



Capital
Markets

Software & Internet Implications of Generative AI and ChatGPT

Rishi Jaluria
(Analyst)
(415) 633-8798
rishi.jaluria@rbccm.com

Brad Erickson
(Analyst)
(503) 830-9488
brad.erickson@rbccm.com

Matthew Hedberg
(Analyst)
(612) 313-1293
matthew.hedberg@rbccm.com

Matthew Swanson
(Analyst)
(612) 313-1237
matthew.swanson@rbccm.com

Jonathan Atkin
(Analyst)
(415) 633-8589
jonathan.atkin@rbccm.com


RBC ImagineTM

Thinking further forward

EQUITY RESEARCH | MARCH 16, 2023

For Required Conflicts Disclosures, see page 192

Credit: DALL-E, RBC Capital Markets



March 16, 2023

RBC Imagine™: Software & Internet Implications of Generative AI and ChatGPT – Moats and Boats

In this deep dive, we discuss the evolution of the Generative AI (GAI) secular mega-trend and the potentially transformative effect it could have on technology, enterprise software & internet broadly, and MSFT and GOOGL more specifically. As highlighted under our [RBC Imagine Artificial Intelligence Activated](#) theme, we believe that GAI is tech's 4th major exponential capacity creator of the last 20+ years building on top of the transformative effects of bandwidth, SaaS/Cloud, and Mobility and as such, we see the potential for material TAM expansion for cloud & search in particular (directly benefiting datacenters) where fast adoption by companies of all sizes may become mission critical to develop or maintain competitive advantages. Meanwhile, we believe GAI's widespread adoption could also lead to commoditization for smaller companies in particular, allowing the ecosystem enablers like cloud providers & digital ad platforms to widen existing moats and accrue disproportionate value. **Our best plays on the GAI secular mega-trend:** Software — MSFT, HUBS, NICE, VEEV, MDB, CRWD, ESTC, GTLB, NET, NOW, and ADBE; Internet — GOOGL, AMZN, META, SNAP & PINS and to a lesser degree, BKNG, EXPE & ABNB; Datacenters — DLR. For ESG implications, see our [review of Hyperscalers and Datacenters](#).

- **With the release and unprecedented rapid adoption of ChatGPT**, we believe GAI has turned from a long-term bet into a near-term reality that has the potential to upend a broad array of industries. **Most positively**, we believe GAI should drive significant cost efficiencies for software & internet companies and could be meaningful TAM expanders for cloud, search, and datacenter operators that specialize in hyperscale requirements, while also creating deeper & wider moats for the tech arms dealers like cloud providers and digital ad platforms to accrue the majority of the value. **Less positively**, customer-facing companies & use-cases could see initial AI-fueled information asymmetry advantages but GAI's widespread availability & adoption could create medium- & long-term participant commoditization, particularly for the smaller/less well-resourced players, where initial differentiation could erode more quickly than prior tech cycles given the faster pace of innovation.
- **Microsoft has emerged as a significant player in GAI**, owing to its investments in OpenAI, where we think it is leveraging GPT as a means to not only disrupt search, but more importantly, enable potential share gains in Azure, drive further differentiation with O365 and Teams, serve as a meaningful driver for GitHub, and improve Dynamics' competitiveness, among other areas within Microsoft's expansive portfolio. Importantly, we don't believe this is priced into the stock today and view GAI as a "call option" on Microsoft's stock.
- **GOOGL has a play on GAI...and it's not all bad like investors seem to think**. We believe GOOGL's GAI strategy is likely underappreciated with investors concerned about search disruption, search share loss and rising computing cost driving LT margin compression with less consideration for a number of important positive offsets like structural distribution advantages, incremental revenue drivers, vertical integration & differentiated data. We come away from a deeper quantitative downside analysis believing we may be within 10-15% of GOOGL approaching a worst-case disruption scenario (contact RBC sales for our choose-your-own-adventure model).
- **Within Software and Internet specifically, we believe GAI will ultimately be a tailwind for three primary categories:** 1) cloud-native market leaders who are quick to embrace GAI and use their data and market leadership as a competitive advantage; 2) vertical-specific companies who can widen the gap vs. challengers and accelerate the winner-take-most dynamics of the TAM; and 3) disruptive mid-market software going after entrenched, legacy competitors, who can use GAI to narrow the gap. **On the other hand, we believe GAI will be negative for companies who:** 1) are slow to embrace it putting them at a competitive disadvantage (native GAI vs. adapted GAI); 2) on-premise companies, where integrating GAI becomes more difficult; and 3) analytics companies that have previously benefited from AI tailwinds.

RBC Capital Markets, LLC
RBCCM Global Research
(Global Research)
Rishi Jaluria (Analyst)
Brad Erickson (Analyst)
Matthew Hedberg (Analyst)
Matthew Swanson (Analyst)
Jonathan Atkin (Analyst)
Dan Bergstrom (Analyst)
Sara Mahaffy (ESG
Strategist)
Richard Poland (AVP)
Logan Reich (Senior
Associate)
Anushtha Mittal (Senior
Associate)
Simran Biswal (Associate)



RBC Imagine™
Thinking further forward



Table of contents

GAI is likely a Cloud & Search TAM expander..... 6

What is GAI, what is OpenAI, and what is ChatGPT? (according to ChatGPT).....9

Why is ChatGPT so important?..... 9

Ten Generative AI predictions..... 11

Societal implications..... 13

Examining ChatGPT/Generative AI Through An ESG Lens..... 15

Microsoft: A multi-pronged strategy well-positioned for the next wave of compute.....17

Azure's push to become de facto cloud platform for AI..... 18

Differentiation in DevOps and Low-Code..... 23

Nothing to lose in search.....25

Making productivity suites more productive.....26

First-mover advantage in intelligent applications..... 28

GOOGL: Down but certainly not out..... 29

What is GOOGL doing with AI/ML today?.....30

How might search business models to be altered or at least adapt to the inclusion of GAI..... 36

What is ChatGPT's impact on traditional search and social platform engagement, if any?.....37

Side by side: Does ChatGPT replace traditional search use cases?..... 40

Taking it to the model - Impact on GOOGL's '24 estimates..... 43

Implications for software companies..... 54

Software company-specific impacts..... 56

Internet & internet company-specific takeaways.....72

Primary GAI implications/themes across internet..... 72

How internet companies are using AI today..... 72

Implications for internet companies by sub-vertical..... 77

Implications for the cloud wars.....81

Broad vertical impacts from AI & ML adoption..... 87

Large language model (LLM) tech underpinnings..... 95

ChatGPT vs. Google Search across internet verticals..... 97

Bing: Illustrative use cases..... 185



GAI is tech's 4th exponential capacity unlock. In the context of the last 20+ years of technology innovation, we believe AI represents the 4th major capacity unlock which exponentially expands technology's capabilities on top of the previous 3 in aggregate: Internet, SaaS/Cloud, and Mobility. Dating back to well-funded NASA projects in the late 1950s, which created much of the modern semiconductor industry & early mainframe computing, it wasn't until the early telephony-based & eventually, wireless broadband infrastructure came along which could support the bandwidth necessary to connect large networks of computers & users.

Second, this bandwidth then enabled cloud services - a relatively small number of large, scaled companies that could provide shared infrastructure for companies of all sizes allowing for expanded & more flexible enterprise workload capabilities alongside significantly greater capital/cost efficiency. Concurrently, the combination of bandwidth & cloud then led to the obsolescence of on-premise software as SaaS's revolutionary innovation potential & subscription model were born. The subscription model attracted enormous capital driving enormous innovation & new players which continues to revolutionize virtually every vertical end-market in the world. **Critical milestones:** 1) the introduction of Salesforce in 1999, which was the first widely-adopted enterprise application delivered over the Internet; and 2) the introduction of Amazon Web Services (specifically S3 storage and EC2 compute) in 2006, which allowed developers to build applications without worrying about the underlying infrastructure.

Third and finally, the late 2000s saw the rise of mobility which once again exponentially expanded computing capabilities with its ability to tie bandwidth, cloud-based, application-specific software together with location & mobility and created thousands of new utilities not previously possible with desktop computing. **Critical milestone:** the introduction of the Apple iPhone, which turned phones from a device to make calls into mobile computers that were simple to use.

With GAI, software development, education, advertising, content creation, professional services and a myriad of general personal & professional productivity tools will all likely see some amount of transformative change, and as humble observers of the still nascent industry, the ceiling to innovation could well be years if not decades away. The critical milestone with GAI has been the release of ChatGPT in November 2022, which has been the fastest-growing Internet service in history.

GAI could be presenting itself at a moment poised for outsized value creation potential. While the major tech revolutions haven't historically operated on the macro's timetable (smartphone explosion into the teeth of the GFC for example), investors' current hyper-focus on cash generation vs. growth-oriented call options (like GAI) means little is ascribed for this type of exponential value creator at this point, in our view. And so while this may also prompt the debate on whether the technology is even ready to release into the wild, one way or another, we think it's here to stay and will now presumptively lead to a veritable wild west period of digestion, exploration, development and ultimately, implementation which could drive seismic change in how people & businesses both seek information, communicate and operate.

Moats & boats. We see two primary outputs of GAI with 1) internal process efficiencies driving cost savings and 2) external, customer-facing capability expansion (sometimes fueled by internal efficiencies). For the internal efficiencies, we see the GAI's democratization of these powerful tools as a boon for all companies with exposure to leverage efficiency benefits. For external use-cases, after an initial digestion & exploration period, we see a rising probability of commoditization, particularly for smaller companies with fewer resources allowing the so-called 'arms dealers' like cloud providers, digital ad platforms & content creation enablers to potentially widen existing moats and accrue disproportionate value.



One of the hallmarks of technology businesses, particularly in software & internet is that the inherent network effects built on the foundation of bandwidth, Cloud, SaaS & mobility meant that the early colonizers of various vertical end-markets with a superior product and/or go-to-market often obtained outsized market share faster than more traditional/non-network effected-oriented industries which pre-dated widespread mobile & internet adoption. This of course has created massive competitive moats for companies like Apple, Microsoft, Google and certain others which short of regulation, would seem highly unlikely to be adversely affected for the foreseeable future.




GAI is likely a Cloud & Search TAM expander

With early ChatGPT/Bing usage data appearing incremental, we believe GAI could be meaningfully additive to not only the search market but also cloud driven specifically by a combination of incremental functionality driving incremental usage, net new workloads as native GAI businesses form and finally, higher compute cost for that incremental consumption.

Search

Specifically for search, we leverage early data from downloads and page views (shown below) as a reflection of GAI expanding use cases and improving functionality & UX. Bing page views are up ~30% vs. pre-ChatGPT release averages while GOOGL traffic is largely flat vs. pre ChatGPT levels (all while ChatGPT traffic is 30% above Bing's total pre-ChatGPT levels).

Figure 1 - Looking at page views, it appears as though Bing did get a noticeable uptick in traffic, coinciding with the rise of ChatGPT, while the impact to Google's traffic appears non-existent, implying Bing/ChatGPT traffic is incremental



Source: Similar Web, RBC Capital Markets. Note: Bing and ChatGPT are on right axis and Google is on left axis

Given the incremental traffic over the past few months, we estimate the potential impact on total search TAM using page views as a proxy. On Bing, traffic is effectively up 160% since the public release of ChatGPT (30% from Bing itself and 130% from ChatGPT) while GOOGL has remained flat. We think the incremental search share gains should continue based on a) GOOGL has yet to implement a GAI solution into Search, b) ChatGPT is still in very early stages where we expect meaningful improvements going forward and c) the integration with Bing is still not rolled out to most users. Assuming GOOGL page views grows a modest 5% and Bing doubles its current incremental growth of 160%, we estimate Search TAM could grow by as much as 26% over the next couple of years as a result of GAI implementation and improvements.



Figure 2 - Given the advent of ChatGPT and integration with Bing has been driving incremental traffic, we estimate the total search TAM expansion as a result of GAI could be ~26%

Search Engine	Current search share	2022 search revenue (\$M)	TAM expansion related to GAI	Total search share post GAI
GOOGL	87.2%	\$162,450	5%	91.6%
Bing	6.9%	\$9,918	320%	28.9%
Yahoo!	3.1%	\$5,755	0%	3.1%
DuckDuckGo	2.4%	\$100	0%	2.4%
Other	0.4%		0%	0.4%
Total	100.0%	\$178,223		126.4%

Source: StatCounter, Similar Web, RBC Capital Markets

Cloud

Along the lines of GAI's additive functionality and incremental traffic in search so far, we also believe there will accordingly be both organically and inorganically higher usage tailwinds translating to cloud spending over time. Further, higher compute cost also should layer on top of this as a driver. While it's difficult to unpack the three separate drivers (GAI driving incremental usage/development, GAI driving net new enterprises & workloads and GAI driving higher computing intensity & cost), we use the initial search uplift as an indirect proxy (admittedly part art/part science) as a way to estimate the uplift to the overall cloud market. Worth noting, we expect double-digit percentage annual cost improvements associated with scale and future (first-party and third-party) GPU generations, making GAI applications more affordable in the future.

Taking the TAM expansion to the models to give an example of the upside potential from increasing cloud adoption/use cases, we estimate an approximate 21% upside to Street's C'23 estimates. Admittedly, TAM expansion may not be meaningful in the near term, but we find it worthwhile as a way to size the potential eventual cloud expansion.

Expanding on our thesis, we see three distinct growth drivers for cloud as a result of generative AI:

1) GAI workloads are more storage and compute intensive. As the usage of GAI continues to grow (especially with enterprise adoption, such as GPT4 APIs), most of that flows through the hyperscale cloud providers (e.g. OpenAI workloads all go through Azure), this drives incremental cloud revenue.

2) GAI drives incremental migration of existing workloads to the cloud. This comes at an important time, when investors are worried about cloud saturation. GAI services are generally cloud-only, which means the urgency for enterprises to adopt GAI should accelerate the pace of migrating workloads.

3) GAI platforms will lead to the next generation of startups. As we outline later in this note, we believe there will be an entire generation of startups built on GAI platforms that natively leverage the technology. Much like the introduction of AWS (and later Azure and GCP) created an entire ecosystem of cloud-native startups and the introduction of the iPhone (and later Android) created an entire ecosystem of mobile applications, we expect the same with GAI. These are all incremental new cloud workloads that do not exist today.



Figure 3 - If we assume cloud TAM expands 26% there's an implied ~21% upside to Street's '23 cloud revenue estimates, though worth noting it's unlikely for the full TAM expansion to be realized this year

Cloud Provider	2022 Revenue (\$M)	Total TAM expansion	Incremental \$ revenue from TAM expansion	Street's '23 estimate	Street's '23 estimate + incremental \$ from TAM expansion	% implied upside to Street's '23 estimates
GCP	\$26,280	26%	\$6,833	\$33,211	\$40,043	20.6%
AWS	\$80,096	26%	\$20,825	\$93,660	\$114,485	22.2%
Azure	\$50,468	26%	\$13,122	\$64,347	\$77,469	20.4%
Total	\$156,844.4	26%	\$40,779.6	\$191,217.5	\$231,997.1	21.3%

Source: Company Reports, Factset, RBC Capital Markets



What is GAI, what is OpenAI, and what is ChatGPT? (according to ChatGPT)

Generative AI, or artificial intelligence that can create new content, is a rapidly developing field with enormous potential to revolutionize industries from creative writing to music composition to product design. Unlike traditional AI models that use pre-existing data to generate predictions or responses, GAI models are trained on vast amounts of data in order to create original content that is similar in style or form to the original.

One of the leaders in the development of generative AI is OpenAI, a non-profit research organization founded in 2015 by a group of tech luminaries including Elon Musk, Sam Altman, Greg Brockman, Ilya Sutskever, and John Schulman. OpenAI's mission is to "ensure that artificial intelligence benefits humanity", and the organization's work spans a wide range of applications, from natural language processing and robotics to climate modeling and bioinformatics.

One of OpenAI's most well-known achievements is the development of GPT-3, a generative AI language model that is capable of generating human-like language with remarkable accuracy and coherence. GPT-3, which was released in 2020, has already been used in a wide range of applications, from chatbots and virtual assistants to content creation tools and language translation software.

ChatGPT is a specific implementation of the GPT-3 language model that has been designed specifically for conversational applications. ChatGPT is capable of engaging in free-form conversations with users, providing responses that are both contextually relevant and grammatically correct. As a result, ChatGPT has become one of the most popular conversational AI tools available, with applications in customer service, marketing, and education.

The importance of generative AI, OpenAI, and ChatGPT lies in their potential to transform the way we interact with machines and with each other. As generative AI models become more sophisticated and more widely deployed, they have the potential to replace or augment many of the routine tasks that humans perform, freeing us up to focus on more complex and creative work. At the same time, they can help us to communicate more effectively with each other and with the machines that surround us, making it easier to access information, solve problems, and connect with others across the globe. Ultimately, generative AI has the potential to change the way we live and work, and OpenAI and ChatGPT are at the forefront of this exciting new field.

(This entire section was written by ChatGPT).

Why is ChatGPT so important?

With ChatGPT's boilerplate intro out of the way, the big questions are: 1) why is ChatGPT so important; and 2) is ChatGPT worth the hype?

Why is ChatGPT so important?

Before ChatGPT, we saw two distinct trajectories for AI systems.

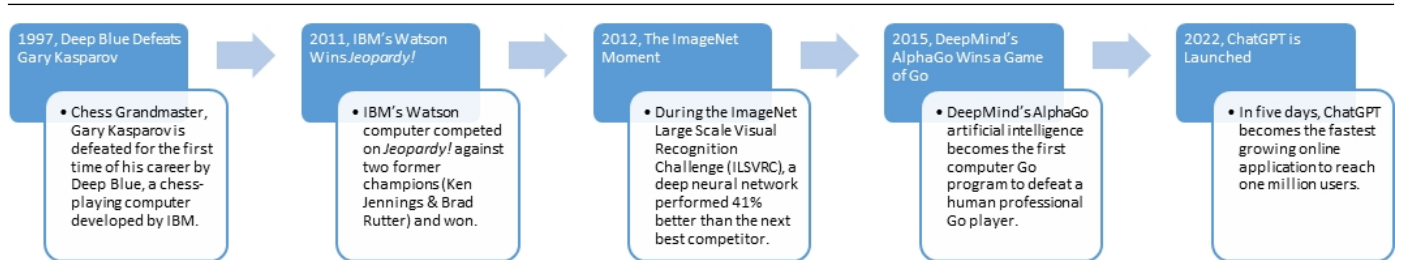
The first is enterprise-grade AI, which includes offerings by companies such as Google, IBM, and Microsoft, among others. These systems have shown impressive milestones over time, such as IBM Watson winning Jeopardy and Google DeepMind defeating a professional human Go player. These systems, although powerful, generally require heavy customization (AlphaGo, for example, is its own computer program, built by DeepMind) and targeting data scientists.

The second is consumer-grade intelligent agents, such as Amazon's Alexa and Apple's Siri, which targeted the consumer and were easy to use (both relying on voice). These systems are intuitive and make use of natural conversational language, but also provide very basic output, such as

simple Internet searches or controlling connected devices.

In many ways, ChatGPT brings together the best of both of these worlds. Underlying ChatGPT is powerful AI, but the actual interface is conversational and iterative and very simple to use. In other words, ChatGPT has democratized AI in a way that other AI systems have not done before. Importantly, the conversational component means users have to react to answers with follow-ups in order to get closer to the answers they desire. ChatGPT can also be trained on specific datasets to provide user-specific or company-specific insights, versus just general insights from the broad Internet.

Figure 4 - Timeline of Major AI Milestones



Source: RBC Capital Markets

Is ChatGPT worth the hype?

We believe ChatGPT is to generative AI what Netscape was to the Internet and what the iPhone was to mobile.

However, the importance of generative AI goes beyond just the consumerization of AI. As the name implies, generative AI can create brand-new content, versus providing answers to existing questions or insights on existing data. This content isn't just limited to text, but also includes unstructured content like images (OpenAI's DALL-E 2 can create original images from a text description), code (we discuss GitHub Copilot later in this note), and video.

While there certainly are limitations to ChatGPT today, it has meaningfully surpassed our expectations of where AI would be today. We also note the public version of ChatGPT is a "free research preview" and likely behind ChatGPT version 3 (ChatGPT-3) in capabilities, while version 4 (ChatGPT-4) has recently been announced. And as we see companies embrace ChatGPT for their own use cases, opening up vast volumes of data, we may end up even more impressed than we currently are.

So is it worth the hype? We absolutely believe so. Generative AI itself represents a major advancement in AI systems beyond deep learning and has flipped the switch on the purpose of AI from analyzing and learning to creating. ChatGPT and the consumer-like interface only magnifies this exponentially, as it opens up significantly more use cases that were previously unheard of (some of which we explore later in this note). Finally, the quality of output from these systems has surpassed our expectations meaningfully and has real societal implications, which we discuss in the following section.



Ten Generative AI predictions

- 1. Companies who don't truly embrace GAI will see their multiple compress by 50% over the next five years.** We believe every technology company needs a strategy to truly embrace generative AI (and not just pay lip service to the technology), otherwise they will be left behind by those that do. Not only will these companies see market share losses over time, but they will also see multiple compression as investors lose confidence in the ability of those companies to be "future-proof". We would draw the analogy to the cloud -- on-premise companies that were late to embrace the cloud (e.g. Cloudera, Tableau) saw meaningful multiple compression as a result.
- 2. Generative AI will be negative for software companies and stocks that have previously benefited from AI tailwinds.** We've seen a number of companies whose businesses have benefited from being associated with AI in the past, including Palantir, and Alteryx. We believe many enterprises will prioritize their spending towards generative AI and view some of those previous solutions as a 'stop-gap' measure. We expect the same to apply to stocks that have outperformed as a result of investor interest in AI, including names like C3.ai, BigBear.ai, and SoundHound.ai, as public investors realize there are few ways to directly invest in generative AI today.
- 3. Accelerated adoption of AI drives next leg of public cloud growth.** One of the big worries investors have is around cloud saturation and how much of the slowdown in hyperscale cloud vendors (AWS, Azure, GCP) is because of macro versus a general slowdown in migration to the cloud. Given the resources necessary for generative AI, we believe that accelerated adoption will drive a major increase in public cloud workloads. In addition, as the hyperscale cloud vendors look to add generative AI services on top of their existing offerings, we believe this also drives up overall growth in the hyperscale cloud vendors.
- 4. Governments & regulatory bodies will increasingly scrutinize GAI.** The rise of generative AI brings about a number of ethical and legal concerns, some of which we discuss in our section on societal implications. We believe governments throughout the world will create legislation around generative AI, including the use of generative AI systems, preventing malicious use, and the use of customer data.
- 5. Generative AI will lead to the next generation of start-ups.** Much like the introduction of the iPhone led to an entire economy of mobile applications and the rise of AWS created entire enterprise software companies made to be cloud native, we believe ChatGPT and generative AI will be the foundation for the next generation of tech start-ups. We also believe the cloud workloads associated with these new start-ups will also be additive to overall cloud growth.
- 6. Security attacks will become more complex.** We believe malicious groups will be able to use generative AI to create more complex and creative cyber-attacks, which can serve as a tailwind for security companies. On the flip side, we think security companies that embrace both AI/ML and GAI more specifically will be in a better position to stop advanced threats. Critical to this thesis is large cloud dataset to train AI and ML models. In particular we think this benefits both CRWD and PANW.
- 7. GOOGL changes the narrative** that GAI prompts an existential crisis over the next 6-12 months as it a) releases new tools, b) demonstrates incremental revenue associated with those tools and c) capex doesn't ramp more than expected from current levels as it does all of this as the market sees further evidence of scaled cost improvements with iterative rollouts of new models.
- 8. AMZN begins publicizing more of its own GAI progress (which has been integrated with its cloud capabilities far longer than GOOGL or MSFT),** with its go-to-market messaging more geared through its developer ecosystem distribution channel which could demonstrate its own but more platform-centric approach to enabling more efficient model training and implementation vs. competitors as opposed to creating the models themselves.
- 9. META brings new GAI tools to market** enabling a differentiated level of automated campaign



creation - and driving advertiser content cost savings that could be incrementally allocated to advertising. These tools become especially important for Reels as advertisers increasingly move to the structurally higher content creation costs in the short-form video.

- 10. The OTAs (ABNB, BKNG & EXPE) get more vocal on their leveraging of GAI solutions,** potentially to get out in front of any negative narrative that may arise surrounding customer acquisition headwinds if and as search discoverability becomes more challenging. While we've never touted loyalty as being a meaningful driver in travel (ie the connected trip or the like), we actually could be turned by a prospective introduction of GAI onto an OTA site where the interactiveness and incremental data capture opportunities relative to any other non-vertical specific discovery engine could be significantly differentiated and drive an acceleration in organic loyalty, if not in actuality.



Societal implications

We've already seen landmark use cases of generative AI and ChatGPT that illustrate just how far the technology has come and its potential to change society. [We wrote on this topic](#) in November 2021 and we are already starting to see some of this play out:

- ChatGPT has passed exams for both law school (four courses at the University of Minnesota) and business school (at Wharton).
- Congressman Jake Auchincloss delivered a speech written by ChatGPT on the U.S.-Israel Artificial Intelligence Center Act.
- Insilico Medicine is beginning clinical trials for a COVID-19 drug designed entirely by generative AI (ISM3312).

From here, we see several major societal implications:

First, we expect disruption in every major industry. Traditionally, automation has been focused on replacing manual labor, but we believe generative AI can do tasks that traditionally fall to 'white collar' jobs and, in many ways, has the potential to displace some of those jobs. For example, generative AI can write articles, legal briefs, blog posts, and scientific research papers, which potentially impacts jobs done by journalists, lawyers, marketers, and scientists.

Second, most discussions around AI impacting white collar jobs have traditionally focused around tasks done by those at more entry level positions and work that is already being outsourced. Examples include analyzing x-rays, researching prior legal opinions, and summarizing events. However, with the advent of generative AI, we see room for even more senior positions to be impacted – for example, not only can generative AI research prior Supreme Court rulings, but has the ability to write legal arguments based on those prior rulings.

Third, having said that, we do not view generative AI as a zero sum game. In many examples we can think of, generative AI can be used to automate more routine tasks, allowing the user to focus on higher value-add actions – for example, a journalist could use generative AI to summarize the background to a piece, spending more time on original research and generating unique insights. We also believe generative AI can be yet another tool for high-skilled workers to have in their arsenal and to be more productive. Instead of a marketer writing every single blog post and starting from scratch, the marketer can use ChatGPT to create a basic template and then build on top of that, allowing them to meaningfully increase the volume of blog posts, thus generating more new business and greater pipeline.

Fourth, relationship building becomes even more important. Arguably, at the top of many professions, be it law, banking, or consulting, relationships are one of the most important ways to distinguish from the competition. With many other areas potentially getting commoditized by generative AI, the ability of senior people to build relationships with others and better understand the unique problems they're trying to solve becomes even more paramount.

Fifth, we believe generative AI will soon become tablestakes and every company will have to have a strategy around utilizing generative AI. Those companies that don't will be left behind, while those that do will likely use generative AI to enhance the business itself and customer relationships, but may struggle to directly monetize the technology. Generative AI, however, will make the economic moat around data even more powerful.

Sixth, generative AI can narrow the gap between traditional players and digital natives. As we've discussed in the past, it's increasingly important to differentiate what is a technology company from what isn't – for example, a mattress company doesn't become a technology company just by selling goods online and having a mobile application. Similarly, we've seen many traditional industries where a startup disrupts the landscape with more technology-enabled functionality,



as well as self-service capabilities – for example, insurtech companies like Lemonade, Hippo, and Root, have disrupted traditional property and casualty insurers with their easy-to-use mobile apps, but the capabilities of generative AI could soon leapfrog what those insurtech companies offer. This would allow traditional insurers to quickly narrow the gap and hurt the differentiation the insurtech players offer, even if they claim to have AI and ML capabilities.

Seventh, widespread adoption of generative AI brings about myriad ethical and legal concerns. If a developer builds an application using code generated by ChatGPT, who owns that code? If a student submits an essay written by ChatGPT, is it their own work? These are tough questions to answer, and we believe there will be endless debate on these topics.

Eighth, in response to the legal side of the equation, we expect governments around the world to take action around generative AI, potentially restricting its use in certain areas and putting guard rails around the use of personal data. In addition, if someone uses generative AI for nefarious reasons, does the creator of that AI hold any liability for not having enough guardrails around safety and compliance? If a sick person uses ChatGPT for medical advice, does OpenAI bear any liability if the advice is wrong or are the current disclaimers enough (or is ChatGPT compared to a website like WebMD instead)?


Ninth, we expect generative AI to change the very nature of work. In addition to changing what is important or differentiated in each job, generative AI has the ability to make each worker more productive. On one hand, this could mean an individual is just completing more tasks and doing more work in the same amount of hours, and potentially mean there are fewer jobs available, resulting in higher unemployment. On the other hand, generative AI could become truly additive and accelerate growth in the economy and even potentially result in shorter workweeks for the average knowledge worker.



Examining ChatGPT/Generative AI Through An ESG Lens

In the figure below, RBC’s ESG Strategy team has partnered with the RBC Elements™ team to examine which SASB topics have been coming up the most in ESG related news activity specific to ChatGPT/Generative AI. Their search included news articles, industry publications, and NGO documents from November 2022 through mid-March 2023. While much of the focus has been on “S” related implications, there are also “E” and “G” issues to consider. In this section of the report, we dive into emerging ESG risks and opportunities as it relates to Customer Welfare, Data Security & Customer Privacy, Human Capital Management, and Resource Use.

Figure 5 - ESG Related News Activity Specific To ChatGPT/Generative AI By SASB Topic (# Articles)



Source: RBC ESG Strategy, RBC Elements, Factset

Customer Welfare: The SASB category that has been the biggest source of ESG related news activity is Customer Welfare. For ChatGPT/Generative AI, this could include impacts from potentially accelerating the spread of false/misinformation (particularly in this current era, as companies are rushing to compete and integrate AI technology) and perpetuating biases. From a financial materiality perspective, associated risks could potentially open companies up to regulatory scrutiny (and costs of compliance). Some examples of recent regulatory initiatives focused on digital rights and content moderation more broadly include, the EU’s Digital Services Act and Digital Markets Act which were adopted in July 2022 (and applies starting in early 2024), as well as the EU’s AI Act which is expected to be voted on later this year.

Data Security & Customer Privacy: Another consideration is around how ChatGPT/Generative AI could impact data security and customer privacy through its collection of data in training models, as well as how the tool could potentially be leveraged for malicious use (e.g. more sophisticated cybersecurity attacks). We have already seen a number of companies such as AMZN, VZ warn employees on their use of ChatGPT due to concerns that sensitive data could be exposed. Similarly, the potential for this tool to harvest and leverage customer data brings up customer privacy concerns. From a financial materiality perspective, these issues could potentially open companies up to reputational impacts and costs if they aren’t properly managing heightened security risks.

Human Capital Management: As highlighted earlier in the report, use of ChatGPT/Generative AI has the potential to transform the nature of work. On the one hand, there have been concerns



about the automation of certain tasks related to legal professions, research, coding, that have raised concerns about job losses and higher unemployment and could lead to reputational impacts for companies engaging in layoffs, particularly if viewed to be mismanaged. Improved productivity from the automation of repetitive tasks could also free up more time for employees to focus on higher value tasks, potentially improving employee engagement and wellbeing at a time when burnout is high across many industries.

Resource Use: In terms of environmental impact, concerns have been raised around the resource consumption (energy, water usage) and emissions needed to fuel, train and utilize Generative AI tools. Some of these impacts can be reduced through making the models/processors more efficient and increasing usage of renewable energy. Additionally, a number of benefits could be realized from the application of this technology. For example, generative AI could be leveraged to aid climate scenario analysis, as well as better manage and measure a company's resource use and carbon footprint in their operations and supply chains.



Microsoft: A multi-pronged strategy well-positioned for the next wave of compute

First, let's discuss our over-arching thesis of what we believe Microsoft's strategy is with OpenAI and publicly going after search. We believe Microsoft's primary goal is to have Azure become the de facto standard cloud platform for AI workloads, in order to help close the gap with Amazon Web Services and create further distance from other hyperscale cloud vendors, like Google Cloud Platform and Oracle Cloud Infrastructure.

We believe Microsoft's messaging around gaining share in search, to an extent, is a deliberate effort to pressure Google. Microsoft's share in search today is small, with Microsoft's Search and News Advertising (includes all advertising revenue, except LinkedIn) revenue of \$12.0B in CY22, or roughly 7% the size of Google's Search revenue of \$185.2B. Therefore, even if Microsoft significantly degrades the Search businesses' gross margins, reflecting the higher compute intensity of AI-based search (according to various sources and conversations, estimates range from 2-10x), the downside impact if share gains don't materialize is negligible, while the upside impact if shares gains do materialize could be meaningfully EPS accretive, in our view. Also, regardless of the EPS impact, if successful, this could lead to multiple expansion given investors have historically not given Microsoft much credit, if any, for Bing. At the same time, this could to a certain extent distract Google from more aggressively competing in cloud infrastructure.

