**What was I made for: Large Language Models in the Real World**

On markets, not much has changed since our August **Rasputin** piece. While most leading indicators point to weaker US growth by Q1, the expected decline is modest as potential recessions go. Tighter Fed policy is partially offset by large fiscal deficits, US industrial policy (incentive-driven spending on infrastructure, energy and semiconductors), strong corporate and household balance sheets, private sector actions to lock in low borrowing rates before 2022 and tight housing/labor markets. Equities have been sustained by the anomaly of equity valuations rising at a time of muted earnings growth, and the AI catalyst (more on that shortly).

If there's anything new, **it's the rise in oil prices**. Three things to keep in mind: (a) OPEC spare capacity is at one of its highest levels outside recessions, (b) for all the obvious reasons, publicly traded energy companies are spending a small share of cash flow on future oil & gas projects despite no decline (yet) in global fossil fuel use, and (c) the US strategic petroleum reserve is already at its lowest level in decades. I don't think the Biden Administration has any serious intention to refill the SPR despite the national security implications, and there are questions about the SPR's structural integrity given more frequent drawdowns than originally intended.
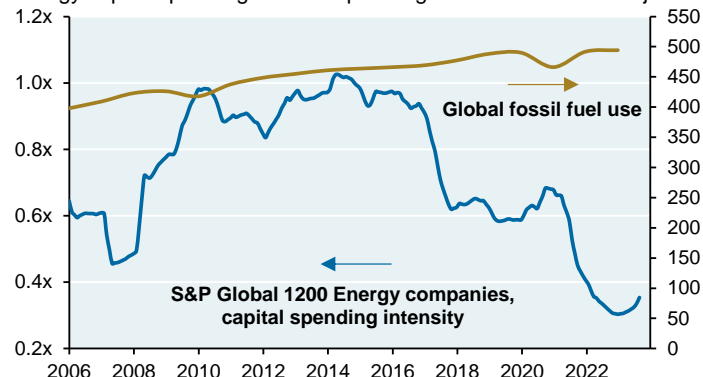
Estimated OPEC spare capacity
Million barrels per day
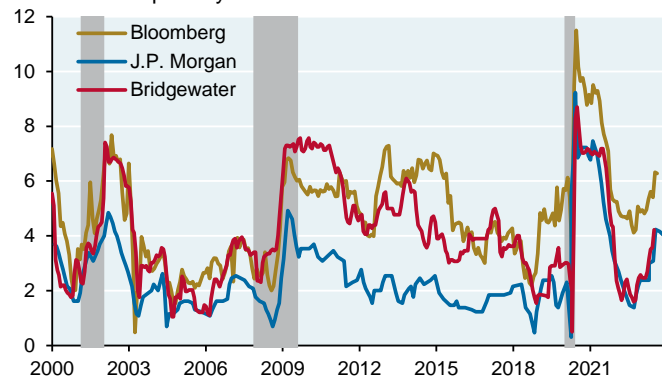
Source: Bridgewater, Bloomberg, J.P. Morgan. August 2023.

Fossil fuel consumption vs energy capital spending
Energy capital spending as % of operating cash flow   Exajoules

Source: Bloomberg, Energy Institute, JPMAM. August 2023.

Higher oil prices often feed into core inflation within a few months. So, a dose of reality: given rising energy prices and sticky wage inflation, futures markets are no longer pricing in a decline in Fed policy rates early next year as they were a few months ago, which made no sense to us at the time.

Fed funds target rate: current vs prior market expectations
Percent
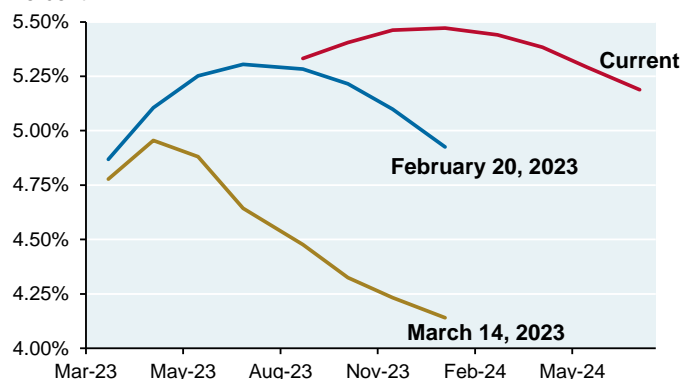
Source: Bloomberg, JPMAM. September 20, 2023.

Strategic petroleum reserve total inventory
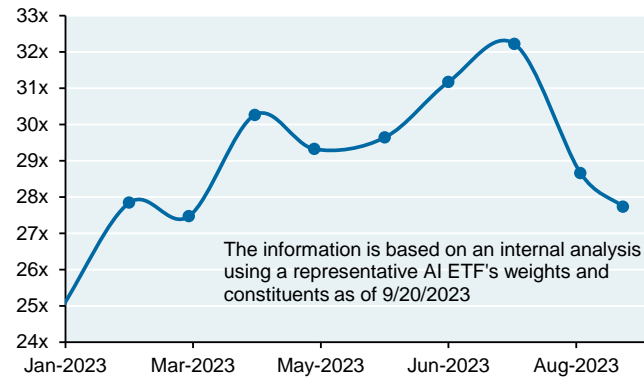Barrels (millions)

Source: US Department of Energy, JPMAM. September 15, 2023.

The major US equity catalyst this year has been the rise in AI-linked stocks. They've come off the boil since July, but there's still a lot of optimism regarding AI's impact on growth, profits and productivity.
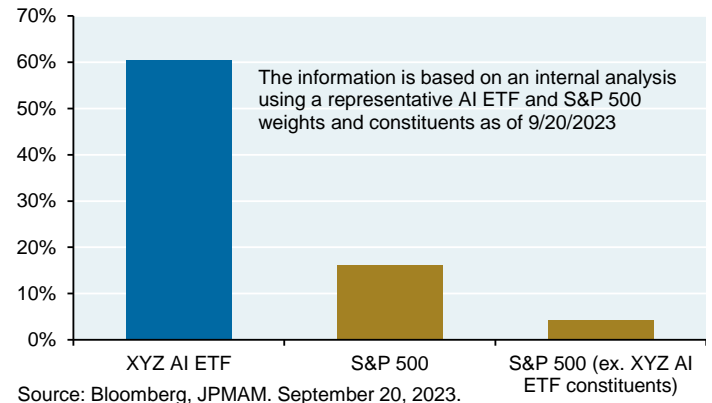
XYZ Generative AI ETF P/E ratio
Price / 12m forward earnings

The information is based on an internal analysis using a representative AI ETF's weights and constituents as of 9/20/2023

Source: Bloomberg, JPMAM. September 20, 2023.

