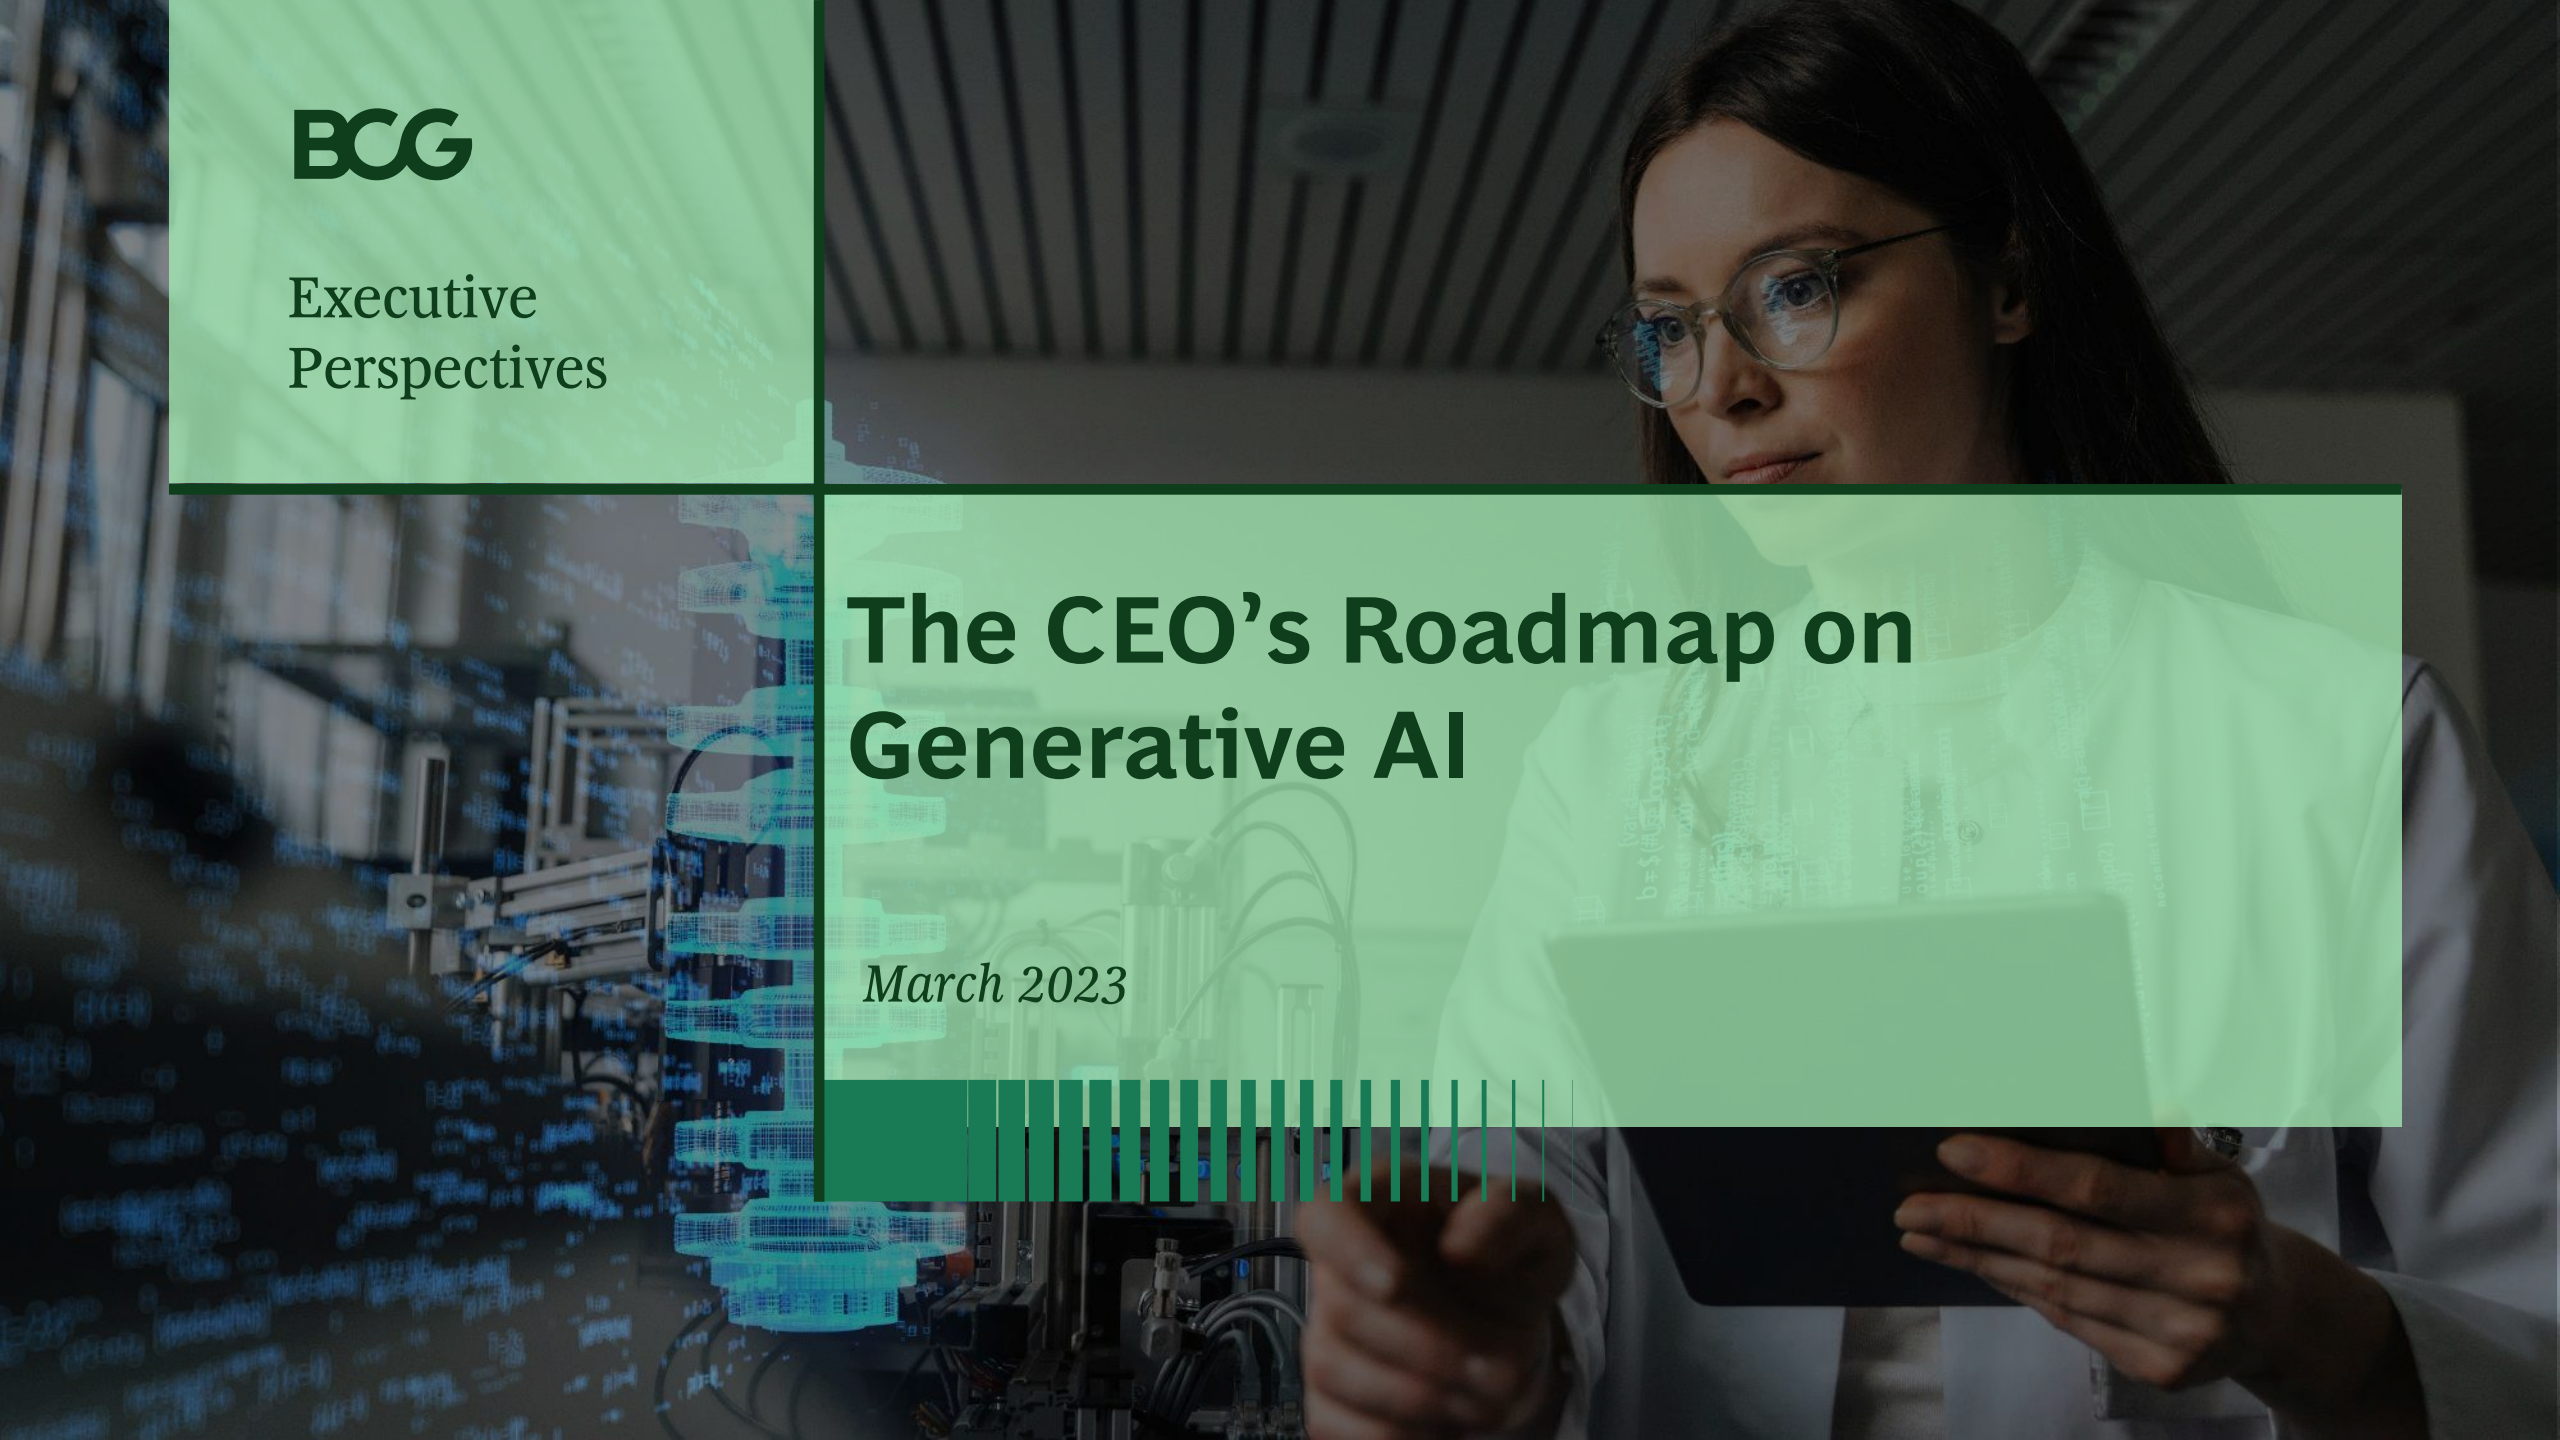

BCG

Executive Perspectives

The CEO's Roadmap on Generative AI

March 2023

Introduction to this document

The release of ChatGPT in late 2022 is analogous to Mosaic's launch three decades prior. In 1993, it was clear that the internet would bring a major revolution across all businesses in less than a decade.

The most focused of business models, and the strongest of brands can be blown to bits by new information technology

– Philip Evans, in his book "Blown to Bits"

Similarly, it is clear that Generative AI will bring another major revolution across all businesses. Today companies are focused on productivity gains and technical limitations, but CEOs need to move the focus to business model innovation.

This is no small task, and CEOs—who are likely several steps removed from the technology itself—may feel they are at a crossroads. But from our perspective, the priority for CEOs is not to be fully immersed in the technology. It is to understand how Generative AI will impact their organization and their industry, and what strategic choices will enable them to exploit opportunities and manage challenges.