In the worst case scenario, Microsoft gains minimal share of search and absorbs some margin impact, while defending share dominance in productivity suites (from Google Workspace and potential AI-based disrupters). In the best case scenario, Microsoft gains a few points of share in search, gains share in CRM/ERP with Dynamics (by using large language models and generative AI to enhance the product suite), defends share dominance in productivity suites, creates further differentiation in cybersecurity, and, most importantly, Azure becomes known as the de facto standard cloud platform for AI workloads.

Background of Microsoft's relationship with OpenAI

In 2019, Microsoft invested \$1B in OpenAI. Leading up to this initial investment, OpenAI transitioned from a non-profit to a unique capital structure with a capped-profit element, whereby economic returns for investors and employees are capped at 100x. Perhaps more importantly, this opened the door for OpenAI to seek out investment from venture capital funds and public companies. We believe there are a few dynamics restricting Microsoft from outright acquiring OpenAI. Aside from the unique capital structure, we believe the consortium of founders are likely unwilling to sell, especially since they would then lose control of the company's founding mission to build safe and beneficial artificial general intelligence.

We believe OpenAI chose Microsoft as a partner for several reasons. First and foremost, Azure provides AI-scale core compute and storage, with an exhaustive list of up-the-stack infrastructure solutions, including security and governance, data and analytics (database, data warehouse, data lake, data governance, etc.), and developer tools. Second, Microsoft has arguably the broadest portfolio of enterprise application software, spanning productivity suites, collaboration, communication, CRM, ERP, and more. In addition to this giving OpenAI a route to commercialization/monetization in the enterprise (Office alone touches nearly 400M users), it provides OpenAI with the corresponding enterprise datasets to train models. Third, Microsoft had already developed its own family of large AI models, the Microsoft Turing models, which had been used to improve language understanding tasks across Bing, Office, Dynamics, and more.

Figure 6 - Timeline of Microsoft and OpenAI's combined innovation



Source: Company website and press releases, RBC Capital Markets

Azure's push to become de facto cloud platform for AI

On one hand, we believe Microsoft's intentions are clear. In our view, Microsoft established a first-mover advantage in turning generative AI into commercial products with the ambition of becoming the de facto cloud platform for AI workloads. Following viral adoption, ChatGPT is not only a household name, but we believe has provided Microsoft-OpenAI with early mindshare in the enterprise. Several companies (Spotify and Snap to name a couple) are not only experimenting with OpenAI's foundation models, but are releasing OpenAI-powered applications into production. To us, the strong early adoption of OpenAI indicates fast followers likely need to show strong differentiation, in order to convince enterprises already experimenting on OpenAI to look elsewhere. Furthermore, Microsoft-OpenAI have already enacted preemptive price reductions to quickly establish a large ecosystem of partners and customers, as well as to limit the potential for competition on price alone.

On the other hand, we do subscribe to the idea of generative AI as a rising tide lifting all boats, in

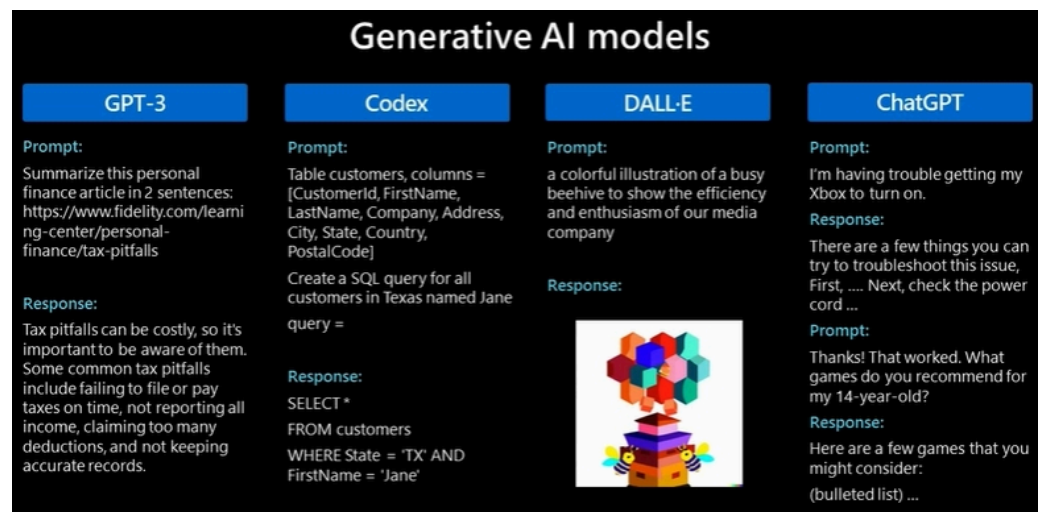
reference to the hyperscale public cloud vendors. We acknowledge AI as a market opportunity remains early and dynamics, despite progressing at an exponential rate. In terms of competition, we also would not discount the technology of close followers who, similar to Microsoft, have been investing in AI for more than a decade.

All together, we believe Microsoft Azure is currently in a strong position to win an above-average share of AI workloads running in the cloud and has been executing well strategically to sustain its early advantage.

Azure OpenAI Service

Azure OpenAI Service became generally available in January and already has 1,000+ enterprise customers. Azure OpenAI Service provides professional developers with API access to OpenAI's foundation models, including GPT, DALL-E, Codex, and Embeddings. Developers can use OpenAI's algorithms out-of-the-box to build new applications or augment existing applications with generative AI functionality. In addition, developers can fine-tune OpenAI's models by training them on their own datasets, in order to tailor the technology to specific use cases, similar to how Microsoft molded GPT-3.5 into Prometheus to make the technology more useful as a search engine alternative. As shown in the Figure below, there are the three model families available in the Azure OpenAI Service today.

Figure 7 - GAI models available in Azure OpenAI Service



Source: Company website

On March 9th, Microsoft announced the more out-of-the-box ChatGPT model is available in public preview. Using ChatGPT, developers can easily augment applications with AI functionality to summarize content, generate email copy, and act as an intelligent virtual assistant. Similar to the other models, developers can adapt the technology to better serve business-specific or vertical-specific use cases by training the model on their own datasets.

Azure can monetize OpenAI in several ways. As mentioned, Azure can sell API access and SDKs for OpenAI's large language models. More importantly, each workload using an OpenAI model (in testing and production, or to train a model) requires compute and storage, generating infrastructure-as-a-service (IaaS) revenue for Azure. This not only refers to the IaaS needed to support OpenAI's own offerings, such as the production version of ChatGPT, but all the workloads being powered by OpenAI's technology, such as Spotify's DJ, Snap's My AI, Salesforce's Einstein GPT, etc.



OpenAI Service's pricing is based on a pay-as-you-go consumption model with a cost per inference, which varies depending on the underlying model. For reference, an inference refers to the process of calculating a response to the query (i.e., for simplicity, think of each response generated by ChatGPT as an inference). We note this inference pricing is associated with AI models in testing or production. Also, in order to further train models, which is more compute intensive process than inferencing, Microsoft charges an hourly rate based on compute units.

OpenAI plans to sell subscriptions for solutions built on top of the company's large language model, such as the recently introduced \$20 monthly subscription to ChatGPT. While it's unclear how Microsoft earns a share of revenue from the subscription plan, the more important takeaway is all cloud compute and storage associated with supporting ChatGPT usage translates to Azure revenue.

Figure 8 - Azure OpenAI Service Pricing

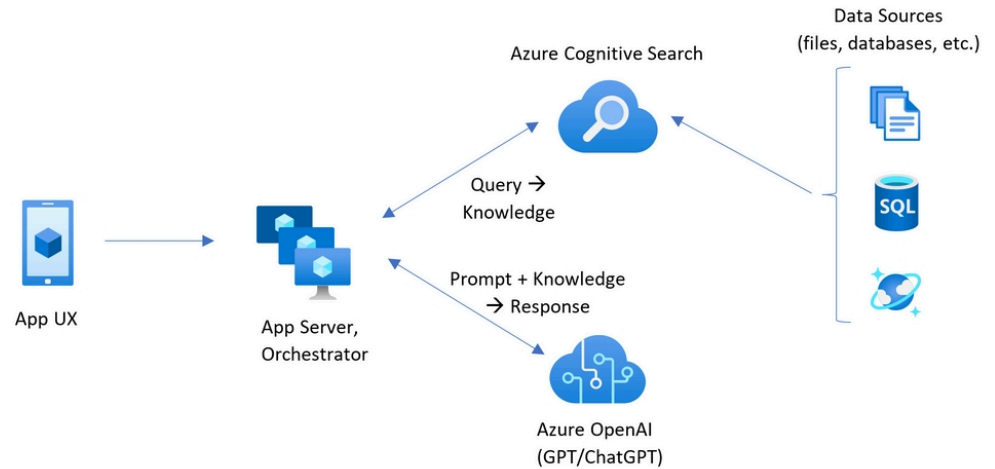
Azure Open AI Service: Pricing			
Series	Models	Inferencing (per 1K tokens)	Hosting (per hour)
Base Series	Ada	0.0004	N/A
	Babbage	0.0005	N/A
	Curie	0.002	N/A
	Davinci	0.02	N/A
Base Series Fine-Tuned	Ada	0.0004	0.05
	Babbage	0.0005	0.08
	Curie	0.002	0.24
	Davinci	0.02	3.00
Codex Series	Code-Cushman	0.024	N/A
	Code-Davinci	0.10	N/A
Codex Series Fine-Tuned	Code-Cushman	0.024	0.54
Embeddings Series	Ada	0.0004	N/A
	Babbage	0.0005	N/A
	Curie	0.002	N/A
	Davinci	0.02	N/A
Azure Open AI Service: Pricing, fine-tuning training			
Series	Models	Inferencing (per 1K tokens)	Training (per compute hour)
Base Series Fine-Tuned	Ada	N/A	20.00
	Babbage	N/A	22.00
	Curie	N/A	24.00
	Davinci	N/A	84.00
Codex Series Fine-Tuned	Code-Cushman	N/A	26.00

Source: Company website, RBC Capital Markets

Translates to more revenue for Azure's up-the-stack offerings too

Increased adoption of Azure OpenAI Service not only generates more Azure IaaS revenue, but should lead to increased interest in Azure's broader data and analytics platform. Microsoft has an extensive portfolio of artificial intelligence, machine learning, data, and analytics offerings. As customers increasingly adopt Azure OpenAI Service, we believe there is a strong value proposition around leveraging Microsoft's broader data and analytics portfolio, given out-of-the-box integrations and connectors for authentication (Azure Active Directory), security and compliance (Defender, Sentinel), and data stores (CosmosDB, Synapse, Data Lake), to name a few.

Figure 9 - Example of how Azure OpenAI opens the door to the broader Azure data and analytics portfolio



Source: Company website

- Azure Cognitive Services:** On top of OpenAI's foundation models, developers can purchase API access and software development kits (SDKs) for Microsoft's broader set of machine-learning models within Azure Cognitive Services. Microsoft is already showcasing efforts to cross-sell OpenAI's models with Microsoft's in-house technology, such as Azure Cognitive Search. As an example, developers can combine Cognitive Search to index, understand, and retrieve data from their own datasets (e.g., a company's knowledge base) and ChatGPT to understand leverage this foundation to answer questions or have a human-like conversation. Azure Cognitive Services includes APIs and services spanning vision (computer vision, custom vision, face recognition), speech (speech-to-text, text-to-speech, speech translation, speech recognition), language (natural language processing, machine-based text translation, conversational AI), decision (anomaly detector, content moderator, personalizer), and of course, OpenAI's models.

Figure 10 - Gartner Magic Quadrant for Cloud AI Developer Services




Source: Gartner, Magic Quadrant for Cloud AI Developer Services 23 May 2022 Baker et al.

All statements in this report attributable to Gartner represent RBC Capital Markets' interpretation of data, research opinion or viewpoints published as part of a syndicated subscription service by Gartner, Inc., and have not been reviewed by Gartner. Each Gartner publication speaks as of its original publication date (and not as of the date of this report). The opinions expressed in Gartner publications are not representations of fact, and are subject to change without notice.

- Data stores:** Not always appreciated is the fact that Microsoft is actually among the top providers in the aggregated database, data warehouse, and data lake landscape. Based on market share data from Gartner, Microsoft was actually the overall share leader in 2021, above Oracle and AWS. This is important, because AI is compute-intensive, requires large datasets train models, and generates equally large data volumes once in production. This data needs to be stored, queried, manipulated, and analyzed, hence we believe the database, data warehouse, and data lake software markets are likely to be clear secular beneficiaries of the AI wave. Microsoft's flagship offerings are in the database side of the market, with Azure SQL and SQL Server for relational workloads and CosmosDB for non-relational workloads. On the data warehouse side, Azure Synapse and Azure Data Lake are less established, but certainly appear to be growing in popularity.

Figure 11 - Gartner Magic Quadrant for Cloud Database Management Systems



Source: Gartner, Magic Quadrant for Cloud Database Management Systems, 13 December 2022, Cook et al.

Differentiation in DevOps and Low-Code

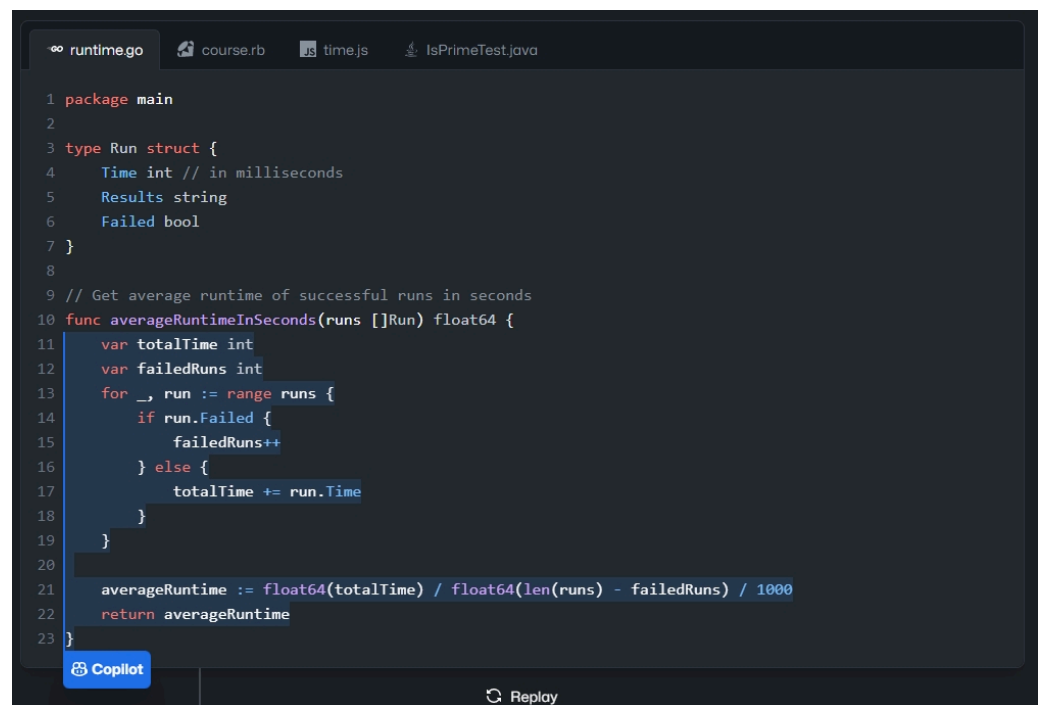
Initial collaboration between Microsoft and OpenAI was focused on enhancing Microsoft’s developer toolset, including GitHub and Visual Studio for professional developers and Power Platform for low-code, no-code tools. GitHub Copilot was the most notable innovation, in our view. While more niche in focus, GitHub Copilot saw viral adoption similar to ChatGPT, as Copilot gained 400,000+ subscribers just one month after becoming generally available in June 2022.

GitHub Copilot

So, how does GitHub Copilot work? When a developer begins writing lines of code, GitHub Copilot will recommend code that can complete the operation (similar to how Outlook will finish your sentence for you when you start typing). Also, a developer can prompt GitHub Copilot to generate code by describing the task in plain English. For example, a developer could prompt GitHub Copilot with, “Write a function to remove all duplicates from a list” and Copilot will then generate code to best match this description. In terms of the underlying technology, GitHub Copilot is powered by OpenAI Codex, a machine learning model based on GPT-3 and trained

using the millions of lines of code in open-source source projects hosted on GitHub. Copilot can also use previous code written by a developer to adapt to their individual programming style over time. GitHub Copilot supports a number of programming languages, including popular ones like Python, JavaScript, TypeScript, Ruby, PHP, and Go, as well as more niche languages like Rust.

Figure 12 - Example prompt of how Github Copilot writes code based on natural language prompts



```
runtime.go course.rb time.js IsPrimeTest.java
1 package main
2
3 type Run struct {
4     Time int // in milliseconds
5     Results string
6     Failed bool
7 }
8
9 // Get average runtime of successful runs in seconds
10 func averageRuntimeInSeconds(runs []Run) float64 {
11     var totalTime int
12     var failedRuns int
13     for _, run := range runs {
14         if run.Failed {
15             failedRuns++
16         } else {
17             totalTime += run.Time
18         }
19     }
20
21     averageRuntime := float64(totalTime) / float64(len(runs) - failedRuns) / 1000
22     return averageRuntime
23 }
```

Source: Github website

A legitimate differentiator in the developer toolchain market. GitHub Copilot could be a game-changer for developer efficiency, which provides Microsoft with a source of competitive differentiation in the developer toolchain market. While we believe competitors could eventually build a rivaling solution (AWS has a similar solution named Whisper), no other companies have the dataset of GitHub’s large repository of open-source code upon which Copilot/Codex has been trained. Also, Copilot is already being monetized today as a seat-based subscription offering for individuals (free trial then \$10/month or \$100/year) and business (\$19 per user per month).

Power Platform

Turning low-code into no-code with PowerFx. During Build in May 2021, Microsoft announced the first features based on GPT-3 as part of Power Apps, the company’s low-code application platform. The feature, called Power App Ideas, allows business users to create applications by turning conversational language in Power Fx, an open-source low-code programming language similar to Excel formulas/functions. In other words, this feature can effectively transform Power Apps from low-code to no-code, enabling a wider audience to build apps and workflows on the platform. Since then, Microsoft has introduced express design in Power Apps, which turns design into code. Using large AI models including DALL-E image recognition and Codex, users can turn design or Figma files into a fully functioning application.

Figure 13 - Gartner Magic Quadrant for Enterprise Low-Code Application Platforms



Source: Gartner, Magic Quadrant for Enterprise Low-Code Application Platforms, 31 December 2022, Vincent et al.

Nothing to lose in search

One of the areas where Microsoft hopes to gain share with generative AI is search, where even if the share gain is modest, it can be meaningful for the stock. During the investor session on the search opportunity, management noted each incremental point of share would equate to \$2B in revenue.

Can Microsoft monetize chat-based search? Currently, Bing’s chat-based search function lies on the same webpage as Bing’s traditional search, but on a separate tab. Over time, we would not be surprised to see future versions of Bing include both traditional search and chat search in one unified tab, where you can simultaneously get a response to your query in both formats. Most importantly, Bing provides embedded links (unclear whether these are sponsored or monetized today) to dive deeper into the answers to your query. So, yes, chat-based search is monetizable, and while the advertising opportunity might take some time to gain traction (based on our conversations, most advertisers seem to be in wait-and-see mode), Microsoft seems to already be making significant strides around the potential to make embedded links into sponsored ads

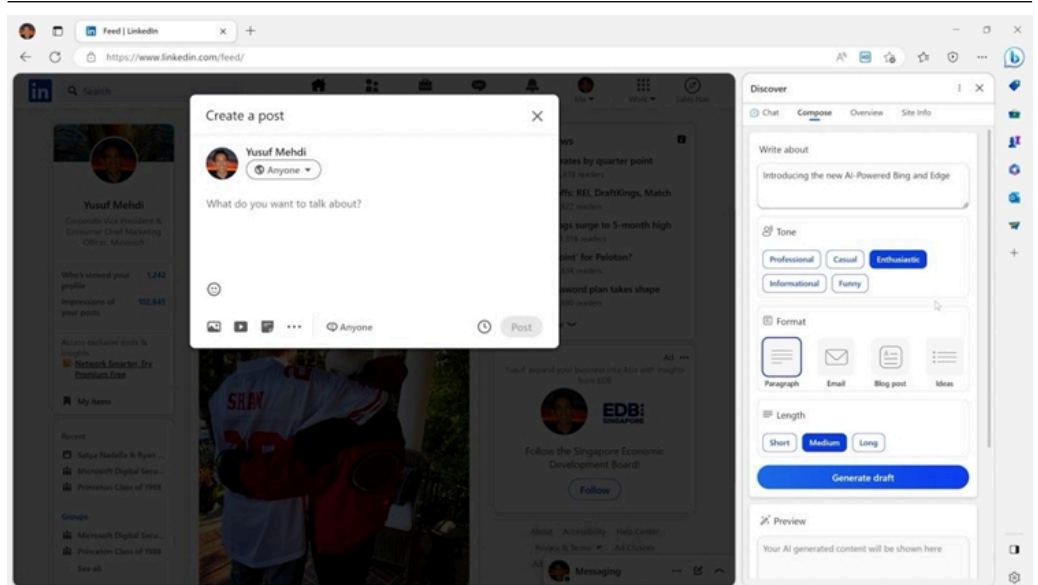
(eventually we expect display ads on the Bing Chat landing page).

In the Bing illustrative use cases section starting on page 185, we highlight a few use cases that illustrate the potential of ChatGPT for a typical consumer. Overall, we find the use cases to be really impressive, even if there are some slight inaccuracies and a learning curve to identify the right questions (known as "prompt engineering") to get the most useful responses. Some of these types of prompts are available on Github today.

The new Bing is powered by a large language model developed with OpenAI, called Prometheus, which management noted is more powerful than ChatGPT and tailored specifically for search. As part of the new Bing experience, Bing has a new sidebar with more detailed responses as well as an interactive chat experience (on the same webpage) for more complex searches with embedded links and citations (to the sites where Bing found the most relevant answers to the search). Also, Microsoft applied Prometheus to improve users' traditional search experience (i.e., the basic search engine and ranking algorithm).

The new Edge now has new functionalities in the sidebar – Chat and Compose. This allows users to leverage Bing's interactive chat experience to summarize an article or webpage and explore the content further, without opening a separate tab. In addition, users can prompt for AI-generated content in multiple tones, formats, and lengths, as shown in the Figure below. Importantly, the new sidebar adapts to the webpage a user is on for additional context.

Figure 14 - Edge Sidebar with Chat and Compose AI



Source: Company website

Making productivity suites more productive

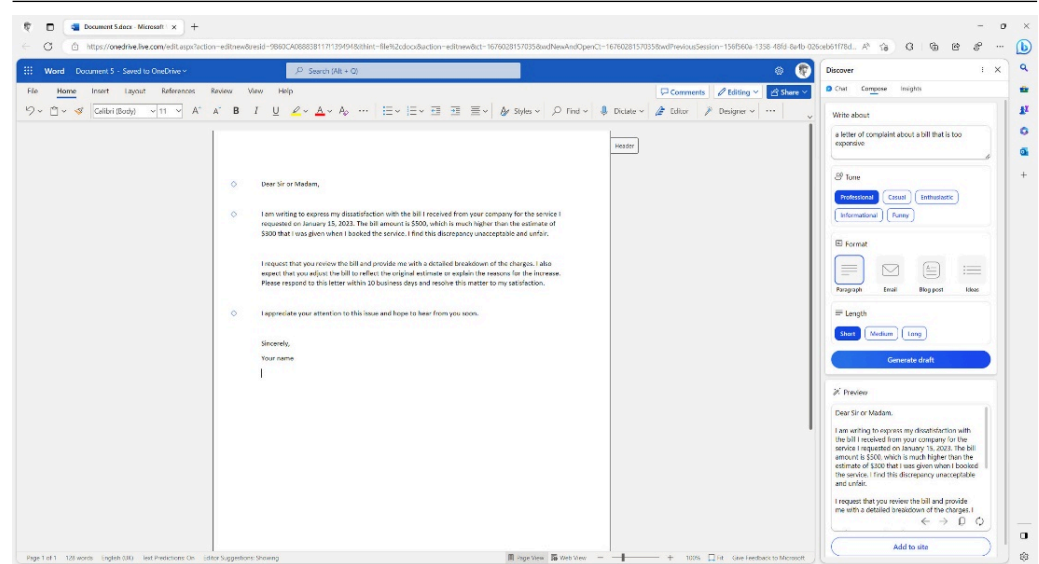
We believe Microsoft can leverage OpenAI's large language models to drive more innovation around the core productivity and collaboration suite of apps, including Word, Excel, PowerPoint, Outlook, and Teams, further extending its share dominance in the space.

Frustrated with Outlook's search experience? So are most people, in our view, but large language models should help improve it. Given Microsoft's enhancements to Bing's search engine, we believe a logical next step would be to take the learnings from Bing with Prometheus to give Outlook's search experience a much-needed face lift. We would not be surprised to see

Microsoft introduce chatbots to help search through emails and calendar invites (so rather than typing in keywords, you could ask the bot “find my most recent email conversation with Rishi”).

In addition, we believe other low-hanging fruit where AI can enhance Outlook is by helping write, finish, and edit emails. Microsoft has already announced Copilot for Viva Sales, which uses Azure OpenAI Service and GPT to help sellers respond to customers or prospects. We believe the logical next step would be to extend this functionality to all Outlook users, by using the context of conversation history, org charts, etc. to curate a personalized or role-based experience. Technically, this functionality is already available in the browser-based version of Office applications, as users can use the new Edge sidebar outlined above in order to compose an email or a paragraph as shown below.

Figure 15 - ChatGPT Auto Generates Emails



Source: Company website

Could we see a new Clippy, but useful? In our view, it is not hard to imagine how a real AI assistant could contribute meaningful value to Outlook (as described above), Word (AI-generated paragraphs, sentence completion), Excel (formula suggestions, automating macros), and PowerPoint (auto-formatting, AI-generated slides based on a Word document).

What has been announced so far? Microsoft is already incorporating generative AI in Teams Premium. We believe it is important to mention Teams Premium has several incremental features unrelated to OpenAI’s technologies but for the purpose of this report we will only mention features that leverage OpenAI’s technologies.

Teams Premium became generally available in February 2023. Teams Premium features leveraging OpenAI’s large language model include:

- **Meetings:** AI enhancements to standard Meetings in Teams include post-meeting transcription translations, AI-generated chapters (analogous to Spark Notes for each meeting, tailored to each meeting participant), time stamps for when a participant joins or leaves a meeting and when a participant’s name is mentioned, recommended tasks, and more.
- **Virtual appointments:** AI enhancements to virtual appointments includes auto-generated post-appointment follow-ups and organizational / departmental analytics.
- **Premium events:** In addition to the features outlined above, AI enhancements to Events



include full transcripts.

- **Live Transcription and Translation:** Participants can focus on paying attention to the meeting (and not taking notes word-for-word) given the meeting will be fully transcribed in live time (and then summarized using GPT-3) and participants will be provided with to-do suggestions. Teams Premium can translate 40 spoken languages (only the meeting organizer needs Teams Premium for all users to see the live transcription and translation features).

In terms of monetization, Teams Premium is available as an add-on for users with an Office 365 or Microsoft 365 subscription and costs \$10 per user per month. For a limited time, users can experiment with a 30-day free trial or lock in a 30% discount or \$7 per user per month.

First-mover advantage in intelligent applications

On March 6th, Microsoft announced Dynamics 365 Copilot, which brings capabilities from generative AI into Microsoft's portfolio of CRM and ERP apps. While monetization routes remain unclear, we do believe this helps improve Microsoft's value proposition in the space and can drive increased interest in the platform with a first-movers advantage around generative AI.

- **Dynamics 365 Sales and Viva Sales:** Here Copilot helps sellers write email responses to customers, creates email summaries of Teams meetings in Outlook, and pulls in details from the CRM system, like product and pricing information.
- **Dynamics 365 Customer Service:** Here Copilot helps customer service agents respond to customer inquiries in both chat and email (i.e., drafts answers based on the context) as well as provides context data from the knowledge base, case history, and more. In addition, Microsoft is now offering more out-of-the-box Power Virtual Agents, in order to build custom bots faster and to leverage Azure OpenAI Service and Bing to provide answers from company websites and integrated internal knowledge bases.
- **Dynamics CDP and Marketing:** Here, Copilot helps marketers segment customers using data from Customer Insights (Microsoft's CDP product). In addition, marketers can leverage generative AI to generate ideas for marketing campaigns.
- **Dynamics 365 Business Central:** For commerce, Copilot helps streamline the process of product listings for ecommerce, product descriptions can be auto generated using GPT prompts, and there are integrations to Shopify.
- **Dynamics 365 Supply Chain Management:** Here, Copilot helps flag external factors such as weather, financials, and geography that might cause supply chain issues. Supply chain planners can use Copilot to auto generate email alerts to send to the appropriate colleagues and partners when there is a detected issue with materials, inventory, carriers, distribution, etc.



GOOGL: Down but certainly not out

Along the lines of our thinking that the GAI strategy for hyperscalers must include best-in-class LLMs driving a user-facing GAI chatbot, MSFT has taken the lead with the rollout of ChatGPT, in the minds of investors. This has driven an unprecedented, almost existential crisis for GOOGL shareholders as this advance hits at both its core search business as well as its budding cloud business, which are the two primary sources of the company's cash generation and the stock's call optionality, respectively.

While this is the perception today, we think much of the difference relates to the marketing & aggressive messaging out of MSFT to date. As such, we see GOOGL having plenty of artillery to help shape a less destructive narrative for itself related to GAI over time through a few different means. 1) We expect Bard (which is based on GOOGL's LaMDA model) to get increasing play in mainstream media (as well as Google's other several larger models) which may help ease search share losses which is the easy-to-arrive-at conclusion. 2) We expect it to become clearer in the coming months that GOOGL is implementing many leading-edge GAI capabilities across its entire workspace suite that are intended to drive functionality that already mirrors the next generation of ChatGPT, which has not yet been released (ChatGPT4), and should also drive improving news flow and public awareness of GOOGL's own industry leadership. 3) Along these lines we'd expect enterprises & developers to take increasing notice, which should drive better partnership/buy-in and accordingly API adoption into using GOOGL's models & tools; this would be similar to the MSFT strategy and thus could/should become its own cloud adoption tailwind.

The four main debates around GAI's implications for GOOGL

1) Search disruption. The first key question is whether GAI displaces a significant amount of existing search activity by responding to queries with more relevant, more intended and/or easier to understand responses. Later in the note, we lay out several dozen examples of search next to the GAI where some level of displacement seems likely, but whether it would displace or reduce the monetization of search appears much more of an open question. Our work suggests that Google or any other search engine in control of their own solution would have meaningful latitude to manipulate the taxonomy of the UI such that it could still be solving for monetization with adequate conversion to support the economics in-line if not better than what's available today.

2) Google search share loss. With ChatGPT, OpenAI's MSFT partnership and the imminent Bing launch with ChatGPT integration, this has driven enormous buzz in both the enterprise & startup community along with mainstream media. Along these lines, it has driven an order of magnitude increase to Bing/Edge downloads and a significant increase already in Bing page views, which is obviously concerning for GOOGL given Bing's LSD share today. The share loss scenario becomes quite simple in that to the degree Bing's salesforce or self-serve platform can monetize the incremental traffic, we'd expect advertisers to look to move more budget over. With that said, we don't think this is occurring today and given the Google traffic & download data is essentially unchanged since the launch of ChatGPT, the early data suggests that much or all of the usage is incremental at this point. All that said, we're reserving further judgment until a full integration of GAI at Bing and we're better able to observe Bing's competitive advances in the wild.

3) GAI tools driving cloud share shift. Based on ChatGPT's first-mover advantage and strong early momentum, an important longer-term concern for GCP is that many of its customers will or already have begun leveraging the ChatGPT model through APIs and in doing so, those customers will be forced to run those workloads through Azure which could open up broader cloud service discussions leading to GCP being displaced. The opposing argument to this would be that GOOGL's already rolling out tools with bigger models & more capability vs. the current



version of ChatGPT, and as enterprises & developers become more aware of this (GOOGL not exactly doing great at marketing at this point), enterprises & developers will suddenly have two significant generative AI tools to work with. And while we think MSFT has the strategy & resources to price these tools aggressively, given its lack of vertical integration for both software & its foundation models, we'd think that compute costs could be meaningfully more attractive on GCP (also AWS) vs. Azure while it can leverage that into pricing advantages on other application layers within the broader cloud service offering (again, AWS likely has the best overall value prop here given its broader array). Sustaining first-mover advantages tends to work best when the advantage itself allows the technology to maintain the advantage and, given that's already not the case, it's unclear to us how much impact MSFT's Trojan horse strategy can drive if it has a finite shelf life.