YTD return of XYZ AI ETF vs S&P 500
Percent, YTD return

The information is based on an internal analysis using a representative AI ETF and S&P 500 weights and constituents as of 9/20/2023

Source: Bloomberg, JPMAM. September 20, 2023.

Is this justified? Let's start with some artwork from multimodal AI image generation models. I asked them to create an image showing two workers looking nervously at a robot at a table with them. I also asked that the robot be labeled "Strategy Team Trainee". The model outputs combine impressive interpretive proficiency with strange errors. Bing and Starry AI ignored the robot label request entirely while Dall-E splattered gibberish letters on the table; Bing and Dall-E ignored the request for the workers to be looking at the robot; Starry AI left out one of the people; and Bing added a third person that looks like they're in a horror film.

This mix of good, bad and bizarre is a good place to start. What follows is an assessment of generative AI in the real world, including my own: I graded GPT 4.0 based on its responses to 71 questions from the Eye on the Market (see pages 9-10 for the question list). Results were mixed, and that's how I view AI's long-run impact on equity markets: better than recent investment fads but not the transformational productivity shock that its most vocal adherents expect.

Bing

Starry AI

Dall-E

2
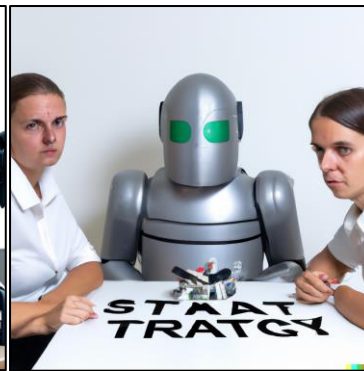
What was I made for: Large Language Models in the Real World

Earlier this year I sent a note to colleagues showing questions I asked of large language models (LLMs). These questions were designed to trick them and sure enough, many LLMs got some answers wrong; see box below for my favorite wrong answer. But Manuela Veloso who runs JP Morgan's AI Research group took me to task for this since my questions were related to math that LLMs can solve with additional prompting[1], obscure US presidential trivia and other mostly useless information[2].

Manuela walked me through generative AI use cases that her group is currently exploring. These projects seem capable of improving productivity when used for software development, charting and information checking, particularly when run internally on corporate data. I was very impressed by the end of the discussion.

Manuela then challenged me to think about how a large language model could help me in my own role. So, that's what I did: I asked Chat GPT-4 questions that my analysts and I worked on over the last two years.   This piece reviews the results, along with the latest achievements and stumbles of generative AI models in the real world and comments on the changing relationship between innovation, productivity and employment.

Results: Chat GPT-4 did a good job on some tasks but was a waste of time on others. Its GPA on the 71 questions: 2.50 out of 4.00. Would it improve my research process? It could speed certain things up, such as locating data in the vastness of the internet. But its computational and reasoning mistakes weigh heavily since I had to double-check everything it did to make sure no errors got through, and it took a time to find and fix them when they did. How real are my productivity gains if I have to check every answer anyway?

A large language model can process massive amounts of text efficiently; that's what it's made for. But it cannot think or reason; it's just something I paid for. At just $20 per month for GPT-4, the price was right.

Michael Cembalest
Chairman of Market and Investment Strategy
JP Morgan Asset Management

My favorite wrong LLM answer

One LLM cited Franklin Delano Roosevelt, creator of the New Deal, as a **Republican**. Republicans of FDR's era would have objected strenuously to that. According to FDR scholar William Leuchtenburg, J.P. Morgan Jr's family kept newspapers with pictures of FDR out of his sight to avoid angering him, and in one Connecticut country club, mention of FDR's name was forbidden as a "**health measure against apoplexy**". Even Democrat Al Smith whom FDR succeeded as governor of New York described FDR's New Deal policies as socialist. So no, FDR was not a Republican.

"*The FDR Years: On Roosevelt and His Legacy*", William Leuchtenburg (UNC/Chapel Hill), April 3,1987

---

[1] Math problems can be more reliably solved by LLMs when allowing access to an external Python interpreter that has access to mathematical functions, or to a variety of plug-ins such as Wolfram.

[2] Chat GPT-4 got these two questions wrong repeatedly: "Can you give me a six-letter word with only one consonant?" and "If I want to take a time machine back to when the Winter Olympics were held in a year whose digits add to 23, which year should I travel to?" **I concede that these questions have no real purpose other than to torture the LLM into getting them wrong.**

Something I'm made for: Generative AI use cases expand

There's been some progress in real-life LLM applications. Notable examples:

- A September 2023 paper from Harvard Business School and Boston Consulting Group found substantial improvements in "performance quality" when arming **management consultants** with GPT-4. The degree to which you are impressed by this will depend on how impressed you are generally with consultants[3]

Impact of GPT-4 on consultant task response quality
Density (normalized frequency)
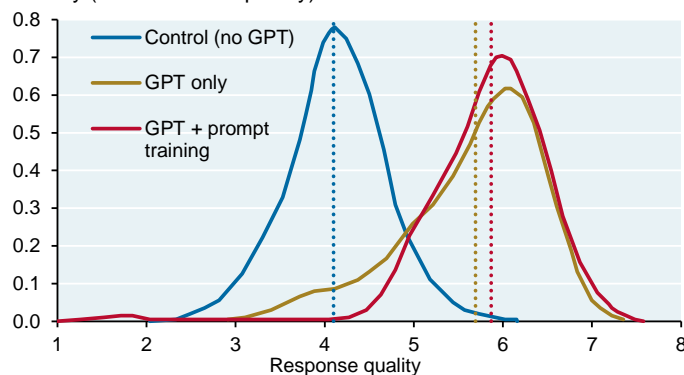
Source: HBS / BCG, "Navigating the Jagged Technological Frontier." 2023.