In this BCG Executive Perspectives edition, we explore how CEOs can take full advantage of the coming revolution with Generative AI

Human-AI augmentation of the future

Prior to Traditional ML	• Use cases focus around automation • Humans as passive recipients of technology tools • Humans as operators of processes	**Focus on standardization and routinization to reduce costs and replace human effort**
With Traditional ML	• Use cases around making decisions with data • Humans actively using technology with data • Humans as operators of processes	**Focus on augmenting decision making to create most efficient systems and processes**
With Generative AI/Foundation Models	• Use cases around augmenting human creativity • Humans supervising AI on first drafts • Humans as designers of content and auditors of AI • Making decisions based on statistics and sequencing	**Focus on enabling greater productivity and creativity, to solve unsolved problems / Might augment decision making in some cases**

CEOs don't need to understand the technology behind Generative AI to create business model innovation; instead, they need to understand its key features

No Code / Low Code

With a convenient chatbot-like interface, Generative AI democratizes access for all including those not well versed in tech. "English is the hottest new programming language" according to Andrej Karpathy[1]

"Infinite Memory"

Generative AI, trained on vast amounts of data, offers users access to an automated system that provides seemingly infinite memory and acts as a knowledgeable personal aide[2]

Lack of Truth Function

As a probabilistic model, Generative AI generates the most likely output to a query. This can sometimes create hallucinations i.e., outputs completely separated from objective truth

Defining features that will drive **Business Model Innovation**

NOTE: 1. Andrej Karpathy is a premier computer scientist and one of the founding members of OpenAI. 2. While not technically infinite, GPT-3 was trained on ~500 billion tokens, which gives users an impression of an "infinite memory" database

Executive Summary | CEOs must make choices across three key pillars

POTENTIAL

Which use cases will differentiate your organization?

1

Discover your strategic advantage through experimentation

a. Generative AI is accelerating across every industry, **it is time to act now** or be left behind
b. Use cases that rely on existing large language model (LLM) applications will be important to stay competitive, but they won't offer differentiation – CEOs need to **discover the company's golden use case**
c. When use cases are identified plan the right implementation approach: **fine-tune or train**
d. **Plan for long-term advantage** through investment in talent and infrastructure

PEOPLE

How should CEOs adapt org structures and prepare employees for deployment?

2

Prepare your workforce with strategic workforce planning and transforming op models

a. CEOs will need to **address key org questions** for change management, talent and operating models
b. Generative AI will **redefine roles and responsibilities** across the organization
c. As AI adoption accelerates, CEOs need to **develop a strategic workforce plan**
d. CEOs will need to **consider new operating models,** however we expect that agile (or bionic) models will remain the most effective and scalable in the long term

POLICIES

How will the company ensure ethical guardrails and legal protections are in place?

3

Protect your business with clear policies that address the limitations of Generative AI

a. **Generative AI presents critical risks** for which companies will need to be prepared
b. **Prepare for risk** through clear policies and training that define roles and responsibilities on how to use Generative AI with a measure of confidence
c. CEOs should ensure the organization **adapts responsible AI norms** for long term risk mitigation

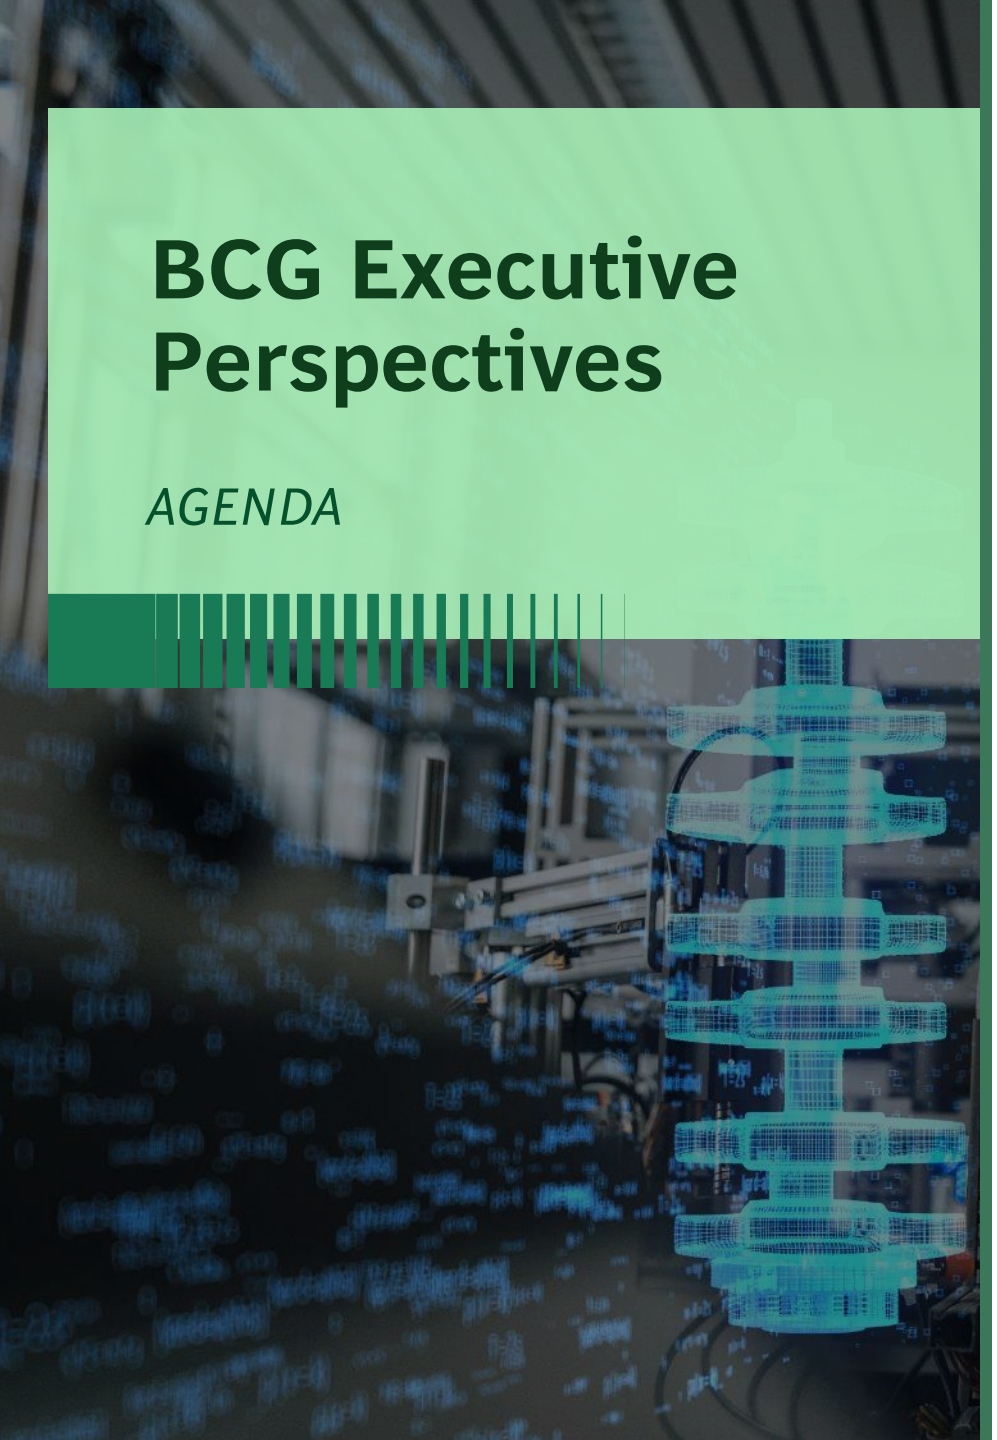

BCG Executive Perspectives

AGENDA

- **Potential: Discover your strategic advantage**
- People: Prepare your workforce
- Policies: Protect your business

Interest in Generative AI is exploding, fueled by the launch of ChatGPT

1a | It is time to act now

Interest in Generative AI has grown exponentially since Q4 2022

Google Search Interest (100 = max interest)

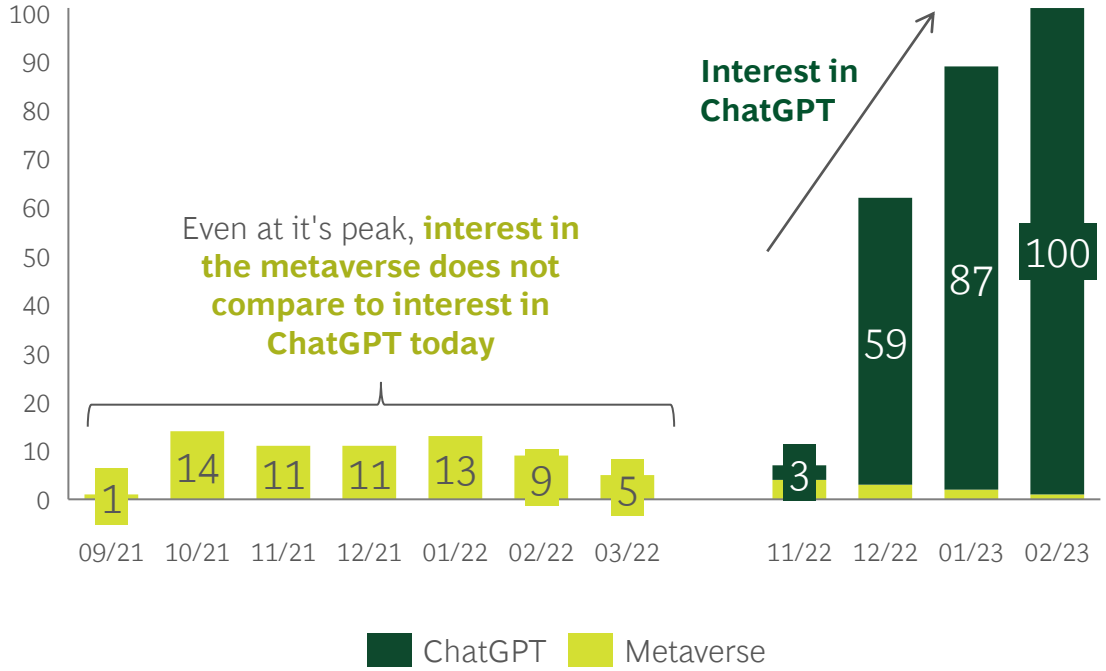

This is driven by the release of ChatGPT, which has taken the world by storm

INSIDE VIEW

Opinion: Can ChatGPT Write This Column?

Not yet, which points to the big question about any technology: Can it scale?

By Andy Kessler January 22, 2023 03:32 pm ET

Appeared in the Jan 23, 2023 , print edition as 'Can ChatGPT Write This Column?'

Wall Street Journal

TECH · CHATGPT

ChatGPT passed a Wharton MBA exam and it's still in its infancy. One professor is sounding the alarm

BY STEVE MOLLMAN

January 21, 2023 at 4:44 PM PST

Fortune

Robotics & AI

OpenAI begins piloting ChatGPT Professional, a premium version of its viral chatbot

OpenAI this week signaled it'll soon begin charging for ChatGPT, its viral AI-powered chatbot that can write essays, emails, poems and even computer code. In an announcement on the company□...

Kyle Wiggers

2:09 AM PST • January 11, 2023

TechCrunch

..and many more

Companies are already seeing a transformative effect from using Generative AI

1a | It is time to act now

Technology

~88%

Of **software developers** reported higher productivity when using a generative AI code assistant[1]

Consumer

Automated on-model **fashion image generation** resulted in

1.5X

Increase in retailer conversion rate[2]

Biopharma

Generative AI Identified a **novel drug candidate** for the treatment of Idiopathic Pulmonary Fibrosis in

21 days

(vs. years with traditional methods)[3]

Financial Institutions

Synthetic GAN-enhance training set for fraud detection achieved a

~98%

accuracy rate

(vs. 97% with unprocessed original data)[4]

Entertainment

Generate novel animated motions from a single training motion sequence with

~97.2%

quality score on natural movements

(vs. 84.6% with traditional methods)[5]

Insurance

InsureTech platforms leveraging generative AI to reduce up to

~30%

of customer service costs[6]

1. Quantifying GitHub Copilot's impact on developer productivity and happiness 2. Vue.ai helps fashion retailers create high-quality on-model product photos 3.Deep learning enables rapid identification of potent DDR1 kinase inhibitors 4. Using generative adversarial networks for improving classification effectiveness in credit card fraud detection 5. GANimator: Neural Motion Synthesis from a Single Sequence 6. Insurtech COVU Leverages OpenAI to Streamline Insurance Agency Operations