4) Rising compute cost drives persistent margin compression. While we've yet to get clarity on what GAI queries (as an example) may cost relative to traditional search queries, we've heard estimates ranging between 2-10x which is obviously concerning relative to GOOGL's gross margins and capex intensity and is a key reason for the stock's recent underperformance after MSFT's CEO commented publically that search gross margins would decline 'forever'. Further, as larger, higher-parameter models come out, we believe there's a strong correlation between parameter count & compute intensity which only bolsters the margin compression bear case (which the market is well aware of at this point). While we certainly have our concerns on this topic, there are solid reasons why this may not be as impactful as it might appear and secondarily, we do think this is at least partially baked into the stock at this point given MSFT's prolific commentary to the markets 3 weeks ago.

a) On its Q4 earnings call, GOOGL management was asked about the cost of rising compute intensity associated with GAI and basically dismissed it as a concern highlighting that it has always approached its capital intensity for all of its innovation keeping a keen eye on ROIC. Given the company announced it is rolling out a fully capable GAI solution within many parts of workspace earlier in the week (6 weeks after the earnings conference call), we believe the company may already be displaying signs of leveraging the technology with virtually no discernible impact on capex intensity.

b) We believe having its own hardware and vertically integrating with the LLM may be a longer-term competitive advantage. We believe the use of 1p GPU vs. merchant silicon can be as much as 50% cheaper where GPU vendors may actually have significant pricing power for leading edge chips to non-vertically integrated players given most hyperscalers can use their own.

c) Further, as future generations of 1P GPUs are achieved, we think the scale and cost improvements relative to current levels could be substantial where even if GAI query costs are on the order of 5-10x vs. traditional search queries today for example, it seems very unlikely that that gap wouldn't significantly narrow over a time frame which could have a substantial upward effect on GOOGL's terminal margin that is very likely not the case today.

d) The simplest and most obvious offset not discussed would be revenue where GAI aspires to drive additional functionality & intelligence for which we think GOOGL can extract more revenue. Performance Max has been the best example of this (ML/DL) where while there likely has been higher computing costs associated with it, we've consistently heard that not only does Pmax extract more revenue on an apples-to-apples basis but there's also a widely held view that it is driving better conversion by automating campaign creation, audience building and ultimately, performance. We think things like targeting, productivity tools and the potential business model adaptation based on differentiated value could all serve to offset any compute cost difference.

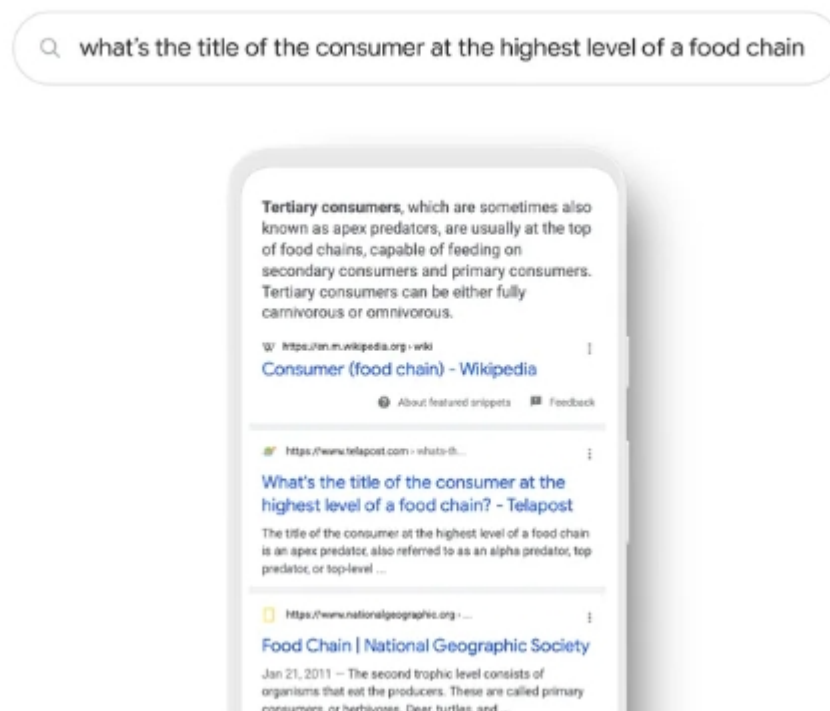
What is GOOGL doing with AI/ML today?

While OpenAI and MSFT have gotten the lion's share of AI attention recently, it's worth remembering that GOOGL has been developing and implementing AI/ML in their business for years. Over six years ago CEO Sundar Pichai spoke to GOOGL being an AI first company and the company has pioneered several innovations in the space, particularly their 2017 Transformer research paper. Given some of the ML/AI innovations are less well-publicized than ChatGPT, we highlight some of the work GOOGL has been implementing across Search, Advertising and Cloud below.

Search

RankBrain, launched in 2015, was the first AI implementation in Search and helps the engine relate words to concepts. This innovation broadened the capabilities of Search where it can better understand less explicit queries and improve the relevance of results. As an example, with RankBrain the search engine is able to produce more relevant results by recognizing the context and concepts relevant to the query where previously it may have not recognized the user is looking for the term "apex predator" vs. results related to human consumers.

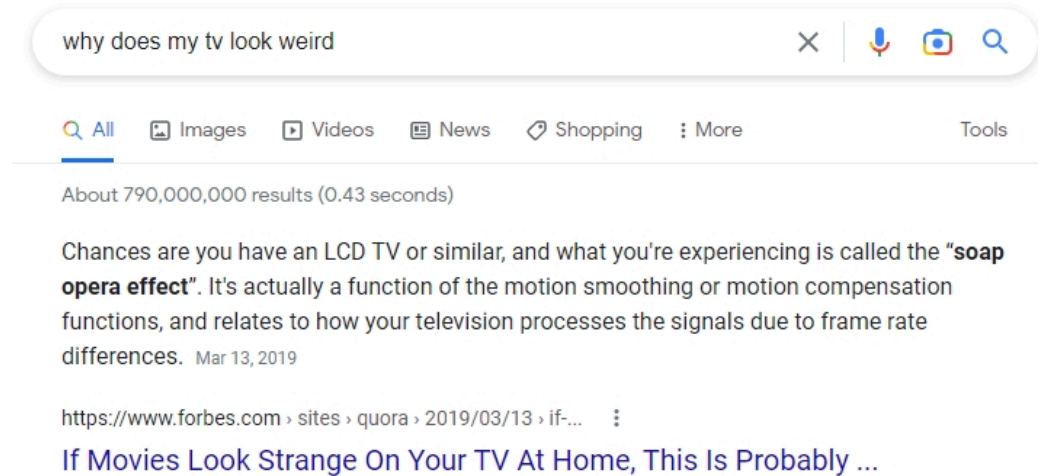
Figure 16 - With RankBrain, the search engine is better able to identify the context of a user's query and produce more relevant results



Source: Google Blog

Neural Matching was implemented in Search in 2018 and better understands how queries relate to pages. This innovation builds upon RankBrain's foundation to enhance the engine's ability to process more ambiguous queries by looking at the entire query vs. keywords and a better comprehension of the underlying concepts, broadening the index search to return the most relevant results. For example, if you search "why does my TV look weird", the engine will recognize the context of the query and return results for "soap opera effect" even though the search text doesn't include any of those keywords.

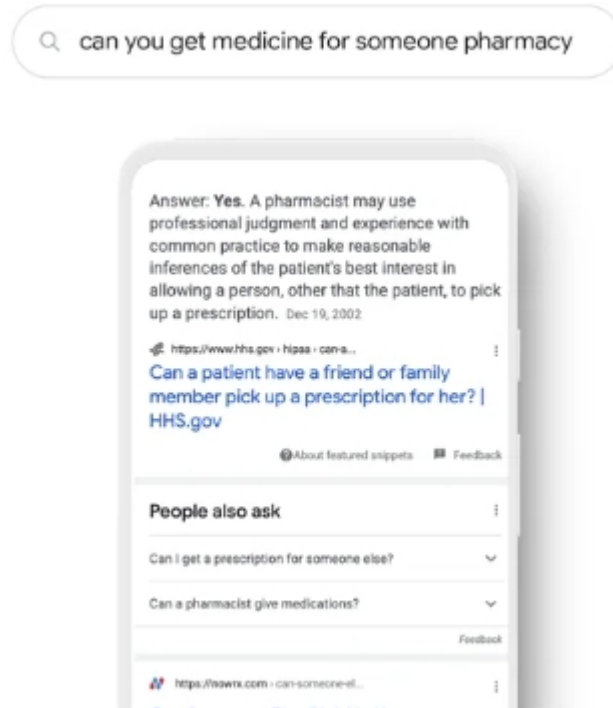
Figure 17 - With Neural Matching, the search engine is able to identify concepts and produce a relevant result despite the search text not including any of the keywords of the correct response



Source: Google.com

Bidirectional Encoder Representations from Transformers (BERT), launched in 2019, understands how different combinations of words have different meanings and intents. Instead of identifying each keyword individually, BERT can comprehend the sequence of words and how that relates to the context of a query, thereby improving web page rankings. In the example below, BERT can recognize the user is looking to pick up medicine for someone else where previously the search engine might have omitted the word "for" and provided results on how to fill a prescription for yourself.

Figure 18 - With BERT, the search engine can better relate combinations of words to user intent where the sequence of words in the query is important to the context of the search



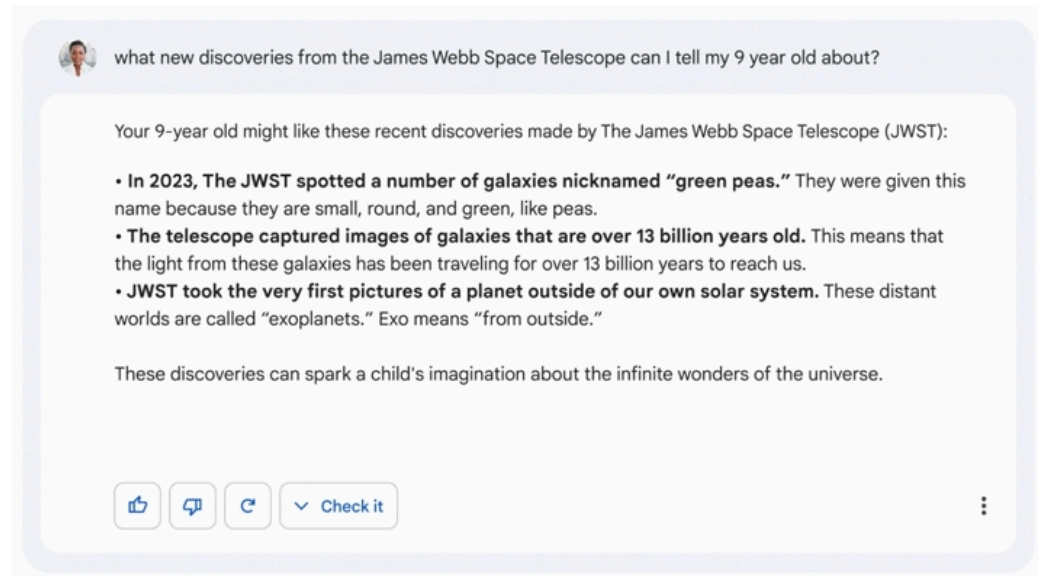
Source: Google Blog

Multitask Unified Model (MUM) was introduced in May '22 and can better understand and generate language (in 75 languages), which Google describes as 1,000 times more powerful than BERT. Beyond text, MUM can also understand images. While still relatively early days, GOOGL plans to further integrate and implement MUM applications in their products going forward.

Language Model for Dialogue Applications (LaMDA) was announced in 2021 and importantly adds conversation capabilities to GOOGL's AI foundation where the AI can engage in conversation with users where previously the AI functionality was largely limited to better language comprehension.

Bard was announced in February '23. Powered by LaMDA (initially on the lightweight version of the model), Bard allows users to engage in conversation with the chatbot similar to ChatGPT. While still in the testing phase GOOGL plans to make it more widely available and integrate within Search, again, similar to Bing integrating ChatGPT.

Figure 19 - An example of Bard's capabilities to generate text responses



Source: Google Blog

Advertising

Performance Max (PMax) launched in 2021 and allows advertisers to access all of GOOGL's ad inventory in a single campaign. PMax uses machine learning models (Smart Bidding) to automatically optimize for an advertiser objective(s) (traffic, leads, sales, etc.) in real time across their ad properties. Creating a PMax campaign is relatively straightforward. Advertisers simply select a goal, budget & bidding strategy and location & languages, add their assets/asset groups (images, videos, logos, etc.) and publish the campaign. A key component of PMax is discovering new customers advertisers might not have been previously targeting in manually created campaigns. Using AI to understand an advertiser's historical customer behavior, GOOGL can better identify converting audiences outside of a perceived target market without necessarily lowering advertiser ROI. Additionally, when advertisers define the value of each type of conversion (phone call, online form fill, etc.), the AI can take that information into account when bidding across their ad properties. With responsive search ads, campaigns can adjust content in real-time based on users' searches to improve performance.

Cloud

GCP has been incorporating AI/ML tools for years for a broad range of applications (see table below). Beyond the software tools/capabilities, GCP also has their own custom TPUs purpose built for training models, which could lower costs vs. to off-the-shelf silicon. Most recently, GOOGL announced (on 3/14) new AI tools within Google Workspace (which appear similar to ChatGPT's capabilities) and importantly a GAI app builder. For example, a "Help Me Write" tool will be integrated within Docs and Gmail to assist users in creating text, editing syntax & tone, adjusting the language based on the intended audience, etc. The tool will be launched in the U.S. for testing in March.

Figure 20 - GCP has an extensive array of AI/ML Cloud API/capabilities which we expect them to build upon going forward

Use Case	Product	Description
Scientists	Vertex AI	Centralized UI to manage all ML workflows. Pretrained APIs to quickly develop GAI applications
Scientists	Vertex AI Workbench	Development environment to manage the entire application/model workflow. Allows for faster building and training, better scale and data analysis
Scientists	Vertex Explainable AI	Set of tools/frameworks to better understand ML outputs to debug and improve model performance.
Developers	AutoML	Simplified ML development for less experienced developers across image, text, video, translation, etc applications
Developers	Natural Language AI	Analyzes unstructured text to provide better insights including sentiment analysis, content classification, syntax analysis, etc
Developers	Dialogflow	Customer service chatbots/voicebots
Developers	Media Translation	Real-time audio translation
Developers	Speech-to-Text	Converts audio to text in 125+ languages. Available on-premise or via cloud.
Developers	Text-to-Speech	Converts text to speech at near human quality levels in 40+ languages in 220+ voices. Ability to train custom voice models via customer audio recordings.
Developers	Recommendations AI	Analyzes customer data to maximize metric business is optimizing for at scale to improve conversions and ROAS
Developers	Translation AI	Translate text in 100+ languages
Developers	Video AI	Automatically recognizes 20k+ objects/places/actions. Generates video metadata. Builds video apps.
Developers	Vision AI	ML models to analyze videos and images. Detects objects and handwriting. Pretrained APIs to extract metadata.
Infrastructure	Deep Learning Containers	Environment to develop, test and deploy AI applications
Infrastructure	Deep Learning VM Image	Access to VMs with pre-installed applications
Infrastructure	Cloud GPUs	Access to GPUs to improve job speed with a wide range of performance and prices
Infrastructure	Cloud TPUs	Access to TPUs for larger workloads and ML model training/interating
Infrastructure	TensorFlow Enterprise	Enterprise-grade CPU, GPU, TPU resources
Solutions	Contact Center AI	Human-like AI for customer service
Solutions	Document AI	Provides structured data from documents
Solutions	Intelligent Products	Customer insights to develop and monetize new products
Solutions	Product Discovery	Implement Google-quality search on websites

Source: Google, RBC Capital Markets



Google's investment in Anthropic

Google invested ~\$400M in Anthropic for a ~10% stake in the company as reported by Bloomberg, and the money will be used in part to acquire additional computing resources. Anthropic was created in 2021 by Dario Amodei and other executives who were colleagues at OpenAI. The founders left due to a disagreement with OpenAI pertaining to the original mission of the company vs. the current direction. Prior to taking money from Google, the startup had raised \$700M+. Anthropic has developed a chatbot called Claude, that has similar capabilities to ChatGPT.

How might search business models to be altered or at least adapt to the inclusion of GAI

Given the inevitably higher compute costs associated with GAI, we think it's possible to see the business model evolve over time - not necessarily just for search per se but for higher frequency GAI users where we'd think there could be several different ways to at least partially or fully cover the higher associated compute costs.

Traditional ad load but with better economics due to more clicks above the fold etc: Depending on taxonomy fueled by a likely whole new realm of A/B testing, to the degree that user engagement for above-the-fold GAI responses is better than traditional search (which seems plausible), conversion would inevitably be better as well, which would drive CPMs higher. What's unclear to us is whether the showing of fewer ad impressions in general would be at least somewhat of a headwind of CPM inflation as a function of theoretically same intent (or more in the event that monetizable usage is truly incremental). Directionally, we're reasonably confident that CPM inflation will result here not only as a function of GAI, but more broadly as AI increasingly drives targeting & performance, however, the magnitude of that improvement is what is likely to remain an open question.

Throttling. Within 2 months following launch, driven by server overload and a host of bad actors using the product in a variety of manipulative & sometimes nefarious ways, OpenAI & MSFT have already put caps on how many questions can be asked on a daily basis from an initial 6 chats/day which was then expanded to 60 with an expected increase to 100 chats per day in the future. A common initial reason which drove this change, in our view was that the bot was being drawn into acrimonious conversations which were then publicized and obviously carried reputational risk with potential to both indict the technology as doing more harm than good which, of course, could then threaten its virality & adoption. Underneath this, we have to think OpenAI (& maybe MSFT) decided to contain the likely rapid compute cost expansion as they worked through the initial kinks & bugs phase of making the product & underlying tech more commercially viable. Over time, we also think throttling could be an important mechanism to limit cost as just like traditional search, LLMs will have an enormous responsibility to index as much of the web as possible on a daily/weekly/monthly periodic basis in order to maintain the quality & relevance to provide a durably differentiated experience both vs. its own search engine as well as competitors.

Subscription/pay-to-play. Between the likely large group of individual GAI super-users combined with companies wanting enterprise access to both the bot as well as any of the underlying LLM's, we have to think there'll be a price to pay for this type of access & associated elevated compute cost. We'd think a normal consumption model would be sufficient here and while this would at some level manifest as a price increase, ie impediment to adoption, we think this should align well with companies (both new GAI native as well as those migrators) which are drawing increasing competitive advantages & associated incremental revenue.

Ad paywall. The viability of an ad-supported model is less clear given that there's already likely to be a hefty amount of ad load in the core use-case to a certain threshold. Beyond that point,




the Netflix analog could hold true in that subscription would entirely support plenty of adoption for elevated usage and then as more of user-behavior & habits became known over time and adoption slowed down, then a plausible incremental layer of adoption & usage could be driven by an ad-supported tier. We'd think the economic contribution from this would likely be dilutive due to relevance & monetization potential while also would likely precede something of a more mature state and therefore could be many years if not decades away.

What is ChatGPT's impact on traditional search and social platform engagement, if any?

The increasing interest in ChatGPT doesn't appear to have had any/a material impact on Google or engagement on social platforms so far. With the meteoric rise of ChatGPT, reaching 100M users in ~2 months, according to Similar Web, investors are likely contemplating the impact on Bing and Google search engines (particularly post ChatGPT integration with Bing) and social media usage. That said, the traffic, download and time spent data suggests the interest in Bing/ChatGPT is incremental.

As background for our analysis, we look at historical search engine market share in the U.S. and globally where GOOGL is the dominant player and has been gaining share over the past decade+.


Figure 21 - Google Search vs. Bing market share in the U.S. since 2009 and find Google has gained ~800bps of share



Source: StatCounter, RBC Capital Markets



Figure 22 - Globally, Google has steadily gained share, though less than in the U.S., since 2009 while Bing share has been flat/slightly down




Source: StatCounter, RBC Capital Markets

One of the more intuitive conclusions investors could draw is that Bing could take search share from Google as it integrates ChatGPT into the search engine. However, the download and web traffic data would generally refute that notion (though worth noting ChatGPT integration in Bing isn't available to all users yet and traffic is something we'll monitor going forward). While Google had been on a downtrend, downloads largely troughed before the spike in Bing interest and had since risen from the lows (worth noting these are 7 day trailing average daily downloads so some volatility is to be expected). Additionally, Bing downloads are down 70% from the peak which may be an indication the recent interest in Bing is temporary vs. durable share gains.




Figure 23 - As shown in the U.S. app download data, Bing downloads spiked in February which may have been a marginal headwind to Google downloads at most, though downloads have actually risen since the lows in February and Bing downloads declined significantly after



Source: Sensor Tower, RBC Capital Markets

On web traffic, the uptick in Bing in the beginning of the year is noticeable but doesn't appear to be at the expense of Google, implying the Bing traffic is incremental.


Figure 24 - Looking at page views, it appears as though Bing did get a noticeable uptick in traffic, coinciding with the rise of ChatGPT, while the impact to Google's traffic appears non-existent, implying Bing traffic is incremental



Source: Similar Web, RBC Capital Markets. Note: Bing and ChatGPT are on right axis and Google is on left axis


On the social side, download data suggests no/minimal impact from ChatGPT as all five apps' downloads are higher now than they were when ChatGPT was released. Engagement (total time spent) on social apps also appears to continue following seasonal trends despite the viral interest in ChatGPT.

Figure 25 - ChatGPT doesn't appear to have a downward impact on social downloads as all five apps' downloads are up since the release of ChatGPT



Source: Sensor Tower, RBC Capital Markets

Figure 26 - While total U.S. time spent on popular socials did come down in February, it doesn't appear materially different than normal seasonality and we wouldn't infer the m/m decline as driven by ChatGPT taking mindshare from social platforms



Source: Sensor Tower, RBC Capital Markets

Side by side: Does ChatGPT replace traditional search use cases?

In our analysis of ChatGPT and the potential ramifications for search advertising, we did a side-by-side comparison of general searches in different verticals on both Google and ChatGPT. The screenshots of each vertical are shown on page 97-184. Given the closed nature of ChatGPT currently, some searches are not perfectly analogous, but we wanted to demonstrate the potential and where ChatGPT could disrupt the search business.

Big brands vs. the long tail (SMBs). We see a top-of-funnel shift to GAI having the potential to



widen the moat for big brands & agencies vs. the long-tail of internet/Ecommerce participants (which are the majority of the economy - and critical to digital ad players in particular like GOOGL & META). Specifically, we expect those with greater resources to leverage both internal & external tools to create the most compelling content, thus creating both an SEM as well as an SEO discoverability headwind for those smaller companies lacking the resources to leverage GAI to their advantage. An area of significant intrigue for us lies in whether the purveyors of LLM's will allow developers to build solutions which effectively can replicate near-perfect relevance (think dynamic content algorithmically optimized) to drive outsized discoverability. Can anyone become a search engine? If digital ads were the new rent, does GAI become the new digital ad?

Key themes that stood out as we observed & analyzed side-by-side comparisons of search vs. ChatGPT (@ also within Bing):

1. In terms of ad load, we found that most cases produce results that are largely similar and thus, we believe a search operator should have a similar ability to manipulate the taxonomy on the page to build in efficient ad load within the GAI UI. With that said, to the degree that GAI's responses hit the mark on the first page more than traditional search, CPM's would likely rise further given the higher conversion which would net out any ad impression weakness with less searching required.
2. GAI results in significantly lower SEO organic results in many commercially monetizable search categories like travel, professional services, auto retail & real estate as some examples. Unclear here which force is stronger though - either a) organic search scarcity & the higher odds of conversion off the first page or two of GAI responses driving potentially significant upward CPM pressure OR b) do the organic headwinds created by the reduced organic page space simply weaken the long tail of sites seeking to be discovered which in turn, drives reduced demand by those same entities which are frequently also SEM customers. Our view is that the former likely outweighs the latter given we think SEO discoverability for the long tail is already a significant minority of traffic, but it's simply too early in our understanding of likely user behaviors to say for certain.
3. The back and forth potential of GAI should create better recommendation engines for paid ads (as well as organic) which should push conversion higher and accordingly, CPMs & CPCs. Traditionally, the goal of any search engine is to find the most relevant response in working with only one question or query. With GAI enabling follow-up questions, this should raise the odds of hitting the most relevant results which should also yield more relevant ads - hence, the likely potential of higher CPC/CPMs.
4. On the one hand, a consistent GAI UI over will likely provide growing comfort levels for digesting & hopefully trusting content - on the other hand, we believe that many searches ultimately require a primary source to ensure a user is comfortable with the validity/accuracy of what they found. Conceding we're fighting an extreme amount of dogma given our own two decades of exclusive Google usage, to the degree that GAI's UI would suffice for a response, we'd think a similar outcome would happen in traditional search where Google often returns a similar, singular response (though coming from a single website referenced below it). We'd think the inclusion of linked sources seems critical to augment that comfort level which again, could challenge the taxonomy on the page to the degree that ad load is increasingly incorporated.
5. GAI may be the closest approximation to a truly unbiased response mankind has ever seen, however, given it is often a single response to a question with multiple answers, it is by definition not possible to be biased in at least some direction (should you visit the Colosseum in Rome on the first day of your trip or third for example, or which movie theater has the best air conditioning). In lieu of providing simple lists of primary sources as responses which will ultimately likely require further clicks, we don't see a way around giving users the choices of primary sources which most users would usually desire. When it comes to anything subjective, (even those searches which must subjectively select a set of objective responses),



we see GAI facing a virtually impossible task of satisfying all users without the prospect of offering adequate permutations of correct responses.

6. Categories where visual-based search is important or necessary seem better protected from disruption - it's difficult to come up with an Ecommerce category of goods or even services where at least some level of visual shopping isn't involved before moving further down the funnel in a targeted way on the way to conversion/purchase. Over time, we'd expect GAI UI's to reflect a similar image/text mix where applicable for different types of searches but again, given the ease of manipulating taxonomy to maximize ad load, we'd think a GAI wouldn't have any trouble delivering an optimal image/text mix to drive shopping & down-funnel conversion.
7. SEO professionals need to become GAI experts asap. To the degree that GAI greatly reduces organic discoverability, we'd imagine we'll see a significant education period by SEO professional followed by a substantial increase in demand for GAI SEO discoverability.

Figure 27 - Past 10 years of top Google searches

Search Rank	Trending Searches									
	2022	2021	2020	2019	2018	2017	2016	2015	2014	2013
1	Wordle	NBA	Election results	Disney Plus	World Cup	Hurricane Irma	Powerball	Lamar Odom	Robin Williams	Black Friday
2	Election results	DMX	Coronavirus	Cameron Boyce	Hurricane Florence	Matt Lauer	Prince	Jurassic World	World Cup	Cyber Monday
3	Betty White	Gabby Petito	Kobe Bryant	Nipsey Hussle	Mac Miller	Tom Petty	Hurricane Matthew	American Sniper	Ebola	Daytona 500
4	Queen Elizabeth	Kyle Rittenhouse	Coronavirus update	Hurricane Dorian	Kate Spade	Super Bowl	Pokemon Go	Caitlyn Jenner	Malaysia Airlines	Chinese New Year
5	Bob Saget	Brian Laundrie	Coronavirus symptoms	Antonio Brown	Anthony Bourdain	Las Vegas shooting	Slither.io	Ronda Rousey	Flappy Bird	Tour de France
6	Ukraine	Mega Millions	Zoom	Luke Perry	Black Panther	Mayweather vs McGregor fight	Olympics	Paris	ALS Ice Bucket Challenge	AMAs
7	Mega Millions	AMC Stock	Who is winning the election	Avengers: Endgame	Mega Millions Results	Solar eclipse	David Bowie	Agar.io	ISIS	French Open
8	Powerball numbers	Stimulus Check	Naya Rivera	Game of Thrones	Stan Lee	Hurricane Harvey	Trump	Chris Kyle	Ferguson	Cinco de Mayo
9	Anne Heche	Georgia Senate Race	Chadwick Boseman	iPhone 11	Demi Lovato	Aaron Hernandez	Election	Fallout 4	Frozen	Yom Kippur
10	Jeffrey Dahmer	Squid Game	Playstation 5	Jussie Smollett	Election Results	Fidget Spinner	Hillary Clinton	Straight Outta Compton	Ukraine	LA Auto Show

Source: Google Trends, RBC Capital Markets

Figure 28 - Average GOOGL search CPC by vertical

Industry	Average CPC
Legal	\$6.75
Consumer Services	\$6.40
Technology	\$3.80
Finance & Insurance	\$3.44
B2B	\$3.33
Home Goods	\$2.94
Dating & Personals	\$2.78
Health & Medical	\$2.62
Industrial Services	\$2.56
Auto	\$2.46
Education	\$2.40
Real Estate	\$2.37
Employment Services	\$2.04
Travel & Hospitality	\$1.53
Advocacy	\$1.43
E-Commerce	\$1.16

Source: WordStream.com, RBC Capital Markets



Figure 29 - Average GOOGLE search CPA by vertical

Industry	Average CPA
Technology	\$133.52
Real Estate	\$116.61
B2B	\$116.13
Advocacy	\$96.55
Consumer Services	\$90.70
Home Goods	\$87.13
Legal	\$86.02
Finance & Insurance	\$81.93
Industrial Services	\$79.28
Health & Medical	\$78.09
Dating & Personals	\$76.76
Education	\$72.70
Employment Services	\$48.04
E-Commerce	\$45.27
Travel & Hospitality	\$44.73
Auto	\$33.52


Source: WordStream.com, RBC Capital Markets

Taking it to the model - Impact on GOOGL's '24 estimates

But first, a look at CapEx over the past several years. As GOOGL continues to build out AI, ongoing technical infrastructure (data centers) investments will likely be required. That in mind, we look at historical CapEx as % of revenue vs. when GOOGL rolled out their major AI advancements (worth noting GOOGL does have other CapEx not related to AI but is at least directionally consistent, in our view).

2018 was the year with highest CapEx intensity which preceded the launch of BERT in 2019 and 2014 is the second highest which preceded the launch of RankBrain, both of which were significant innovations in Search. Since 2018, CapEx intensity has come down consistently with the exception of 2022 which notably featured the launch of MUM and the rise of LaMDA & Bard. Going forward, significant CapEx intensity increases seem unlikely, in our view, given a) revenue has grown so much over the past few years and b) the company has already been investing in AI for nearly a decade already. Directionally higher CapEx intensity in the next few years is certainly possible, but we don't expect going back to '18 levels.


Figure 30 - GOOGL's capex as % of revenue is near historical lows



Source: Company Reports, RBC Capital Markets

Looking across our internet coverage more broadly (excl. AMZN, GOOGL, META), CapEx intensity years has fallen over the past several years even as the adoption of AI has grown over that period. And while some may argue the CapEx leverage is a result of cloud adoption where the costs were largely shifted from CapEx to COGS, gross margins increased significantly over the same time period suggesting that the increasing compute intensity associated with historical shifts to AI has not translated into margin pressures as scale and processor innovation has been able to offset. Admittedly, GAI's compute intensity is likely altogether higher than legacy AI workloads but to the extent that computing becomes increasingly more efficient as in past generations, we'd expect investor fears on the matter to abate over time.


Figure 31 - Across our Internet coverage, CapEx as % of revenue has been trending downwards over the past several years, coinciding with the rise of AI



Source: Company Reports, Factset, RBC Capital Markets



Figure 32 - And at the same time, GMs have been increasing, implying the rise of AI has not driven margins downwards



Source: Company Reports, Factset, RBC Capital Markets

The potential impact of GAI going forward is unknown. Given investor attention on the topic, we attempt to size the revenue/gross margin impact on the Street's '24 estimates under three separate scenarios: 1) Impact on GOOGL's search revenue due to share loss to Bing. 2) Impact on GOOGL's search revenue based on GOOGL searches fulfilled by their own GAI chatbot. 3) GOOGL search GM impact due to compute intensity delta between GAI and traditional search. While these scenarios likely won't occur in a vacuum, we think sizing the range of potential impacts of each scenario separately will help investors dimensionalize the factors that could affect GOOGL's search business. For purposes of this analysis, we don't assume any Search TAM expansion as a result of GAI.