HBS/BCG: effects of AI on management consulting
- 12% more tasks completed, tasks completed 25% faster, quality improved by 40%
- Tasks: idea generation, segment analysis, marketing
- GPT-4 was as a skill leveler: lowest ranked consultants had the biggest improvement in their performance
- In some tasks, excessive reliance on AI backfired since workers did not apply enough judgment and checking

"_Navigating the Jagged Technological Frontier_", HBS/BCG, Fabrizio Dell'Aqua et al, September 2023

- GitHub's **AI programming tool** Co-pilot already has 1 million users for whom it writes 45% of their code. GitHub's CEO believes that will eventually rise to 80%[4]. GPT-4 can translate code from Python to Java and C++, and to obscure languages like Julia that are used in computational science. However, a lot of code used in enterprise applications is copy-paste anyway, so the productivity benefits need careful measurement

- Data scientists evaluated ChatGPT on **statistical exams** which involve coding and conceptual problems. Without any prompt engineering, ChatGPT scored 104 points out of a possible 116[5]

- MIT analysts examined LLMs in the context of **professional writing**. They assigned occupation-specific writing tasks to college-educated professionals and randomly exposed half of them to ChatGPT. Access to LLMs raised productivity: time required for each task declined and output quality improved. Inequality between workers decreased as ChatGPT benefited low-ability workers more[6]

- LLMs for **customer support agents** improved productivity by ~15% when measured as issues resolved per hour, improved customer sentiment and employee retention, and helped lowest skilled workers the most[7]

- LLM testing in **medical research.** A 2023 medRxiv pre-print analysis found that an AI tool that summarized doctor-patient interactions improved skin condition diagnoses[8]. Separately, a 2023 study in JAMA fed 70 notoriously difficult-to-diagnose medical cases into GPT-4[9]. The results: GPT-4 scored 4.2 out of 5.0, and for 64% of cases the correct diagnosis was included in GPT's response.

---

[3] There aren't many empirical studies on the success rate of management consulting services. **One such study found that using management consulting services simply leads to greater use of management consultants**. "_The management consultancy effect: Demand inflation and its consequences in the sourcing of external knowledge_", Sturdy et al, University of Bristol, 2020

[4] "_GitHub CEO says Copilot will write 80% of code sooner than later_", Freethink.com, July 7, 2023

[5] "_What Should Data Science Education Do with Large Language Models?_", James Zou (Stanford) et al, July 2023

[6] "_Evidence on the Productivity Effects of Generative Artificial Intelligence_", Noy and Zhang (MIT), March 2023

[7] "_Generative AI at Work_", Brynjolfsson (Stanford), Li and Raymond, 2023

[8] "_Testing the Limits of Language Models_", Johri et al (Harvard Medical School), September 12, 2023

[9] "_Accuracy of a Generative AI Model in a Complex Diagnostic Challenge_", Kanjee et al, JAMA, June 15, 2023

Warning: you're not going to like all LLM use cases. People are using LLMs to generate reams of fake content, fake web pages, fake product reviews, fake eBooks and phishing emails[10]. This spam gets passed off as the real thing to content aggregators, eBook buyers and other consumers of information who unwittingly pay directly or indirectly (via advertising) for this junk. Some of it seems designed solely to profit from Google's automated advertising process. One fake news site posted 1,200 AI generated articles in one day and ended up receiving ads from 141 different corporate advertisers[11]. **Can LLMs be used to spot fake news sites that other people created using the very same LLM tools?** The news is not great: on small sample sizes they do fine, but exhibit unpredictable hallucinations on larger ones. Four LLMs tested at the University of Wisconsin were unable to spot fake news around one third of the time and performed worse than human fact-checkers[12].

All of these use cases have created a frenzy of analysts comparing large language models and other generative AI to 20ᵗʰ century milestones such as the electrification of farms, the interstate highway system and the internet itself. One example: Goldman Sachs has gone all-in regarding the potential impact of large language models on corporate profits. Goldman's analysts estimated the potential earnings boost due to AI adoption and increased labor productivity for each company in the Russell 1000, concluding that the median company's earnings could rise 19% vs current levels and that ~8% of companies could see an average rise of 60% or more.[13]

If Goldman is right about AI, what's the consequence for employment? Most AI firms, economists and investment banks argue that while productivity displaces some workers, it creates enough new jobs to more than offset the losses. Unfortunately, this simple axiom is not always true as we illustrate in the Appendix based on work from Daron Acemoglu and David Autor at MIT.

Why didn't I include this OpenAI chart on GPT-4 vs GPT-3.5 as evidence of LLM improved real-world impact? That's what we will discuss next.

GPT-4 improvements vs GPT-3.5
GPT percentile vs human test takers
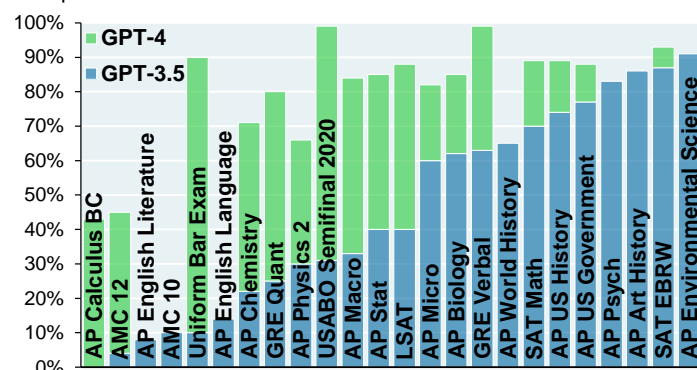
Source: OpenAI. 2023.

Is GPT-4 really this good at math and calculus if it also sometimes struggles with middle school geometry? Or is this chart measuring something else?

