DE GRUYTER

ARTIFICIAL INTELLIGENCE

Rethinking Business Strategy



Teresa Martín-Retortillo

DE G

AI ESSENTIALS FOR LEADERS

Teresa Martín-Retortillo
Artificial Intelligence

AI Essentials for Leaders

Edited by
Alfred Essa and Teresa Martín-Retortillo

Teresa Martín-Retortillo

Artificial Intelligence

Rethinking Business Strategy

DE GRUYTER

Author's proceeds from the sale of this book are going to the not-for-profit organization Fundación Valle de Tena (Spain).

ISBN 978-3-11-158240-5
e-ISBN (PDF) 978-3-11-158349-5
e-ISBN (EPUB) 978-3-11-158380-8
ISSN 2944-2540

Library of Congress Control Number: 2026930339

Bibliographic information published by the Deutsche Nationalbibliothek
The Deutsche Nationalbibliothek lists this publication in the Deutsche Nationalbibliografie; detailed bibliographic data are available on the internet at http://dnb.dnb.de.

© 2026 Walter de Gruyter GmbH, Berlin/Boston, Genthiner Straße 13, 10785 Berlin
Cover image: Icon: bgblue/DigitalVision Vectors/Getty Images; Chess: Kyuubisa/AdobeStock #1075371925
Typesetting: Integra Software Services Pvt. Ltd.
Printing and binding: CPI books GmbH, Leck

www.degruyterbrill.com
Questions about General Product Safety Regulation:
productsafety@degruyterbrill.com

To the person who taught me the concept of a *“petit précis”* and to his wonderful wife.

To Marcos, with deep gratitude for his endless support.

Contents

Preface —— IX

Section 1: **Business Foresight in the Era of AI**

Chapter 1
Approach to Leading into the Future —— 3
Starting with “What Ifs” —— **3**
Bias for Action —— **7**

Chapter 2
Identifying Known-Knowns —— 9
Known-Knowns —— **9**
AI: The Queen in the Geopolitical Chess Board —— **9**
Infrastructure for AI will Continue to Improve —— **13**
Capital Will Not be Free —— **21**
AI Safety Will Prevail —— **23**

Chapter 3
Known-Unknowns: Navigating Polarities —— 27
Trust in AI —— **27**
Human Augmentation Versus Automation —— **29**
Role of Big Tech —— **33**
Artificial General Intelligence —— **39**

Chapter 4
Plausible Challenging Future Scenarios —— 41
The AI-Run Ecosystem —— **43**
Digital Feudalism —— **48**
The Augmented Enterprise —— **52**
The Personalized Tech-Mosaic —— **55**

Section 2: **Monday 8 a.m.**

Chapter 5
Dare Re-Coding Your Business Model — 65
Business Model Innovation — **65**
Dynamic Value Chains — **73**
AI as a Catalyst for Business Model Innovation — **77**

Chapter 6
Re-Design for Optionality, Not Perfection — 79
Optionality in Light of Growing Uncertainty — **79**
Managing the Tension of Investing in the Core Versus Creating Options Beyond the Core — **83**
Much Beyond the Core — **86**
Managing the Portfolio of Options — **87**

Chapter 7
Re-Wire your Operating Model for the AI Factory — 89
The Burning Platform — **89**
Building your AI Factory (Not to be Confused with Data Centers) — **92**
Break Silos, Focus on Decisions — **95**

Chapter 8
Learn, Learn, Re-Learn — 97
Setting your GPS — **97**
Demanded Skills — **98**
How to Learn: Developing and Training your Workforce — **104**
Engage and Mobilize the Team: There is Value in the Process — **105**

Conclusion — 107

List of Figures — 109

Notes — 111

Selected Bibliography — 117

Index — 121

Preface

Writing about AI in the 2020's feels like composing a symphony with a keyboard that keeps rearranging its keys. AI developments could take many different paths. Yet the future is not completely unpredictable.

My hope is that this book helps recalibrate your perspective on AI with more grounded expectations and fewer myths. Recent academic research[1] has found something counterintuitive: when we consume information in constant little bites (headlines, news feeds), our beliefs become less accurate than when we take it in all at once. Thus a short book can help better navigate the strategy questions arising from the recent AI revolution.

I have taken on the challenge of writing this book because I believe the unique perspectives I have gathered during my professional career of 30+ years may offer valuable insights for executives competing in the Era of AI. This book is aimed at C-suite leaders and practitioners responsible for driving their businesses and for shaping the business strategy of their companies. It is targeted mostly at non-tech leaders in resource-conscious organizations hungry to take advantage of moments of disruption to move their companies to enduring competitive positions. The book cateres to those readers frustrated with the literature that comes in the form of recipes for success from a variety of product and service vendors trying to push their own solutions. It aims to provide a deeper understanding of what is required to rethink business strategy in the Era of AI.