Total addressable market is expected to reach ~$120B by 2027

1a | It is time to act now

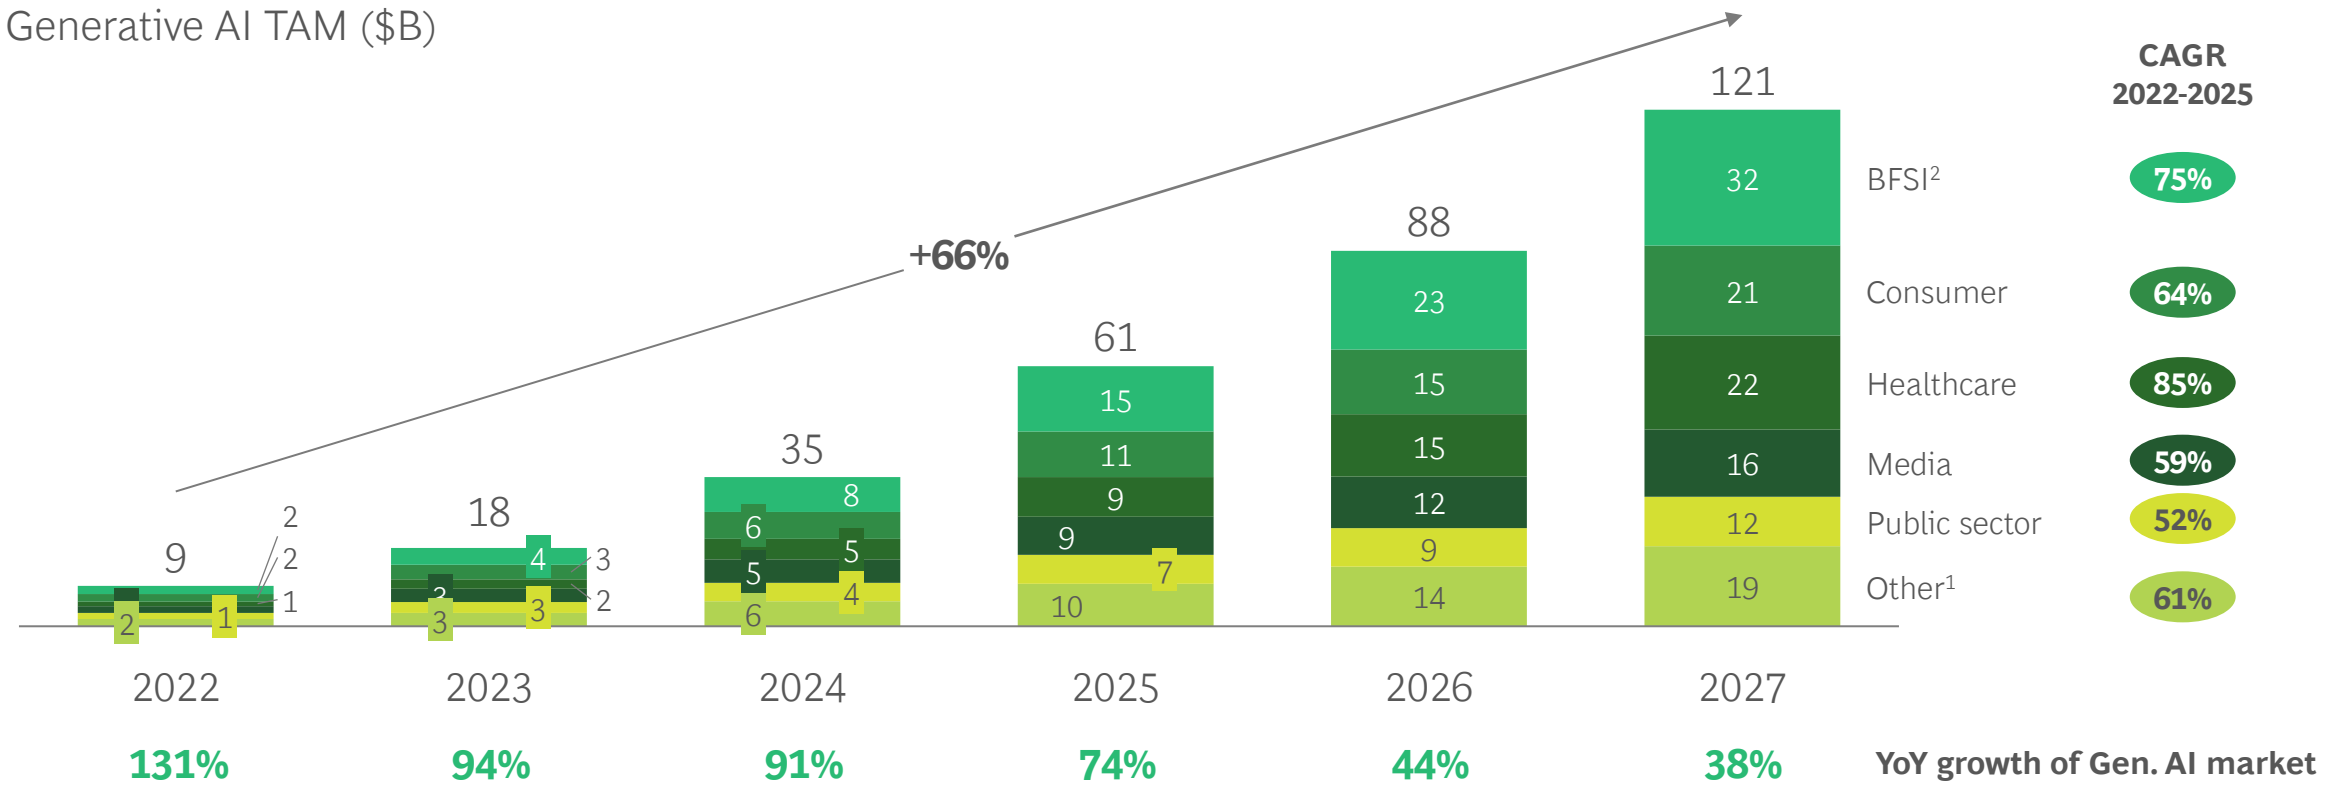

NOTE: 1. Other includes Industrial Goods, Energy, and Telecom markets; 2. BFSI includes Insurance (~$2B 2025) and Financial Institutions (~$13B 2025) including retail and wholesale banking, asset and wealth management, and private equity
Source: AI TAM research; Expert interviews; BCG analysis

For the CEO, the key is to identify the company's so-called golden use cases that drive competitive advantage

Non-core use cases are table stakes, everyone will adopt them

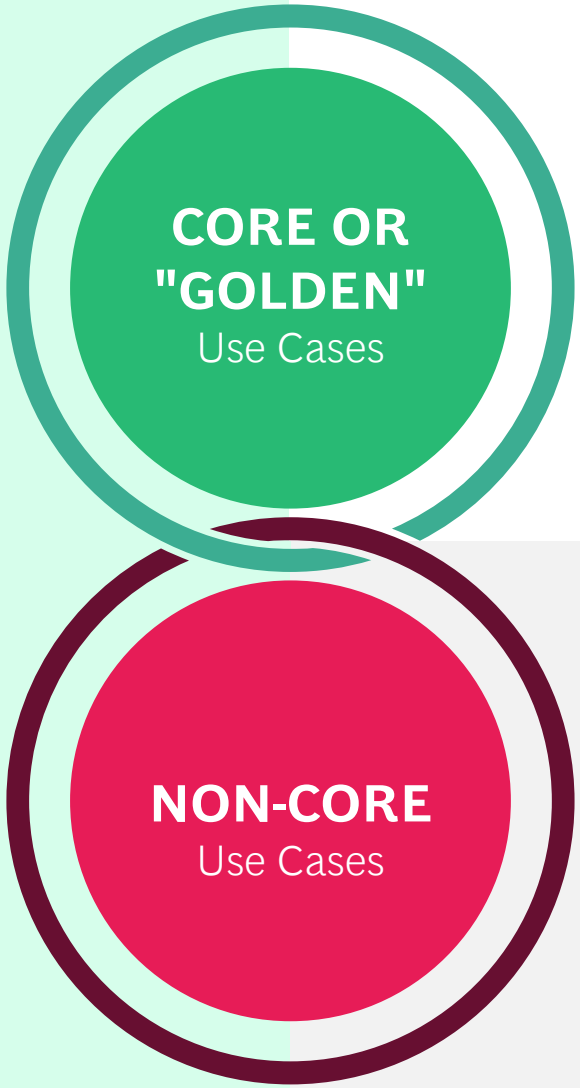

Focus on Building Long-Term Advantage

Strengthen competitive positioning with truly unique use cases that both drive value and are challenging to adopt (i.e., have a barrier to entry for competitors)

- For example, in pharmaceuticals companies, Generative AI can drive core R&D to produce new drugs/molecules at record pace

There is low barrier to adopting use cases that rely on existing LLM applications, but they will be important to keep pace with other organizations

- For example, purchasing Generative AI tools that create automatic summaries of meeting notes

Table-Stakes to use cases will improve efficiency

Golden use cases will add to a company's unique competitive advantage in the marketplace, while non-core use cases are readily adopted by all

1b | Discover the company's golden use case

Productivity Gains

First drafts
with Jasper AI

What is Jasper Doing?
Web-based application for businesses powered by Generative AI that helps teams create tailored content up to 10x faster

How is Jasper Doing it?
Built a model on top of OpenAI's GPT-3, fine-tuned on 50+ use-cases such as writing, copyediting, advertising, and content creation

Why is generative AI better vs traditional ML?
Traditional ML incapable for such a task. It does not have any "generative" capabilities for new text adapted to use-case

Efficiency Gains

Predictive maintenance
with an Equipment Manufacturer

What is the Equipment Manufacturer Doing?
Building proof-of-concept for global end-to-end predictive maintenance of fleet with IoT sensors powered by Generative AI

How is the Equipment Manufacturer Doing it?
IoT sensors constantly monitor key indications of performance through signals from parts, and relay that information back to a Generative AI powered back-end software

Why is generative AI better vs traditional ML?
Identification of anomalies in sensor data is difficult since failure data is rare in real-world. Generative AI can generate synthetic data, and better predict failures before occurrence

Innovation

Building novel proteins
with ProFluent

What is ProFluent Doing?
Creating novel proteins that do not exist in nature, aimed at advancing drug treatment. Proof-of-concept shown with creation of novel proteins with anti-microbial properties

How is ProFluent Doing it?
Using "inverse design", i.e., working backwards from desired properties to create proteins. Gartner believes that by 2030, 30% of new drugs will be discovered using this method

Why is generative AI better vs traditional ML?
Similar to Jasper, traditional ML does not have "generative" capabilities and thus is not great at creating never before seen protein structures by self-learning from training dataset

Non-core use case

Productivity improvement will be table stakes since all businesses will adopt

Golden use cases

For the equipment manufacturer, high quality of maintenance is a core part of their business model. Similarly for ProFluent, protein synthesis is at the heart of their business. Generative AI strengthens competitive positioning for both companies in their core business activities

While foundation models today are used for generative use cases, this may expand to include discriminative use cases as well in the future

1b | Discover the company's golden use case

Not exhaustive

Discriminative uses of AI

Currently in domain of Traditional ML

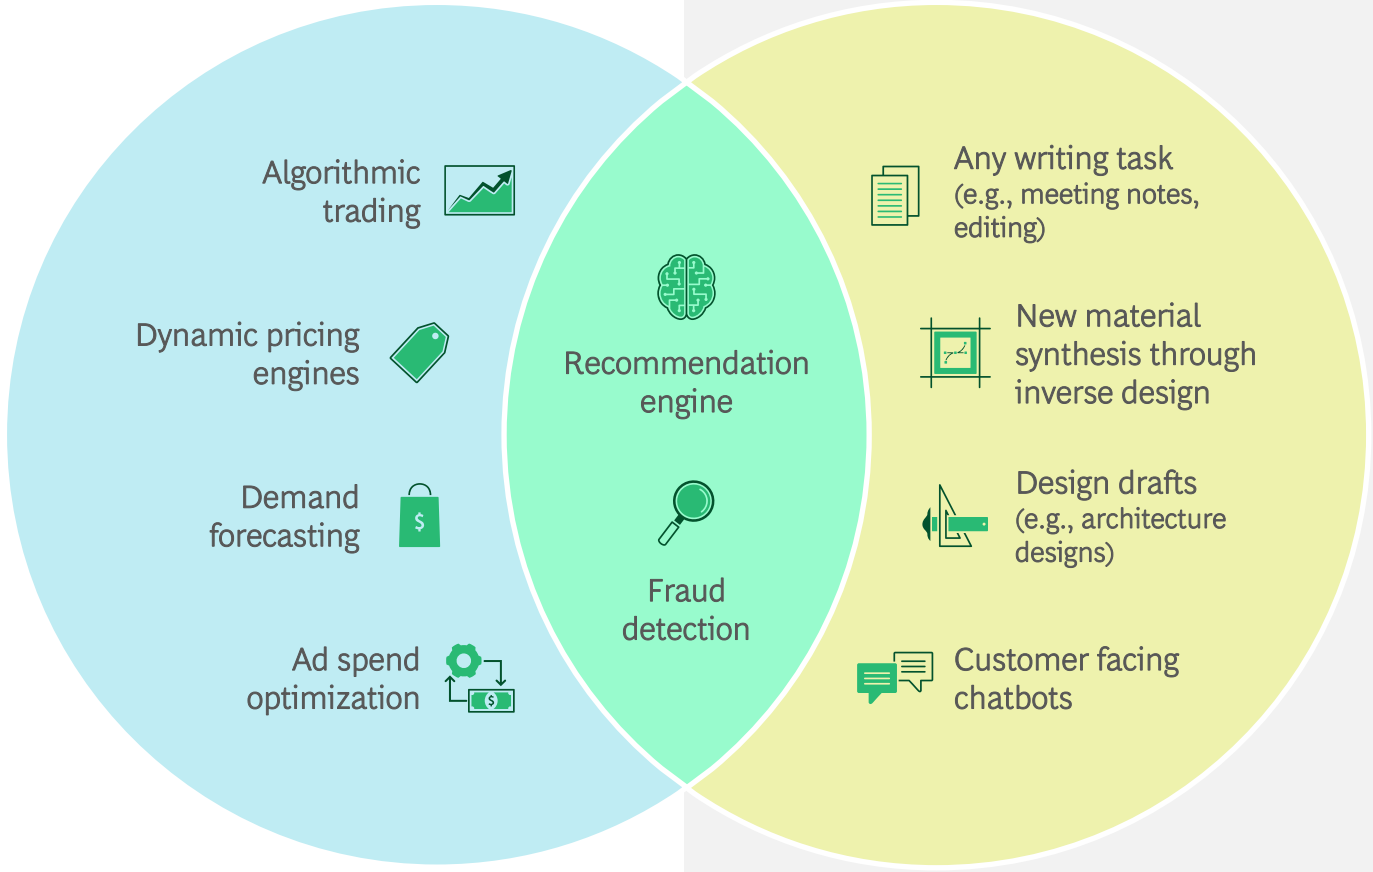

Generative uses of AI

Currently in domain of Foundation models

Foundation models are currently being used for **generative use cases;** however, this may expand in the future to cover certain **discriminative use cases** as well