---

[10] "*Pivot to AI*", Amy Castor and David Gerard, September 12, 2023

[11] "*AI-generative models are driving a surge in content on fake news sites*", The Week, July 18, 2023

[12] **Accuracy in detecting fake news: Bard 64%, Bing 64%, GPT-3.5 62% and GPT-4.0 71%.** See "*A Comparative Performance Evaluation of ChatGPT 3.5, ChatGPT 4.0, Bing AI, and Bard in News Fact-Checking*", Kevin Caramancion, University of Wisconsin Department of Math, Statistics and Computer Science, 2023

[13] "*Identifying potential long-term EPS beneficiaries of AI adoption*", Hammond et al (GS), August 21, 2023

It's not what I'm made for: large language models and their limitations

That OpenAI chart bars shows GPT's performance on primarily multiple choice exams (average multiple choice share of exam score 68%). But…many LLMs can be trained on datasets which contain prior exam questions and answers, something referred to as "contamination". According to one analysis, many datasets are contaminated while others are highly suspected of being so[14]. **So, what's really measured is GPT's ability to respond to a question whose answer it has probably already seen or one that is highly similar to it**. The chart may simply be telling us that GPT 4.0 is getting better at memorization given 10x more parameters than GPT 3.5.

> **Testing for LLM contamination.** "Simple contamination tests based on substring checks such as those carried out by OpenAI in their March 2023 GPT-4 Technical Report are not sufficient to guarantee lack of contamination. Substring checks carried out by OpenAI were not applied on the entire problem instance, only on 3 randomly selected substrings of 50 characters each. This is not enough to ensure disjointness for long (or even moderately long) problems, which are quite common in tests like the Uniform Bar Exam". "_GPT-4 can't reason_", Konstantine Arkoudas (RPI/MIT), August 2023

There's also a more fundamental issue to consider: the real-world value of a lawyer is not the ability to answer multiple choice bar exam questions all day. Many professional disciplines require reasoning and synthesis of new information. **And when tested on more complex synthesis of ideas, many LLMs still struggle:**

- *Law*. GPT-4 got a C in Constitutional Law and a C- in Criminal Law from professors at the University of Maryland who noted that "GPT-4 produced smoothly written answers that failed to spot many important issues, much like a bright student **who had neither attended class nor thought deeply about the material**"[15]

- *Actuaries*. GPT-4 failed an actuarial exam this summer, registering a 19.75 out of a possible 52.50 score[16]

- *Economics*. Professor Steve Landsberg at the University of Rochester had GPT-4 take his college sophomore economics exam; GPT-4 scored just 4 out of a possible 90[17]

- *Taxes*. GPT-4 gets tax liabilities exactly right only 1/3 of the time and miscalculates tax liabilities by over 10% nearly a quarter of the time due to a misreading of the statutes[18]; and performed "terribly" on graduate level tax and trust & estates exams[19]

- *Math teaching*. GPT-4 can be a poor teacher. It botched explanations of Pythagoras' theorem, instructed users that if you know the hypotenuse of a right triangle, that's enough info to determine the length of both sides; and provided misinformation on geometric concepts until it was informed that its answers were wrong. GPT-4 also got stuck in a "death loop" of nonsense when provided with mathematically impossible dimensions of a triangle that it should have been able to recognize and point out upfront[20]

- *Journalism*. Online editors cite a growing amount of AI-generated content that is so far beneath their standards that they consider it a "new kind of spam". As a result, they reject all AI-written submissions since they have perfect spelling and grammar but lack a coherent story, and are useless to them[21]

- *The Big Bench project*. The most comprehensive LLM assessment that we've seen is "BIG-bench"[22]. This project, encompassing 204 tasks compiled by 400+ researchers, still finds substantial underperformance of LLMs compared to the average human, and well below highly performing humans

---

[14] "_LM Contamination Index_", Agirre et al, University of the Basque Country UPV/EHU, May 30, 2023

[15] "_GPT-4's Law School Grades_", Blair-Stanek et al, University of Maryland, May 25, 2023

[16] "_GPT-4 Fails Final Actuarial Exam_", March 29, 2023, David Wright, Market Solutions

[17] "_GPT-4 Fails Economics_", Steve Landsburg, University of Rochester, April 5, 2023

[18] "_Can GPT-4 really do tax?_", Blair-Stanek (U Maryland), Holzenberger and Van Durme (Johns Hopkins), August 14, 2023

[19] James Hines at the University of Michigan (Professor of Law & Economics, Director of Office of Tax Policy Research)

[20] "_Chat GPT-4 is still not ready to teach geometry_", Paul Von Hippel, UT Austin, April 2023

[21] "_AI Junk is Starting to Pollute the Internet_", WSJ, July 12, 2023

[22] "_Beyond the Imitation Game_", BIG-bench (Google et al), June 2022, revised June 2023. Tasks include linguistics, math, common-sense reasoning, biology, physics, social bias, software development, etc

Is this what I'm made for? Large language models and the Eye on the Market

I took Manuela's advice and evaluated GPT-4 for myself. The ground rules:

- I asked 71 questions (see pages 9-10) that we worked on over the last two years for the Eye on the Market
- I graded GPT-4's speed and accuracy vs our normal process. **In other words, GPT-4 was evaluated against a process that works for us without yielding hallucinations and errors, even if our research process takes time**. Speed isn't worth it if it comes with too many mistakes
- **We enabled GPT-4 features to upload data files[23] which we prepared for it when it was unable to source the data itself. We also enabled several plug-ins that allow web browsing of PDFs and excel files when necessary. When testing different plug-ins, we used GPT-4's best answer. As a result, the end of GPT-4's parameter training in 2021 was not a constraint on its ability to answer our questions**
- Grading was affected by the consistency of GPT-4's response (lower grades for less consistency)
- Wrong answers were more heavily penalized than no answer given extra work needed to find and fix them
- We provided follow-up prompts when initial answers were too broad or did not answer the question

Results. GPT-4's performance was bimodal: a lot of good grades but a lot of bad grades as well, averaging out to a ~2.50 GPA which is a between a C+ and a B-

Chat GPT-4 grading

Grade	# of questions	Grade	# of questions
A	26	C	7
A-	5	C-	2
B+	0	D+	0
B	3	D	13
B-	3	D-	2
C+	2	F	8

Chat GPT-4 GPA by Subject	
Economics	2.62
Markets	2.31
Energy	2.70
Politics	2.06
Overall GPA	2.47

Source: Cembalest assessments

Examples of GPT-4 mistakes[24]:

- it would hallucinate numbers and then refuse to provide a source for where it found them
- it would outline the correct steps to solve a problem and then execute them incorrectly
- it would misread data files it found or that we provided
- it didn't notice that subtotals should be excluded from summation calculations
- it used the wrong constants for certain energy conversions
- it asserted certain facts that are contradicted by other readily available information
- GPT-4 complained about one unanswered question by saying "this is a complex analysis that requires data gathering, data cleaning and statistics". OK but by the time we've done all that, what do we need GPT for?