My sources of inspiration come from the intersection of four key roles I have had in my professional career as well as a deep-rooted passion. Starting with my passion: learning is the pillar that brings it all together for me. AI is nothing other than a new technology based on a learning process. Strategy and innovation have learning at the core; foresight itself can be characterized as a way to channel the curiosity to learn and prepare for the future.

The experience that I have relied upon in authoring this book stems from four key roles:

- advising corporate executives and investors on challenging strategic issues across the globe as a partner at one of the leading global strategy consulting firms.
- supporting global leaders in their leadership and technical development, as head of Executive Education and Lifelong Learning at a business school.
- being part of the C-suite driving digital transformation at a global publisher and at a global university.
- serving on the board of directors of both private and listed companies.

https://doi.org/10.1515/9783111583495-203

Stating my Biases

Every source of inspiration carries its own bias. I've done my best to correct for these with third-party references but I invite you, the reader, to keep this context in mind. I also need to be transparent about one of my own biases: I am an incurable optimist. I believe that every organization has the power to craft its own future—if it has the courage and discipline to do so.

That belief is not naïve; it is hard-earned. Over nearly 20 years in management consulting, advising leaders in corporations and private equity firms, I've seen how daunting disruption can be. As a junior consultant, I debated the fate of Altavista – one of the early and most powerful search tools – at Digital Equipment. I wrestled with how Kodak should respond to the rise of digital photography. Later, I helped clients design strategies to compete with Amazon in categories where the odds looked impossible. These experiences shaped how I see digital transformations: they demand both rigorous analysis and an imaginative leap into the art of the possible. Unfortunately, imagination and creativity are often in short supply in boardrooms—yet in the Era of AI, they may be the most vital skills of all.

My perspective has also been forged inside organizations, not just as an advisor. Serving on executive boards, I've helped design and implement digital transformations. I know the scars and the exhilaration that come from "rebuilding the plane while flying it." Transformations are no longer rare events—they are business as usual. The real art lies in managing a portfolio of change: some projects require structural, long-term shifts, while others look incremental but can prove transformational. A single new data source can open up new revenue streams. A subtle shift in the technology stack can demand a full redesign of business processes.

In recent years, my role as an educator—leading Lifelong Learning and Executive Education at a major European Business School with global reach—has widened my lens. From Europe to the United States, from Latin America to the Middle East and Asia, I've seen executives grapple with the same challenge: how to steer digital change when the pace of disruption outstrips the capacity of traditional playbooks. The lesson is clear: learning—both individual and organizational—must be seen as a multi-year journey, drawing equally on formal training and informal discovery. And above all, learning and experimenting needs to be in full alignment with the enterprise current and future value drivers.

Finally, as a board member involved in high-stakes decisions on acquisitions and technology investments, I've seen the power of clarity. A well-stated strategy and purpose, anchored in a deep understanding of the firm's capabilities, can tilt the balance between value creation and risk.

I bring these four perspectives together with one goal: to make your job as a leader a little easier. My hope is that an integrated view of the forces at play in the AI Era will help you separate noise from signal, sharpen your choices, and set a more confident direction for your enterprise.

Approach

Artificial Intelligence as a General-Purpose Technology

AI started in the 1950s with Alan Turing developing the first intelligent machine. AI became more prevalent in the business world in the early 2000s during the first digital revolution, which brought the unique combination of robust communications infrastructure, the Internet, and exponential data creation and exchanges. However, it was not until the late 2010s and early 2020s that we witnessed the large-scale adoption of AI in the business world.

I believe that it is worth authoring a book—and not only a blog or a post in social media—because artificial intelligence is likely to become a general-purpose technology (GPT).[2] AI is likely to be as transformational as electricity, or the Internet. Understanding and anticipating the implications of a new general-purpose technology is critical for business leaders.

What are the characteristics of a general-purpose technology? First, these types of technologies have an impact across many sectors in the economy. Like electricity or the Internet, AI is likely to transform your business whether you are in manufacturing or in services. Second, these technologies have ample room for improvement and further development. As we will see in this book, not only is the computing infrastructure improving every year, but the accuracy and breadth of models is in constant evolution. Finally, general-purpose technologies have the potential to foster the emergence of related new businesses. In the same way that electricity opened the door for a new industry in household appliances or the Internet on mobile devices opened the door for new businesses based on geolocation capabilities like hail rides, AI is opening the door for the development of entirely new businesses like AI assistants.

Despite the speed of innovation in AI, full deployment of its potential in business will take years, and so there is time for traditional companies to respond and pursue the upcoming opportunities. This response is likely to be even more powerful if it comes from a fundamental understanding of business strategy in the Era of AI—not just from piloting selected use cases in different areas of the business.

Focus on Superior and Sustainable Value Creation

A successful business strategy has often been defined by a set of choices that ensure a competitive advantage able to generate superior and sustainable value relative to peers. I will focus on how a company generates value but also how it captures it and distributes the value generated.