At a high level - we think investors are already contemplating much of the potential adverse impact to GOOGL's earnings associated with GAI & ChatGPT. As is summarily shown below in Figures 48-50, we estimate a base case scenario where there's a theoretical 12-13% downside risk to '24 EPS under what we'd argue are somewhat generous assumptions around share losses to Bing (5% points of share), the portion of Bing's share-gain searches which are monetizable (7.5%) and the compute cost relative to traditional search (base case assumes 4x).

So with the stock having dislocated by 12% YTD and virtually sitting on top of its 9 year trough at just below 15x '24 EPS, assumptions of more than 5% share gains, double digit monetizable searches from share gains and 6x+ relative compute costs would have to be the assumption to arrive at a downside scenario closer to 25-30%.

Variable 1 - Impact on GOOGL's search revenue due to share loss to Bing

Bing taking share from Google Search will likely not be a significant revenue headwind. The integration of ChatGPT in Bing has been widely followed by the media and appears to be a strategic move by MSFT as CEO Satya Nadella recently noted in an interview that he's willing to demonetize Bing (accept higher compute costs associated with GAI) in order to take share from GOOGL given Bing margins are incremental to the business, even at higher CoR levels. We conduct an analysis to determine the range of outcomes of GOOGL share losses to Bing and the % of share losses that would have otherwise been monetized by GOOGL (GOOGL currently monetizes ~20% of total search queries), and calculate the implied % revenue impact. We find



that if Bing took 10% share (highly unlikely, in our view) and 30% of those searches were monetizeable by GOOGL, the estimated total revenue headwind be under 2%.

Figure 33 - We sensitize the potential impact on GOOGL search and total revenue due to search share loss to Bing and the % of the share loss that was monetizeable by GOOGL

GOOGL search revenue share loss to Bing sensitivity analysis (\$M) - Current '24 search revenue estimate = \$185.2B						
		GOOGL % share loss to Bing				
		1.0%	2.5%	5.0%	7.5%	10.0%
% of Bing share gain as monetizeable searches	15%	\$184,930	\$184,513	\$183,818	\$183,124	\$182,429
	30%	\$184,652	\$183,818	\$182,429	\$181,040	\$179,651
	45%	\$184,374	\$183,124	\$181,040	\$178,957	\$176,873
	60%	\$184,096	\$182,429	\$179,651	\$176,873	\$174,095
	75%	\$183,818	\$181,735	\$178,262	\$174,790	\$171,317

GOOGL search % change in revenue sensitivity analysis						
		GOOGL % share loss to Bing				
		1.0%	2.5%	5.0%	7.5%	10.0%
% of Bing share gain as monetizeable searches	15%	-0.1%	-0.4%	-0.7%	-1.1%	-1.5%
	30%	-0.3%	-0.7%	-1.5%	-2.3%	-3.0%
	45%	-0.4%	-1.1%	-2.3%	-3.4%	-4.5%
	60%	-0.6%	-1.5%	-3.0%	-4.5%	-6.0%
	75%	-0.7%	-1.9%	-3.8%	-5.6%	-7.5%

GOOGL total revenue share loss to Bing sensitivity analysis (\$M) - Current '24 total revenue estimate = \$333.6B						
		GOOGL % share loss to Bing				
		1.0%	2.5%	5.0%	7.5%	10.0%
% of Bing share gain as monetizeable searches	15%	\$333,290	\$332,873	\$332,179	\$331,484	\$330,790
	30%	\$333,012	\$332,179	\$330,790	\$329,401	\$328,012
	45%	\$332,735	\$331,484	\$329,401	\$327,317	\$325,234
	60%	\$332,457	\$330,790	\$328,012	\$325,234	\$322,456
	75%	\$332,179	\$330,095	\$326,623	\$323,150	\$319,677

GOOGL total % change in revenue sensitivity analysis						
		GOOGL % share loss to Bing				
		1.0%	2.5%	5.0%	7.5%	10.0%
% of Bing share gain as monetizeable searches	15%	-0.1%	-0.2%	-0.4%	-0.6%	-0.8%
	30%	-0.2%	-0.4%	-0.8%	-1.2%	-1.7%
	45%	-0.2%	-0.6%	-1.2%	-1.9%	-2.5%
	60%	-0.3%	-0.8%	-1.7%	-2.5%	-3.3%
	75%	-0.4%	-1.0%	-2.1%	-3.1%	-4.2%

Source: Company Reports, Factset, RBC Capital Markets

Variable 2 - Impact on GOOGL's search revenue based on GOOGL searches fulfilled by their own GAI chatbot

From a revenue perspective, if/when GOOGL implements Bard into Search, we estimate the revenue headwind associated with Bard taking search share from traditional Search will likely not be a significant revenue headwind. With the recent Bard announcement and Bing adding ChatGPT functionality directly in search interface, we assume GOOGL will introduce similar GAI chatbot functionality into search and estimate the impact on search revenue as GAI fulfills a portion of total searches. Additionally, GOOGL currently only monetizes ~20% of searches which introduces another variable of what percentage monetizeable search use cases can be fulfilled by GAI. We sensitize those two variables in the below tables and calculate the implied Search revenue % headwind. If GAI ends up fulfilling 15% of total searches but only 20% of those searches were monetizeable (likely will be lower than 20%, in our view), the total revenue headwind would be under 2%.



Figure 34 - We sensitize the potential impact on GOOGL search revenue due to GAI fulfilling a % of total searches and the % of GAI searches that were previously monetizeable

GOOGL search revenue sensitivity analysis (\$M) - Current '24 search revenue estimate						
		% of monetizeable search use cases fulfilled by GAI				
		10%	20%	30%	40%	50%
% of searches fulfilled by GAI	5.0%	\$184,281	\$183,355	\$182,429	\$181,503	\$180,577
	7.5%	\$183,818	\$182,429	\$181,040	\$179,651	\$178,262
	10.0%	\$183,355	\$181,503	\$179,651	\$177,799	\$175,947
	12.5%	\$182,892	\$180,577	\$178,262	\$175,947	\$173,632
	15.0%	\$182,429	\$179,651	\$176,873	\$174,095	\$171,317

GOOGL search % change in revenue sensitivity analysis						
		% of monetizeable search use cases fulfilled by GAI				
		10%	20%	30%	40%	50%
% of searches fulfilled by GAI	5.0%	-0.5%	-1.0%	-1.5%	-2.0%	-2.5%
	7.5%	-0.7%	-1.5%	-2.3%	-3.0%	-3.8%
	10.0%	-1.0%	-2.0%	-3.0%	-4.0%	-5.0%
	12.5%	-1.2%	-2.5%	-3.8%	-5.0%	-6.3%
	15.0%	-1.5%	-3.0%	-4.5%	-6.0%	-7.5%

GOOGL total revenue sensitivity analysis (\$M) - Current '24 search revenue estimate = \$333.6B						
		% of monetizeable search use cases fulfilled by GAI				
		10%	20%	30%	40%	50%
% of searches fulfilled by GAI	5.0%	\$332,642	\$331,716	\$330,790	\$329,864	\$328,938
	7.5%	\$332,179	\$330,790	\$329,401	\$328,012	\$326,623
	10.0%	\$331,716	\$329,864	\$328,012	\$326,160	\$324,308
	12.5%	\$331,253	\$328,938	\$326,623	\$324,308	\$321,993
	15.0%	\$330,790	\$328,012	\$325,234	\$322,456	\$319,677

GOOGL total % change in revenue sensitivity analysis						
		% of monetizeable search use cases fulfilled by GAI				
		10%	20%	30%	40%	50%
% of searches fulfilled by GAI	5.0%	-0.3%	-0.6%	-0.8%	-1.1%	-1.4%
	7.5%	-0.4%	-0.8%	-1.2%	-1.7%	-2.1%
	10.0%	-0.6%	-1.1%	-1.7%	-2.2%	-2.8%
	12.5%	-0.7%	-1.4%	-2.1%	-2.8%	-3.5%
	15.0%	-0.8%	-1.7%	-2.5%	-3.3%	-4.2%

Source: Company Reports, Factset, RBC Capital Markets


Variable 3 - GOOGL search GM impact due to compute intensity delta between GAI and traditional search

From a gross margin perspective, the incremental compute intensity of GAI fulfilled searches could be a significant headwind. Given the compute intensity delta between GAI and standard search queries, we estimate that the gross margin dilution could be significant depending on the relative compute intensity and % of total searches fulfilled by GAI which we sensitize in the below tables. We assume that Search's GM ex TAC is approximately 85% and calculate the GM on Search and the entire company. While a likely draconian scenario, if relative compute costs are 10x and only 10% of searches are fulfilled by GAI, gross margins would fall 13%.

It appears as though Street estimates are not contemplating incremental compute intensity related to GAI as GMs are expected to expand over the next several years. While it's likely fair to say GAI compute intensity will fall on an absolute basis with scale and GPU innovation, we think GAI compute intensity relative to legacy search is still significantly higher and the transition to GAI-fulfilled queries will likely happen faster than costs can fall, at least in the near-term, where gross margins could see degradation over the next couple of years.



Figure 35 - Gross margins have been declining relatively steadily over the past several years where Street estimates imply expanding margins over the next several years



Source: Company Reports, Factset, RBC Capital Markets

Figure 36 - We sensitize the potential impact on GOOGL search GM due to increased compute intensity associated with GAI searches relative to the % of total searches done by GOOGL's GAI and the relative compute intensity vs. traditional search

GOOGL search gross margin sensitivity analysis - assume 85% GM ex TAC						
GAI compute cost intensity vs. traditional search						
		2x	4x	6x	8x	10x
% of searches fulfilled by GAI	5.0%	84.3%	82.8%	81.3%	79.8%	78.3%
	7.5%	83.9%	81.6%	79.4%	77.1%	74.9%
	10.0%	83.5%	80.5%	77.5%	74.5%	71.5%
	12.5%	83.1%	79.4%	75.6%	71.9%	68.1%
	15.0%	82.8%	78.3%	73.8%	69.3%	64.8%

GOOGL % change in search gross margin ex. TAC sensitivity analysis						
GAI compute cost intensity vs. traditional search						
		2x	4x	6x	8x	10x
% of searches fulfilled by GAI	5.0%	-0.9%	-2.6%	-4.4%	-6.2%	-7.9%
	7.5%	-1.3%	-4.0%	-6.6%	-9.3%	-11.9%
	10.0%	-1.8%	-5.3%	-8.8%	-12.4%	-15.9%
	12.5%	-2.2%	-6.6%	-11.0%	-15.4%	-19.9%
	15.0%	-2.6%	-7.9%	-13.2%	-18.5%	-23.8%

GOOGL total GM sensitivity analysis - Street estimates imply 56%						
GAI compute cost intensity vs. traditional search						
		2x	4x	6x	8x	10x
% of searches fulfilled by GAI	5.0%	55.4%	54.5%	53.7%	52.9%	52.0%
	7.5%	55.2%	53.9%	52.7%	51.4%	50.2%
	10.0%	54.9%	53.3%	51.6%	49.9%	48.3%
	12.5%	54.7%	52.7%	50.6%	48.5%	46.4%
	15.0%	54.5%	52.0%	49.5%	47.0%	44.5%

GOOGL % change in total gross margin sensitivity analysis						
GAI compute cost intensity vs. traditional search						
		2x	4x	6x	8x	10x
% of searches fulfilled by GAI	5.0%	-0.7%	-2.2%	-3.7%	-5.2%	-6.7%
	7.5%	-1.1%	-3.4%	-5.6%	-7.8%	-10.1%
	10.0%	-1.5%	-4.5%	-7.5%	-10.5%	-13.4%
	12.5%	-1.9%	-5.6%	-9.3%	-13.1%	-16.8%
	15.0%	-2.2%	-6.7%	-11.2%	-15.7%	-20.2%

Source: Company Reports, Factset, RBC Capital Markets

Impact of all three scenarios on the full P&L

Putting all three scenarios together, we estimate bull, base and bear case scenarios to size the impact to the Street's implied EPS estimate. While the combined revenue headwind from Scenarios 1 & 2 may not be significant even in a bear case scenario, the potential impact on EPS is meaningful, in our view. Note: in the below tables, green borders denote the variables driving our EPS downside estimates.



Figure 37 - In a bull case, we estimate only a marginal 4% downside to Street's implied '24 EPS estimate for GOOGL

GAI sensitivity analysis to '24 Street revenue and gross profit estimates		Impact of GAI on Street's '24 full PnL			
Line Item	Value	In M except for EPS	Street estimates	Adjusted estimate	Delta
Total revenue	\$333,568.0	Total revenue	\$333,568.0	\$332,388.1	-0.4%
Search revenue	\$185,207.4	CoR	\$147,506.8	\$149,568.0	1.4%
Total CoR	\$147,506.8	Total gross profit	\$186,061.2	\$182,820.1	-1.7%
Search CoR	\$27,781.11	S&M	\$30,497.7	\$30,497.7	
Total gross profit	\$186,061.2	G&A	\$17,336.9	\$17,336.9	
Search gross profit	\$157,426.3	R&D	\$45,702.6	\$45,702.6	
Search GM (ex TAC)	85.0%	EBIT	\$92,524.1	\$89,283.0	-3.5%
Total GM	55.8%	Interest expense	\$1,600.0	\$1,600.0	
		Tax expense	\$15,755.0	\$15,193.4	
		Tax rate	17%	17%	
		Net Income	\$75,169.0	\$72,489.5	-3.6%
		Diluted shares outstanding (4Q22)	12,947	12,947	
		EPS	\$5.81	\$5.60	-3.6%

Key Variables	Value
% share loss to Bing	1.0%
% of searches Bing gains are monetizeable	7.5%
% of searches fulfilled by GOOGL GAI	7.5%
% of monetizeable search use cases fulfilled by GOOGL GAI	7.5%
GAI compute cost intensity vs. traditional search	2.0x

Scenario 1 - Share loss going to Bing

% share loss to Bing	1.0%
% of searches Bing gains are monetizeable	7.5%
Search revenue % headwind	0.1%
Search revenue post share loss to Bing	\$185,068.5
Total revenue % headwind	0.0%
Total revenue post share loss to Bing	\$333,429.1

Scenario 2 - GAI fulfills a % of GOOGL searches

% of searches fulfilled by GOOGL GAI	7.5%
% of monetizeable search use cases fulfilled by GOOGL GAI	7.5%
Search revenue % headwind	0.6%
Search revenue post shift to GOOGL GAI	\$184,027.5
Total revenue % headwind	0.3%
Total revenue post shift to GOOGL GAI	\$332,388.1

Scenario 3 - Compute costs on CoR

Search CoR post share loss to Bing	\$27,760.3
Total CoR post share loss to Bing	\$147,485.9
% of searches fulfilled by GOOGL GAI	7.5%
GAI compute cost intensity vs. traditional search	2.0x
Increase in compute costs impact on search CoR	8%
Search CoR including impact of GAI higher compute intensity	\$29,842.3
Search GM	83.9%
Search GM bps impact	-113 bps
Total % increase in CoR	1.4%
Total \$ CoR	\$149,568.0
Total GM	55.1%
Total GM bps impact	-64 bps

Source: Company Reports, Factset, RBC Capital Markets



Figure 38 - In a more base case scenario, our assumptions imply 13% downside to Street's implied '24 EPS estimate for GOOGL

GAI sensitivity analysis to '24 Street revenue and gross profit estimates		Impact of GAI on Street's '24 full PnL			
Line Item	Value	In M except for EPS	Street estimates	Adjusted estimate	Delta
Total revenue	\$333,568.0	Total revenue	\$333,568.0	\$329,421.7	-1.2%
Search revenue	\$185,207.4	CoR	\$147,506.8	\$155,570.2	5.5%
Total CoR	\$147,506.8	Total gross profit	\$186,061.2	\$173,851.4	-6.6%
Search CoR	\$27,781.11	S&M	\$30,497.7	\$30,497.7	
Total gross profit	\$186,061.2	G&A	\$17,336.9	\$17,336.9	
Search gross profit	\$157,426.3	R&D	\$45,702.6	\$45,702.6	
Search GM (ex TAC)	85.0%	EBIT	\$92,524.1	\$80,314.3	-13.2%
Total GM	55.8%	Interest expense	\$1,600.0	\$1,600.0	
		Tax expense	\$15,755.0	\$13,639.4	
		Tax rate	17%	17%	
		Net Income	\$75,169.0	\$65,074.9	-13.4%
		Diluted shares outstanding (4Q22)	12,947	12,947	
		EPS	\$5.81	\$5.03	-13.4%

Key Variables	Value
% share loss to Bing	5.0%
% of searches Bing gains are monetizeable	15.0%
% of searches fulfilled by GOOGL GAI	10.0%
% of monetizeable search use cases fulfilled by GOOGL GAI	15.0%
GAI compute cost intensity vs. traditional search	4.0x

Scenario 1 - Share loss going to Bing

% share loss to Bing	5.0%
% of searches Bing gains are monetizeable	15.0%
Search revenue % headwind	0.8%
Search revenue post share loss to Bing	\$183,818.4
Total revenue % headwind	0.4%
Total revenue post share loss to Bing	\$332,178.9

Scenario 2 - GAI fulfills a % of GOOGL searches

% of searches fulfilled by GOOGL GAI	10.0%
% of monetizeable search use cases fulfilled by GOOGL GAI	15.0%
Search revenue % headwind	1.5%
Search revenue post shift to GOOGL GAI	\$181,061.1
Total revenue % headwind	0.8%
Total revenue post shift to GOOGL GAI	\$329,421.7

Scenario 3 - Compute costs on CoR

Search CoR post share loss to Bing	\$27,572.8
Total CoR post share loss to Bing	\$147,298.4
% of searches fulfilled by GOOGL GAI	10.0%
GAI compute cost intensity vs. traditional search	4.0x
Increase in compute costs impact on search CoR	30%
Search CoR including impact of GAI higher compute intensity	\$35,844.6
Search GM	80.5%
Search GM bps impact	-450 bps
Total % increase in CoR	5.6%
Total \$ CoR	\$155,570.2
Total GM	53.2%
Total GM bps impact	-261 bps

Source: Company Reports, Factset, RBC Capital Markets



Figure 39 - In a bear case scenario, our assumptions imply a 26% downside to Street's implied '24 EPS estimate for GOOGL

GAI sensitivity analysis to '24 Street revenue and gross profit estimates		Impact of GAI on Street's '24 full PnL			
Line Item	Value	In M except for EPS	Street estimates	Adjusted estimate	Delta
Total revenue	\$333,568.0	Total revenue	\$333,568.0	\$326,229.2	-2.2%
Search revenue	\$185,207.4	CoR	\$147,506.8	\$164,192.8	11.3%
Total CoR	\$147,506.8	Total gross profit	\$186,061.2	\$162,036.3	-12.9%
Search CoR	\$27,781.11	S&M	\$30,497.7	\$30,497.7	
Total gross profit	\$186,061.2	G&A	\$17,336.9	\$17,336.9	
Search gross profit	\$157,426.3	R&D	\$45,702.6	\$45,702.6	
Search GM (ex TAC)	85.0%	EBIT	\$92,524.1	\$68,499.2	-26.0%
Total GM	55.8%	Interest expense	\$1,600.0	\$1,600.0	
Key Variables		Tax expense	\$15,755.0	\$11,592.1	
% share loss to Bing	7.5%	Tax rate	17%	17%	
% of searches Bing gains are monetizeable	20.0%	Net Income	\$75,169.0	\$55,307.1	-26.4%
% of searches fulfilled by GOOGL GAI	12.5%	Diluted shares outstanding (4Q22)	12,947	12,947	
% of monetizeable search use cases fulfilled by GOOGL GAI	20.0%	EPS	\$5.81	\$4.27	-26.4%
GAI compute cost intensity vs. traditional search	6.0x				
Scenario 1 - Share loss going to Bing					
% share loss to Bing	7.5%				
% of searches Bing gains are monetizeable	20.0%				
Search revenue % headwind	1.5%				
Search revenue post share loss to Bing	\$182,429.3				
Total revenue % headwind	0.8%				
Total revenue post share loss to Bing	\$330,789.9				
Scenario 2 - GAI fulfills a % of GOOGL searches					
% of searches fulfilled by GOOGL GAI	12.5%				
% of monetizeable search use cases fulfilled by GOOGL GAI	20.0%				
Search revenue % headwind	2.5%				
Search revenue post shift to GOOGL GAI	\$177,868.6				
Total revenue % headwind	1.4%				
Total revenue post shift to GOOGL GAI	\$326,229.2				
Scenario 3 - Compute costs on CoR					
Search CoR post share loss to Bing	\$27,364.4				
Total CoR post share loss to Bing	\$147,090.1				
% of searches fulfilled by GOOGL GAI	12.5%				
GAI compute cost intensity vs. traditional search	6.0x				
Increase in compute costs impact on search CoR	63%				
Search CoR including impact of GAI higher compute intensity	\$44,467.1				
Search GM	75.6%				
Search GM bps impact	-938 bps				
Total % increase in CoR	11.6%				
Total \$ CoR	\$164,192.8				
Total GM	50.4%				
Total GM bps impact	-542 bps				

Source: Company Reports, Factset, RBC Capital Markets



Given the model's sensitivity to relative compute costs of GAI fulfilled queries, we show the below data tables which calculate the EPS impact of GAI headwinds at different relative compute costs ranging from 2x-10x vs. traditional search. The two variables we sensitize are share loss to Bing and % of searches fulfilled by GOOGL's own GAI while holding both % of searches Bing gains that are monetizeable and % of monetizeable search use cases fulfilled by GOOGL GAI constant at 15%. In a best case scenario, EPS would only fall by 3%, where in a worst case EPS falls nearly 48%.

Figure 40 - Given the incremental compute costs associated with GAI queries are still a debate amongst investors, we show several sensitivity tables, each based on a different compute intensity level, to calculate the impact on the Street's '24 EPS estimates for GOOGL

Impact on EPS assuming 2x compute cost intensity for GAI fulfilled searches						
		GOOGL % share loss to Bing				
		1.0%	2.5%	5.0%	7.5%	10.0%
% of searches fulfilled by GAI	5.0%	\$5.61	\$5.59	\$5.55	\$5.52	\$5.48
	7.5%	\$5.53	\$5.50	\$5.47	\$5.43	\$5.39
	10.0%	\$5.44	\$5.41	\$5.38	\$5.34	\$5.31
	12.5%	\$5.35	\$5.33	\$5.29	\$5.25	\$5.22
	15.0%	\$5.26	\$5.24	\$5.20	\$5.17	\$5.13

EPS % impact vs. current '24 consensus at 2x compute costs						
		GOOGL % share loss to Bing				
		1.0%	2.5%	5.0%	7.5%	10.0%
% of searches fulfilled by GAI	5.0%	-3.3%	-3.7%	-4.3%	-5.0%	-5.6%
	7.5%	-4.8%	-5.2%	-5.8%	-6.5%	-7.1%
	10.0%	-6.4%	-6.7%	-7.4%	-8.0%	-8.6%
	12.5%	-7.9%	-8.3%	-8.9%	-9.5%	-10.1%
	15.0%	-9.4%	-9.8%	-10.4%	-11.0%	-11.6%

Impact on EPS assuming 4x compute cost intensity for GAI fulfilled searches						
		GOOGL % share loss to Bing				
		1.0%	2.5%	5.0%	7.5%	10.0%
% of searches fulfilled by GAI	5.0%	\$5.44	\$5.41	\$5.38	\$5.34	\$5.31
	7.5%	\$5.26	\$5.24	\$5.20	\$5.17	\$5.13
	10.0%	\$5.08	\$5.06	\$5.03	\$4.99	\$4.96
	12.5%	\$4.91	\$4.88	\$4.85	\$4.82	\$4.78
	15.0%	\$4.73	\$4.71	\$4.67	\$4.64	\$4.61

EPS % impact vs. current '24 consensus at 4x compute costs						
		GOOGL % share loss to Bing				
		1.0%	2.5%	5.0%	7.5%	10.0%
% of searches fulfilled by GAI	5.0%	-6.4%	-6.7%	-7.4%	-8.0%	-8.6%
	7.5%	-9.4%	-9.8%	-10.4%	-11.0%	-11.6%
	10.0%	-12.5%	-12.8%	-13.4%	-14.0%	-14.6%
	12.5%	-15.5%	-15.9%	-16.5%	-17.1%	-17.6%
	15.0%	-18.6%	-18.9%	-19.5%	-20.1%	-20.7%

Impact on EPS assuming 6x compute cost intensity for GAI fulfilled searches						
		GOOGL % share loss to Bing				
		1.0%	2.5%	5.0%	7.5%	10.0%
% of searches fulfilled by GAI	5.0%	\$5.26	\$5.24	\$5.20	\$5.17	\$5.13
	7.5%	\$4.99	\$4.97	\$4.94	\$4.90	\$4.87
	10.0%	\$4.73	\$4.71	\$4.67	\$4.64	\$4.61
	12.5%	\$4.46	\$4.44	\$4.41	\$4.38	\$4.34
	15.0%	\$4.20	\$4.18	\$4.15	\$4.11	\$4.08

EPS % impact vs. current '24 consensus at 6x compute costs						
		GOOGL % share loss to Bing				
		1.0%	2.5%	5.0%	7.5%	10.0%
% of searches fulfilled by GAI	5.0%	-9.4%	-9.8%	-10.4%	-11.0%	-11.6%
	7.5%	-14.0%	-14.3%	-14.9%	-15.5%	-16.1%
	10.0%	-18.6%	-18.9%	-19.5%	-20.1%	-20.7%
	12.5%	-23.1%	-23.5%	-24.0%	-24.6%	-25.2%
	15.0%	-27.7%	-28.0%	-28.6%	-29.1%	-29.7%

Impact on EPS assuming 8x compute cost intensity for GAI fulfilled searches						
		GOOGL % share loss to Bing				
		1.0%	2.5%	5.0%	7.5%	10.0%
% of searches fulfilled by GAI	5.0%	\$5.08	\$5.06	\$5.03	\$4.99	\$4.96
	7.5%	\$4.73	\$4.71	\$4.67	\$4.64	\$4.61
	10.0%	\$4.37	\$4.35	\$4.32	\$4.29	\$4.26
	12.5%	\$4.02	\$4.00	\$3.97	\$3.94	\$3.91
	15.0%	\$3.67	\$3.65	\$3.62	\$3.59	\$3.56

EPS % impact vs. current '24 consensus at 8x compute costs						
		GOOGL % share loss to Bing				
		1.0%	2.5%	5.0%	7.5%	10.0%
% of searches fulfilled by GAI	5.0%	-12.5%	-12.8%	-13.4%	-14.0%	-14.6%
	7.5%	-18.6%	-18.9%	-19.5%	-20.1%	-20.7%
	10.0%	-24.7%	-25.0%	-25.6%	-26.1%	-26.7%
	12.5%	-30.8%	-31.1%	-31.6%	-32.2%	-32.7%
	15.0%	-36.9%	-37.2%	-37.7%	-38.2%	-38.7%

Impact on EPS assuming 10x compute cost intensity for GAI fulfilled searches						
		GOOGL % share loss to Bing				
		1.0%	2.5%	5.0%	7.5%	10.0%
% of searches fulfilled by GAI	5.0%	\$4.91	\$4.88	\$4.85	\$4.82	\$4.78
	7.5%	\$4.46	\$4.44	\$4.41	\$4.38	\$4.34
	10.0%	\$4.02	\$4.00	\$3.97	\$3.94	\$3.91
	12.5%	\$3.58	\$3.56	\$3.53	\$3.50	\$3.47
	15.0%	\$3.13	\$3.12	\$3.09	\$3.06	\$3.03

EPS % impact vs. current '24 consensus at 10x compute costs						
		GOOGL % share loss to Bing				
		1.0%	2.5%	5.0%	7.5%	10.0%
% of searches fulfilled by GAI	5.0%	-15.5%	-15.9%	-16.5%	-17.1%	-17.6%
	7.5%	-23.1%	-23.5%	-24.0%	-24.6%	-25.2%
	10.0%	-30.8%	-31.1%	-31.6%	-32.2%	-32.7%
	12.5%	-38.4%	-38.7%	-39.2%	-39.7%	-40.2%
	15.0%	-46.0%	-46.3%	-46.8%	-47.3%	-47.7%


Source: Company Reports, Factset, RBC Capital Markets. Note: we assume % of searches Bing gains that are monetizeable and % of monetizeable search use cases fulfilled by GOOGL GAI are both 15%



Our base case implies ~13% downside to '24 EPS

Leveraging our base case above which assumes Google loses 5% market share to Bing and compute costs for GAI runs 4x vs. traditional searches implies an estimate ~13% downside to '24 EPS. Given the stock's valuation is trading near trough levels over the past decade and 30% below the mean/median, it would appear the market would have to bake in more draconian assumptions to warrant more than ~10% further downside, in our view.

Figure 41 - GOOGL is trading near historical trough P/E multiple and 30% below historical average



Source: Factset, RBC Capital Markets

Figure 42 - GOOGL is trading near trough P/E levels

	Value	Current multiple delta
Current multiple	14.9x	N/A
Peak	29.4x	-49.4%
Trough	14.6x	1.5%
Mean	21.2x	-30.1%
Median	21.3x	-30.4%

Source: Factset, RBC Capital Markets



Implications for software companies

As we think across the broader software landscape, we see several puts and takes behind the impact of generative AI and ChatGPT on enterprise software companies. To a certain extent, we have to separate the impact of ChatGPT into the back end (how companies build software) and the front end (how companies utilize ChatGPT in their products). In other words, how does ChatGPT benefit the company and how does it benefit the actual software?

On the back end, ChatGPT can speed up development of new products and features by helping write and debug code and help companies deal with the shortage of talented developers. To a certain extent, this can even be a margin tailwind, reducing the amount of R&D spend necessary, but we would prefer that companies reinvest those savings in further innovation. These companies can respond to customer needs faster and the time between customers demanding new features and bringing those features to market could decrease.

We expect this trend to be a tailwind for innovative challengers – for example, we believe this can help HubSpot narrow the gap versus Salesforce on enterprise-grade functionality in Sales Hub (while also acknowledging that Salesforce has a major data advantage) and a headwind for complacent market leaders that already underinvest in innovation. We would also highlight some impressive use cases we've seen with OpenAI integrations with private companies, including Coda, where customers can use OpenAI to create personalized emails for sales prospects and tweet ideas, and Notion, where customers can use OpenAI to draft up blog posts.

We also see a potential tailwind for market-leading vertical software companies (e.g. Guidewire, Veeva) to leverage the data from their market leadership position to build more differentiated and customized products, which would increase the separation between them and their competitors. In addition, this could help reduce the services attach rate, as customization could become easier with generative AI.

On the front end, broadly speaking, we believe generative AI will be a net positive for enterprise software companies that embrace the technology. This technology can make enterprise software applications easier to use and navigate, more bespoke (since the experience can be easily customized to the specific needs of the enterprise), and “smarter”, in that the software will adapt over time. For example, we can imagine generative AI underlying CRM marketing campaigns and helping marketers iterate so each campaign is more targeted and effective than the last. Or ChatGPT making HR software easier to navigate so that all employees can more easily log their time cards and find resources for specific questions (instead of a canned FAQ section), generating meaningful time and cost savings.

On the other hand, we feel generative AI will be a net negative for two categories of enterprise software.

The first is those companies that don't quickly embrace generative AI, as it becomes a competitive disadvantage against those companies that do. If two companies are selling competing data warehousing solutions and one embeds ChatGPT, making it easier for users to find and analyze specific data sets, while the other doesn't, the former is more likely to win out. We arguably may see this when comparing Microsoft's GitHub, which has GitHub Copilot, making it easier for developers to write and debug code, versus GitLab, which currently has no such capabilities (although we can imagine GitLab embracing generative AI sooner than later to narrow that gap).

The second is companies whose entire value proposition is selling solutions that offer AI-like capabilities or where the “secret sauce” is their own proprietary AI, especially if these companies are benefiting from enterprises with AI mandates. We believe generative AI has



leapfrogged many of these technologies, such as offerings from Palantir, and Alteryx. We would have particular caution on Palantir, as those technologies require so much manual customization (and Palantir collects meaningful revenue from professional services, per our estimates) and ChatGPT can make that level of manual customization unnecessary. In addition, solutions utilizing generative AI are more likely to gain share of budget associated with AI mandates versus companies selling AI-like solutions.



Software company-specific impacts

In this section, we weigh the tailwinds and headwinds for select companies in our software coverage universe.