Once use cases are selected, CEOs should make strategic choices about whether to fine-tune existing LLMs or to train a custom model

1c | Fine tune or train

Decision Tree for Foundation Model Choice

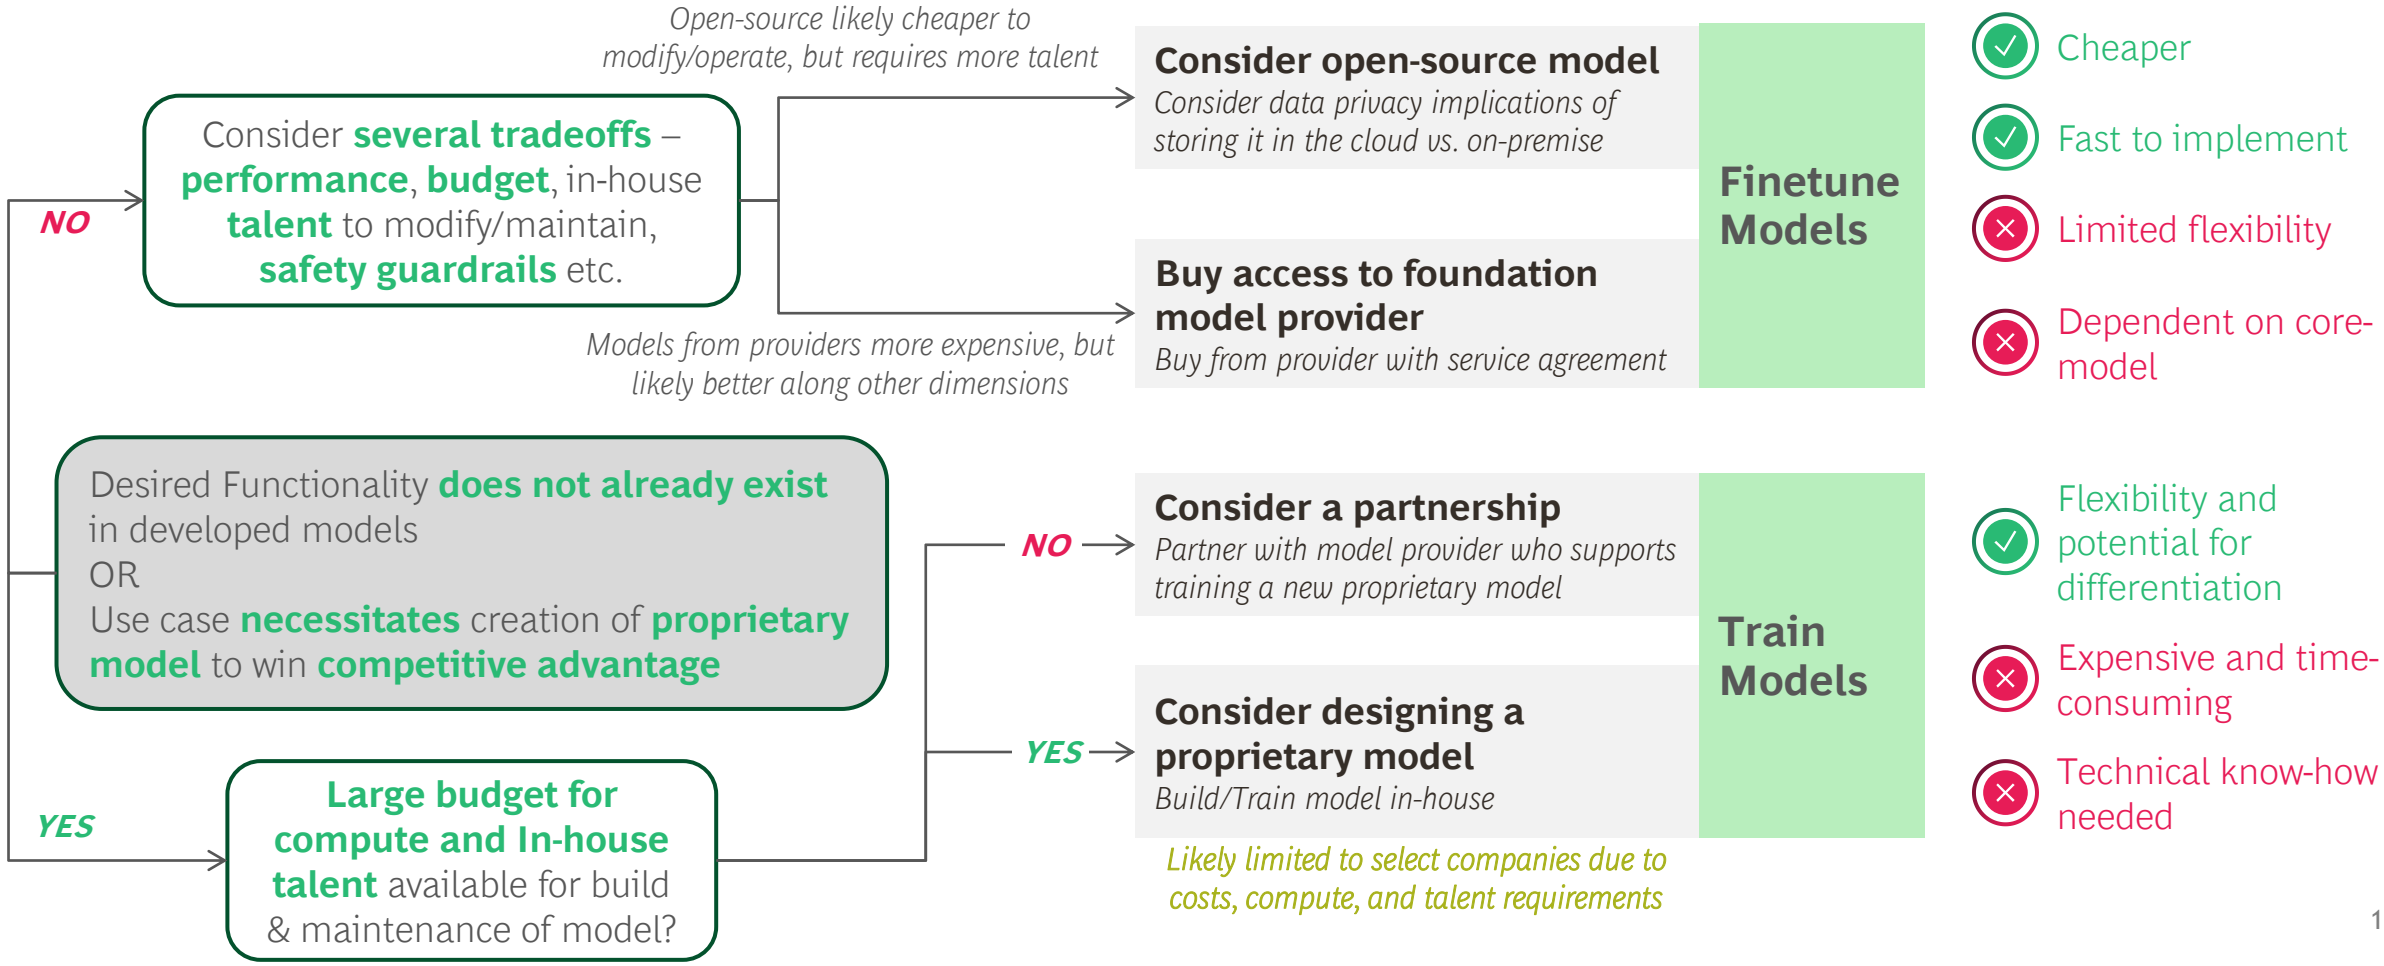

Training a custom LLM will offer greater flexibility, but that comes with high costs and capability requirements

1c | Fine tune or train

1 | Develop New, Cutting-edge foundation model

Create a new foundation model in-house from scratch. Costs scale with model complexity

$50 - $90M+

Estimated cost for complex models

Main drivers of cost:

- Hardware (i.e. GPUs or TPUs): $30M[1]
- Training runs: $10M+[2]
- People and R&D costs: variable

2 | Enhance Existing foundation model

Partner with LLM provider to significantly enhance existing model (e.g., feeding complex company-proprietary data)

$1 - $10M

Estimated cost

Main drivers of cost:

- Training runs: $1M -$5M[3]
- Partnership costs: variable

3 | Fine-tune Existing foundation model

Fine-tune existing foundation model for related tasks (e.g., fine-tuning ChatGPT for legal memo writing)

$10 - $100k+

Estimated cost

Main drivers of cost:

- Data gathering and labelling: $10k+[4]
- Computational costs: minimal

Usage Costs – $7M to $15M yearly *(costs 30x to 50x lower if not using the most advanced model)*

GPT4 costs $0.06 for ~750 words. 5k to 10k employees each using the technology 100 times a day costs ~$7M to $15M

1. Meta's LLaMA used 2048 A100 GPUs for training, each of which can cost ~$20k. See https://wandb.ai/vincenttu/blog_posts/reports/Meta-AI-Released-LLaMA--VmlldzozNjM5MTAz?galleryTag=ml-news. 2. A single training run for GPT-3 is projected to cost $12M. See https://venturebeat.com/ai/ai-machine-learning-openai-gpt-3-size-isnt-everything/. 3. Training runs here likely less intensive than full-scale model training, leading to lower costs.

Carefully assess the timing of Generative AI investments considering tech and talent; move too soon and risk wasting money, too slow and risk falling behind

1d | Plan for long-term advantage

It could take 5+ years for low error tolerance use cases to be feasible[1] …

Error tolerance

Key metric to evaluate readiness of Generative AI is the error tolerance of chosen use-case

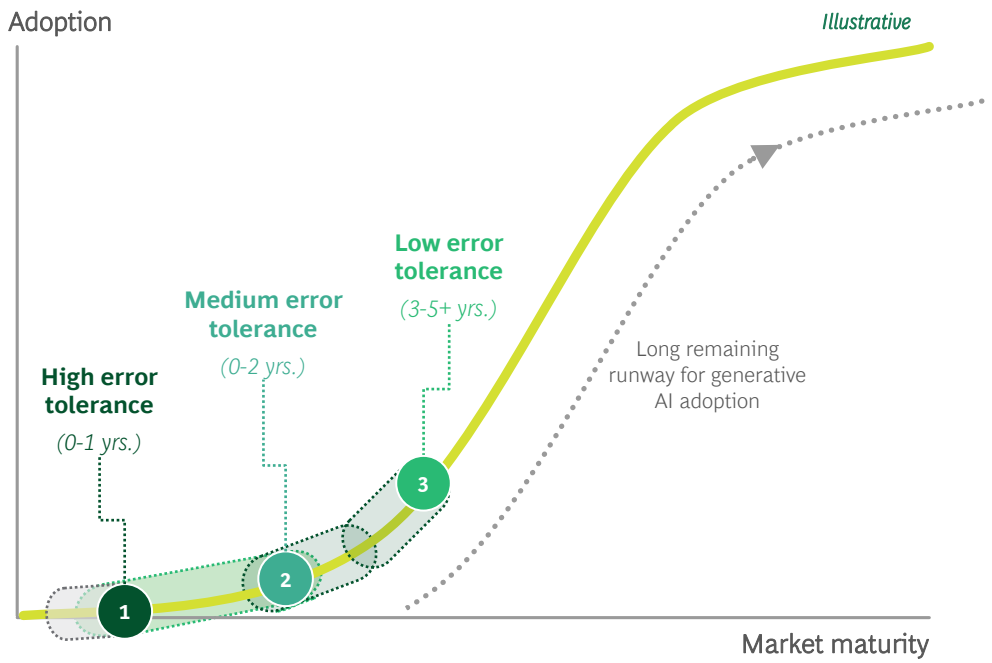

1 Near-term | **HIGH error tolerance**

Use cases where errors are OK

- e.g., drug development since scientists review every molecule suggested by AI for safety and efficacy

3 Longer-term | **LOW error tolerance**

Use cases with low room for error

- e.g., doctors using chatbots to retrieve and query a patient's medical history for easy access

…but research is becoming proprietary

Open-source: OpenAI's GPT-2

Proprietary: OpenAI's GPT-3; Meta's LLaMA

Research is also moving very quickly:

Meta's LLaMA released 2/24/23, outperforming GPT-3 on many tasks

GPT-4 released on 3/14/23

Waiting too long to invest into Generative AI today may mean that businesses risk falling behind. Research into high-performing foundation models is increasingly proprietary and guarded as a source of competitive advantage.