The 71 questions appear on pages 9-10 with my grading assessments. The **supplemental attachment** contains my questions, GPT-4's answers and the correct responses if you're interested in all the details.

---

[23] Data files are uploaded using the **Advanced Data Analysis** feature in Chat-GPT. This is a Python environment that can perform a wide range of mathematical operations and analysis, including basic arithmetic, algebra, calculus, statistics, linear algebra and numerical methods

[24] Other LLMs have problems of their own. **Google claims to have integrated Bard with other Google tools like its email and calendar**. Yet when my son asked Bard for the last time that I had sent him an email via Google, it answered wrong (!!). When he told Bard it was wrong, it gave an even worse answer

How can we make sense of GPT's mixed success? I was pleased at how well GPT-4 did on many questions, but the time spent finding and fixing the wrong answers was exhausting. And I'm not the only one:

- Meta Chief AI scientist Yann LeCun: LLMs are useful writing aids, they don't reason, they don't plan, they don't really understand and make stuff up. Why are they so good at writing code? "Unlike the real world, the universe that a program manipulates is limited, discrete, deterministic and fully observable"[25]

- NYU's Gary Marcus: "Large Language Models are trained to produce outputs that are *plausible*, but not necessarily *correct*. Although their abilities are often uncanny, they are lacking in aspects of reasoning, leading LLMs to be less than completely trustworthy. Furthermore, their results tend to be both unpredictable and uninterpretable."[26]

- Konstantine Arkoudas: "LLMs are going to be helpful mostly because they understand language very well, they've memorized a tremendous amount and can process text at scale". Even so, "GPT-4 is plagued by internal inconsistency, an inability to correctly apply elementary reasoning techniques and a lack of understanding of concepts that play a fundamental role in reasoning…Use of generative AI in software development, science and engineering for anything other than tedious tasks (as a sort of turbo-charged auto-complete for knowledge-heavy coding questions) is fraught with serious risks. Normative standards of correctness are of paramount importance in these fields, and current LLMs cannot meet them"[27]

There are many LLM applications that are clearly thriving despite the caveats above. Even so, GPT-4 will have a bigger impact in Manuela's world than mine since her tasks conform more to what LLMs are made for, at least right now. I'm still unsure how to use a tool whose answers to complex questions have to be checked whether they're right or wrong given the frequency of the latter. I guess I will just use it for the simplest questions where it performs well; for me, that's what it's made for. At just $20 a month for GPT-4, I got what I paid for.

What about AI's long-run impact on equity markets? It should be much more durable than other recent investment themes, but the productivity shock that its most vocal adherents expect seems exaggerated.

It's a fad, fad, fad, fad world
Google search interest over time (100 = peak interest)
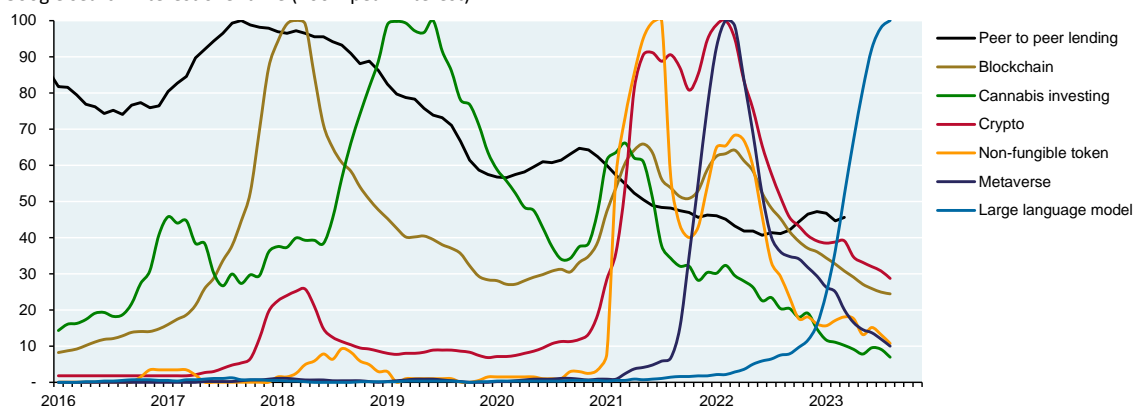
Source: Google Trends, JPMAM. September 2023.

Michael Cembalest, JP Morgan Asset Management

Title reference
"*What Was I Made For*", Billie Eilish O'Connell and Finneas O'Connell, Atlantic/Interscope Records, July 2023

---

[25] Meta Chief AI Scientist Yann Le Cun on LinkedIn, 2023

[26] "*Getting from Generative AI to Trustworthy AI*", Gary Marcus and Doug Lenat, July 31, 2023

[27] "*GPT-4 can't reason*", Konstantine Arkoudas (RPI/MIT PhD, CTO Dyania Health) on Medium, August 7, 2023

Appendix: What are the implications of generative AI for employment? Some analysis from MIT

Many AI firms, economists and investment banks argue that while productivity displaces some workers, it creates enough new jobs to more than offset the losses. **Unfortunately, this simple axiom is not always true**.

Daron Acemoglu and Pascual Restrepo found that overall measures of labor demand can mask substantial job creation and job destruction effects that accompany productivity gains[28]. They split historical data into two time frames: 1947-1987 and 1987-2017. The first chart shows that job displacement effects were almost identically offset by job creation effects from 1947-1987. **However, this trend did not persist into the 1990's and beyond**. The second chart shows that during the latter period, job displacement effects were much higher than job creation effects; and the third chart shows that for manufacturing, the net effects were even worse for workers.

David Autor also found a shift in the 1980's in which innovation shifted towards worker replacement rather than worker augmentation[29]. The table shows that from 1940-1980, employment and wage gains in occupations in which workers benefit from exposure to productivity ("augmentation") were roughly equal to employment and wage losses in occupations whose workers suffer from exposure to productivity ("automation"). **But after 1980, this changed**: the pace of employment erosion in automation-exposed occupations accelerated relative to the pace of employment gains in augmentation-exposed occupations.