Adobe (ADBE) — Covered by Matt Swanson

Adobe's current product set: Adobe's most recent Max event demonstrated its advancements within the generative AI space with both AI and ML being major parts of the creative cloud roadmap. On its first quarter call, Adobe highlighted generative AI as a focal point of its Summit user event starting 3/20 noting that they are excited to unveil a differentiated product set with a "creative co-pilot". Previous demonstrations of utilizing AI within photo shop focuses both on generative capabilities but also enhancing the variety of tools utilized to edit images, enhancing the potential of creativity not replacing it. In other areas like Adobe Express, AI helps users with less experience move beyond pre-made templates to create something more unique to them. Express is another area where we expect a product announcement from Summit where AI and generative AI could be major components. While currently the company has focused on generative AI within images, investments are currently being made for video and 3D design. Within video, Sensei already has the ability to help edit video, with Substance 3D bringing AI to the 3D design space. While Adobe should benefit on a standalone basis, being a central hub of enterprise creativity should also create an opportunity for Adobe to partner with point solutions that will need work flows built around their tools for use in the enterprise. We feel Adobe should benefit from an expanded partner eco-system around generative.

Tailwinds:

- **Creativity for the people:** Expanding the total addressable market for creativity has been a focal point for Adobe as well as newer entrants to the market including Canva and Figma. The idea of democratizing creativity has historically come from simplifying user interface, growing detail and prevalence of templates, more drag and drop features and of course entry level price points. Generative AI is the next iteration of opening up the creative market, building of products like Adobe Express that make basic tools accessible to consumers to democratizing visually professional content.
- **Market sizing:** At the company's 2022 analyst day, Adobe sized its market as 68 million creative professionals, 900 million communicators and 4 billion consumers. The ability to expand this market over time will depend on expanding these groups but also bringing current users up the curve. Within enterprises, this is a license expansion opportunity. Much in the way that Tableau opened up the business intelligence market to knowledge workers, the thought process was those who have the questions should be able to find the answers. We feel in creative, those who have the ideas, should be able to create the content.
- **Augmented intelligence:** The new age, augmented intelligence is the idea of meshing the strengths of humans and machines to see compounding value. Within creativity, the idea is not to replace creative professionals but to emphasize them. Making sure the work they are doing is worthy of their skill set, allowing AI to help with drafts not just for their own work but to be able to receive drafts from non-creative professionals reducing reworks. Generative AI editing will speed the iteration process and allow workers to create more content and more expansive content.

Potential headwinds:

- **Image issues:** Presenting generative AI within creative has come with some challenges, attending user events Adobe customers generally seem excited, with products demos for things like Sensei getting large ovations but also apprehension. A tale as old as technology, the fear of being replaced by robots is real, personally we would not be overly excited about an



AI industry around writing equity research. Working with and not against change, disrupting to not be disrupted, is not new and concerns will pass. Much as low-code, no-code has not reduced the need for developers, data scientists lived through the democratization of business intelligence and security architects now couldn't live without AI for threat detection. The digital environment continues to expand at faster rate than the workforce with shortages creating a choke point in digital transformation. Concerns around alienating the creative community we feel will ultimately be short-lived. Adobe has also taken proactive steps working as a leader of the Content Authenticity Initiative or CAI which has standards to help creators maintain control over their work.

- **Market impact:** The impact on the creative market could be massive, in our view. As we mentioned above, we don't see this as a disruptive force to displace creative professionals, but it should increase the efficiency of creatives allowing for work to be done faster with fewer iterations. Where we do see changes are on the low-end of the market. In the consumer segment, generative AI opens up a new user base of non-creatives to do more on their own with the ability to create professional quality images. We feel this could fragment the low-end of the market with easier entry for application based solutions to the consumer market which may have a high bar for quality but lower for specificity. For example, using DALL-E, we tried to make a cover for this report. With a prompt of "Royal Bank of Canada robot lion writing a research report" this was the result. While the result looks professional, particularly given the speed from ideation to the result, it is not professional in the sense that it does not match our company brand or image. This would serve as a good jumping off point if we were trying to explain to someone what we wanted the cover to look like, but at the same time, it is not fine-tuned to meet specific parameters particularly trying to reflect an existing brand. In Adobe's view, unlocking this creativity in non-creative workers could end up as a top-of-funnel marketing tool.

Figure 43 - DALL-E Generative Image



Source: DALL-E, RBC Capital Markets

Net/net: Adobe is well positioned to be a beneficiary of generative AI within the digital media landscape, but more importantly we believe the investments will de-risk and future proof the model against disruption from players within the space. The greatest areas of risk come from the potential fragmentation of the consumer market, which could also work as a top-of-funnel entry point to broader platforms like Adobe, while the greatest areas of upside are likely to come from expanding the total addressable market for creativity including expanded enterprise seat counts, as well as premium tools to deal with the challenges of more expansive digital environments, such as the metaverse, increasing the efficiency of limited creative professionals.

Autodesk (ADSK) — Covered by Matt Hedberg

Tailwinds: We think Autodesk is in a good position to leverage AI capabilities, initially through generative design capabilities, which is used to provide engineers the ability to quickly explore, optimize, and make informed decisions to complex design problems. To an engineer, generative design software is like an assistant that helps with creating, testing, and evaluating design options. However, we found it interesting that Microsoft and Satya Nadella demoed a natural language processing tool for generating Maya scripts for both direct-to-Maya and for Bifrost.



Over time, we think doing real-time training on customer data is where the real innovation and tailwinds to Autodesk is likely to emerge.

Headwinds: We don't see any real headwinds to generative AI for Autodesk at the moment.

Bottom line: With a growing portfolio of cloud-based design tools, we believe Autodesk is in a good position to leverage large customer data sets and generative AI models to give users a better design experience.

Akamai (AKAM) — Covered by Rishi Jaluria

Tailwinds: Generative AI does pose increasing security risks, which could be a tailwind for Akamai's security business. We could also see a potential tailwind for the delivery business, although we do not believe OpenAI is a meaningful Akamai customer today. We would also note one of OpenAI's co-founders, Jack Hughes, was also a co-founder of Akamai. Finally, we could see OpenAI (and other competitors) utilize Akamai for delivery and performance acceleration as the volume of traffic grows exponentially (we note Microsoft will be OpenAI's exclusive cloud provider going forward and Akamai is one of Azure CDN's partners).

Headwinds: We note OpenAI utilizes Cloudflare, an Akamai competitor, as a security vendor today. If Cloudflare meaningfully benefits, it could worsen Akamai's competitive positioning against Cloudflare.

Bottom line: We don't see many meaningful tailwinds or headwinds for Akamai from ChatGPT.

Asana (ASAN) — Covered by Rishi Jaluria

Tailwinds: As generative AI becomes more embedded in business processes, we believe there will be a greater focus on collaboration, as human-to-human collaboration becomes a more differentiated component of work (while siloed tasks like data entry get outsourced to ChatGPT). This could serve as a tailwind for Asana, which helps manage collaborative work and projects.

Headwinds: We expect ChatGPT to be able to replace some of the basic task management capabilities of Asana. Our due diligence suggests there are paid users utilizing Asana for basic assigning and following up of tasks and if that use case can be automated and replaced by ChatGPT, this could be a meaningful headwind.

Bottom line: We see generative AI as being neither a strong tailwind, nor a strong headwind for Asana, as we believe the headwind to basic use cases can be balanced out by tailwinds towards adoption of collaboration software.

Couchbase (BASE) — Covered by Matt Hedberg

Tailwinds: We think generative AI could ultimately become a bit of a tailwind to Couchbase. The company has spent a lot of time thinking about AI and ways that it could be useful to help developers expedite the application build process. More broadly, efforts continue around leveraging AI and ML as a core part of the value proposition. Additionally, we believe there could be a tailwind around AI as far as helping customers store massive amounts of data to train AI and ML models built on the Couchbase platform.

Headwinds: The headwind could be that Couchbase is small relative to other database peers and could easily be outspent regarding generative AI.

Bottom line: Overall we see generative AI as more neutral to Couchbase today, but over time, it could become a tailwind.

Box (BOX) — Covered by Rishi Jaluria

Tailwinds: Generative AI could lead to exponential growth in the volume of content (including



images and videos), serving as a tailwind for Box to serve as a system of record or repository for all this content. In addition, given the vast volumes of data Box has, the company could figure out how to leverage ChatGPT as a monetization vector (e.g. deeper analytics on all the user's content).

Headwinds: We don't see any major headwinds for Box in the near-term from generative AI.

Bottom line: We see generative AI as generating slight tailwinds for Box.

Coursera (COUR) — Covered by Rishi Jaluria

Tailwinds: With generative AI becoming a new skill for people to learn, this can serve as a tailwind for the Consumer business, as learners take courses around generative AI and ChatGPT (we can imagine OpenAI partnering for ChatGPT certifications). In addition, if there is, in fact, meaningful disruption to many careers as a result, this could serve as a tailwind for all three businesses (Consumer as people look to upskill/reskill, Degrees as people look to make a pivot in their careers with graduate degrees, and Enterprise, especially on the government side, as governments look to reskill/upskill their citizenry). In addition, we can see Coursera leverage ChatGPT technology in its courses to help provide more "hands-on" education to learners.

Headwinds: ChatGPT can itself be used as an educational tool and lead to commoditization for some of Coursera's more basic offerings. For example, instead of a learner taking a course in financial modeling on Coursera, ChatGPT could serve as the instructor and teach how to build a DCF model in an interactive manner.

Bottom line: While ChatGPT can disrupt learning for single course offerings, we note Coursera offers a greater breadth of materials than others in the edtech space and, importantly, the Consumer business comes with professional certifications (which can be used on resumes or LinkedIn profiles) that ChatGPT cannot offer. Ultimately, while we understand investor concerns around the impact of ChatGPT on edtech – ChatGPT certainly makes it harder for edtech companies focused on helping students cheat on exams – we see more tailwinds than headwinds and view Coursera as a net beneficiary of ChatGPT.

Salesforce (CRM) — Covered by Rishi Jaluria

Tailwinds: For a business as diversified as Salesforce, we have to think of a few different angles. From an overarching perspective, however, we believe Salesforce has vast volumes of data over time, which can be a major competitive advantage, if utilized properly.

Within the core sales, service, and marketing offerings, we see a number of tailwinds that can make Salesforce's products better. Across these portfolios, ChatGPT can reduce the amount of customization required and even reduce the need for Salesforce Administrators. This can help reduce the TCO for Salesforce solutions, which remains one of the biggest sources of pushback, based on our due diligence.

Service Cloud is likely the biggest beneficiary, as integrating ChatGPT makes the product more effective for end users and reduces the reliance on human intervention to tackle customer service issues. For example, if customers submit support tickets through Service Cloud, ChatGPT can bring immediate resolution to the problem based on customer data and behavior, without any human intervention required. In addition, allowing Service Cloud customers to offer a white labeled, integrated ChatGPT would allow real-time interactions before having to engage with a contact center, resulting in major cost savings.

Sales Cloud can also benefit from ChatGPT making the CRM smarter. This can be used to help users more intelligently target sales prospects with higher propensity to spend and analyze behavior of salespeople to help determine the most appropriate channel, depending on the



type of customers they're going after. In addition, ChatGPT could be used to verify contacts in the database and ensure there's less stale contact information (we note the public version of ChatGPT refuses to find contact information for specific people, even business email addresses). We wouldn't be surprised if some Sales Cloud users try to use ChatGPT to automate the outreach process and draft emails to prospects, although we could caution that, as intelligent as ChatGPT is, its content (so far) can be distinguished from human-generated content and recipients of automated emails that pretend to be human-generated may push back.

Marketing Cloud can see benefits from helping users create more intelligent marketing campaigns and better tweak language to drive better outcomes. In addition, AI can go back and analyze prior marketing campaigns and suggest improvements for future campaigns.

Even in recently-acquired parts of the business, we see tailwinds. Adding AI on top of Tableau (along with all the new data being generated by generative AI) can make the product and its analytics more effective. Even Slack can benefit from ChatGPT bringing about better search capabilities (Slack's weak spot, based on our conversations, has been its weak search – ironic for a company whose name stands for Searchable Library of All Conversation and Knowledge).

We note that Salesforce has already made some major announcements around generative AI, which we view rather positively. The first is Einstein GPT to create personalized content across Salesforce clouds, including Sales, Service, and Marketing Clouds (which incorporates data from Salesforce Data Cloud). The second is a ChatGPT app for Slack for conversation summaries, research tools, and writing assistance. The third is Salesforce Ventures launching a \$250MM Generative AI Fund. While Einstein GPT is currently in closed pilot, we are encouraged to see Salesforce moving quickly.

Headwinds: On the flip side, we do see numerous potential headwinds throughout Salesforce's portfolio. From an overarching company perspective, we do note Salesforce is a direct competitor with Microsoft and has historically had a somewhat contentious relationship with the company, which only got worse with the purchase of Slack, even though the companies try to play nice publicly (witness Microsoft CEO Satya Nadella on stage with Salesforce CEO Marc Benioff at Dreamforce 2015). Salesforce embracing OpenAI and ChatGPT may be a tough pill for management to swallow, as it directly benefits Microsoft, which is seen as a meaningful competitor.

In addition, as Microsoft more rapidly embraces OpenAI and integrates the technology throughout its suite of products, Dynamics could become a more significant threat to Salesforce on the core CRM side of the equation. We expect this impact to be even worse with Slack. Microsoft Teams has already been a major competitive threat to Slack and is already embedding ChatGPT in Teams Premium, which may increase the competitive advantages versus Slack and further hurt that business.

Finally, while Salesforce does have the advantage of its vast amount of data, we also believe that competitors utilizing ChatGPT for product development can help narrow the gap in features and functionality. For example, we've spoken with users who find HubSpot's Marketing Hub superior to Salesforce's Marketing Cloud, but stay with Salesforce because Marketing Hub is lacking certain enterprise-grade features. ChatGPT can allow HubSpot to narrow that gap and make its product suite more competitive against Salesforce at the mid-market and even the enterprise level sooner than expected.

Bottom line: It's difficult to make a call right now, given the diversity of Salesforce's portfolio and the competitive dynamics with Microsoft. As things stand right now, we see a number of tailwinds and are encouraged by Salesforce launching offerings around generative AI so quickly. On the flip side, we do believe the competitive environment could get tougher from both



Microsoft and HubSpot. Putting this all together, we think Salesforce will ultimately see more tailwinds than headwinds from generative AI.

CrowdStrike (CRWD) — Covered by Matt Hedberg

Tailwinds: We think generative AI will give hackers new ways to write and deliver malicious code. As such, we think there are several potential tailwinds AI can provide to CrowdStrike to help customers prevent AI-based attacks. CrowdStrike was built with AI and ML at its core as their models are trained on 2 trillion data points collected every day by the CrowdStrike Security Cloud. What this translates to is that the Falcon platform achieves 100% ransomware prevention with zero false positives. Recently the company launched ExpRT.AI for Falcon Spotlight. ExpRT.AI or the Expert Prediction Rating artificial intelligence model leverage vulnerability and threat-based telemetry provide a dynamic, responsive ExpRT Rating within Falcon Spotlight’s console which has shown a 3x improvement in proactive vulnerability prioritization over other methods. With the importance of SecOps, the company also introduced AI-powered Indicators of Attack (IoAs), a new innovation for fileless attack prevention. In the future, we believe the company is in a good position to leverage its large and growing data lake to stay ahead of hackers that also are in a position to leverage generative AI.

Headwinds: We don’t see any immediate headwinds to CrowdStrike from generative AI.

Bottom line: We see AI as a positive for CrowdStrike given all their customer data is cloud-based and not on premise. Ultimately we believe AI will work best with large data sets, something CrowdStrike is in a good position to leverage.

Clearwater Analytics (CWAN) — Covered by Rishi Jaluria

Tailwinds: One of Clearwater’s competitive advantages is the amount of data it has, especially related to the pricing of assets. Clearwater could integrate that data into ChatGPT and use it to further its competitive advantage. In addition, Clearwater could leverage generative AI to provide faster insights, improving its value proposition.

Headwinds: ChatGPT may potentially put pressure on the investment management industry (in ways that so-called “robo advisors” like Betterment and Wealthfront could not), creating fee pressure and limiting budget for Clearwater’s solutions.

Bottom line: We see a slight benefit for Clearwater from generative AI.

Dropbox (DBX) — Covered by Rishi Jaluria

Tailwinds: Generative AI could lead to exponential growth in the volume of content (including images and videos), serving as a tailwind for Dropbox to serve as a system of record or repository for all this content. In addition, given the vast volumes of data Dropbox has, the company could figure out how to leverage ChatGPT as a monetization vector (e.g. deeper analytics on all the user’s content).

Headwinds: We don’t see any major headwinds for Dropbox in the near-term from generative AI.

Bottom line: We see generative AI as generating slight tailwinds for Dropbox.

Datadog (DDOG) — Covered by Matt Hedberg

Tailwinds: Modern applications are constantly in flux, cycling through ephemeral cloud and container infrastructure as they scale up and down in response to usage. Datadog provides built-in AI and ML logic to simplify the complexity of managing and monitoring environments. With anomaly detection, forecasting, outlier detection, and automatic metric correlation, engineers can leverage AI and ML to intelligently monitor larger numbers of objects more effectively. Watchdog is a great example of how the company is helping customers in areas such as anomaly



detection, root cause analysis and enhanced forecasting to prevent application performance bottlenecks. As such, we believe AI and ML could increasingly act as tailwinds to Datadog as it integrates additional features such as natural language query processing that could help customers uncover trends faster and more intuitively.

Headwinds: We don't see any immediate headwinds for generative AI for Datadog but believe they have to continue to invest beyond AI/ML into GAI.

Bottom line: We see AI becoming increasingly relevant for observability and given Datadog's cloud-based model, we believe it will be in a good position to leverage advanced generative AI capabilities in the future.

DocuSign (DOCU) — Covered by Rishi Jaluria

Tailwinds: While we don't see any major tailwinds on the core e-signature side, we could see generative AI benefiting the broader agreement cloud, especially CLM. ChatGPT can improve search capabilities in finding historical contracts (building on the Seal Software acquisition). In addition, we could see ChatGPT helping with negotiations of contracts and helping users better understand what is embedded in contracts.

Headwinds: We don't see any meaningful headwinds in the near-term from generative AI.

Bottom line: We see generative AI as a slight tailwind for DocuSign.

Dynatrace (DT) — Covered by Matt Hedberg

Tailwinds: Dynatrace could be in a good position to leverage AI through Davis, their underlying AI engine. Davis uses deterministic AI which we believe is a different approach to traditional ML techniques that focuses on causation not correlation. With data being a differentiator, the platform analyzes over 368 billion dependencies per second and pulls data from mainframes, core infrastructure, cloud platforms, and CI/CD pipelines. What this translates to is better, more intuitive monitoring for use cases such as AIOps, anomaly detection and application security. As such, we believe AI and ML could increasingly act as tailwinds to Dynatrace as it integrates additional features such as natural language query processing that could help customers uncover trends faster and more intuitively.

Headwinds: We don't see any immediate headwinds for generative AI for Dynatrace but believe they have to continue to invest beyond AI/ML into GAI.

Bottom line: We see AI becoming increasingly relevant for observability and given Dynatrace's platform with AI built into the core, we believe it will be in a good position to leverage advanced AI capabilities in the future.

Elastic (ESTC) — Covered by Matt Hedberg

Tailwinds: Elastic is a search engine, so intuitively we believe generative AI could be a nice tailwind for Elastic. Over the past 18 months, the company has made investments to enable running large language models directly on its platform with Vector Search. With large transformer models like ChatGPT coming up, management believes this opens up many possibilities for customers to build search centric applications on Elastic given it would be cost prohibitive to run these models on its own infrastructure. Additionally, we think the company is in a good position to leverage generative AI in its enterprise and site search products to provide a more intuitive search process based on natural language processing. We also think Elastic could benefit from AI increasingly built into its security offerings to provide real-time threat intelligence. Finally, we think its observability platform could also benefit from AI to help uncover anomalies faster with less human interaction. That said, management believes it remains in the early innings of the opportunity.



Headwinds: At the moment, we don't see any immediate headwinds for generative AI for Elastic. That said, we do worry they can be outspent by larger peers including they hyperscalers. As such, a potential headwind could emerge if Elastic isn't able to properly integrate generative AI features or that it sees pressure from ChatGPT search capabilities.

Bottom line: While still early, we believe Elastic is in a good position to leverage generative AI and advanced ML capabilities across the entire platform including enterprise search, site search, observability and security. Over time, we wouldn't be surprised to see them as one of the bigger beneficiaries within our coverage.

Fastly (FSLY) — Covered by Rishi Jaluria

Tailwinds: While we do not believe OpenAI is a current Fastly customer, we could see other companies use Fastly for delivery and performance, as companies need to utilize CDN and edge cloud to drive faster response times, as traffic continues to grow exponentially.

Headwinds: We note OpenAI utilizes Cloudflare, a Fastly competitor, as a security vendor today. If Cloudflare meaningfully benefits, it could worsen Fastly's competitive positioning against Cloudflare.

Bottom line: We don't see any major headwinds or tailwinds for Fastly from generative AI.

GitLab (GTLB) — Covered by Matt Hedberg

Tailwinds: We believe generative AI could transform the GitLab platform. The obvious alternative in the space is GitHub Copilot, but we believe GitLab is in a good position to harness generative AI in several ways beyond code suggestions (in closed beta currently and expected to go GA soon) to areas such as how to classify issues, summarize suggestions, hand off and automate work as well as tangential areas such as observability, security and search. The company has focused on AI for years including internal development as well as a March 2022 partnerships with Tabnine, an alternative to Copilot that uses AI to improve code development. That said, the code suggestion functionality in closed beta was developed organically by GitLab. Essentially management sees peers that focus on AI relative to DevOps and code suggestions while it believes GitLab is taking a more holistic and extensive view for DevSecOps. Over time, we believe future internal development and API integrations make a lot of sense given the open source nature of GitLab. That said, we think it will take time to integrate features and positively impact growth.

Headwinds: The obvious potential headwind is Microsoft GitHub Copilot and if GitLab can't keep up with innovation or large cloud data sets.

Bottom line: We believe that GitLab could be one of the better positioned companies in our coverage to benefit from generative AI given its platform approach to DevSecOps and the ability to consolidate spend vs point-based vendors.

Guidewire (GWRE) — Covered by Rishi Jaluria

Tailwinds: We could see Guidewire leverage generative AI to make its systems smarter and more customized (while also reducing the services attach rate). In addition, Guidewire's large volume of data and insights on customer behavior can serve as a major competitive advantage versus others in the P&C insurance software space. In addition, we could see insurance companies leveraging ChatGPT to better serve customers, which frees up budget for Guidewire's solutions and accelerates digital transformation.

Headwinds: We don't see any meaningful headwinds for Guidewire from generative AI.

Bottom line: We see generative AI as a real potential tailwind for Guidewire.



HubSpot (HUBS) — Covered by Rishi Jaluria

Tailwinds: We see numerous tailwinds for each of HubSpot’s core products. HubSpot has already demonstrated some of the use cases we highlight below in the Alpha version of ChatSpot.ai, which is a conversational UX powered by ChatGPT and DALL-E 2 and fine-tuned to operate within HubSpot CRM.

Marketing Hub can benefit from utilizing ChatGPT to generate content to drive inbound marketing or to insert in outbound email marketing campaigns, doing research on potential prospects, and helping to build target customer segments.

Sales Hub can benefit from bringing ChatGPT to automate entry of new contacts, as well as automatically verify contact information. In addition, we can see salespeople use ChatGPT to generate drafts for outbound messages, although we would caution against automating the reach out altogether, as ChatGPT content can be discerned from human-generated content. In addition, ChatGPT could be used to analyze the behavior of salespeople and identify prospects with the highest potential conversion rate or highest propensity to spend.

Service Hub can benefit from having ChatGPT as a “front-end” for customers to interact with service and help lead to resolution of issues without requiring human customer service agents, which generates meaningful cost savings.

Very importantly, if HubSpot embraces generative AI for writing code, we could see HubSpot begin to narrow the gap with its enterprise competitors at the CRM level, most notably Salesforce and Adobe. In our conversations, we find plenty of customers who want to switch from Salesforce or Adobe to HubSpot, but find HubSpot lacks certain enterprise-grade features. Generative AI can make it easier for HubSpot to come up with these features quickly.

Tailwinds: If ChatGPT replaces some of the work of marketers, salespeople, or support/service, it could lead to reduction in potential seats for HubSpot.

Bottom line: We see real potential benefits to HubSpot from generative AI, especially given how quickly the company is moving.

MongoDB (MDB) — Covered by Rishi Jaluria

Tailwinds: We see the massive amount of new content being generated with generative AI as being a tailwind for MongoDB, especially if a meaningful portion of that content contains unstructured data, as MongoDB is built to work with unstructured data (in contrast to traditional relational systems). In addition, as MongoDB is already a tool well-adopted by developers, the use of generative AI for coding can make it easier and faster to build new applications on MongoDB, serving as a tailwind for workloads. Finally, we can see developers using generative AI with MongoDB to automate much of the maintenance and updating of the database (fulfilling the “autonomous database” vision that others have suggested).

Headwinds: We could see a possibility of the rise of generative AI bringing about a brand new category of next-generation database systems that are optimized for generative AI, creating some longer-term competition for MongoDB. We believe it is important that MongoDB moves quickly to optimize its systems for generative AI to defend against potential future competitive threats.

Bottom line: We believe MongoDB will see meaningful tailwinds from generative AI.

Microsoft (MSFT) — Covered by Rishi Jaluria

Please see the dedicated deep dive on Microsoft on pages 17-28.



Cloudflare (NET) — Covered by Matt Hedberg

Tailwinds: We believe generative AI could be a tailwind for Cloudflare for several reasons. OpenAI runs a portion of its network on Cloudflare and leverage technology such as CDN, DDoS, Bot Management and Gateway DNS with future opportunities that could include Workers. So one reason Cloudflare could benefit from generative AI is by hosting other providers of AI technology on the Cloudflare platform. Cloudflare is building out an extensive edge network layer of infrastructure, compute, storage and database as we believe it could see generative AI momentum similar to the hyperscalers. The other way we think Cloudflare could benefit from GAI is through cyber-security solutions that leverage advanced AI capabilities for stopping zero-day threats to more intelligent chatbots. While these trends will likely take time to show up in results, we believe Cloudflare is in a good position to leverage long-term GAI trends.

Headwinds: The headwind could be hyperscalers with larger networks and deeper R&D budgets. The other headwind from a security perspective is competition from larger platform vendors.

Bottom line: Overall we think Cloudflare could be one of the best positioned names in our coverage to benefit from generative AI.

New Relic (NEWR) — Covered by Rishi Jaluria

Tailwinds: If generative AI is, in fact, a tailwind for new application development and cloud workloads, then New Relic, as a player in observability, can benefit. In addition, integration of ChatGPT can make monitoring more effective and easier for users to deploy and analyze underlying data. We note New Relic has announced MLOps capabilities to monitor applications built with OpenAI GPT APIs, with the goal of “instant observability.”

Headwinds: We don’t see any meaningful headwinds from generative AI.

Bottom line: We see slight tailwinds from generative AI for New Relic.

NICE (NICE) — Covered by Rishi Jaluria

Tailwinds: AI is already a big part of NICE’s value proposition, selling IVAs and conversational chatbots alongside the CXone CCaaS platform. NICE has already announced integration of CXone Expert with ChatGPT to create conversational customer experience software. In our view, this creates a major market opportunity for NICE (and other CCaaS market leaders) for several reasons. First, customers already have existing spend on chatbots, most of which have technology that is inferior (from an AI perspective) to ChatGPT, so NICE doesn’t have to generate new spending. Second, the use of IVAs (interactive virtual agents) and chatbots can generate meaningful cost savings for contact centers – for example, an IVA responding to a customer inquiry checking a bank balance is significantly cheaper than having a human agent answer that same question (naturally, the customer checking the bank balance online is the cheapest). Third, NICE has a major data advantage, given it is the market leader in CCaaS and can derive insights from its operational history and customer base to make ChatGPT more effective for the contact center use case, whereas customers using ChatGPT on their own wouldn’t). Finally, given the moving off legacy contact center solutions (e.g. Cisco, Avaya) is a heavy lift, we believe generative AI can help make that move and the customization of CCaaS solutions easier, potentially accelerating migrations to CCaaS.

Outside of the core CCaaS, we also believe generative AI can benefit the financial crime and compliance business, making tasks like AML (anti money laundering) more effective.

Headwinds: To a certain extent, ChatGPT does compete against NICE’s own chatbots and there could be some cannibalization. In addition, some customers could choose to build their own chatbots with ChatGPT, bypassing NICE altogether.



Bottom line: We believe generative AI and especially ChatGPT will be a meaningful tailwind for NICE within the contact center business and further increase its value proposition.

ServiceNow (NOW) — Covered by Matt Hedberg

Tailwinds: We believe generative AI could be a significant tailwind to ServiceNow. We see potential tailwinds including AIOps, advanced chatbots through ServiceNow's customer service solutions, low-code/no-code application development on the ServiceNow platform and general workflow and employee productivity enhancements. ServiceNow has focused on AI for many years including helping customers deliver better self-service, detect incidents, route and prioritize work, discover hidden patterns, optimize knowledge bases, quickly identify workflow that can be automated, recommend action and allow for a natural search experience. In the future we think there are additional areas of expansion opportunities through machine learning frameworks, natural language processing, search and automation. As a cloud-based platform build for workflow, we think the company is in a good position to further leverage generative AI in the future.

Headwinds: We don't see any real headwinds to generative AI for ServiceNow unless other platforms such as Salesforce leapfrog it from a technology perspective.

Bottom line: We believe that ServiceNow could see some of the strongest tailwinds from generative AI over time as we think it could provide a nice benefit to growth for years.

UiPath (PATH) — Covered by Matt Hedberg

Tailwinds: We believe generative AI could represent several tailwinds for UiPath in the form of new use cases customers can utilize. For example, we believe there are several business applications for ChatGPT when combined with a UiPath automation robot including 1) analyzing customer feedback, 2) creating a custom response email, 3) prescreening resumes, 4) creating job interview questions and 5) evaluating customer service conversations. Over time, we think other use cases could emerge for the use of both attended and unattended bots combined with a ChatGPT application. In the near-term, the company plans to release a preview of its GPT connector to help users leverage GPT to generate content in automation as management believes generative AI can be applied to many other use cases. Additionally, the company currently leverages a combination of its own AI models as well as Microsoft/GPT, Google, and Amazon to help accelerate productivity. Overall, management remains bullish on GAI as they believe it can both accelerate the adoption of its platform and democratize access for automation.

Headwinds: One potential headwind that could emerge is if generative AI moves in and replaces the traditional RPA market. We also worry that with a lot of UiPath customer data on premise, it could limit their ability to leverage GAI trends.

Bottom line: Overall we believe the jury is still out on UiPath's ability to leverage GAI as we see it as more of a neutral at this point that could turn into either a bigger headwind or tailwind.

PagerDuty (PD) — Covered by Matt Hedberg

Tailwinds: We believe AI could become a tailwind for PagerDuty over time as they have spent a lot of time talking about foundational AI. AI for IT operations or AIOps is a broad category that comprises the use of big data, ML, and AI to automate and accelerate the identification and resolution of IT issues. In a cloud based world, customers generate millions of events per day. At this scale, it is not viable for IT to manually parse through all that data to detect and remediate issues. As such, we believe customers are increasingly leveraging PagerDuty's AIOps capabilities to help make sense of massive amounts of hybrid-cloud data. Primary products that could see increased customer interest in our view include PagerDuty Event Intelligence and PagerDuty Automation Actions. And while GAI use are interesting from a consumerization perspective, we



believe it would be nearly impossible to leverage internet data for what PagerDuty solves for enterprise customers.

Headwinds: The biggest potential headwind in our view is that PagerDuty could be outspent by larger tech peers.

Bottom line: We see AI as a tailwind for PagerDuty given their focus on enterprise AIOps as well as overall workflow as we believe the company could increasingly leverage aspects of generative AI for additional search capabilities.

Palo Alto Networks (PANW) — Covered by Matt Hedberg

Tailwinds: We think generative AI will give hackers new ways to write and deliver malicious code. As such, we think there are several potential tailwinds AI can provide to Palo Alto Networks to help customers prevent AI-based attacks. From a product perspective, the company launched autonomous digital experience management in FY/22 followed by AIOps and SaaS Security Posture Management this year as management believes AI has the power to transform SASE. Its integrated security services are now all powered by AI to detect and prevent zero-day attacks, and they'll soon be introducing additional AI-driven capabilities. Fundamentally management believes AI is a data problem, which requires clean and comprehensive data sets. Unlike consumer-based AI, we believe the enterprise AI market is more complex and needs to be real-time, especially when it comes to cyber-attacks. Within newer products, we believe AI can also benefit Cortex and XSIAM as well as Prisma Cloud and overall SOC transformations.

Headwinds: We don't see any real headwinds other than some of its customer data is on premise vs in the cloud which makes it more of a challenge to harness customer data sets vs a pure cloud platform.

Bottom line: Overall we think Palo Alto could be one of the biggest benefits of GAI within the cyber-security market given its platform approach that captures petabytes of data.

