverse effects on the value or price of, or income derived from, certain investments.

Bridgewater research utilizes data and information from public, private, and internal sources, including data from actual Bridgewater trades. Sources include BCA, Bloomberg Finance L.P., Bond Radar, Candeal, Calderwood, CBRE, Inc., CEIC Data Company Ltd., Clarus Financial Technology, Conference Board of Canada, Consensus Economics Inc., Corelogic, Inc., Cornerstone Macro, Dealogic, DTCC Data Repository, Ecoanalitica, Empirical Research Partners, Entis (Axioma Qontigo), EPFR Global, ESG Book, Eurasia Group, Evercore ISI, FactSet Research Systems, The Financial Times Limited, FINRA, GaveKal Research Ltd., Global Financial Data, Inc., Harvard Business Review, Haver Analytics, Inc., Institutional Shareholder Services (ISS), The Investment Funds Institute of Canada, ICE Data, ICE Derived Data (UK), Investment Company Institute, International Institute of Finance, JP Morgan, JSTA Advisors, MarketAxess, Medley Global Advisors, Metals Focus Ltd, Moody's ESG Solutions, MSCI, Inc., National Bureau of Economic Research, Organisation for Economic Cooperation and Development, Pensions & Investments Research Center, Refinitiv, Rhodium Group, RP Data, Rubinson Research, Rystad Energy, S&P Global Market Intelligence, Sentix GmbH, Shanghai Wind Information, Sustainalytics, Swaps Monitor, Totem Macro, Tradeweb, United Nations, US Department of Commerce, Verisk Maplecroft, Visible Alpha, Wells Bay, Wind Financial Information LLC, Wood Mackenzie Limited, World Bureau of Metal Statistics, World Economic Forum, YieldBook. While we consider information from external sources to be reliable, we do not assume responsibility for its accuracy.

This information is not directed at or intended for distribution to or use by any person or entity located in any jurisdiction where such distribution, publication, availability, or use would be contrary to applicable law or regulation, or which would subject Bridgewater to any registration or licensing requirements within such jurisdiction. 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 Bridgewater® Associates, LP.

The views expressed herein are solely those of Bridgewater as of the date of this report and are subject to change without notice. Bridgewater may have a significant financial interest in one or more of the positions and/or securities or derivatives discussed. Those responsible for preparing this report receive compensation based upon various factors, including, among other things, the quality of their work and firm revenues.