1. Sequoia expects first drafts produced by Generative AI in certain domains to be better than human professionals by 2030
See https://www.sequoiacap.com/article/generative-ai-a-creative-new-world/

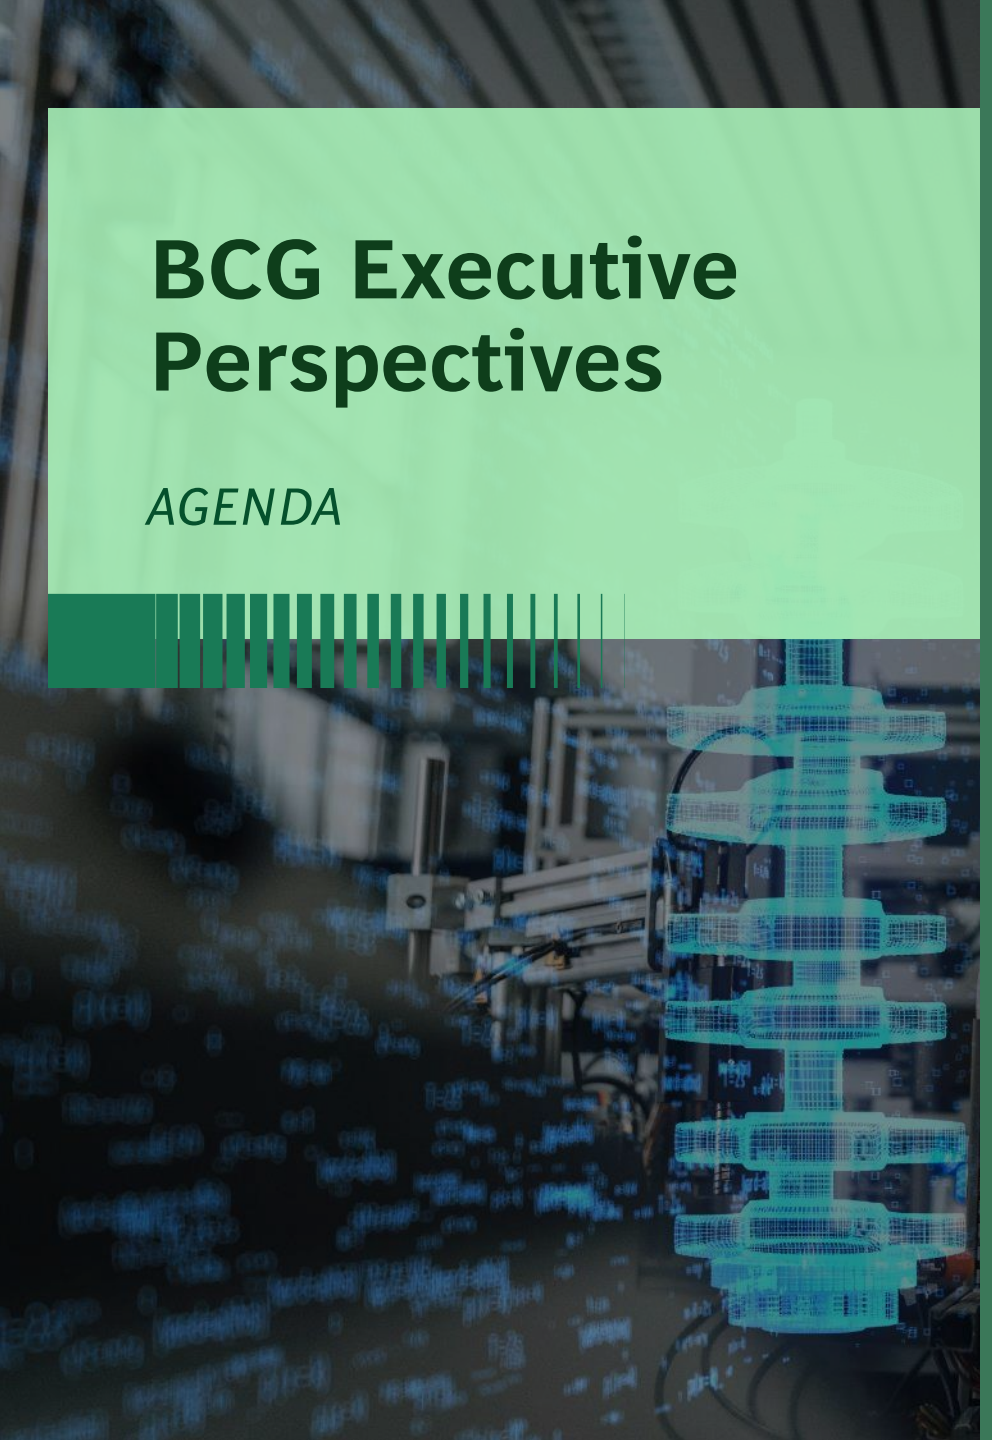

BCG Executive Perspectives

AGENDA

- Potential: Discover your strategic advantage
- **People: Prepare your workforce**
- Policies: Protect your business

To achieve the Human-AI augmentation of the future, CEOs should answer questions for change management, workforce planning, and op model design

2a | Address key organizational questions

Key considerations to craft a Generative AI adoption plan

	Managing Culture and Change in Company	Strategic Workforce Plan	Organization and Operating Model Design
Overarching Goal	*Cultivate a culture that embraces AI like another coworker*	*Build a workforce that will be competitive 10 years from now*	*Create an efficient operating model that balances scale and agility*
Key Questions Addressed:	• How can professional identity concerns be managed to encourage AI adoption? • How can a culture of human and AI collaboration be fostered? • How can management communication create positive momentum	• What new skills and talent will be crucial for long-term advantage? • What new competencies will managers need to lead an AI-augmented workforce? • How should training/recruiting be adjusted to build a high-performing workforce?	• What existing roles and responsibilities will change because of Generative AI? • How should I organize my departments for efficient collaboration with AI • Where should LLMs and data scientists sit within the organization?

A successful Generative AI adoption plan is **customized to each organization**, driven by the **industry** the company operates in, its current **AI readiness**, and the **golden use cases** it selects

While traditional AI has augmented the capabilities of managers and decision makers, Generative AI will augment the capabilities of individual contributors

2b | Redefine roles and responsibilities

Traditional AI/ML empowers individuals to make decisions, changing the role of managers

Traditional AI and ML algorithms augments decision making

Lower-level individuals can now make data-driven decisions without management support

This changes the role of the manager from decision maker to a manager of teaming and relationship dynamics

- For e.g., at ExxonMobil, geoscientists use ML algorithms to decide where and how to extract oil at maximum efficiency with limited guidance of managers

VS.

Generative AI creates first draft content, changing the role of individual contributors

Generative AI augments content creation

Individuals will spend less time creating first-drafts and more time revising or supervising AI generated content

This changes job tasks of individual contributors to include auditor or supervisor of Generative AI

- For e.g., Andrej Karpathy, a founding member of OpenAI, said "Copilot has dramatically accelerated my coding… I don't even really code [anymore], I prompt & edit"

> *"This is the first time that a technology developed in Silicon Valley benefits the lives of everyday people so quickly and so tangibly*
>
> – Satya Nadella, CEO of Microsoft

FROM: Key roles today

A role centered around creating marketing content and executing campaigns

- Creating marketing content and ideas from scratch
- Managing social media accounts, scheduling and uploading posts
- Writing creative briefs to interface with advertisement agencies
- Tracking ad campaign performance metrics
- Creating brand guidelines to drive alignment across all stakeholders

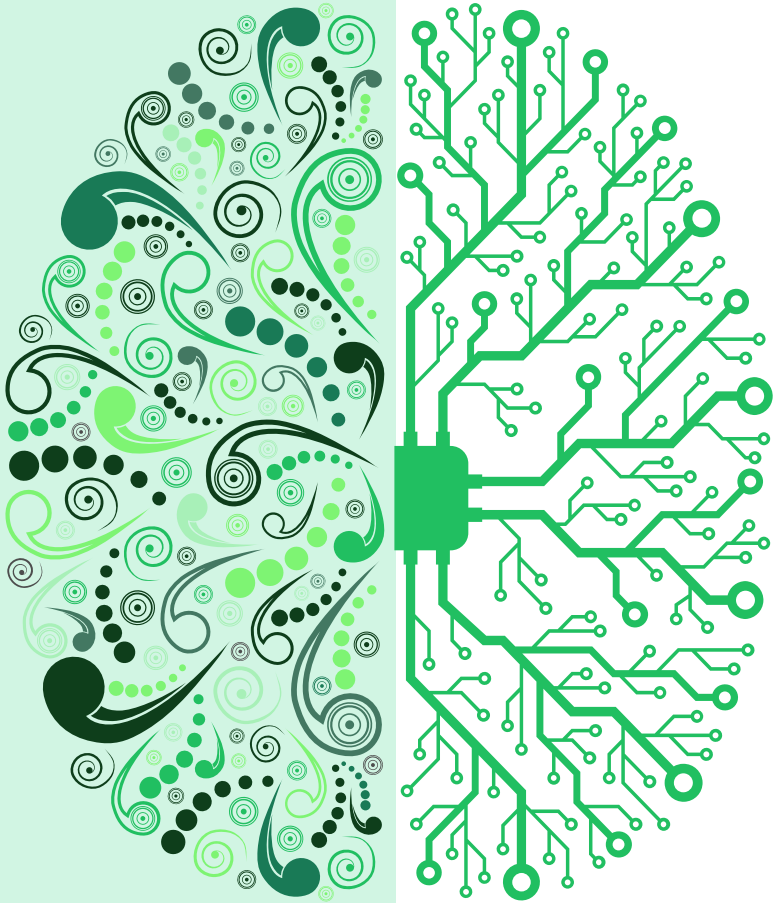

TO: New roles tomorrow

A supervisor role with AI on content, with increased time devoted to strategic thinking

- Supervising AI for first drafts of creative briefs and brand guidelines and overall better and faster marketing content
- Building deeper relationships with customers, suppliers, and brand ambassadors
- Increased focus on brand strategy, positioning, and target audience identification
- Increased focus on personalized marketing campaigns using Generative AI-powered tools

Core role changes for a marketer

2b | Redefine roles and responsibilities

Generative AI will redefine roles across the organization

Carefully consider the professional identity of your employees when making changes to role definitions

Marketing

Finance

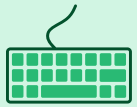

IT

Sales

Sample roles	Tasks today that Generative AI can provide first drafts for	Future tasks *(in addition to verifying first drafts)*
Social Media Specialist	Creating social media content, scheduling and uploading posts	Building relationships with customers and followers
Advertisers	Developing creative material (e.g., videos)	Exploring new advertising channels and opportunities
Accountant	Preparing and maintaining financial accounts	Identifying and implementing new accounting policies and programs
Payroll Specialist	Processing employee payroll and taxes	Ensuring compliance with labor laws and regulations, providing guidance and support to employees
Software Engineers	Low-value coding and debugging, code translation	Reviewing code safety, designing new complex algorithms (e.g., better recommendation engines)
Help Desk Support	Troubleshooting common issues	Resolving system-wide problems, supporting complex technical issues
Sales Rep	Lead generation, follow-ups, logging customer interactions in CRM systems	Build relationships with customers, understand their needs and pain-points
Deals Desk Support	Log quotes, and request sales approvals	Develop complex pricing models, customized deals for customers

Employees are expressing concern about the impact to their professional identity

IDEAS

How ChatGPT Will Destabilize White-Collar Work

No technology in modern memory has caused mass job loss among highly educated workers. Will generative AI be an exception?