Pegasystems (PEGA) — Covered by Rishi Jaluria

Tailwinds: On the DPA (digital process automation) side, we can see generative AI helping further automate these processes and integrating systems more intelligently. In addition, Pega does require a high level of customization and generative AI can help reduce services attach. On the CRM side, we can see generative AI making the systems smarter and ChatGPT helping out in contact center use cases. Finally, we expect Pega to utilize generative AI to make its developers more effective. We expect further announcements at Pega's user conference, PegaWorld iNspire, in June (noting the tagline for the event is "The AI landscape is complex. Unravel it here").

Headwinds: We have concerns that some of Pega's automation could be replaced with generative AI, especially for more basic automation use cases.

Bottom line: We believe generative AI can be a slight tailwind to Pega's business, but are closely monitoring the landscape to see if the headwinds are bigger than we expect.

Palantir (PLTR) — Covered by Rishi Jaluria

Tailwinds: Generative AI can make Palantir's solutions smarter and more effective, especially for advanced analytics use cases. In addition, given the high services attach rate, generative AI can make it easier for Palantir to customize software for customers.

Headwinds: We believe Palantir has been benefiting from positioning itself as an AI company and benefiting from AI mandates at companies. We also believe the insights from generative AI can replace the need for some of Palantir's solutions, especially within the commercial segment.



Bottom line: We see more headwinds than tailwinds and, as such, believe generative AI will be net negative for Palantir.

Smartsheet (SMAR) — Covered by Rishi Jaluria

Tailwinds: As generative AI becomes more embedded in business processes, we believe there will be a greater focus on collaboration, as human-to-human collaboration becomes a more differentiated component of work (while siloed tasks like data entry get outsourced to ChatGPT). This could serve as a tailwind for Smartsheet, which helps manage collaborative work and projects. In addition, with Smartsheet's focus on advanced functionality and workflow capabilities, we could see integration of ChatGPT with Smartsheet to make the systems more effective and easier to use. Finally, ChatGPT can serve as an accelerator on its own, providing custom templates depending on specific use cases (e.g. sales pipeline forecasting).

Headwinds: We expect ChatGPT to be able to replace some of the basic task management capabilities of Smartsheet. Having said that, we don't see many paying users using Smartsheet for basic task management and find limited shelfware.

Bottom line: We see generative AI as a slight tailwind for Smartsheet.

Splunk (SPLK) — Covered by Matt Hedberg

Tailwinds: The power of AI resides in data. Splunk customers have a lot of data within the platform. The issue is that most of this data sits on premise and isn't open to large scale data processing models that reside in the cloud. Over time and as Splunk Cloud gains more traction, we think Splunk could be in a better position to leverage AI, but we think others could be better positioned as of now.

Headwinds: The biggest issue is that a majority of Splunk data sits on premise, which in our view prevents GAI scale.

Bottom line: Over time, we could see Splunk in a better position to see GAI tailwinds, but at the moment, we think other cloud-based security and monitoring vendors are in a better position.

Twilio (TWLO) — Covered by Rishi Jaluria

Tailwinds: Twilio is already leveraging OpenAI technology and offers users the ability to build a chatbot using OpenAI technology (along with Twilio Programmable Messaging). In other words, Twilio can help users have a ChatGPT-like experience, but over SMS, which could be helpful for customer service type use cases. In addition, the ability of generative AI to write code can make it easier for companies to adopt Twilio, especially non-digitally native companies. We have fielded investor questions if a ChatGPT mobile application would use Twilio for 2FA, but we believe it would be more likely to use Microsoft for 2FA, given OpenAI's deep partnership with Microsoft.

Outside of the core messaging platform, we believe generative AI can enhance Segment by making a true AI-based CDP. On the flip side, given the limited traction that Twilio has seen with Flex, its CCaaS solution, we don't see meaningful tailwinds for ChatGPT within the Flex business.

Headwinds: Given the limited traction that Flex has seen, we actually believe ChatGPT will increase the gap between the CCaaS market leaders (NICE, Five9, Genesys) and Flex.

Bottom line: We see generative AI as a slight tailwind for Twilio.

Veeva (VEEV) — Covered by Rishi Jaluria

Tailwinds: We see a number of meaningful potential tailwinds for Veeva.

On the R&D side, we believe generative AI has meaningful potential to change the way drugs are designed and brought to market. For example, we note there are a number of start-ups



focused on designing new drugs with generative AI and, in fact, Insilico Medicine is beginning clinical trials on a COVID-19 drug (ISM3312) designed entirely by generative AI. If companies begin to use generative AI to help design drugs, this could be a meaningful tailwind for Veeva's R&D business with higher volume of drugs going through the R&D and regulatory processes. In addition, Veeva can use generative AI within its own R&D products to make them more effective (e.g. digital trials). Finally, given the vast volumes of data Veeva has from its market leadership position, Veeva can utilize generative AI to give unique insights (which can also extend to the commercial business).

On the commercial side, Veeva can utilize ChatGPT to make its CRM more effective, from updating and testing contact information to helping pharmaceutical reps better target the right healthcare professionals. We also expect ChatGPT to be a tailwind for some of the alternative communication channels, such as Approved Email and for improving marketing campaigns utilizing Crossix's data set. Finally, we can see generative AI improving the Data Cloud business, with greater insights derived from the data, as well as helping Veeva potentially catch up faster with IQVIA (wherever the products are competitive).

Headwinds: We could see ChatGPT making pharmaceutical reps more effective and, as a result, leading to further cuts to pharma reps, creating a headwind for Veeva's CRM business.

Bottom line: On balance, while there could be some challenges for the core CRM offering, we see more tailwinds, especially on the R&D business, and believe Veeva will be a net beneficiary from generative AI.

Ziff-Davis (ZD) — Covered by Rishi Jaluria

Tailwinds: We could potentially see Ziff-Davis utilizing ChatGPT for its Internet properties to more efficiently generate content, serving as a tailwind to the digital advertising business (and improving margins). In addition, given Ziff-Davis is a programmatic acquirer, generative AI could help better identify targets that would be a fit within the Ziff-Davis portfolio.

On the software side, we believe it is mixed for marketing (Moz group). If the nature of search does, in fact, change, as we posit throughout this whitepaper, it could create headwinds for SEO solutions. Generative AI, however, can also help accelerate the need for SEO software, as well as help with marketing campaigns. Within security, we could see a tailwind as cyber-attacks become more complex, although we don't expect this to be a meaningful tailwind, given the more downstream and commoditized nature of Ziff-Davis' security solutions.

Headwinds: We could see headwinds to the digital media business from traffic going directly to ChatGPT for "expert answers" instead of Ziff-Davis properties. For example, instead of researching a health topic on Everyday Health, users may instead ask ChatGPT for answers to a topic. Similarly, we could see users asking ChatGPT for a comparison between different smartphones, versus navigating to one of Ziff-Davis' tech properties.

Bottom line: We see several puts and takes here, but do ultimately believe Ziff-Davis will face a slight headwind from ChatGPT as things stand now. Having said that, if Ziff-Davis is able to utilize ChatGPT in ways we can't imagine right now, that headwind could be overcome.

ZoomInfo (ZI) — Covered by Rishi Jaluria

Tailwinds: We believe most users utilize ZoomInfo for the core data, not the up-the-stack offerings, but ChatGPT can certainly change that. For example, ChatGPT can help salespeople automate outreach, so salespeople can spend more time engaging with customers and prospects, versus manually drafting cold outreach. In addition, generative AI can help better identify context and run testing on prior outreach to create better email templates over time. We would caution against outright using ChatGPT to automatically send emails to targets as



communication generated from ChatGPT is distinguishable from human generated content today.

Headwinds: While ChatGPT today refuses to answer questions as to business emails or phone numbers today, other generative AI systems may be willing to serve up that information. Over time, generative AI may be able to provide the same data that ZoomInfo customers pay for today (e.g. “what is the contact information for the head of infrastructure at company xx and who is their boss?”), which would put pressure on paying users.

Bottom line: Given the puts and take, it’s hard to gauge a net impact, but we believe generative AI could be a slight negative for ZoomInfo, if the company is not able to expand beyond the core data offering fast enough.

Zoom Video Communications (ZM) — Covered by Rishi Jaluria

Tailwinds: We see a number of potential tailwinds for Zoom from generative AI and note that CEO Eric Yuan has spoken publicly about Zoom becoming an “AI-first company”.

Zoom is currently utilizing generative AI (including GPT3) throughout the portfolio of solutions. For example, Zoom offers Zoom Smart, which summarizes meetings. Zoom will also use generative AI alongside Zoom IQ for sales.

We see room for Zoom to further expand the use of generative AI, including improving the video technology, search, and chatbots, which can enhance capabilities within chat, as well as offering a cross-sell opportunity within the nascent contact center business (although we would caution that we were skeptical on Zoom’s ambitions within contact center, beyond the low end of the market).

Headwinds: Given Microsoft is a major competitor with Teams and has already incorporated ChatGPT in Teams Premium, the competitive environment could remain intense against Zoom (and perhaps Zoom is already playing catch-up on AI).

Bottom line: We see numerous puts and takes and have concerns around competition from Microsoft, but also note Zoom is willing to embrace the technology and is willing to work with Microsoft (in contrast to Slack). As a result, we believe generative AI can be a slight tailwind for Zoom.



Internet & internet company-specific takeaways

Primary GAI implications/themes across internet

Software engineering efficiency. We see GAI as having potential to super-charge the underlying engineering output for product development as well as adding meaningful automation to maintenance engineering that is a part of the fabric of every company in our coverage universe. With R&D typically running anywhere from the HSD to mid-teens of revenue, we see there being opportunities to recapture at least a few points of fixed cost leverage over and above what's already generally occurring across the space as a function of revenue growth & generally higher opex discipline.

Not all guns are right for this war. The common characteristic in calling out Cloud and Search as likely being outsized beneficiaries of GAI is that both are providing the foundational tools for virtually all companies to leverage (guns for the war so to speak) as opposed to leveraging GAI for more specific customer-facing tools (though social media and certain types of professional services might be the exceptions to this rule). Much in the same way that there's been enormous profit in making chips for smartphones (ie multiple chip players garnering solid profits vs. Apple being the only hardware maker with real operating profit), we believe the majority of the incremental value creation from GAI will accrue to these GAI arms dealers where a lack of competitive differentiation over the longer-term could erode much of the profit pool for those players.

Innovation/disruption could come from unexpected places - admittedly, search seems the most exposed though it could also largely represent TAM expansion vs. risk to Google. With the potential to re-create others' capabilities by using 3p LLM's for a cost and with GAI tools technically having to provide generally unbiased information in order to maintain relevance, it's especially hard to think of certain, better-resourced players in a box and say they couldn't innovate their way into disrupting an entirely different company or even industry. Just brainstorming but some examples would include Apple implementing GAI into its core iOS UI which would then bypass Google entirely. Amazon implements consumer-facing GAI on its site that places more pressure on search but also risks displacing demand as results could well reflect superior responses & content which drive users off Amazon's platform.

Regulatory risk seems highly likely. Within the first few months of its rollout, ChatGPT has already altered its operating & business model to reflect many different negatively-bent manipulations of the service which we view as a harbinger given we're only scratching the surface on users discovering all the ways GAI can be leveraged. Historically, regulators have taken issue with basic search algorithms in particular where user data & the internet's content drove a virtually uncontrollable level of discovery and underlying ability of the provider (Google primarily) to steer results in its own best interests. We'd view the underlying LLM's which power GAI as exponentially more complex and difficult to unpack vs. Google's organic search algorithms for example and as such, are likely to invite both significant regulatory scrutiny along with a complete lack of ability to ultimately control or regulate it short of shutting it down completely. As such, we see Google's more measured approach to rolling it out as trying to demonstrate prudence around user safety whereas MSFT seems more bent on going as fast as possible which we'd view as carrying much more reputational and regulatory risk going forward.

How internet companies are using AI today

Note: All Internet companies are covered by Brad Erickson

Airbnb (ABNB)

AI and machine learning have been incorporated in ABNB since the early days of the business, and they employ it across a number of different parts of the business. ABNB utilizes AI for ranking



properties, predictive search, predictive pricing, user profile personalization, customer service, and analyzing guest behavior all to create a platform that gives both guests and hosts a better experience so they come back more frequently.

Amazon (AMZN)

AMZN uses AI and machine learning not only to improve the customer experience, but also to improve operational efficiency. They employ AI broadly for customer product recommendations, search relevancy, supply chain management, cloud tools, fraud detection, personalization, Amazon Go stores, image/voice recognition, and customer service.

Angi (ANGI)

ANGI to date has not been very public about their AI efforts. Presumably, they are using AI tools to assist in their search and user recommendations. On the Q4 2022 earnings call, CEO Joseph Levin said regarding AI, "Every new technology is a threat and an opportunity. And we certainly think about them in both ways." Besides that, there has been very little discussion publicly.

Booking (BKNG)

BKNG deploys AI across almost every part of the business. For customer service, they have AI chatbots for English speakers that are able to handle 60%+ of all requests. Marketing decisions are driven by models built with learning algorithms to optimize spend. BKNG also uses AI throughout the purchasing funnel, making personalized and connected trip recommendations for users. Within the company, there is also a task force for generative AI which is testing a variety of products to improve the user experience.

Bumble (BMBL)

AI and machine learning are a backbone of BMBL's business. AI is built into the recommendation engine, used in safety, and customer service. In the future, BMBL is hoping AI will help users portray their best self on their profiles and ultimately create the best matches. The company also plans to use AI's predictive capabilities to detect bad actors, making the platform safer. The CEO, Whitney Wolfe Heard, has emphasized focusing on deeper integration of AI moving forward.

CarGurus (CARG)

As an online car platform, CARG has incorporated AI and machine learning into its core business. One of the products is called Lead AI, which analyzes how shoppers engage with the platform to rank them by quality of leads. Another product is Digital Deal, which allows customers to build their perfect deal and connect with dealers across the country. CARG also uses AI to support its supply chain and backend operations.

Carvana (CVNA)

Similar to CARG, CVNA uses AI to streamline operations throughout the business. CVNA uses AI to determine real-time offers for customers from robust data on the automotive market, customer credit score, and financing. AI and other technology is also used for virtual inspections and other aspects of the buying experience. They deploy AI for customer support throughout the purchasing process as well.

DoorDash (DASH)

DASH leverages AI and machine learning to create a more efficient marketplace while also supporting customer service. Their AI capabilities allow them to predict when an order will be ready, picked up, and delivered, so they can optimize each step of the process. Within the delivery process, DASH even tries to predict minute details like how long it will take a dasher to park their car, deliver the item to the doorstep, and the likelihood a dasher will accept an upcoming order. Further, AI is also incorporated into batching, where a single dasher will pick up multiple orders at one time. DASH also uses AI for purchase recommendations and customer service requests.



Expedia (EXPE)

EXPE is deploying AI-driven solutions to create a great customer experience and better personalization to drive high LTV per customer. It's important to callout, Sam Altman, the CEO of OpenAI, is on the board of EXPE. The company uses AI to improve the search experience for consumers in recommendations and discovery, and ranking properties on the platform based off massive data sets. For example, EXPE is able to adjust pictures it displays to users of a specific property based on their interests and prior experiences to increase conversion. EXPE also offers an AI price prediction tool, predicting the best time to book and sending alerts on price drops. They use it for support requests as well. Further, the company uses AI to improve the experience for users in other languages.

Meta (META)

While META rebranded in the fall of 2021 to focus on bringing the metaverse to life, AI has always been an important part of the business strategy. Further, following significant losses from the Reality Labs division, META has shifted some of its resources to AI and created a more centralized AI team to unlock the immense opportunity as it plays a key role in the advertising and social media businesses. AI has helped improved Facebook and Instagram's recommendation engines to increase the quantity of plays for Reels and drive further user engagement across multiple surfaces. Machine learning algorithms, automation, and better measurement have also helped create more relevant and timely advertising, increasing conversion for advertisers at lower CPAs. META has also been creating its own LLMs, like LLaMA and Galactica, to further assist in its business and to expand the customer base.

Fiverr (FVRR)

While AI will likely disrupt the freelancing business, FVRR is committed to embracing the AI revolution. FVRR is offering freelancers a suite of AI tools to improve their efficiency and customer satisfaction. FVRR is also focused on offering businesses connections to AI professionals for a variety of tasks. On the website, FVRR offers AI services for AI Applications, AI Artists, AI Models, AI Music Videos, Fact Checking, AI Content Editing. Lastly, FVRR is using AI and machine learning to increase the quality of search, recommendations, and matching to improve the platform.

Google (GOOGL)

While detailed above, Google has been focused on AI and machine learning for the larger part of a decade, and it is incorporated into almost every part of their business. For search, Google deploys RankBrain to understand the context of a specific search, helping to produce more relevant results. The advertising and cloud businesses have also seen significant increases in efficiency from AI tools. Further, the Transformer paper Google published in 2017 has served as a foundational resource for LLMs. They have utilized this research to develop and train several LLMs like BERT, MUM, and LaMDA to assist in numerous aspects of their business. Bard, which is still in a testing phase, is an AI chatbot powered by LaMDA which will be incorporated into search as well.

Lyft (LYFT)

While LYFT has focused on providing a quality transportation experience, AI and machine learning is deeply integrated into the tech stack to improve the efficiency of the platform for a better experience for both riders and drivers. The company leverages data from over a billion rides for pricing and predictive algorithms. LYFT also deploys AI to improve matching capabilities and optimize driver engagement spend for better efficiency. Lastly, LYFT is using AI and other technology to develop autonomous driving capabilities.

Match (MTCH)

Similarly to BMBL, AI and machine learning is key to MTCH's business. It is used to enhance



the overall user experience to create deeper engagement and higher paying users. Machine learning references data based on interests, prior matches, engagement, and positive outcomes to create higher quality recommendations and matches. It is also used to enhance safety and detect scammers on the platform. MTCH management has also spoken about AI's potential to "coach" users when creating their profiles or engaging with others for more positive outcomes.

Pinterest (PINS)

PINS uses a variety of AI tools to enhance the experience on their platform. For image search, PINS deploys AI to more accurately categorize and rank photos, resulting in higher quality results given a specific search. Because many of the photos on the PINS platform are high resolution photos as opposed to selfies or mobile uploads, PINS is able to use representation learning, an algorithm related to neural networks, to categorize and easily extract billions of photos. Further, PINS uses AI and machine learning to power its recommendation system to drive engagement. The combination of sorting, ranking, and recommending given the robust dataset allows PINS to establish highly personalized user profiles.

Redfin (RDFN)

For RDFN, AI and machine learning is a vital part of its mission to make the home buying and selling process more efficient. While many AI tools are built into RDFN's offering, the main aspects to highlight are estimate and recommendation. Redfin Estimate calculates the market value of a home leveraging datasets, taking into account over 500 data points. Within its recommendation engine, algorithms track user engagement like clicks, interests, and contacting real estate agents to understand individual user preferences. RDFN then employs this data to provide users a list of recommended houses and sends personalized messages and alerts about new listings or price drops.

Vivid Seats (SEAT)

SEAT has discussed very little publicly about its utilization of AI and machine learning on its platform. Despite this, the website most likely utilizes algorithms to set ticket prices and adjusts them accordingly depending on timing and demand of the event. The company is also one of the few in the industry to offer a loyalty program, and they most likely use AI for ticket recommendations based on a user's past purchases or recent interest. Thirdly, SEAT boasts world-class customer service, which more than likely includes some aspect of a chatbot or AI-related care.

Snap (SNAP)

SNAP employs AI to power many of the content creation capabilities users enjoy daily as well as to create a more powerful platform. SNAP creates filters using AI-powered lenses with the help of image processing. Further, the company is also using AI for scene understanding, scene interaction, and scene rendering so machine learning models can recognize names and shapes of objects. In advertisements, SNAP uses targeted ads with goal-based bidding to optimize advertisement conversion. Further, SNAP deploys AI to increase the safety on the platform and to detect unscrupulous actors or interactions. Lastly, in March 2023, SNAP released My AI for Snapchat Plus subscribers, which is a chatbot designed to answer everyday questions.

Squarespace (SQSP)

SQSP has been following AI innovation for the last decade and views it as a tool to augment its core web design capabilities rather than totally replace it. SQSP's platform has Content Transformation Engine, which offers suggestions for content improvements. SQSP also has a customer service chatbot to decrease response times and improve customer satisfaction. Further, SQSP has a partnership with Grammarly, which is an AI-powered grammar checker. SQSP has also integrated Google images onto its platform. In the future, SQSP is focusing on assisting customers in creating content like text and pictures, so they can create websites faster.



Uber (UBER)

UBER deploys AI across many aspects of its platform to enhance the user experience, increase efficiency, and create higher retention. The company uses machine learning to optimize pricing as well as to power upfront pricing, one of the biggest changes to UBER's business, which decouples rider price and driver pay. UBER also leverages its platform data to cross-sell product offerings. Further, UBER employs AI for better matching between riders and drivers and predictive capabilities for a more efficient platform. AI also helps detect fraud and improve customer support. Lastly, like LYFT, UBER is using AI to test autonomous vehicles.

Upwork (UPWK)

UPWK has not revealed much publicly surrounding its AI capabilities and while they recognize the potential headwinds, they are bullish on AI for their business. UPWK is likely deploying AI in search and recommendations to power better matches. UPWK has freelancers in over 90 categories of work and has acknowledged AI and machine learning tools can enhance and increase the efficiency of work, which will make the platform more appealing to current and potential clients. UPWK also expects the AI revolution to increase the need for technical skills for development, training, and support, which is currently one of the biggest use cases for the platform, increasing the overall TAM.

Wix.com (WIX)

WIX employs AI and machine learning in several aspects of its business. WIX offers Wix Artificial Design Intelligence, which is a custom website builder using algorithms based on user's inputs and preferences without requiring code or detailed design. In February 2023, WIX also announced an AI Text Creator tool utilizing GPT-3 to generate human-like content for websites in English with more languages rolling out later on. Further, WIX has an AI-powered logo maker tool with customization capabilities. In 2023, management has highlighted an AI product pipeline simplifying user creation of websites.

Zillow (ZG)

Zillow started testing AI and machine learning in the tech's nascent days and now deploys it in several aspects of their business. Zillow offers Zestimates for real time home valuation based off current market information, listing date, and pictures, among other data. In January 2023, Zillow announced a new AI-powered natural language search for real estate instead of traditionally starting the home search with location or bed/bath count. Further, Zillow's 3D Home app also allows users to create estimated floor plans and virtual tours based on upload photos or scans. Zillow also leverages their datasets and algorithms to send alerts and prospective houses to users.



Implications for internet companies by sub-vertical

Best positioned companies: GOOGL, META, AMZN, PINS, SNAP, and to a lesser degree, BKNG, EXPE and ABNB.

Companies with mixed exposure or at risk: GOOGL, FVRR, WIX, SQSP, UPWK, BKNG, EXPE, MTCH, BMBL

Cloud and Associated Infrastructure: GOOGL, AMZN & MSFT

The one vertical where it is most difficult to build an argument of being negatively impacted aside from competitive concerns. Higher compute intensity driving higher consumption, new revenue models driving TAM expansion and lastly, all sorts of hardware & software opportunities to drive cost & performance efficiencies over time which should drive structural profitability improvements from initial levels barring new, more aggressive competitive entrants. On the infrastructure front, generative AI represents a material, multi-year demand driver for third-party datacenter developers with a focus on hyperscale campus deployments. This includes most notably DLR amongst the listed providers, given its global scale and solid relationships with the major U.S. and Asian players, as well as numerous unlisted operators that are active across the globe. See our most recent [Global Datacenter Review from December 2022](#).

Search Advertising: GOOGL, AMZN

Positively, we see GAI as a likely TAM expander for search in particular, and potentially for social as well for those with the best enabling tools. We believe GAI will likely induce incremental search activity as users become aware of its greater functionality and for many types of enterprise or super-users, there should be incremental revenue opportunities given the cost of compute for many types of these use-cases would be prohibitive. Further, given the likely coming tighter integration with GAI will product more user inputs alongside improving AI being used for targeting and measurement, we'd expect search's ability to target users should also incrementally improve which could drive structural higher CPM's. Further, we think SEM conversion could also improve as a function of potentially smaller ad load driven by scarcity which would also be a positive offset to CPMs. Lastly, we think the content creation efficiencies (fleshed out more fully in social media section below) could create significant incremental dollars that could be spent incrementally just as well on search vs. social in the event that search's marginal ROAS is higher (which we think it often is).

Negatively, the biggest risk is that GAI replaces many types of queries which are currently fulfilled through regular search and in some cases, searches that are monetizable - particularly across categories like travel, education, professional services etc. Further, to the degree that vertical specific companies of a size with adequate resources are able to leverage their own proprietary/vertical-specific data on top of 3p foundation models, this could create a superior search experience in that vertical which could drive some portion of top-of-funnel activity away from traditional search, or even GAI within search.

Social media advertising: META, GOOGL, PINS, SNAP

Positively, the content creation & overall advertising creative industry could arguably see the most disruption amongst industries aside from software engineering & education with the advent of GAI. We believe this could free up significantly more dollars to be plugged back into social media spend (also search given in many cases depending on where a particular advertisers' marginal dollar is producing the highest ROAS). Content creation costs is not only a significant expense for companies and agencies (estimate as much as a third of total ad spend in some cases) but given the relatively recent rise of short-form video engagement, we can make a compelling case that the combination stands to structurally increase over time in the absence of any tailwinds from GAI. In our channel checks with SMB ad agencies, we're already hearing



evidence that advertisers are leveraging tools like ChatGPT to create copy that is approaching the brand messaging intent that human content creation would produce. Further, with yesterday's introduction of several new, larger LLM's (ChatGPT4, PaLM etc), which can increasingly work with text to images, video & audio, this exponentially expands the toolkit for advertisers to add significant efficiencies to content creation costs. We think this is especially positive for META which we believe possesses superior AI-related tools and has greater resources vs. smaller competitors to invest in GAI tools likely enabling differentiated content creation efficiency.

Negatively, to the degree that GAI does somehow enable advertisers to shift spending to other channels due to capturing more users given GAI's other functionality, this could become a source of lead gen diversification. Further, to the degree that social media platforms begin leveraging GAI solutions without significant oversight, we'd have to think that bad actors would almost immediately begin using it for such things as bullying, fraud, phishing and any nefarious output stemming from impersonation.

Ecommerce: AMZN

Positively, we'd think the greater interactiveness of GAI could drive more effective discoverability, faster & more accurate price discovery and an overall better shopping experience. Further, we'd think these tools would likely benefit horizontal/3p marketplaces given by definition, a supposedly-unbiased GAI solution would have to be objectively producing responses reflecting the core tenets of any usual organic search algorithm vs. favoring one particular site or brand's products.

Negatively, any Ecommerce platform with significant competition from a large number of players could actually be disincentivized to leverage a GAI tool as it could drive business off-platform in the event that more relevance was elevated from elsewhere. It poses quite the paradox whereby if AMZN for example, were to implement a GAI solution for shopping, how could it drive its recommend-only products on AMZN if the tool is a truly objective GAI solution that is searching across the entire internet. The easy solve there would seem that in our example, AMZN's GAI would only reflect content from within its own ecosystem but isn't that simply what already exists today and therefore, the GAI's value add would only be to make discovery and shopping more conversational - not exactly revolutionary or as significant of a value when considering humans can already read, and the whole point of GAI is to aggregate content from the entire internet, not just one company's viewpoint.

Online travel: ABNB, BKNG, EXPE

Positively, we think the OTAs in particular (including ABNB) could develop a significantly better booking process within their own sites around not only planning the nodes of a trip (flight, hotel, rental car etc) but also planning a more comprehensive itinerary while simultaneously leveraging ratings, reviews and other types of user-generated content from which a GAI tool should be able to provide differentiated color and guidance for a trip that would be difficult if not impossible to obtain anywhere else besides word-of-mouth. Further, we think the OTA's would be best situated (vs. the hotels themselves) to capture incremental content regarding individual travel suppliers to improve the trip planning process. This could not only create higher SEM & SEO relevance and conversion but also could further enhance the on-site shopping experience beyond what any non-chain hotel could produce (which drives most of the profit for EXPE & BKNG, in particular).

Negatively, the OTAs (BKNG, EXPE mostly, less so ABNB) are heavily exposed to search where we estimate travel may represent as much as 10-15% of Google's search revenue in any given qtr. One of the most well-publicized speculations on how GAI could change industries has been travel with the over-arching view that discovery, shopping and maybe even booking could be upended and done within a GAI UI in a way that search could never fulfill. While there are a variety of reasons why this is likely not the case (which are a function of structural industry



nuances vs. the GAI technology itself), the big risk, in our view, is around the long-tail of travel bidders in SEM as well as SEO. Based on screenshots we generated & analyzed (later in the note), we believe there's a decent probability that for a portion of travel related, discovery/shopping type of searches, the GAI could fully circumvent a typical search or at least severely limit the ad load. Overall, we'd think this could be a net headwind, though an important offset/minimization to the ad load headwind would be that a) the OTA's would almost certainly appear higher than most independent hotels or smaller OTA's and b) OTA's tend to have very little SEO exposure (we estimate LSD/MSD as % of room nights for BKNG & EXPE, likely lower for ABNB).

Professional services marketplaces - FVRR, UPWK

Positively, we think we're entering into a veritable wild west of GAI discovery where all companies, particularly SBMs, are largely not equipped to even begin exploring how they could leverage GAI for their businesses on their own. As such, we see a huge opportunity for a new or at least migratory category of what will essentially be GAI consultants to help businesses of all sizes more quickly understand how to leverage these tools in hopes of new capabilities and/or competitive advantages. Certainly for areas around education, content creation, digital marketing & health, just to name a few. Further, we think opportunities for horizontal customer acquisition tools & techniques should also arise as more and more content can be automated where a whole new level of SEO consulting is likely to arise along with the potential to conduct deeper GAI analysis based on SEM performance where companies can more quickly learn what's driving conversion off of the most critical customer acquisition channels.

Negatively, once the initial phases of exploration, digestion and implementation occur, we think it's entirely possible that various industries, particularly information-centric ones, could start to realize benefits that would actually reduce their reliance on outside professional services. In our space, this would likely affect FVRR & UPWK mostly where categories like content creation, copy writing and digital marketing represent a significant portion of the company's business. In the near-term, FVRR in particular has talked about leaning into the GAI secular trend as it created many new categories for freelancers to fulfill demand which FVRR saw spike after the rollout of ChatGPT. With that said, given the tools make these freelancers exponentially more productive, at some level, we'd think there could be a long-term dilution dynamic that plays out there given the skill-set for achieving those efficiencies is highly replicable once the learning curve becomes less steep.

Web design - WIX, SQSP

Positively, we'd expect web design companies to implement GAI tools (WIX has already implemented a ChatGPT API into its workflow) and certainly when we think about all the associated content-related marketing tools associated with website building, we'd expected fully-integrated GAI capabilities to rapidly become table stakes for any of the major players like WIX, SQSP, GDDY (NC), Weebly, Webflow. Further, given the underlying technology that exists with website building, we'd think this also raises the odds of M&A given bringing together an existing builder with tighter integration with the most powerful models would be a compelling integration for MSFT's or GOOGL's corporate customers in particular (just as an example).

Negatively, between text to text as well as text to image, there's already evidence that GAI's ramping capability can all but replace the functionality for creating a website for someone who has no skillset to do so. Pragmatically, there's all kinds of reasons why this wouldn't necessarily be widely adopted like there's still more friction to that, maintaining is not easy, Ecommerce is very difficult vs. static sites as it's been demonstrated and finally, security. With that said, when we think about logo design as a likely top-of-funnel hook that often grabs new customers in the web design space, the availability of text to visual tools could pose funnel headwinds for these customers. Further, all these companies offer a la carte services (or additions to existing subscriptions) related to more intensive types of content creation which could



become marginalized by GAI.

Digital real estate - ZG, RDFN

Positively, we think companies with the largest data sets, much of which is often not publicly available, should be able to put more distance between themselves and smaller competitors to provide richer content related to shopping for real estate. With that said, we've no reason to think one would have significant differentiation over the much longer-term given all have enough scale to reasonably invest in GAI tools.

Negatively, to the degree that GAI allowed traditional top-of-funnel players (primarily search) to shift some real estate shopping over from a site (find me houses on a quiet street not near railroad tracks in a particular city, for example), this would obviously be a headwind for traffic & impressions. We'd think this is low likelihood.

Online car retail - CARG, CVNA

Positively, we'd think the recommendation engines and customer interaction could increase relevance and thus the quality of the experience substantially. Further, the shopping websites working with GAI would prompt capturing more information around vehicle inventory for example, which would raise discoverability and conversion. Inspections will of course, never go away in the industry, but given the holy grail in B&M car retail is to run a totally seamless omnichannel experience, we'd think GAI could significantly reduce friction that exists today while also increasing consumer confidence in buying such a large item online.