Job creation and displacement effects for services and manufacturing, 1947-1987, Index (1947=0)
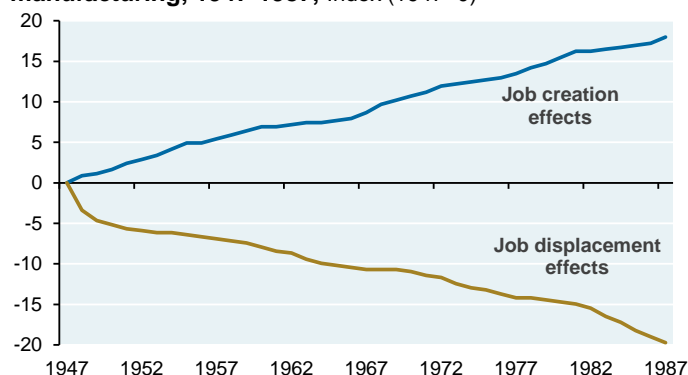
Source: Daron Acemoglu and Pascual Restrepo, 2019.

Job creation and displacement effects for services and manufacturing, 1987-2017, Index (1987=0)
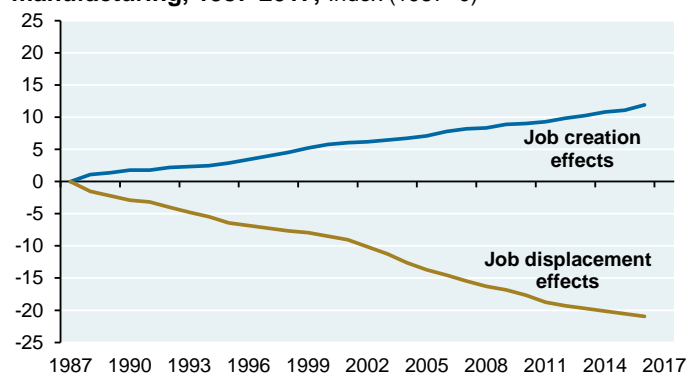
Source: Daron Acemoglu and Pascual Restrepo, 2019.

Job creation and displacement effects for manufacturing only, 1987-2017, Index (1987=0)
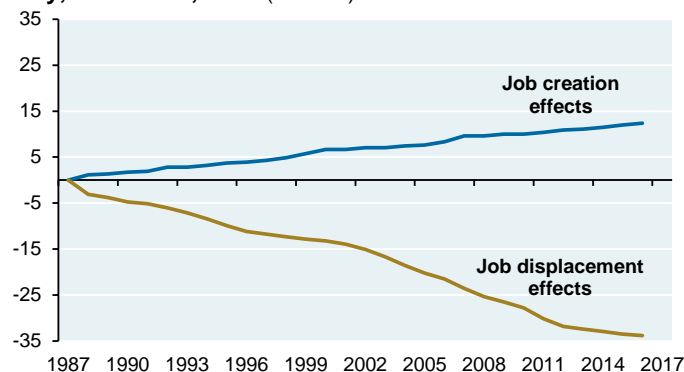
Source: Daron Acemoglu and Pascual Restrepo, 2019.

Change in employment as a function of exposure to augmentation and automation

	1940-1980	1980-2018
Augmentation exposure	1.85	1.31
Automation exposure	-1.52	-3.99

Change in wages as a function of exposure to augmentation and automation

	1940-1980	1980-2018
Wage augmentation	2.05	1.34
Wage displacement	-1.68	-3.93

Source: Autor et al. 2022.

---

[28] "*Automation and New Tasks: How Technology Displaces and Reinstates Labor*", Daron Acemoglu (MIT) and Pascual Restrepo (MIT), 2019

[29] "*New Frontiers: The Origins and Content of New Work: 1940-2018*", David Autor (MIT) et al, 2022

Questions asked of Chat GPT-4 in Sept 2023 with my grading assessments; each was drawn from a prior Eye on the Market

See supplemental attachment for GPT-4's answers, the rationale for our grades and the correct responses

Economics

[1] What was Russia/USSR's share of World GDP in the 19th century? D

[2] How are Economic Complexity Scores computed for US cities? C

[3] Since 1960, has the US yield curve ever inverted without being followed by a recession within 2 years? B

[4] How does the efficiency of LA and Long Beach container ports compare to global container ports? C-

[5] What was the worst famine in the world since 1800 when measured in terms of mortality rate? D

[6] Do corporate interest payments as a share of NIPA profits usually rise when the Fed Funds rate goes up, and is that what is happening in 2023? C+

[7] How much have smaller homes as a share of total completions declined since the 1970's, and what academic studies can you find that discuss the benefits of relaxed zoning restrictions? A-

[8] Where can we find gold holdings as a share of official reserve assets by country going back to 1956? A

[9] Find a source that has calculated the bilateral trade-weighted tariff levels between the US and China compared to tariffs they apply to the rest of the world. Since the Phase One agreement, what is higher: US tariffs on Chinese exports or Chinese tariffs on US exports? Also, what share of US exports are subject to Chinese tariffs A

[10] Where can we find data on male life expectancy in France since the early 1800's? A

[11] Where can we find data on how taxes and energy compare to labor costs as a share of total US corporate sector output? A

[12] In the 3 years following passage of the Affordable Care Act, how much did the share of the US uninsured population decline? A

[13] How have effective federal tax rates on the top 1% of the income distribution changed since the early 1980's? B-

[14] What were the two worst winning percentage years for NY/NJ professional sports teams since 1970? C-

[15] Have any analysts published rebuttals to the Saez/Piketty data on income inequality with their own estimates of income shares for the top 1%? If so, what are the primary reasons for the different findings? A

[16] From 1913 to 2018, which country suffered the largest decline in its per capita GDP relative to other countries? D

[17] In addition to farebox revenue and "other taxes and subsidies", what are the three largest sources of NY MTA revenue? A

[18] Where can we find data comparing the NYC mayor's budget projections with the Comptroller, and what is the financial risk of an increase in asylum seekers? A

[19] What share of municipal interest earned would have been impacted by Obama's proposed Federal tax on municipal bonds? C

[20] When does the CBO project that entitlements, mandatory payments and interest will exceed Federal gov't tax revenues? D