By Annie Lowrey

The Atlantic

TIME Magazine

ARTIFICIAL INTELLIGENCE

How Generative AI Will Change All Knowledge Work

SUCCESS · CHATGPT

Some companies are already replacing workers with ChatGPT, despite warnings it shouldn't be relied on for 'anything important'

BY TREY WILLIAMS

February 25, 2023 at 6:00 AM PST

Fortune

To successfully adopt Generative AI, CEOs must alleviate these concerns

Work with HR to understand how roles will evolve and regularly pulse check employee sentiment as their AI initiatives roll out

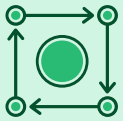

Develop a transparent change management initiative that will both help employees embrace their new AI coworkers and ensure employees retain autonomy

---

While some roles will be adversely impacted by Generative AI, overall Humans aren't going anywhere — and in fact are needed to deploy AI effectively and ethically

As Generative AI adoption accelerates, CEOs need to use their learnings to develop a strategic workforce plan

2c | Develop a strategic workforce plan

DEVELOP

Upskill and reskill talent at speed with high reach and high richness

- What key skills will be needed to work effectively with Generative AI?
- What training programs can upskill the workforce at speed?

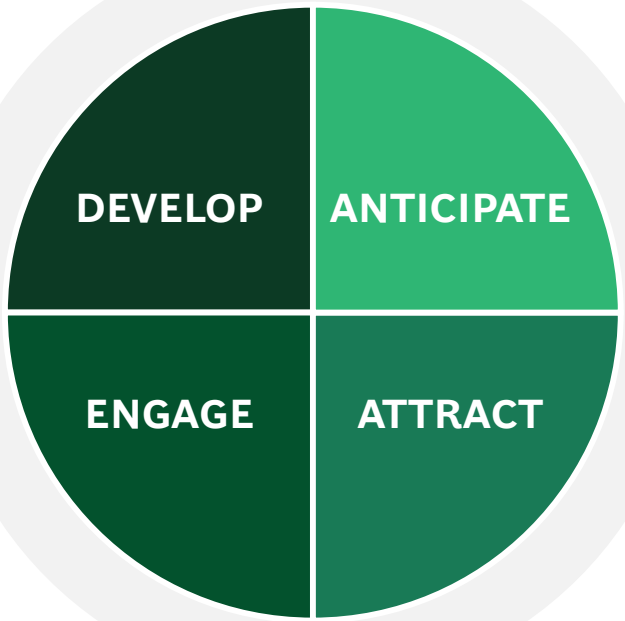

ANTICIPATE

Understand talent and skills needed to deliver on business strategy

- What workforce changes are needed as the company steadily adopts Generative AI?
- What are the risks associated with workforce changes, and how to mitigate them?

ENGAGE

Deliver unmatched talent value proposition and experience

- How to create a culture of continuous learning and development that encourages employees to use Generative AI?
- What is the company's value proposition to employees in a Generative AI world?

ATTRACT

Source creatively securing best-in-class candidate experience

- How should the interviewing process change to surface the talents needed in a Generative AI dominated world?
- How should the sourcing process change to ensure candidates with new skillsets are attracted to the company?

Consider centralizing the IT/R&D function supplying LLMs and data engineers

2d | Consider new operating models

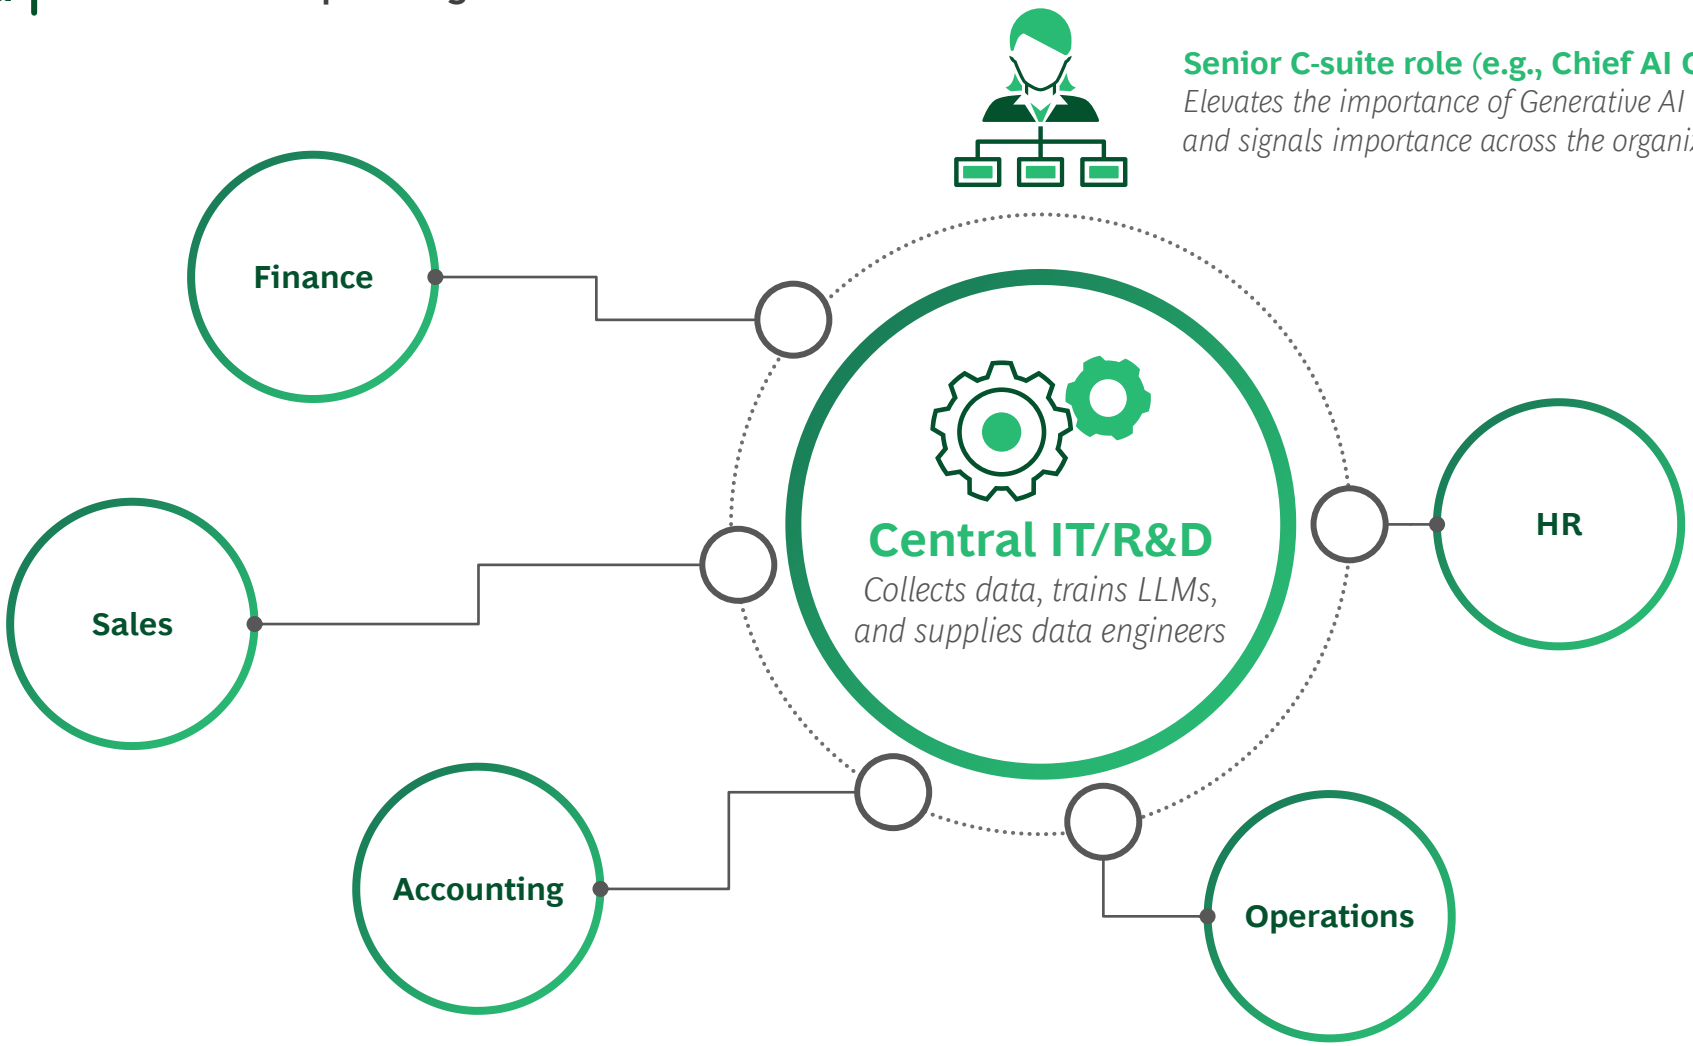

Each functional department **interfaces with the Central IT/R&D** to:

- **Supply** all collected **data** for model training
- **Embed data scientists** within their departments to build functional expertise
- **Request data engineers** to fine-tune LLMs for specific use-cases

This creates a scalable model with a central authority for data ownership and model control

We expect that agile (or platform) models will remain the most effective and scalable in the long term

2d | Consider new operating models

Sample model for a platform organization

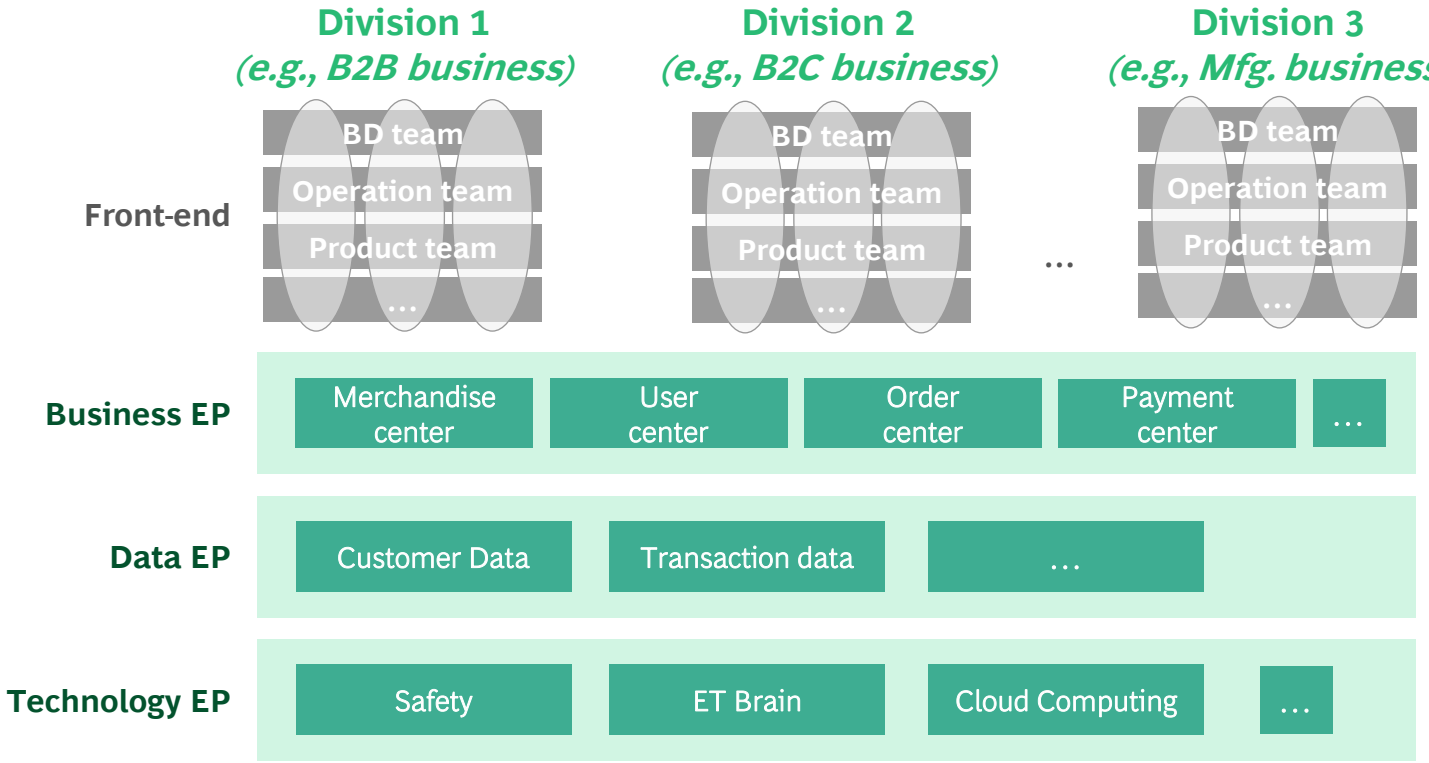

Decentralized
Front end teams have autonomy to serve customers

Scalable
Processes are identified and scaled to serve front end teams and to learn