Negatively, we believe the segment has meaningful SEO exposure where to the degree that search habits in the sector see more of a pronounced mix shift towards GAI vs. traditional search, that could pose more significant traffic headwinds. In 2019, when Google introduced some of the most SEO algo changes in years, CARG, CARS and TRUE all experienced significant slowdowns in traffic which resulted in margin compression for the space as all had to increase paid marketing channels.

Ride-hailing/delivery: UBER, LYFT, DASH

Positively, we see very little if any disruption potential for ride-hailing, while for delivery, we think improving recommendation engines would be the primary benefit. While we don't find what Instacart is doing with its early implementation of ChatGPT as being all that innovative (basically the same as Google today with a smaller set of vendors & options to choose from), over time, we'd think this could provide a differentiated and for some a more palatable (pun intended) user experience.

Online dating: MTCH, BMBL

Positively, we see the potential to increase a user's proverbial shots on goal with the ability to automate content and recommendation engines improving to the degree that the companies capture more information from users with the pitch being that more information is likely to find better matches. Along these lines, we'd think there'd be meaningful payer penetration/monetization opportunities here as a portion of users would likely find outsized value add from this capability.

Negatively, anything that reduces authenticity on a dating site/app generally has a souring effect on the experience and certainly GAI does exactly that. To the degree that companies over-enable GAI tools for user conversations for example, we'd think a stigma around trust would quickly emerge and could become a headwind for a particular brand which would be difficult to come back from.



Implications for the cloud wars

Top-down thesis: We believe the hyperscalers will each bring their own proprietary or exclusively-partnered GAI/chatbot solution to market as a Trojan horse intended to drive user adoption leading to enterprise API usage which is a conduit to locking in cloud services & effective competitive share gains over time.

GAI massive opportunity for the hyperscale cloud vendors. We believe AI is reaching an inflection point as a market opportunity, analogous to public cloud infrastructure five to ten years ago driven by the market acceptance and accelerating innovation around generative AI. Amazon Web Services, Microsoft Azure, and Google Cloud Platform all stand to benefit from widespread adoption of AI. There will likely be millions of applications built on top of these AI platforms, reaching billions of end users — for example, Spotify announced Spotify DJ using OpenAI technology and Snap announced My AI and these two companies alone reach close to a billion consumers. More importantly, the compute intensity of AI workloads and training AI models is several times larger than a normal workload, therefore given the scale elasticity required for AI, these workloads not only necessitate cloud adoption but also likely carry a larger revenue per workload.

Microsoft. We do subscribe to the viewpoint generative AI is likely to be a rising tide lifts all boats type of situation. Having said that, we believe Microsoft is looking to leverage a first-mover advantage to become the de facto cloud platform for AI workloads. ChatGPT is already a household name at this point, providing Microsoft not only with a first-mover advantage but also significant mindshare, though still early. As noted earlier, many companies, such as Spotify and Snap, are not only experimenting with OpenAI but releasing OpenAI-powered applications into production. We believe competitors do have the technology to compete and will soon launch competitive commercialized solutions, but likely need to showcase convincing differentiation in order to attract enterprises that are already experimenting with OpenAI. Furthermore, OpenAI and Microsoft have already done pre-emptive price cuts to establish a large ecosystem of partners and customers as quickly as possible, as well as to limit competitors differentiating on price alone. All in all, while still early days, we do feel like Microsoft Azure has strong competitive positioning at this point and is making the right moves to try and sustain it.

Figure 44 - Small snapshot of Microsoft's AI/ML solutions (not an exhaustive list)

Product	Description
Personalizer	AI service that your applications make smarter decisions at scale using reinforcement learning
Computer Vision	API/SDK providing developers with access to advanced algos for processing images and returning information
Speaker Recognition	API/SDK providing developers with access to AI algorithms that can help determine who is speaking in an audio clip
Content Moderator	API/SDK providing developers with access to AI algorithms that check text, image, and video content for material that is potentially offensive, risky, or otherwise undesirable
Face	API/SDK providing developers with access to AI algorithms that detect, recognize, and analyze human faces in images
Keyword Recognition	API/SDK providing developers with access to AI algorithms that detects a word or short phrase within a stream of audio
Text To Speech	API/SDK providing developers with access to AI algorithms that enable your applications, tools, or devices to convert text into humanlike synthesized speech
Image Analysis	API/SDK providing developers with access to AI algorithms that extract visual features from images, like adult content, brands or objects, or human faces
Speech Translation	API/SDK providing developers with access to AI algorithms that give your applications, tools, and devices access to source transcriptions and translation outputs for the provided audio
Intent Recognition	API/SDK providing developers with access to AI algorithms that perform pattern matching and conversational language understanding
Speech To Text	API/SDK providing developers with access to AI algorithms that perform real-time or offline transcription of audio streams into text
Spatial Analysis	API/SDK providing developers with access to AI algorithms that turn detect the presence and movements of people in video
Azure Form Recognizer	API/SDK providing developers with access to AI algorithms that turn documents into usable data
Anomaly Detector	API/SDK providing developers with access to AI algorithms to easily add anomaly detection into apps
Custom Vision	Build, deploy, and improve owned image classifiers, applying content labels based on visual characteristics
Azure Bot Service	Create bots and connect them across channels
Open Datasets	Curated and cleansed data - including weather, census, and holidays - that you can use with minimal preparation to enrich ML models
Kinect Dk	Developer kit with advanced AI sensors that provide sophisticated computer vision and speech models
Azure Cognitive Search	Infrastructure, API/SDKs, and tools for building AI-powered cloud search into mobile and web apps
Health Bot	Platform for healthcare orgs to build and deploy compliant, AI-powered virtual health assistants
Azure Machine Learning	Platform managing the machine learning project lifecycle, including training, deploying, and managing ML models
Azure OpenAI Service	See sub-section in Microsoft section on this topic
Metrics Advisor	Solution using AI to perform data monitoring and anomaly detection in time series data
Video Indexer	Solutions to extract insights from videos and audio
Immersive Reader	Tool that implements proven techniques to improve reading comprehension for new readers, language learners, and people with learning differences such as dyslexia
Data Science Vms	Virtual machine with pre-loaded data science and machine learning tools
Azure Databricks	Apache Spark-based analytics platform

Source: Company website, RBC Capital Markets

Google. While GOOGL has been playing catch up to MSFT/OpenAI since the release of ChatGPT, the company's deep experience with AI and increasing internal focus on rolling out GAI tools into their suite of products may narrow the gap. Despite the seemingly rushed Bard announcement in February, we think future GAI announcements will be more frequent and could change the narrative for the positive. Just recently (on 3/14), GOOGL announced a new GAI app builder and GAI tools within Google Workspace to write emails and documents, create presentations, take meeting notes, summarize email threads, etc. within Gmail, Documents, Sheets, Meet, etc.



Beyond Workspace, Search is the obvious channel to highlight their GAI capabilities to the largest audience possible. While a Search/GAI integration is not there yet today, an eventual release could be a tailwind to GOOGL in terms of mindshare/marketing. And it's worth noting that the current state of GAI chatbots is far from perfect. While we expect them to improve significantly going forward with the release of newer & larger models, the public's (and potential Cloud customers) confidence in the quality of outputs likely has room to the upside. Lastly, it's worth noting that GOOGL develops their own custom TPUs which drive significantly lower relative compute costs for both model training & utilization.

Figure 45 - GOOGL has an extensive array of AI/ML Cloud API/capabilities which we expect them to build upon going forward

Use Case	Product	Description
Scientists	Vertex AI	Centralized UI to manage all ML workflows. Pretrained APIs to quickly develop GAI applications
Scientists	Vertex AI Workbench	Development environment to manage the entire application/model workflow. Allows for faster building and training, better scale and data analysis
Scientists	Vertex Explainable AI	Set of tools/frameworks to better understand ML outputs to debug and improve model performance.
Developers	AutoML	Simplified ML development for less experienced developers across image, text, video, translation, etc applications
Developers	Natural Language AI	Analyzes unstructured text to provide better insights including sentiment analysis, content classification, syntax analysis, etc
Developers	Dialogflow	Customer service chatbots/voicebots
Developers	Media Translation	Real-time audio translation
Developers	Speech-to-Text	Converts audio to text in 125+ languages. Available on-premise or via cloud.
Developers	Text-to-Speech	Converts text to speech at near human quality levels in 40+ languages in 220+ voices. Ability to train custom voice models via customer audio recordings.
Developers	Recommendations AI	Analyzes customer data to maximize metric business is optimizing for at scale to improve conversions and ROAS
Developers	Translation AI	Translate text in 100+ languages
Developers	Video AI	Automatically recognizes 20k+ objects/places/actions. Generates video metadata. Builds video apps.
Developers	Vision AI	ML models to analyze videos and images. Detects objects and handwriting. Pretrained APIs to extract metadata.
Infrastructure	Deep Learning Containers	Environment to develop, test and deploy AI applications
Infrastructure	Deep Learning VM Image	Access to VMs with pre-installed applications
Infrastructure	Cloud GPUs	Access to GPUs to improve job speed with a wide range of performance and prices
Infrastructure	Cloud TPUs	Access to TPUs for larger workloads and ML model training/interating
Infrastructure	TensorFlow Enterprise	Enterprise-grade CPU, GPU, TPU resources
Solutions	Contact Center AI	Human-like AI for customer service
Solutions	Document AI	Provides structured data from documents
Solutions	Intelligent Products	Customer insights to develop and monetize new products
Solutions	Product Discovery	Implement Google-quality search on websites

Source: Google, RBC Capital Markets



Amazon. While AWS may appear to be at a relative disadvantage at the moment, we think they have the potential to be a fast follower in GAI workloads. Relative to OpenAI/MSFT and GOOGL it's not entirely clear that AMZN has developed its own foundational models similar to the scale of GPT or LaMDA/PaLM. While the partnership with Hugging Face (which is also used by MSFT and GOOGL) gives AMZN access to the models which can be trained and deployed via SageMaker, it could be viewed as a relative disadvantage to GOOGL/MSFT by not owning the models themselves. Additionally, beyond Alexa devices, it's not clear they have an intuitive means of distributing a GAI to get in front of users in a more visible way (we think it's unlikely AMZN would implement GAI functionality into amazon.com's search similar to ChatGPT/Bing).

That said, we think AWS has likely been developing cloud AI capabilities longer than any and have a lot of the same functionality as GCP and Azure, if not more advanced. And while they haven't publicly announced a direct competitor to GPT/LaMDA CEO Andy Jassy recently noted they've been working on large GAI models for a long time and the company does have some GAI they offer to clients (Lex and Polly). Additionally, with the Hugging Face partnership, AWS has access to foundational models which can be trained by AWS's Tranium for half the price of comparable GPUs, like the ones used by OpenAI & MSFT. Finally, given their vast developer relationships, they may not need a consumer distribution model to drive appropriate capture of workloads going forward, though may face some NT headwinds given the rapid pace of ChatGPT's more aggressive marketing & both MSFT & GOOGL's likely GAI innovation/implementation.



Figure 46 - AMZN has an extensive array of AI/ML Cloud API/capabilities which we expect them to build upon going forward

Use Case	Product	Description
Computer Vision	Rekognition	ML to automate image recognition and video analysis
Computer Vision	Lookout for Vision	Identify product defects in real time, automating quality inspection
Computer Vision	Panorama	Process and analyze video at the edge to improve logistics
Automated data extraction and analysis	Textract	Extracts text from scanned documents for finance, healthcare and gov't applications. Understands data forms and structures.
Automated data extraction and analysis	Comprehend	Analyze document text and provide insights
Automated data extraction and analysis	A2I	Human review of ML models/systems to ensure quality/precision
Language AI	Lex	Conversational AI for customer service. Chatbots and voice assistants.
Language AI	Transcribe	Speech-to-text transcription
Language AI	Polly	Text-to-speech transcription
Improve customer experience	Kendra	ML powered enterprise search
Improve customer experience	Personalize	Powers personalized customer experiences. Days to implement. Real-time recommendation updates based on new information.
Improve customer experience	Translate	Batch and real-time translation. Incorporated into chat/email/helpdesk applications.
Business metrics	Forecast	ML business prediction model. Retail/inventory, workforce, travel demand, etc. use cases.
Business metrics	Fraud Detector	Custom fraud detection model. Build without any ML experience. Immediate deployment.
Business metrics	Lookout for Metrics	Identify and attribute outlier datapoints. Automated alerts.
Code and DevOps	DevOps Guru	Insights on application behavior to identify anomalies and recommend necessary action. Identify operational issues early.
Code and DevOps	CodeGuru Reviewer and Profiler	Identify security vulnerabilities. Code cost and efficiency analysis.
Industrial AI	Lookout for Equipment	ML model to process data from machinery sensors to minimize downtime
Industrial AI	Monitron	Monitor machine status. Predict maintenance needs.
Healthcare	Healthlake	Securely analyze unstructured data. Make health predictions. Reduce costs on imaging.
Healthcare	Comprehend Medical	Process and analyze unstructured medical text .

Source: AWS, RBC Capital Markets

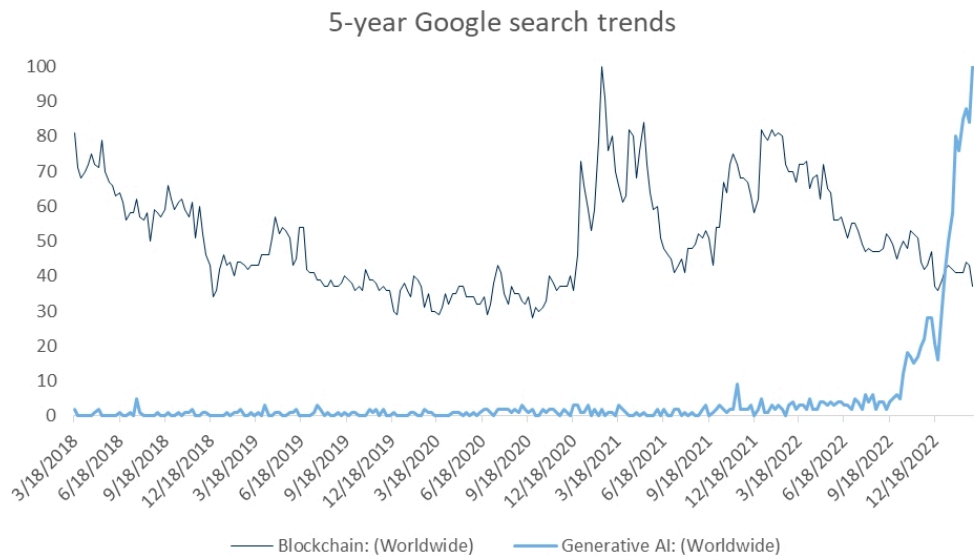
Broad vertical impacts from AI & ML adoption

All aboard the hype train

We feel generative AI is a transformative advancement in technology that will have broad implications across verticals and sectors with applications for a wide range of technology vendors. That said, it has also become a common theme in transcripts across our space as management teams attach themselves to the topic. A focus of this report has been to separate the companies we feel are most likely to have a material impact from the technology and to help investors sort through the noise that is currently being generated. We saw a similar dynamic emerge around blockchain in Q1 2021, and while we feel generative AI has broader applications, we see similarities in the level of investor interest correlating with the number of companies looking to show themselves as well-positioned in the space.

Below, we show Google Search trends for the two topics over the past five years. We can see the peak in blockchain and the drop off afterwards as the hype around the technology transitioned toward the more pragmatic view of applicable use cases. Generative AI has seen a much steeper increase in interest, which has gained the attention of investors and consumers as the accessibility of the technology through applications like ChatGPT and DALL-E has demonstrated the consumerization of tech making the theme much more tangible.

Figure 47 - Google Search trends




Source: Google Trends

Following the money

According to Gartner, over \$1.7B has been invested into generative AI solutions by VC firms over the last 5-years across 255 deals. As seen below, Gartner lays out the investments, number of deals and number of companies receiving VC investments from 2019 until Q3 of 2022.

Figure 48 - Generative AI VC Investment

Generative AI VC Investment by Number of Deals and Companies, 2019-3Q22



Source: adapted from PitchBook investment data.
776430_C


Gartner.

Source: Gartner "Emerging Tech: Venture Capital Growth Insights for Generative AI" By Annette Jump, Vibha Chitkara, Alys Woodward on December 2, 2022

Within those deals, we can see the largest dollar amounts have gone to AI-enabled drug discovery while the AI Avatars have seen the largest number of companies receiving VC investments. From a tech sector perspective, AI Code Generation is the only segment that saw increased funding in 2022 compared to 2021.

Figure 49 - Generative AI VC Investment by Segment

Generative AI VC Investment by Segment, 2019-3Q22



Source: adapted from PitchBook investment data.
776430_C




Source: Gartner "Emerging Tech: Venture Capital Growth Insights for Generative AI" By Annette Jump, Vibha Chitkara, Alys Woodward on December 2, 2022

Breaking these deals down further, later stage startups have seen the most funding while early stage has seen the largest number of companies receiving VC funding.

Figure 50 - Generative AI VC Investment by Stage

Generative AI VC Investment by Stage, 2019-3Q22



Source: adapted from PitchBook investment data.
776430_C



Source: Gartner "Emerging Tech: Venture Capital Growth Insights for Generative AI" By Annette Jump, Vibha Cithara, Alys Woodward on December 2, 2022

Robot revolution

Generative AI is likely to have an impact across a wide range of industries both by enhancing the quantitative skills that computers have always excelled at, which will be an advancement of existing AI and ML technology, but also flowing into new areas like writing and creative processes which have typically seen less machine intervention.

We continue to believe in the long-run augmented intelligence will be the end-result of most of these changes. As seen below Gartner has laid out uses cases for generative AI as well as the industries they are most likely applicable in. Through this we can see more real-world examples of how generative AI is likely to reshape aspects of work across verticals.

Figure 51 - Generative AI Use Cases by Industry

Use Cases by Industry/Capability

	Industries									Business Capabilities				
	Automotive and Vehicle Manufacturing	Media	Architecture and Engineering	Energy and Utilities	Healthcare Providers	Electronic Product Manufacturing	Manufacturing	Pharmaceutical	Human Resources Management	Information Technology Management	Marketing and Sales	Research and Development		
Drug Design								●						
Material Design	●			●		●						●		
Chip Design						●								
Synthetic Data	●		●	●	●	●	●	●		●	●	●		
Generative Design (Parts)	●		●				●					●		
Generative Design (Architecture)			●											
Code Generation										●				
Text Generation		●									●			
Image Generation		●									●			
Video Generation		●							●		●			
Audio Generation		●									●			
Media Content Improvement		●									●			

Source: Gartner
778501_C



Source: Gartner “Innovation Insight for Generative AI” By Brian Burke, Arun Chandrasekaran and Svetlana Sicular on December 15, 2022

How the market may develop

While generative AI will likely impact most verticals in one way or another, as we often see with emerging technology, there will likely be a lag between early adopters and when the technology becomes more pervasive. Investments made within generative AI do not only pertain to the technology companies developing them, but for the customers who use them. Industries will need to retrain employees in many cases to realize the full efficiency gains of generative AI with forward-looking enterprises likely to see the earliest benefits.

Strategic planning assumptions from Gartner’s “Top Strategic Technology Trends for 2022: Generative AI”:

- By 2025, generative AI will account for 10% of all data produced, up from less than 1% today.
- By 2027, 30% of manufacturers will be using generative AI to increase product development efficiency.
- By 2025, 50% of drug development initiatives will use generative AI.
- By 2024, 50% of the largest low-code/no-code development platforms will provide “text to code” functionality in their AI suites.
- By 2025, 20% of all test data for consumer-facing use cases will be synthetically generated.
- By 2025, 30% of outbound marketing messages from large organizations will be synthetically generated.
- By 2025, 90% of the material in quarterly reports will be synthetically generated.



- By 2024, the European Union will pass legislation to mandate the “watermarking” of AI-generated artifacts.
- By 2025, 20% of procedural code professionals will have retrained because generative AI will have cannibalized their core skill set and market value.
- By 2025, stunt doubles in the entertainment industry will no longer need to look like the stars they imitate.

Finding a balance

The largest concern around AI and more particularly generative AI has been around the impact on employees and the potential for jobs being replaced by machines. As seen in the survey results from Gartner below, employees are not opposed to AI assistance, but only AI replacement. The key for successful implementation will be finding the balance that maximizes productivity without alienating employees, maximizing the potential of human capital.

Traditional AI use cases have been widely accepted focusing on automating manual and repetitive tasks and streamlining processes as seen below in the strong preference for AI assistance in data processing. Additionally, most employees are also comfortable with measures for mistake reduction utilizing AI - we are all familiar with tools like spell-check.

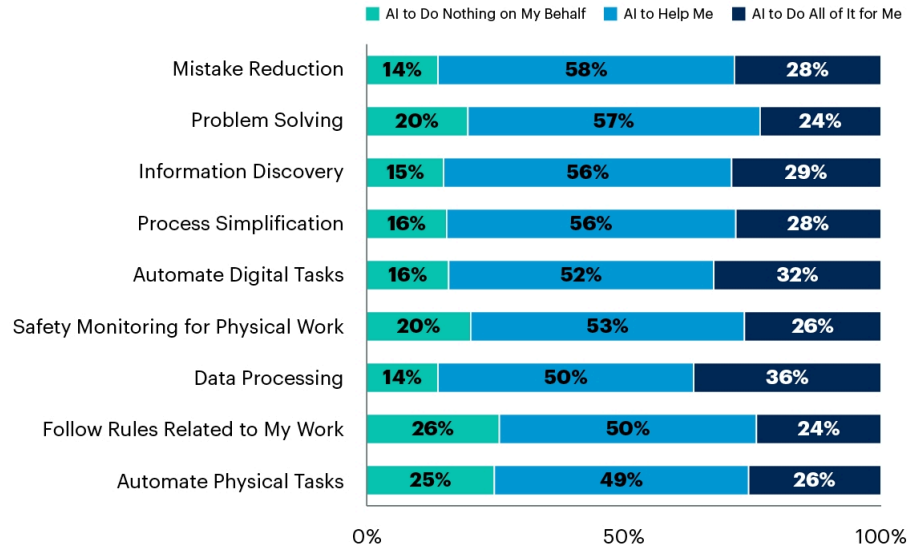
The areas people are less comfortable with are when they move away from machines and into the physical workplace which in our view continues to show a lack of trust in artificial intelligence more generally. Areas where workers stated they wanted the least AI intervention included “automate physical tasks” and “follow rules related to my work” as well as “safety monitoring for physical work”.

The answer that we feel relates the most closely to generative AI is in problem solving. 20% of respondents want AI intervention, 24% want AI to “do it all for me”, while 57% want AI help. This latter segment which has the highest proportion of respondents falls squarely in the realm of augmented intelligence. We feel this is the long-tail of generative AI, with the human role adapting to the changing dynamics within the workplace focusing on areas where their domain specific expertise can enhance the automated processes within their workplace.



Figure 52 - Gartner Survey Results on the Extent Workers Would Like AI to Help

Extent of Help
Percentage of U.S. Workers



n = 1,424 Base: U.S. workers, excludes "not applicable"

Q: To what extent would you like AI to help you achieve the following on your behalf?


Source: 2021 Gartner Worker-Consumer AI Attitudes Survey
767371_C



Source: Gartner "How to Convince Your CEO to Use AI to Augment, Not Replace, Workers" By Leigh McMullen, Whit Andrews, Svetlana Sicular, Dave Aron on May 16, 2022


While still in its nascent days and yet to see true mass adoption, AI and machine learning will disrupt many industries and sectors moving forward, especially given the massive funding it has and will continue to receive. As seen in the two charts below, corporations and venture capital firms have been pouring money into AI and machine learning at an accelerating rate.

Figure 53 - Global AI investment continues to ramp



Source: Statista, RBC Capital Markets

Figure 54 - Venture Capital AI & ML Deal Value and Deal Count



Source: Pitchbook, * Only through Q3 2022

Large language model (LLM) tech underpinnings

What is a large language model?

A large language model is a machine learning model that is trained on a massive amount of data in order to perform natural language processing tasks. Some of the most well-known LLMs are GPT and BERT. LLMs can be used for a variety of natural language processing tasks, including answering questions, language translation, sentiment analysis, and text summarization. They are also used to generate text in applications like chatbots, virtual assistants, and content creation tools.

LLMs are trained on a massive amount of text or a subset of the data which is first cleaned and processed. GPT-3, for example, has 175 billion parameters and is one of the most powerful LLMs to date. LLMs have neural networks, consisting of many layers of interconnected nodes, designed to learn and recognize patterns in the text. LLMs like GPT-3 and LaMDA are transformer neural networks which are very effective at language tasks.

One of the training methods uses self-supervised learning, which teaches the model to predict missing words or phrases after being given a sequence of words. It involves feeding the model many examples of texts with certain words masked or adjusting the weights of the neural network, so it becomes better at predicting the missing words. The process is repeated iteratively until the model reaches an acceptable level of accuracy. One advantage of this training method is that it trains the model on the statistical nature of language, leading to more natural and fluent language responses. On the flip side, because the model is trained to predict the next word, it is not necessarily learning the higher level representation of meaning that human's rely on, creating the need for reinforcement learning from human feedback.

OpenAI and Anthropic are known for pioneering "instruction tuning", which is the method of fine-tuning self-supervised machine-learning models with human feedback. This practice has enabled much smaller models to perform on-par, if not better, than models that are much larger in parameter size.

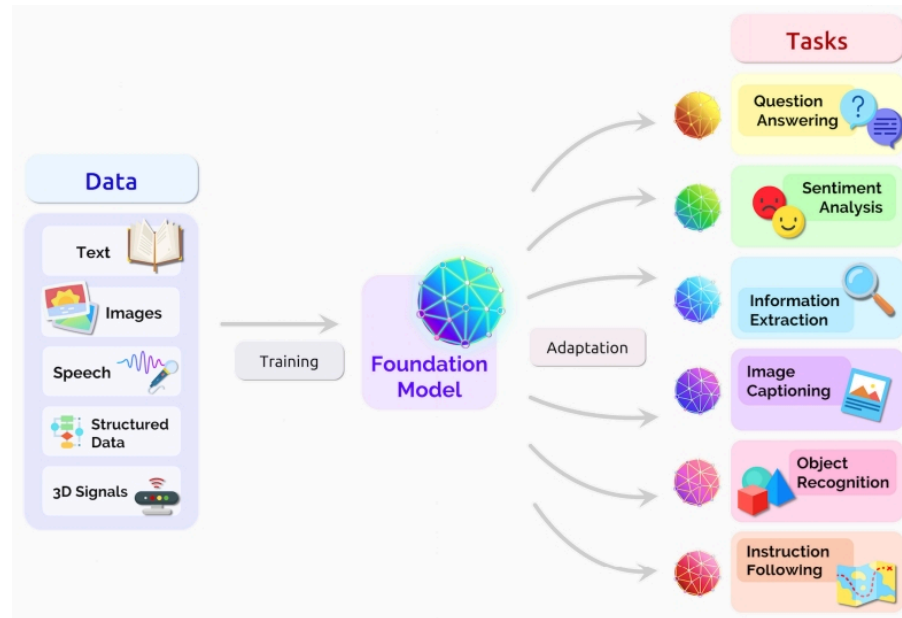
As the model becomes better at predicting the next word or phrase, it also becomes better at generating natural and coherent language. The model then utilizes the patterns it has learned from training and can be fine-tuned to generate outputs based on the desired goal. This allows the model to be tailored for specific requirements.

In terms of training costs for LLMs, a rough estimate is ~\$1/1000 parameters, depending on the operating costs, model complexity, and type of training data. Extrapolating that out, it costs ~100k-\$1.5M to train 1.5B parameters. GPT-3, for example, was built on 175B parameters and is estimated to have cost several million dollars to train according to Capgemini. As these models increase in size and complexity, the training costs will continue to rise.

What is a foundation model?

While large language models as the name suggests refer to machine-learning models that can understand language, foundation model refers to the broader umbrella of models that can understand language, vision, robotics, or reasoning and search. Foundation models as a term was popularized by Stanford's Center for Research on Foundation Models (CRFM). A good example of a non-language foundation model is OpenAI's DALL-E 2, which would be considered a large vision model, as the algorithm not only understands and recognizes images, but generates them too (in response to human prompts, much like ChatGPT works with language).

Figure 55 - Foundational Models



Source: Center for Research on Foundation Models

Google Transformer Paper

Google published a paper in 2017 titled “Attention is All You Need” introducing the Transformer, a novel neural network that was groundbreaking for understanding and completing tasks like question answering, machine translation, and language modeling. Instead of processing sequential inputted data, the algorithm processes the entire input all at once. The paper detailed the Transformer’s outperformance vs. prior models and demonstrated how it required fewer steps compared to prior models, less computation to train, and is an overall better match for machine learning hardware.

In the past, recurrent neural networks (RNN) were the most prominent approaches to language understanding. In general, neural networks process and understand language by creating fixed or variable length vector space representations. They start with single words, or even pieces of words, and they collect data from adjacent words to determine meaning. RNNs process words and language in left to right or vice versa manner, reading one word at a time. This forces RNN to perform multiple steps, like an iterative process, before it can accurately understand the words or phrases. Research has shown, the higher the number of steps a network requires, the harder it is for a network to understand and then ultimately make a decision. Based off the left-to-right or right-to-left nature of RNNs, they are not the best suited for advanced computing hardware like TPUs and GPUs.

Compared to RNNs, the Transformer is a novel neural network and performs a small but constant number of steps to understand language. The Transformer is able to directly model relationships between all the words in the sentence by comparing an individual word to all the other words in the sentence at the same time, regardless of position, severely cutting down the number of steps and allowing for faster computing.

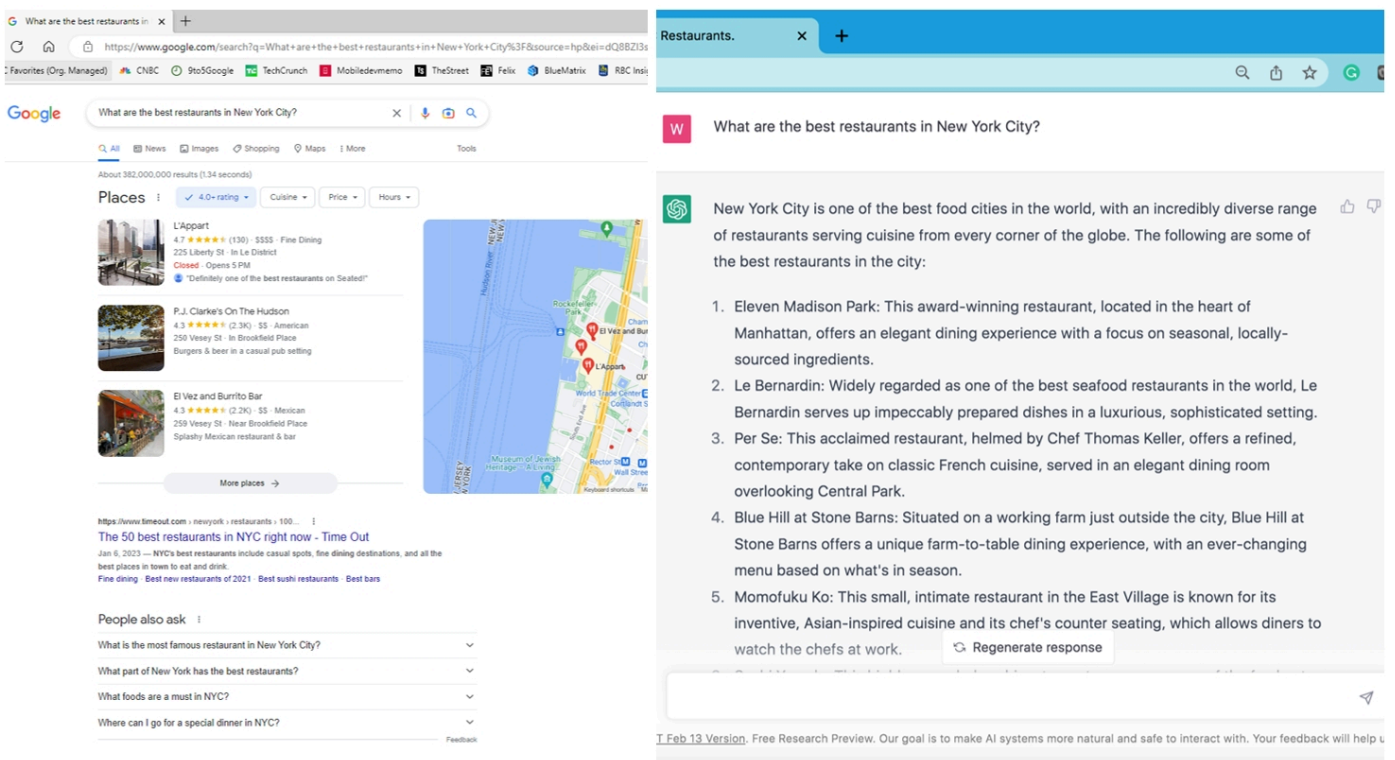
This paper has served as the foundation for many of the AI tools and capabilities that Google has released in the past six years. The current versions of LaMDA, PaLM, and other features would not have been possible without the Transformer research.