[21] What pension and retiree healthcare plans is Washington DC responsible for, and what is the NPL sensitivity of each? F

Markets

[22] For banks that failed in 2008, how frequently were losses imposed on uninsured depositors? D

[23] Do US banks hedge interest rate risk in their hold to maturity portfolios of bonds? D

[24] Can all banks add back losses on available-for-sale bond portfolios to capital? C+

[25] What share of global cross border loans and international debt securities are denominated in US dollars? B-

[26] Since 2006, what was the peak year for foreign purchases of US Treasury securities? C

[27] In Q4 2022, what were unrealized losses on available for sale securities as a share of Tier 1 capital for SVB and FRC? F

[28] Which of the following variables is most closely connected to the bottom in equity markets during a recession: PMI Index, GDP, payrolls, earnings, housing starts and household/corporate delinquencies? A

[29] What has been the outperformance of US vs non-US equities since the year 1900? A

[30] What share of drug sales and drug trials are made up of cell and gene therapy treatments? B

[31] What parts of the body do cell and gene therapies usually deliver treatments to? A-

[32] From 1999 to 2009, what was higher, the GSE loss rate on mortgage loans with private mortgage insurance or the GSE loss rate on mortgage loans without private mortgage insurance? C

[33] How much Chinese semiconductor consumption is met via Chinese production? C

[34] How much money do US companies and their subsidiaries make in China and how does that compare to the trade deficit that the US has with China? A

[35] The following is a list of 31 bank tickers. Which banks disclosed office loans as a % of CRE loans in 2021? [Provided tickers] D

[36] What data source could we use to find the price level of the S&P 500, dividends, inflation, and interest rates going back to the early 20th century? A

[37] Of the following banks and broker-dealers (list provided), which institution had the highest cumulative collateral postings to the Primary Dealer Credit Facility and which had the least? F

[38] What were peak loan to deposit ratios in European banks before the 2009 financial crisis, and how far did they fall by 2017? B-

Energy

[39] We have read that China is adding more coal plants faster than the US and Europe are decommissioning them in 2019, 2020, 2021, and 2022. Is this true? D

[40] What are average capacity credits reported by ISOs with respect to electricity systems and additional wind and solar power? A

[41] What share of US transportation energy consumption is accounted for by cars and light trucks? A

[42] What is the semiconductor content of a BEV compared to an internal combustion engine car? A

[43] What is the grid electricity share of US industrial energy consumption? D

[44] What is the energy density of jet fuel compared to the energy density of lithium ion batteries? A-

[45] What are the largest sources of GHG emissions in Indonesia and Brazil? A

[46] In enhanced oil recovery, how much CO_2 is injected per barrel of oil? A

[47] What US states have the highest and lowest CO_2 emissions per MWh of electricity generation? A

[48] Which three US states produce the greatest dollar value of food, energy and minerals combined? A

[49] Where can we find information on European natural gas storage levels by day and year? A

[50] For all countries in Europe, please make a scatterplot showing wind and solar penetration combined by country on the X axis, and household electricity prices by country on the Y axis. What is the slope of the regression line? C

[51] Where can we find information about how much heat the ocean has been absorbing since the 1950s? A

[52] Where can we get office building emissions by city? D

[53] What are the capacity-weighted wind and solar capacity factors for New York (not including plants in either calculation that have less than 12 months of net generation data)? D

[54] Where can we find data on residential heating oil costs compared to propane and natural gas, and what conversions would I need to make to compare them in the same energy unit? A

[55] If the CO_2 emissions of emerging economies continue to grow at the2014-2022 pace, how much would the CO_2 emissions of developed economies have to decline compared to their growth rate from 2014-2022 in order to keep global emissions flat? F

[56] In 2019, what was the ratio of electricity prices to natural gas prices in common energy terms for industrial users in California? F

[57] What industrial sectors are seen as having high electrification potential and how do they differ from sectors with low projected electrification potential? B

[58] How much did sulfur dioxide emissions decline after passage of the Clean Air Act through 2016? A

[59] What are net fossil fuel imports of the US and Europe in 2021 and 2022, combining fossil fuels into a common energy unit? F

Politics

[60] Name two bills that have passed by the Senate with 100% support from one major party and 0% support from the other party? D

[61] What is the narrowest party margin in the US House of Representatives since 1900? F

[62] If the debt ceiling is not raised, how would the Treasury handle interest on the Federal debt? A

[63] What year after 1945 saw the largest tax increase when measured as new tax revenues collected as a share of GDP? F

[64] In the 1992 Presidential election, what share of Perot voters would George Bush have needed to win the state of Georgia? D-

[65] What happens if no Presidential candidate reaches 270 electoral votes? A-

[66] Since 1976, what is the highest vote share that a non-Democratic and non-Republican Presidential candidate has received in any state in Presidential elections? A

[67] Where can we find information on the actual cost of medical programs compared to original estimates since 1960? C

[68] Using the latest data from the IRS, which taxpayers are audited more frequently: those earning from $200k-$500k, or those in the lowest income category? D

[69] Where can we find data comparing the political polarization of US senators since 1900, and which field should we plot as a proxy for a polarization score? A-

[70] Which two-term president's administration coincided with the largest increase in economically significant regulations? A

[71] Assume three possible Congressional alignments: (1) Unified government, (2) unified Congress with a president in the opposing party and (3) split Congress, and that these three political alignments can co-exist with either a Republican or Democratic President. From the 79th Congress to the end of the 116th Congress, which of the six permutations yielded the highest average monthly S&P 500 return, and what was it? Assume that each political term begins in November of the election year and ends in the October of the next election year D-

IMPORTANT INFORMATION

This report uses rigorous security protocols for selected data sourced from Chase credit and debit card transactions to ensure all information is kept confidential and secure. All selected data is highly aggregated and all unique identifiable information, including names, account numbers, addresses, dates of birth, and Social Security Numbers, is removed from the data before the report's author receives it. The data in this report is not representative of Chase's overall credit and debit cardholder population. The views, opinions and estimates expressed herein constitute Michael Cembalest's judgment based on current market conditions and are subject to change without notice. Information herein may differ from those expressed by other areas of J.P. Morgan. This information in no way constitutes J.P. Morgan Research and should not be treated as such.