Flexible
Technology allows for personalization and localization, to create the pull

Integrated
One source of all data and information

Responsive
Modular technology available to all

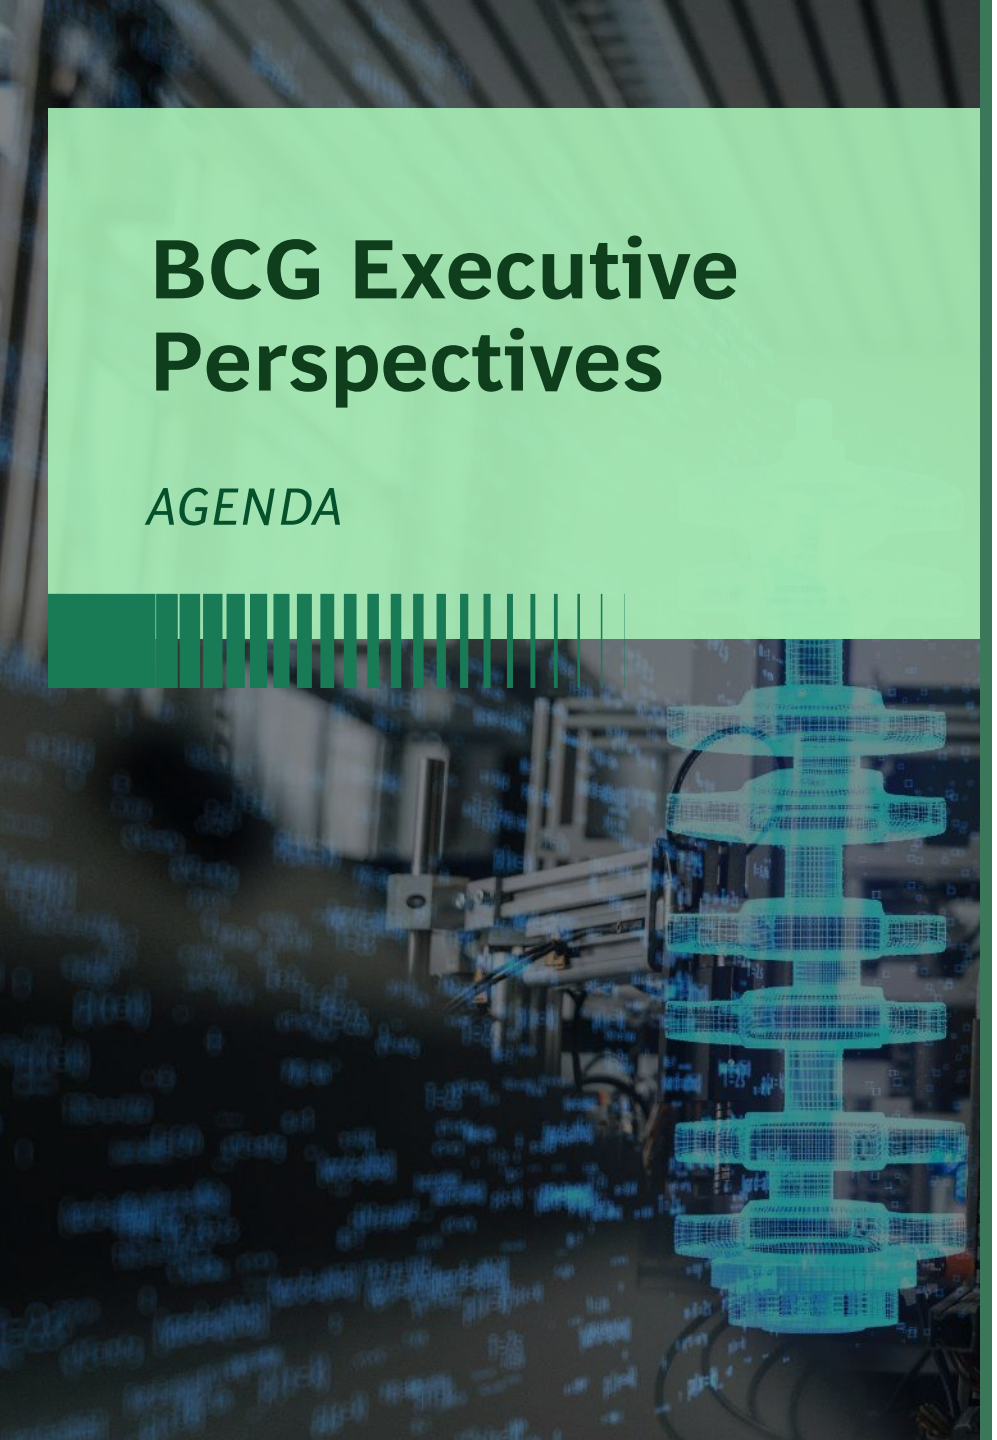

BCG Executive Perspectives

AGENDA

- Potential: Discover your strategic advantage
- People: Prepare your workforce
- **Policies: Protect your business**

Risks associated with Generative AI are showing up in the real world rapidly

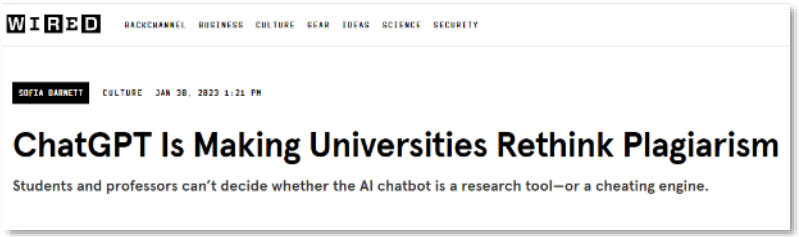

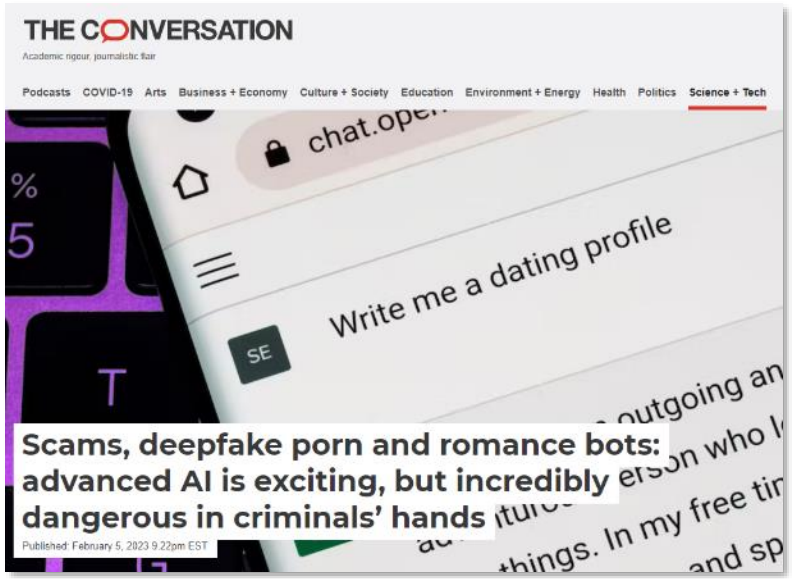

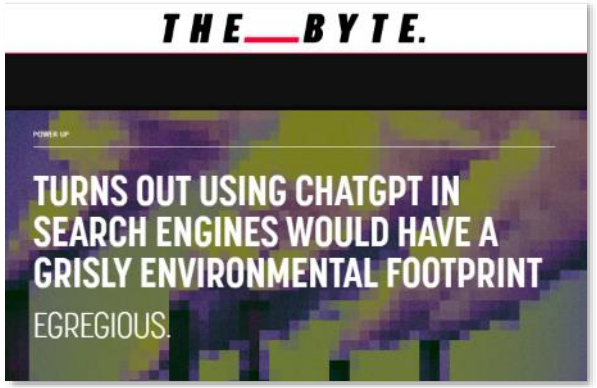

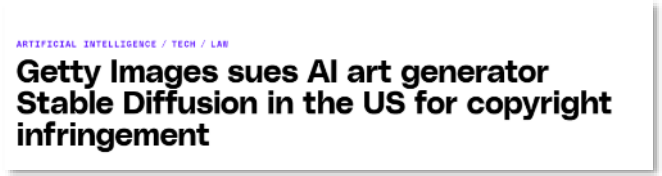

GenAI creates fundamental shifts impacting the Responsible AI (RAI) approach

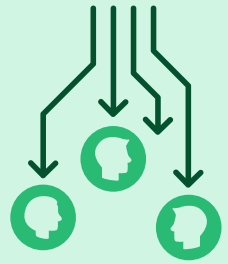

Democratization

Ease of use is much higher now:

- Anybody (even non-technical staff) can use these capabilities with **very few technical resources** (e.g., data, compute, expertise)
- Smaller need for large teams and budgets **limiting visibility** for managers and governance mechanisms

Shadow AI will be on steroids:

- Capability overhang can **emerge in unexpected corners** of the organization (e.g., compared to only technical divisions before)
- Time to detect, resolve, and mitigate incidents is much higher **violating the principle of surprise aversion** in risk management

3PP Reliance

Buying / renting from 3PP:

- Foundation models require a lot of compute, data, and expertise and are **overwhelmingly procured** rather than built in-house
- **Small set of entities** can provide these foundation models

Latent and opaque risks outside of in-house scope:

- **Limited visibility** on data lineage (e.g., copyright infringement) and model training (e.g., using confidential information to upgrade models)
- **Limited control** on functionality changes on the technical roadmap

Companies must be wary of critical risks of Generative AI today before adopting the technology

3a | Generative AI presents critical risks

Not Exhaustive

Energy use and environmental harm

Generative AI uses more energy on compute, both during model training and usage than traditional ML. While more efficient computation techniques are being developed, mitigation today is limited to usage of more environmentally sustainable energy sources

Capability Overhang

Due to its probabilistic nature, Generative AI can sometimes show unexpected capabilities upon deployment (e.g., several users tricked ChatGPT and bypassed its security to access kernel model). This risk is difficult to fully mitigate, but extensive pre-launch testing will help

Biased Outputs

Real world data is often biased. Without oversight, the Generative AI models trained on this data also carry bias. Mitigation techniques include Reinforcement Learning with Human Feedback (RLHF) where the model is taught to be unbiased, yet this method is not perfect

Copyright Infringement

Generative AI is trained on publicly available data, much of which is copyright protected. This can lead to lawsuits by IP holders. Mitigation strategies rely heavily on foundation model providers to obey copyright laws, and for governments to create new laws for Generative AI

Lack of Truth Function

Generative AI can sometimes produce factually incorrect responses presented in a very convincing manner. To mitigate risks from using incorrect information, companies must mandate double checking all Generative AI outputs, and limiting its use to non-critical tasks today

Sophisticated Phishing and Fraud

Generative AI makes cybercrime easier – generating convincing phishing emails or deepfakes instantly. To mitigate this risk, companies must strengthen cybersecurity protocols, train employees on new safety risks, and consider deploying Generative AI themselves to to catch fraud

Leaks of proprietary data

When training Generative AI models in the cloud, companies transmit proprietary data which the data may be leaked in a security breach. To mitigate this risk, companies can instead choose to train models on-prem vs. cloud, although this necessitates other tradeoffs

Shadow AI

Employee application of external generative AI tools without adequate guidance or supervision creating risk and causing harm. To mitigate this risk, companies must create detailed and clear Generative AI use guidelines and policies

Not all risks are created equal, with some posing a higher business risk while also being harder to mitigate