ChatGPT vs. Google Search across internet verticals

Travel

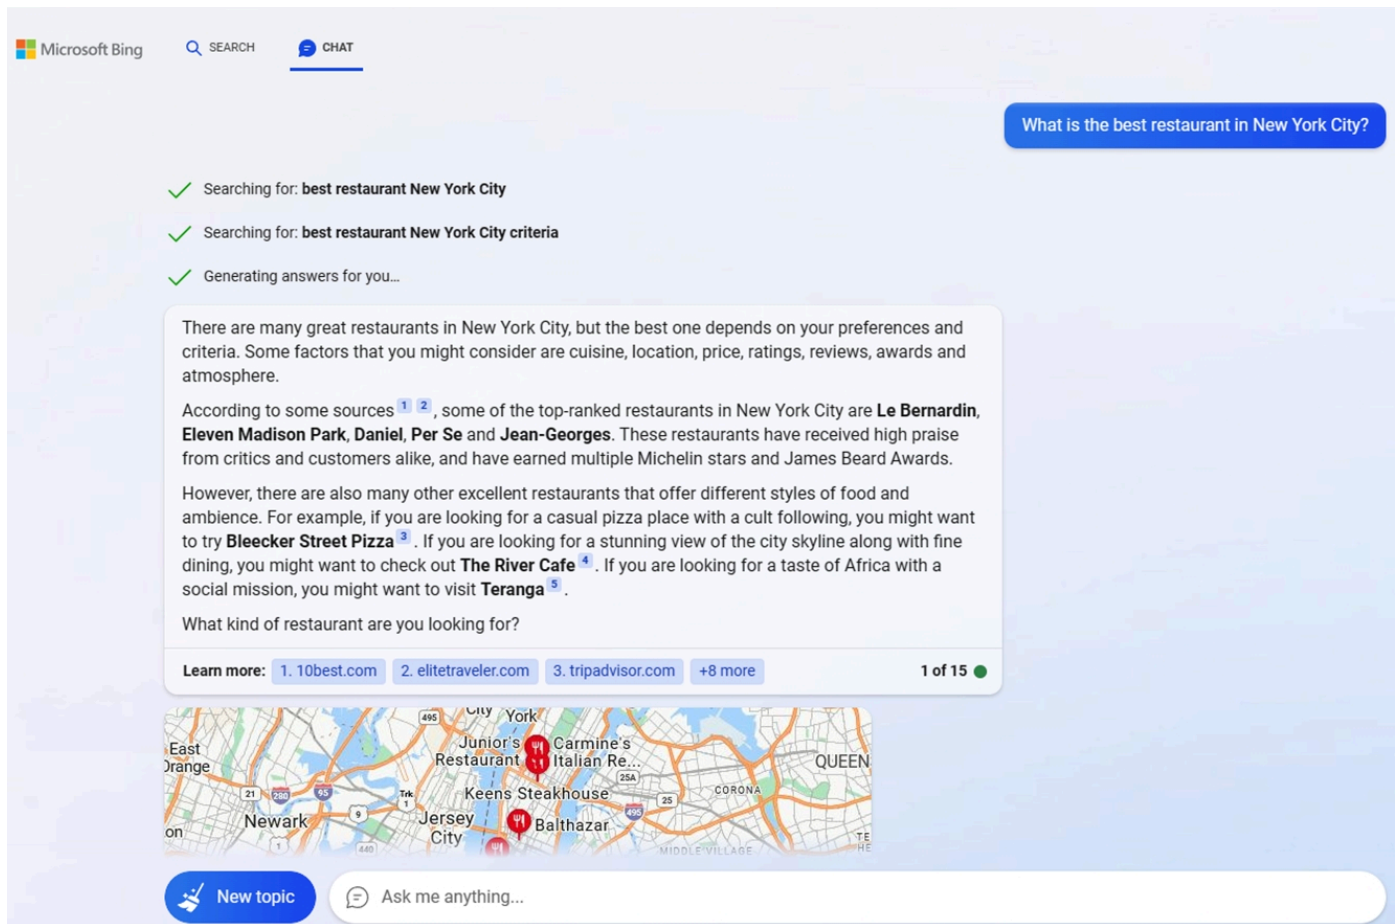
Restaurants in a big city produces a very similar type of output vs. traditional Google and judging by ChatGPT's UI, we think there are likely similar opportunities to apply a similar set of recommendations while changing the taxonomy of the page to reduce the organic results and increasing the paid results. Further, in terms of research, ChatGPT's nicely laid out initial screen, just like Google's, is totally insufficient for showing a menu, reading reviews and providing at least some, merchant-specific information which would be somewhat necessary for going deeper in the funnel (calling, emailing, booking etc).

Figure 56 - Google Search vs. ChatGPT - Restaurants



Source: Google Search, ChatGPT, RBC Capital Markets

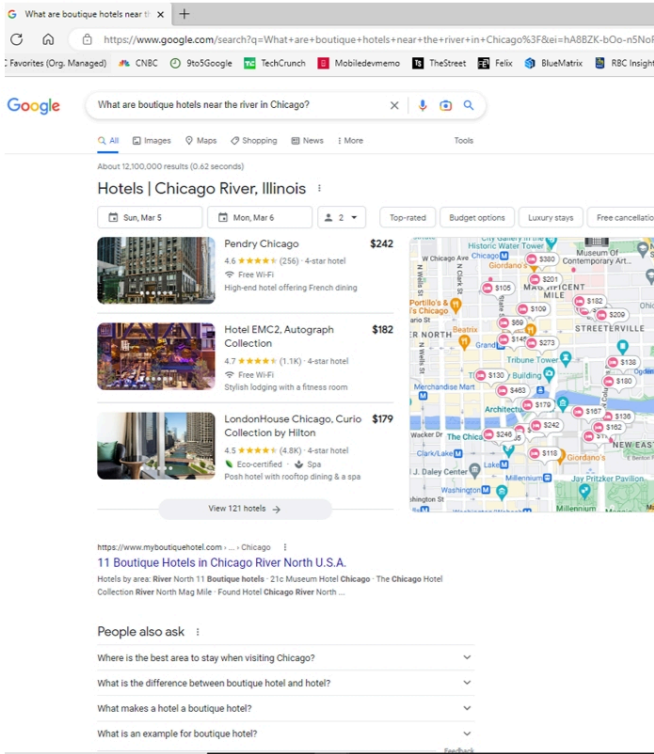
Figure 57 - Bing - Restaurants



Source: Bing, RBC Capital Markets

Boutique hotels near the river in Chicago yields some helpful detail and in this case, the simply constructed explanations feel similar to hearing a recommendation from a friend who'd stayed there which is superior to anything search could yield. With that said, hotels are one category where for discoverability, research & even booking, text is likely not sufficient to drive down-funnel conversion as we believe visual shopping this case seems somewhat of a critical component both within the search experience and on the hotel's website if a user is going to go further or all the way down the funnel to conversion.

Figure 58 - Google Search vs. ChatGPT - Hotel



What are boutique hotels near the river in Chicago?

About 12,100,000 results (0.62 seconds)

Hotels | Chicago River, Illinois

Sun, Mar 5 | Mon, Mar 6 | 2 | Top-rated | Budget options | Luxury stays | Free cancellation

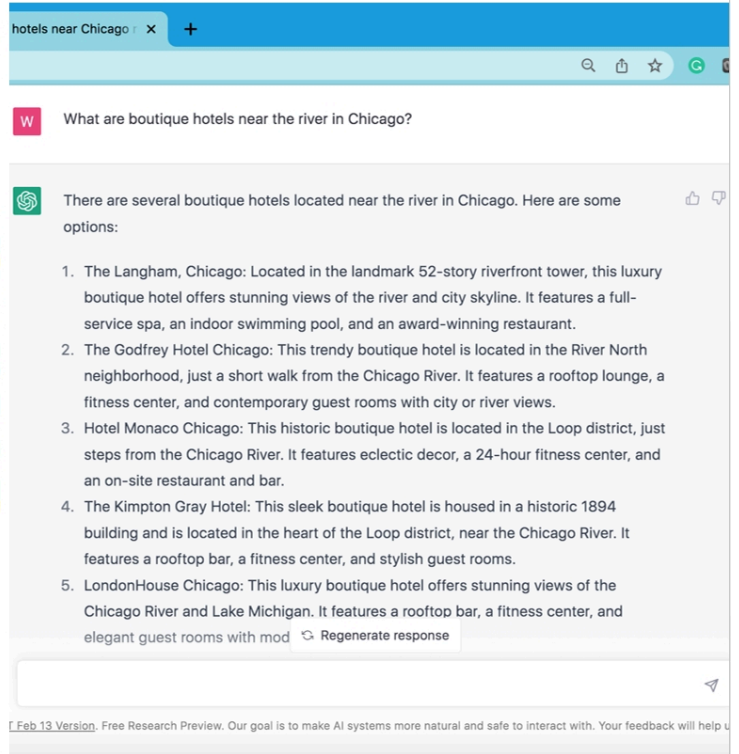
- Pendry Chicago** \$242
4.6 ★★★★★ (256) 4-star hotel
Free WiFi
High-end hotel offering French dining
- Hotel EMC2, Autograph Collection** \$182
4.7 ★★★★★ (1.1K) 4-star hotel
Free WiFi
Stylish lodging with a fitness room
- LondonHouse Chicago, Curio Collection by Hilton** \$179
4.5 ★★★★★ (4.8K) 4-star hotel
Eco-certified | Spa
Posh hotel with rooftop dining & a spa

View 121 hotels →

11 Boutique Hotels in Chicago River North U.S.A.
Hotels by area: River North 11 Boutique hotels - 21c Museum Hotel Chicago - The Chicago Hotel Collection River North Mag Mile - Found Hotel Chicago River North ...

People also ask

- Where is the best area to stay when visiting Chicago?
- What is the difference between boutique hotel and hotel?
- What makes a hotel a boutique hotel?
- What is an example for boutique hotel?



What are boutique hotels near the river in Chicago?

There are several boutique hotels located near the river in Chicago. Here are some options:

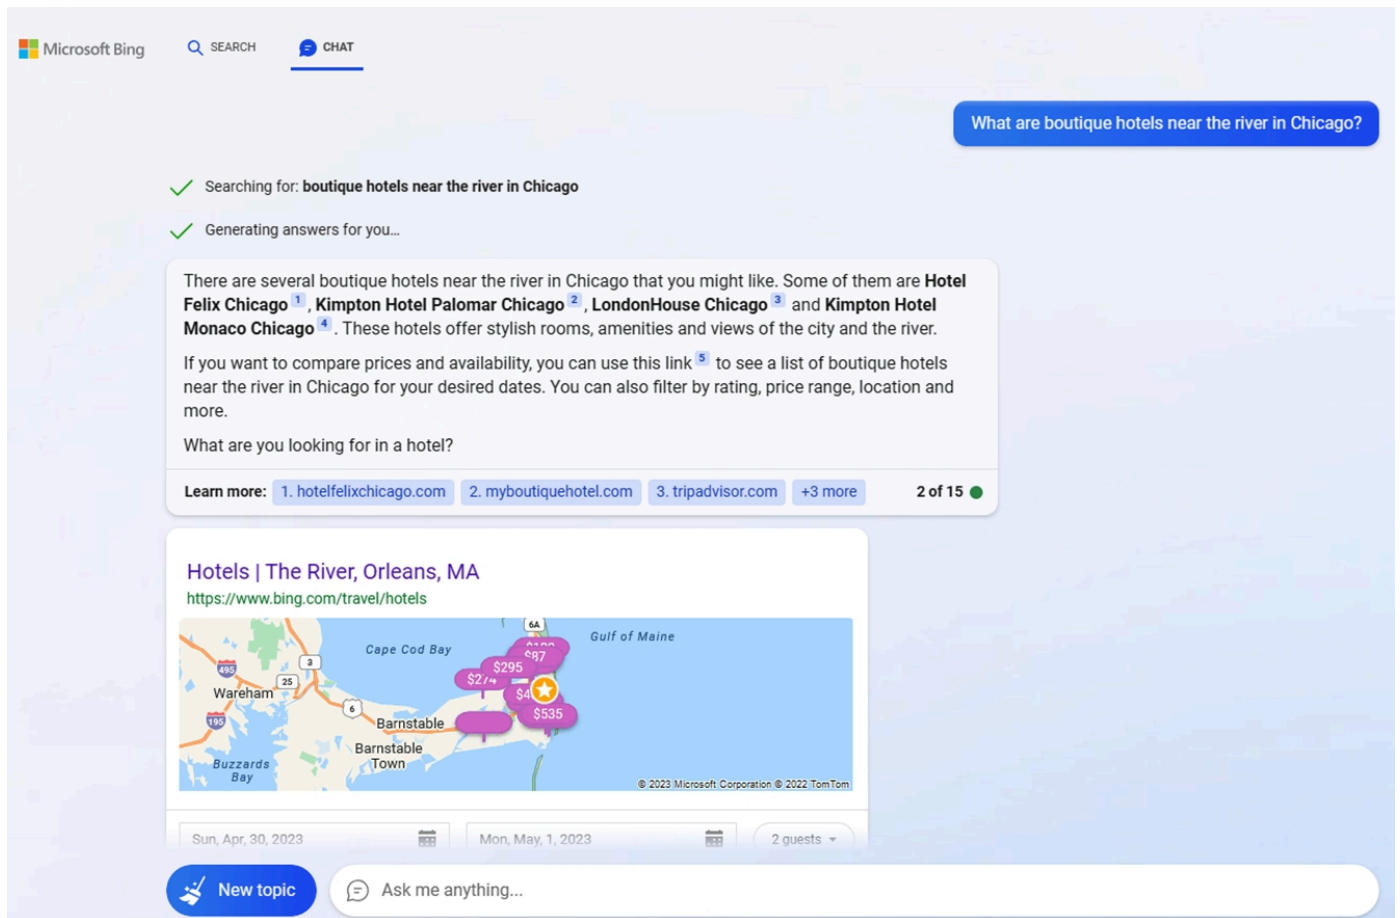
- The Langham, Chicago:** Located in the landmark 52-story riverfront tower, this luxury boutique hotel offers stunning views of the river and city skyline. It features a full-service spa, an indoor swimming pool, and an award-winning restaurant.
- The Godfrey Hotel Chicago:** This trendy boutique hotel is located in the River North neighborhood, just a short walk from the Chicago River. It features a rooftop lounge, a fitness center, and contemporary guest rooms with city or river views.
- Hotel Monaco Chicago:** This historic boutique hotel is located in the Loop district, just steps from the Chicago River. It features eclectic decor, a 24-hour fitness center, and an on-site restaurant and bar.
- The Kimpton Gray Hotel:** This sleek boutique hotel is housed in a historic 1894 building and is located in the heart of the Loop district, near the Chicago River. It features a rooftop bar, a fitness center, and stylish guest rooms.
- LondonHouse Chicago:** This luxury boutique hotel offers stunning views of the Chicago River and Lake Michigan. It features a rooftop bar, a fitness center, and elegant guest rooms with mod

Regenerate response

Feb 13 Version. Free Research Preview. Our goal is to make AI systems more natural and safe to interact with. Your feedback will help u

Source: Google Search, ChatGPT, RBC Capital Markets

Figure 59 - Bing - Hotel

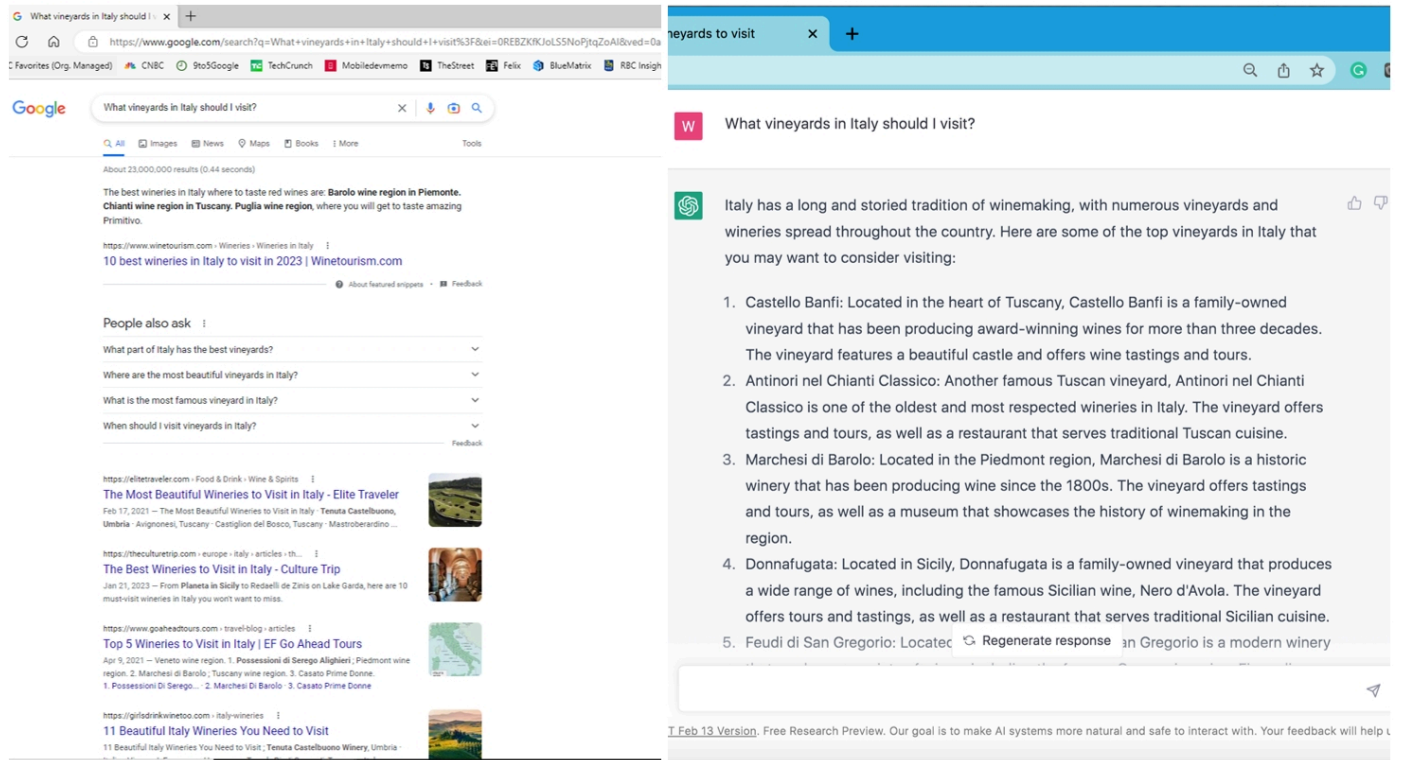


The screenshot shows the Microsoft Bing search interface. At the top, there are navigation options for 'Microsoft Bing', 'SEARCH', and 'CHAT'. A search bar contains the query 'What are boutique hotels near the river in Chicago?'. Below the search bar, there are two green checkmarks indicating search progress: 'Searching for: boutique hotels near the river in Chicago' and 'Generating answers for you...'. The main content area features a text-based answer: 'There are several boutique hotels near the river in Chicago that you might like. Some of them are **Hotel Felix Chicago**¹, **Kimpton Hotel Palomar Chicago**², **LondonHouse Chicago**³ and **Kimpton Hotel Monaco Chicago**⁴. These hotels offer stylish rooms, amenities and views of the city and the river. If you want to compare prices and availability, you can use this link⁵ to see a list of boutique hotels near the river in Chicago for your desired dates. You can also filter by rating, price range, location and more. What are you looking for in a hotel?'. Below the text, there are three links: '1. hotelfelixchicago.com', '2. myboutiquehotel.com', and '3. tripadvisor.com', followed by '+3 more' and '2 of 15'. A secondary search result is visible below, titled 'Hotels | The River, Orleans, MA' with the URL 'https://www.bing.com/travel/hotels'. This result includes a map of the Orleans, MA area showing hotel locations with price tags: '\$277', '\$295', '\$4', and '\$535'. The map also shows 'Cape Cod Bay', 'Wareham', 'Barnstable', and 'Barnstable Town'. At the bottom of the search results, there are date and guest selection options: 'Sun, Apr, 30, 2023', 'Mon, May, 1, 2023', and '2 guests'. A 'New topic' button and an 'Ask me anything...' input field are also present.

Source: Bing, RBC Capital Markets

(No ad impact): Vineyards in Italy to visit yields a solid overview in GAI and then its usual list in rank order. Interestingly, in search, it gives the beginning of a similar list which came from a single organic site which apparently ranked best in its quality/intent/relevance score. Lastly, there were no ads here which seemed interesting given a competitive wine industry (though more rational given its history and cultural institutionalization relative to American wineries).

Figure 60 - Google Search vs. ChatGPT - Italian Vineyard



The figure compares search results for the query "What vineyards in Italy should I visit?".

Google Search Results:

- Search query: "What vineyards in Italy should I visit?"
- Results: About 23,000,000 results (0.44 seconds).
- Top results include:
 - "The best wineries in Italy where to taste red wines are: Barolo wine region in Piemonte. Chianti wine region in Tuscany. Puglia wine region, where you will get to taste amazing Primitivo."
 - "10 best wineries in Italy to visit in 2023 | Winetourism.com"
 - "The Most Beautiful Wineries to Visit in Italy - Elite Traveler"
 - "The Best Wineries to Visit in Italy - Culture Trip"
 - "Top 5 Wineries to Visit in Italy | EF Go Ahead Tours"
 - "11 Beautiful Italy Wineries You Need to Visit"

ChatGPT Results:

- Search query: "What vineyards in Italy should I visit?"
- Response: "Italy has a long and storied tradition of winemaking, with numerous vineyards and wineries spread throughout the country. Here are some of the top vineyards in Italy that you may want to consider visiting:"
- List of vineyards:
 - Castello Banfi: Located in the heart of Tuscany, Castello Banfi is a family-owned vineyard that has been producing award-winning wines for more than three decades. The vineyard features a beautiful castle and offers wine tastings and tours.
 - Antinori nel Chianti Classico: Another famous Tuscan vineyard, Antinori nel Chianti Classico is one of the oldest and most respected wineries in Italy. The vineyard offers tastings and tours, as well as a restaurant that serves traditional Tuscan cuisine.
 - Marchesi di Barolo: Located in the Piedmont region, Marchesi di Barolo is a historic winery that has been producing wine since the 1800s. The vineyard offers tastings and tours, as well as a museum that showcases the history of winemaking in the region.
 - Donnafugata: Located in Sicily, Donnafugata is a family-owned vineyard that produces a wide range of wines, including the famous Sicilian wine, Nero d'Avola. The vineyard offers tours and tastings, as well as a restaurant that serves traditional Sicilian cuisine.
 - Feudi di San Gregorio: Located in Sicily, Feudi di San Gregorio is a modern winery

Source: Google Search, ChatGPT, RBC Capital Markets



Figure 61 - Bing - Italian Vineyard

Microsoft Bing SEARCH CHAT

Let's learn together. Bing is powered by AI, so surprises and mistakes are possible. Make sure to check the facts, and [share feedback](#) so we can learn and improve!

[Terms of Use](#) | [Privacy Statement](#)

Choose a conversation style [Preview](#)

More Creative **More Balanced** More Precise

What vineyards in Italy should I visit?

✓ Searching for: **best vineyards in Italy**

✓ Generating answers for you...

Hello, this is Bing. Italy has many wonderful vineyards to visit, but it depends on your preferences and budget. According to one source ¹, some of the best vineyards in Italy are **Marchesi di Barolo**, **Contucci**, **Azienda Vinicola Contini**, **Possessioni di Serego Alighieri**, and **Savignola**. You can also find local results near you by using your location ². For example, one nearby vineyard is **Winery Fattoria Cantagallo** which has excellent reviews and offers wine tasting and tours ².

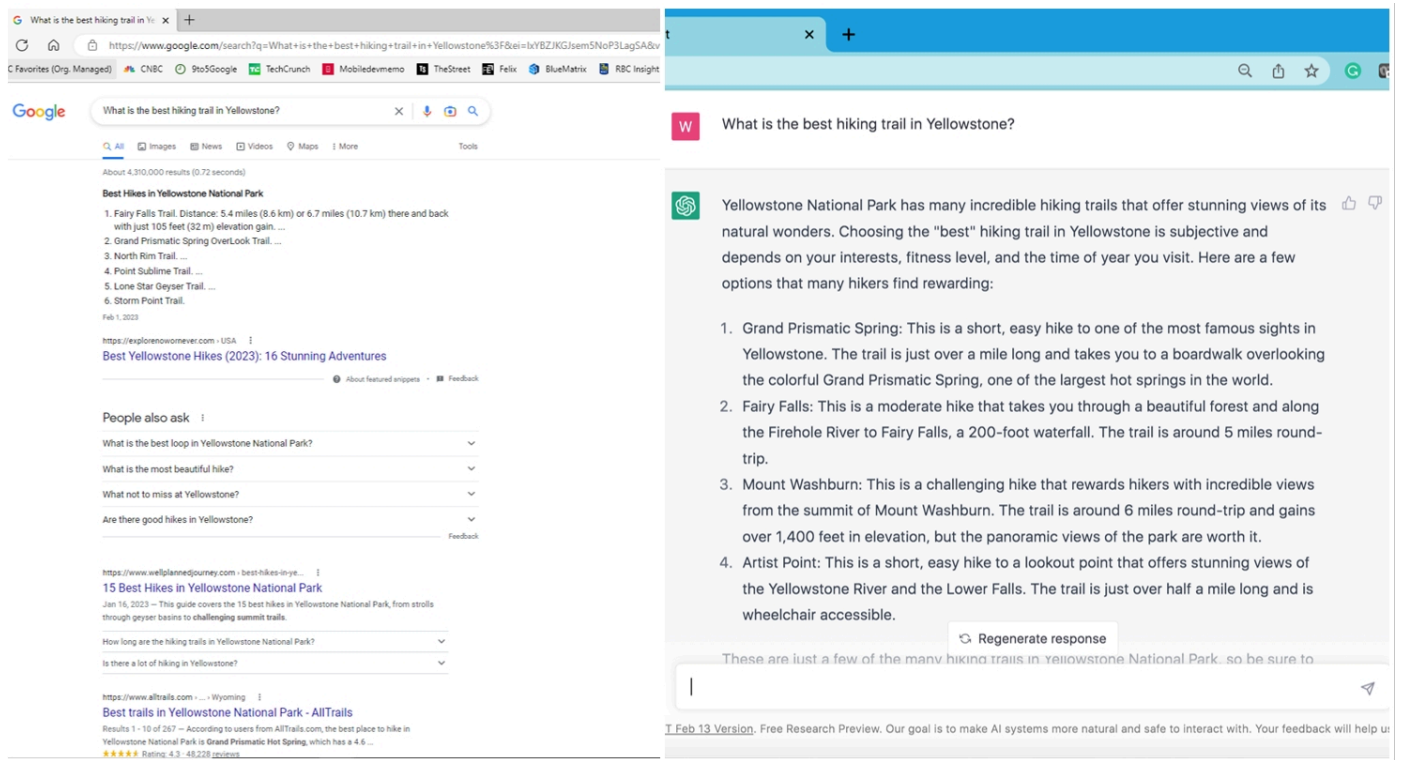
Learn more: [1. bing.com](#) [2. bing.com](#) [3. theworldbucketlist.com](#) +3 more 1 of 15 ●

[?](#) [Tell me more about Marchesi di Barolo.](#) [How can I book a tour at Winery Fattoria Cantagallo?](#) [What are some other wine regions in Italy?](#)

[New topic](#)

Source: Bing, RBC Capital Markets

Figure 62 - Google Search vs. ChatGPT - Hiking in Yellowstone



The figure compares the results of a Google search and a ChatGPT response for the query "What is the best hiking trail in Yellowstone?".

Google Search Results:

- Search query: "What is the best hiking trail in Yellowstone?"
- Results: About 4,310,000 results (0.72 seconds)
- Section: **Best Hikes in Yellowstone National Park**
- List of 6 trails:
 1. Fairy Falls Trail. Distance: 5.4 miles (8.6 km) or 6.7 miles (10.7 km) there and back with just 105 feet (32 m) elevation gain. ...
 2. Grand Prismatic Spring Overlook Trail. ...
 3. North Rim Trail. ...
 4. Point Sublime Trail. ...
 5. Lone Star Geyser Trail. ...
 6. Storm Point Trail.
- Source: <https://explorenowornever.com> - USA
- Section: **Best Yellowstone Hikes (2023): 16 Stunning Adventures**
- Section: **People also ask**
 - What is the best loop in Yellowstone National Park?
 - What is the most beautiful hike?
 - What not to miss at Yellowstone?
 - Are there good hikes in Yellowstone?
- Source: <https://www.wellplannedjourney.com> - best-hikes-in-ye...
- Section: **15 Best Hikes in Yellowstone National Park**
- Date: Jan 16, 2023 — This guide covers the 15 best hikes in Yellowstone National Park, from strolls through geyser basins to challenging summit trails.
- Section: **How long are the hiking trails in Yellowstone National Park?**
- Section: **Is there a lot of hiking in Yellowstone?**
- Source: <https://www.alltrails.com> - Wyoming
- Section: **Best trails in Yellowstone National Park - AllTrails**
- Results: 1 - 10 of 267 — According to users from AllTrails.com, the best place to hike in Yellowstone National Park is **Grand Prismatic Hot Spring**, which has a 4.6 ...
- Rating: 4.3 - 48,228 reviews

ChatGPT Response:

- Query: "What is the best hiking trail in Yellowstone?"
- Response: "Yellowstone National Park has many incredible hiking trails that offer stunning views of its natural wonders. Choosing the 'best' hiking trail in Yellowstone is subjective and depends on your interests, fitness level, and the time of year you visit. Here are a few options that many hikers find rewarding:"
- List of 4 trails:
 1. Grand Prismatic Spring: This is a short, easy hike to one of the most famous sights in Yellowstone. The trail is just over a mile long and takes you to a boardwalk overlooking the colorful Grand Prismatic Spring, one of the largest hot springs in the world.
 2. Fairy Falls: This is a moderate hike that takes you through a beautiful forest and along the Firehole River to Fairy Falls, a 200-foot waterfall. The trail is around 5 miles round-trip.
 3. Mount Washburn: This is a challenging hike that rewards hikers with incredible views from the summit of Mount Washburn. The trail is around 6 miles round-trip and gains over 1,400 feet in elevation, but the panoramic views of the park are worth it.
 4. Artist Point: This is a short, easy hike to a lookout point that offers stunning views of the Yellowstone River and the Lower Falls. The trail is just over half a mile long and is wheelchair accessible.
- Buttons: "Regenerate response"
- Text: "These are just a few of the many hiking trails in Yellowstone National Park, so be sure to..."
- Footer: "Feb 13 Version. Free Research Preview. Our goal is to make AI systems more natural and safe to interact with. Your feedback will help us..."

Source: Google Search, ChatGPT, RBC Capital Markets



Figure 63 - Bing - Hiking in Yellowstone

Microsoft Bing SEARCH CHAT

Creative Balanced Precise

What is the best hiking trail in Yellowstone?

✓ Searching for: **best hiking trail in Yellowstone**

✓ Generating answers for you...

Hello, this is Bing. I'm happy to help you find the best hiking trail in Yellowstone National Park. 😊

There are many great trails to choose from, depending on your preferences and abilities. Some of the most recommended hikes are:

- **Mount Washburn Trail**¹: a moderate hike of 6.4 miles round trip with stunning views of the park and wildlife sightings.
- **North Rim Trail**²: an easy hike of 7.6 miles out and back along the Grand Canyon of Yellowstone with waterfalls and canyon vistas.
- **Beaver Pond Trail**³: a family-friendly loop trail of 5.4 miles starting from Mammoth Hot Springs with beaver ponds, meadows and forests.
- **Grand Prismatic Overlook**⁴: a short but steep hike of 1.2 miles round trip to see the largest hot spring in the park from above.

What kind of hike are you looking for?

Learn more: [1. a-z-animals.com](#) [2. theplanetd.com](#) [3. muchbetteradventures.com](#) +3 more 1 of 15

? I want a challenging hike. I want a scenic hike. I want a short hike.

New topic Ask me anything...

Source: Bing, RBC Capital Markets