The views contained herein are not to be taken as advice or a recommendation to buy or sell any investment in any jurisdiction, nor is it a commitment from J.P. Morgan or any of its subsidiaries to participate in any of the transactions mentioned herein. Any forecasts, figures, opinions or investment techniques and strategies set out are for information purposes only, based on certain assumptions and current market conditions and are subject to change without prior notice. All information presented herein is considered to be accurate at the time of production. This material does not contain sufficient information to support an investment decision and it should not be relied upon by you in evaluating the merits of investing in any securities or products. In addition, users should make an independent assessment of the legal, regulatory, tax, credit and accounting implications and determine, together with their own professional advisers, if any investment mentioned herein is believed to be suitable to their personal goals. Investors should ensure that they obtain all available relevant information before making any investment. It should be noted that investment involves risks, the value of investments and the income from them may fluctuate in accordance with market conditions and taxation agreements and investors may not get back the full amount invested. Both past performance and yields are not reliable indicators of current and future results.

Non-affiliated entities mentioned are for informational purposes only and should not be construed as an endorsement or sponsorship of J.P. Morgan Chase & Co. or its affiliates.   Company names are for illustrative purposes only and may or may not be held in the portfolio at any point in time. The views presented are those of the Portfolio Manager and may differ from the views of other J.P. Morgan employees and affiliates. The examples are not an endorsement, solicitation or recommendation to purchase the security

Key Risks

This material is for information purposes only, and may inform you of certain products and services offered by private banking businesses, part of JPMorgan Chase & Co. ("JPM"). Products and services described, as well as associated fees, charges and interest rates, are subject to change in accordance with the applicable account agreements and may differ among geographic locations. Not all products and services are offered at all locations. If you are a person with a disability and need additional support accessing this material, please contact your J.P. Morgan team or email us at accessibility.support@jpmorgan.com for assistance. **Please read all Important Information.**

GENERAL RISKS & CONSIDERATIONS

Any views, strategies or products discussed in this material may not be appropriate for all individuals and are subject to risks. **Investors may get back less than they invested, and past performance is not a reliable indicator of future results.** Asset allocation/diversification does not guarantee a profit or protect against loss. Nothing in this material should be relied upon in isolation for the purpose of making an investment decision. You are urged to consider carefully whether the services, products, asset classes (e.g. equities, fixed income, alternative investments, commodities, etc.) or strategies discussed are suitable to your needs. You must also consider the objectives, risks, charges, and expenses associated with an investment service, product or strategy prior to making an investment decision. For this and more complete information, including discussion of your goals/situation, contact your J.P. Morgan team.

NON-RELIANCE

Certain information contained in this material is believed to be reliable; however, JPM does not represent or warrant its accuracy, reliability or completeness, or accept any liability for any loss or damage (whether direct or indirect) arising out of the use of all or any part of this material. No representation or warranty should be made with regard to any computations, graphs, tables, diagrams or commentary in this material, which are provided for illustration/ reference purposes only. The views, opinions, estimates and strategies expressed in this material constitute our judgment based on current market conditions and are subject to change without notice. JPM assumes no duty to update any information in this material in the event that such information changes. Views, opinions, estimates and strategies expressed herein may differ from those expressed by other areas of JPM, views expressed for other purposes or in other contexts, and **this material should not be regarded as a research report.** Any projected results and risks are based solely on hypothetical examples cited, and actual results and risks will vary depending on specific circumstances. Forward-looking statements should not be considered as guarantees or predictions of future events.

Nothing in this document shall be construed as giving rise to any duty of care owed to, or advisory relationship with, you or any third party. Nothing in this document shall be regarded as an offer, solicitation, recommendation or advice (whether financial, accounting, legal, tax or other) given by J.P. Morgan and/or its officers or employees, irrespective of whether or not such communication was given at your request. J.P. Morgan and its affiliates and employees do not provide tax, legal or accounting advice. You should consult your own tax, legal and accounting advisors before engaging in any financial transactions.

For J.P. Morgan Asset Management Clients:

J.P. Morgan Asset Management is the brand for the asset management business of JPMorgan Chase & Co. and its affiliates worldwide.

To the extent permitted by applicable law, we may record telephone calls and monitor electronic communications to comply with our legal and regulatory obligations and internal policies. Personal data will be collected, stored and processed by J.P. Morgan Asset Management in accordance with our privacy policies at https://am.jpmorgan.com/global/privacy.

ACCESSIBILITY

For U.S. only: If you are a person with a disability and need additional support in viewing the material, please call us at 1-800-343-1113 for assistance.

This communication is issued by the following entities:

In the United States, by J.P. Morgan Investment Management Inc. or J.P. Morgan Alternative Asset Management, Inc., both regulated by the Securities and Exchange Commission; in Latin America, for intended recipients' use only, by local J.P. Morgan entities, as the case may be.; in Canada, for institutional clients' use only, by JPMorgan Asset Management (Canada) Inc., which is a registered Portfolio Manager and Exempt Market Dealer in all Canadian provinces and territories except the Yukon and is also registered as an Investment Fund Manager in British Columbia, Ontario, Quebec and Newfoundland and Labrador. In the United Kingdom, by JPMorgan Asset Management (UK) Limited, which is authorized and regulated by the Financial Conduct Authority; in other European jurisdictions, by JPMorgan Asset Management (Europe) S.à r.l. In Asia Pacific ("APAC"), by the following issuing entities and in the respective jurisdictions in which they are primarily regulated: JPMorgan Asset Management (Asia Pacific) Limited, or JPMorgan Funds (Asia) Limited, or JPMorgan Asset Management Real Assets (Asia) Limited, each of which is regulated by the Securities and Futures Commission of Hong Kong; JPMorgan Asset Management (Singapore) Limited (Co. Reg. No. 197601586K), which this advertisement or publication has not been reviewed by the Monetary Authority of Singapore; JPMorgan Asset Management (Taiwan) Limited; JPMorgan Asset Management (Japan) Limited, which is a member of the Investment Trusts Association, Japan, the Japan Investment Advisers Association, Type II Financial Instruments Firms Association and the Japan Securities Dealers Association and is regulated by the Financial Services Agency (registration number "Kanto Local Finance Bureau (Financial Instruments Firm) No. 330"); in Australia, to wholesale clients only as defined in section 761A and 761G of the Corporations