3a | Generative AI presents critical risks

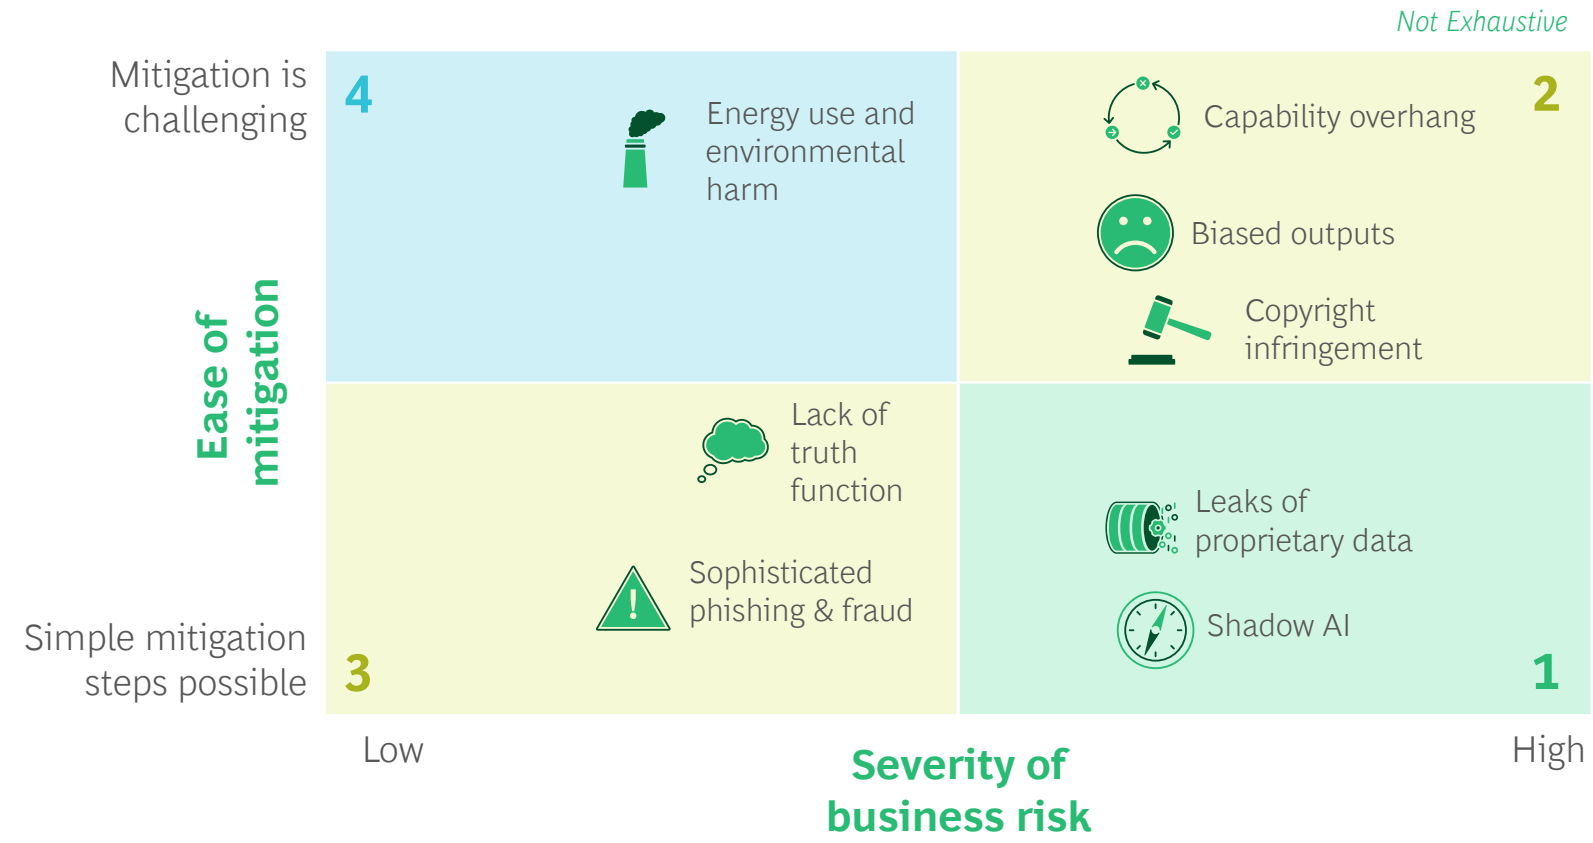

Source: BCG Henderson Institute

1 Companies should focus on risks with high severity of business impact, that are easily **mitigated first**

2 **3** Risks that are either difficult to mitigate, or that have low potential business impact can be **tackled second**

4 Risks with low business impact are toughest to mitigate, but nonetheless important to **tackle after low hanging challenges are addressed**

Timeline to potential resolution of issues varies based on complexity and urgency

3a | Generative AI presents critical risks

Short-term *Medium-term* *Long-term*

	1	2	3	4	5	6	7	8
	Shadow AI	**Sensitive data leaks**	**Enhanced Fraud**	**Lack of Truth Function**	**Copyright challenges**	**Biased outputs**	**Environmental Harm**	**Capability Overhang**
Description	More widespread and unconventional/ non-traditional uses of AI **invisible to top-down governance approaches**	Sensitive data can potentially be extracted from models via prompt engineering or more sophisticated abuse/attacks	Constantly evolving capabilities to **deceive and detect with higher stakes** over time from more pervasiveness	**Dynamically changing information ecosystem** makes it difficult to mitigate this issue without extensive human intervention	**Unclear data provenance** from publicly scraped datasets and quality issues	Generative AI models trained on real world data can carry over bias to outputs in hard to detect ways	SOTA achievable through larger models with more complex architectures with **higher computational needs**	Latent capabilities and how they might be leveraged depends on the **interacting ecosystem** and development of new systems
Potential Resolution	Clear guidelines captured in **enforceable policies accompanied by cultural change** and top-down communication	Sensitive data should be anonymized and/or encrypted Chained models can be used to block leakage of sensitive data	Better detection methods with **verifiable claims** on performance and reduced false negatives and false positives	Fact-checking services and **problematic information fingerprinting** can provide stop-gap measures to stem issues	**Legal precedents** will be set as courts figure out IP and copyright issues which will shape response and development approaches	Reinforcement Learning with Human Feedback (RLHF) can catch some biases but not all. Tech. advancement likely needed for comprehensive solve	**Research work in TinyML** with an emphasis on edge-device deployment, faster inference, and cost + environmental considerations	**Better governance mechanisms** and understanding of how capabilities emerge

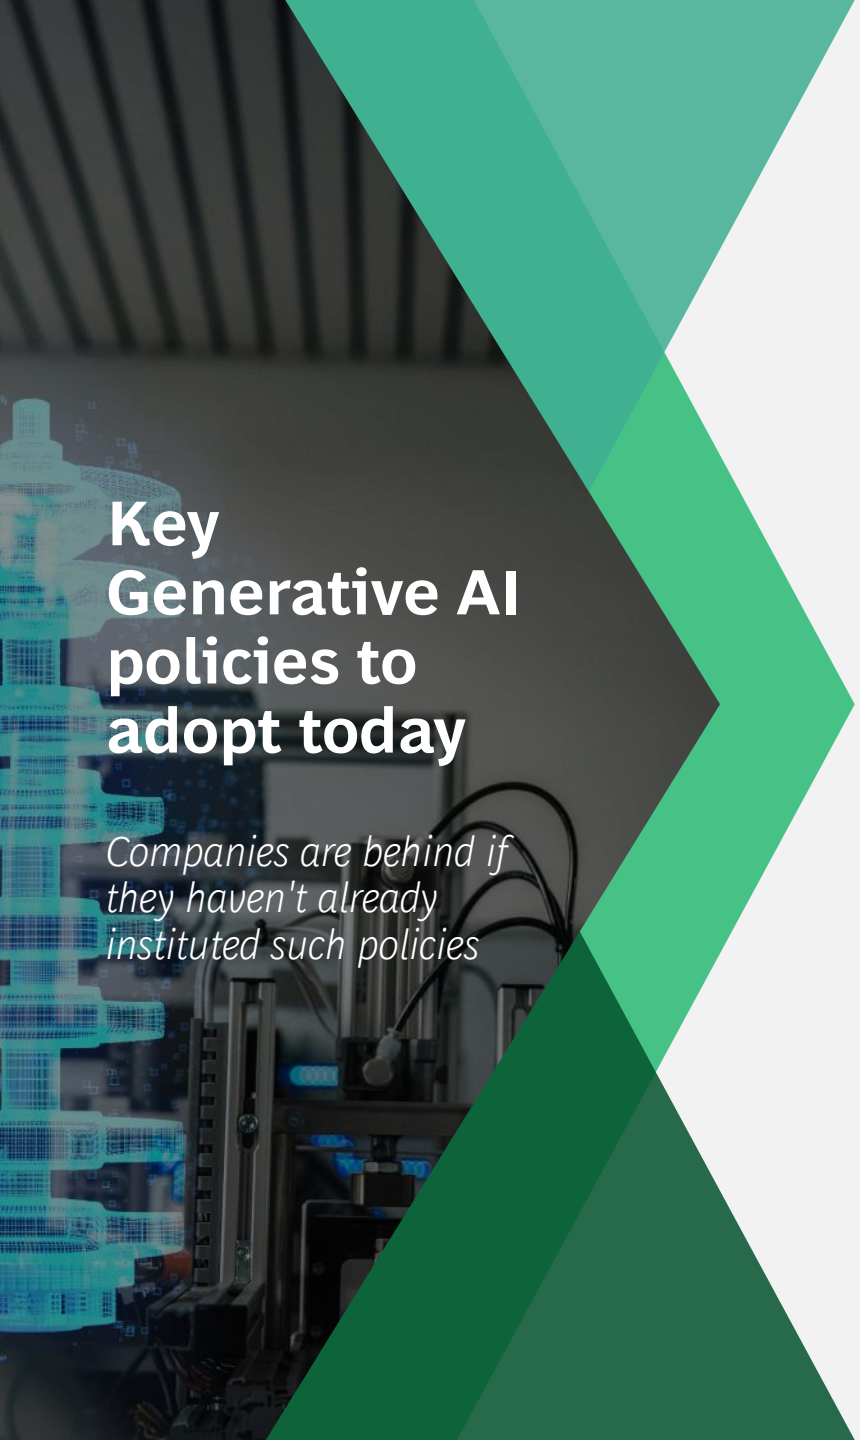

Key Generative AI policies to adopt today

Companies are behind if they haven't already instituted such policies

To manage risk with Generative AI, companies can adopt the following policies today:

Enact Responsible Research/Release Norms

- Like academia, set up an institutional review board to a priori assess impact of any Generative AI use cases

Set and Communicate Clear Generative AI Use Policies

- To manage IP and hallucination risks set clear guidelines on when Generative AI can and cannot be used

Sanitize Sensitive Data Before Training Models

- To minimize losses during data breaches, sanitize sensitive data (such as names and addresses) before training foundation models

Improve Generative AI Risk Assessment Capabilities

- Consider setting up a "red-team" to deliberately find failure models and vulnerabilities with Generative AI applications

Leaders will need to revisit these policies continuously, as the pace of innovation with Generative AI is high and produces new capabilities (and correspondingly new risks)

Connect with our Generative AI leadership team with any questions

Generative AI Team

Abhishek Gupta
Senior Solution Delivery Manager, RAI
Gupta.Abhishek@bcg.com
Montreal

David Martin
Managing Director & Partner
Martin.David@bcg.com
Dallas

Nicolas de Bellefonds
Managing Director & Senior Partner
deBellefonds.Nicolas@bcg.com
Paris

Akash Bhatia
Managing Director & Partner
Bhatia.Akash@bcg.com
San Francisco

François Candelon
Managing Director & Senior Partner
Candelon.Francois@bcg.com
Paris

Steve Mills
Managing Director & Partner
Mills.Steven@bcg.com
Washington D.C.

Amanda Luther
Managing Director & Partner
Luther.Amanda@bcg.com
Austin

Leonid Zhukov
Vice President – Data Science
Zhukov.Leonid@bcg.com
New York

Suchi Srinivasan
Managing Director & Partner
Srinivasan.Suchi@bcg.com
Seattle

Daniel Sack
Managing Director & Partner
Sack.Dan@bcg.com
Stockholm

Matthew Kropp
Managing Director & Partner
Kropp.Matthew@bcg.com
San Francisco – Bay Area

Click here for our latest insights on Generative AI

Disclaimer

The services and materials provided by Boston Consulting Group (BCG) are subject to BCG's Standard Terms (a copy of which is available upon request) or such other agreement as may have been previously executed by BCG. BCG does not provide legal, accounting, or tax advice. The Client is responsible for obtaining independent advice concerning these matters. This advice may affect the guidance given by BCG. Further, BCG has made no undertaking to update these materials after the date hereof, notwithstanding that such information may become outdated or inaccurate.

The materials contained in this presentation are designed for the sole use by the board of directors or senior management of the Client and solely for the limited purposes described in the presentation. The materials shall not be copied or given to any person or entity other than the Client ("Third Party") without the prior written consent of BCG. These materials serve only as the focus for discussion; they are incomplete without the accompanying oral commentary and may not be relied on as a stand-alone document. Further, Third Parties may not, and it is unreasonable for any Third Party to, rely on these materials for any purpose whatsoever. To the fullest extent permitted by law (and except to the extent otherwise agreed in a signed writing by BCG), BCG shall have no liability whatsoever to any Third Party, and any Third Party hereby waives any rights and claims it may have at any time against BCG with regard to the services, this presentation, or other materials, including the accuracy or completeness thereof. Receipt and review of this document shall be deemed agreement with and consideration for the foregoing.

BCG does not provide fairness opinions or valuations of market transactions, and these materials should not be relied on or construed as such. Further, the financial evaluations, projected market and financial information, and conclusions contained in these materials are based upon standard valuation methodologies, are not definitive forecasts, and are not guaranteed by BCG. BCG has used public and/or confidential data and assumptions provided to BCG by the Client. BCG has not independently verified the data and assumptions used in these analyses. Changes in the underlying data or operating assumptions will clearly impact the analyses and conclusions.


BCG