Historically, particularly in the late 20th century, this was commonly measured in terms of shareholder returns, that is, dividends and capital appreciation relative to initial capital invested. However, the definition of success has expanded to encompass broader stakeholder impact, including value delivered to employees, customers, communities, and the environment. Empirical evidence suggests that companies integrating both shareholder and stakeholder objectives can outperform on both financial and non-financial dimensions (e.g., the meta study of more than 2,000 academic papers by Friede, Busch and Bassen, 2015;[3] Serafeim, 2020;[4] Bain & Company, 2024[5]). This is particularly true if the so-called sustainability initiatives are well aligned with strategy and engrained in the core business. There is also general agreement that sustainability type efforts are a multiplier for an otherwise sound business performance; they can deepen the competitive moat but typically they do not offset lagging performance metrics in the core business standards such as Global Reporting Initiative (GRI) and Sustainability Accounting Standards Board (SASB), and integrated reporting frameworks are advancing toward cross-industry benchmarks, though significant variation remains across geographies and sectors. I have taken the point of view that shareholder returns need to match expectations within an industry, and only then can stakeholder capitalism be additive. Danone, the global dairy company, is a good reminder of what happens when you focus on sustainability initiatives at the expense of competitive performance: CEO Emmanuel Faber was ousted by the board of directors in 2021 at the request of activists because in addition to lagging financial performance to peers there was the perception of excessive focus on stakeholder well-being. Ideally, a company strives for both. In rethinking strategy in the Era of AI, I focus on strategies that can deliver competitive advantage while generating superior and sustainable value.

The Value of Business Foresight

It is difficult to predict the future of AI or even design key pillars of successful strategies in the Era of AI that would be valid for all sectors and companies. However, in addition to rigor and tested strategic principles, a combination of creativity, imagination, and ambition are needed. We would not be discussing AI if some business

leaders had not used creativity, imagination and ambition to create some of the AI products and services we have today. In an environment of growing dynamism but also uncertainty, looking creatively into the future is not a "nice to have," it is a "must have." Business foresight has become increasingly important, and I would argue is becoming a core component of strategy.

By foresight I do not mean trend watching, but the specific articulation of options that allow businesses to prepare for uncertain yet plausible and relevant futures. In and of themselves, these options can be a source of value creation, they can ensure agility without overcommitment, and they can allow experimentation while avoiding chaos. Creating a small portfolio of well-selected real options is the desirable outcome of a strategic foresight exercise looking into a horizon of +/- 10 years. Applying foresight to the Era of AI is at the core of this book.

How this Book is Organized

The book will cover two sections.

In the 2020s, the rules of business strategy are being rewritten in real time. The AI Era is still taking shape, and no one yet knows which formulas will guarantee success. That uncertainty, however, should not lead to paralysis—it's an invitation to use foresight. History shows that the companies that thrive in times of upheaval are those that imagine multiple futures, prepare for them, and build competitive moats that endure.

In this first section, I will lay out plausible, relevant, and deliberately challenging scenarios for the AI Era. Think of them as wind tunnels for your strategy: safe places to test how strong your business is under different futures. The process begins with the known-knowns—the elements we can be confident about. By naming them, we strip away false uncertainty and distraction. From there, we turn to the known-unknowns: the key polarities and tipping points that could disrupt incumbents and crown new winners. Spotting these polarities early—and building strategic options around them—is what separates survivors from casualties. Shell, the Dutch oil giant and the corporation that made strategic foresight a core capability, offers a famous example: in the 1970s, it used scenario planning to anticipate a collapse in oil demand, sparing billions in wasted refinery investments. In the same spirit, this section aims not to forecast the future but to prepare you for it. My goal is to stretch your imagination, sharpen your awareness, provide prompts that help you test—and stress-test—the resilience of your strategy and to cognitively prime you for relevant signs and prevent any distraction from noise.

The second section moves from non-linear futures, from imagination to action. Here, I describe the strategic superpowers every organization needs in the AI Era. These superpowers emerge by looking across scenarios, finding the common threads, and turning them into durable capabilities. Along the way, I'll give you prompts to ask your strategic co-pilot the questions that will make your plans future-proof. For context, remember how the Digital Era, beginning in about 2000, transformed business models. Subscription and platform models redefined economics: high fixed costs, near-zero marginal costs, powerful network effects, and the race to maximize lifetime customer value. Companies built moats not just with capital, but with data and user experience. Now, the AI Era challenges these foundations. Its economics are different and so must be your operating models. By critically reassessing how your business creates value, you can unlock entirely new opportunities—not just survive the disruption, but use it to your advantage.

I want this to be a provocative, energizing journey. Rather than overwhelm you with endless footnotes and references, I'll cite only the most essential sources. The focus is not on cataloguing research, but on sparking reflection and equipping you with practical tools to reimagine your business in the AI Era.

Additional Resources

The speed of innovation in this field, as well as the practitioner orientation, will require updating the conceptual framework used in the book. I will provide regular updates on the web page of the book (https://ai-dialogos.com) or on the social media channels I tend to use.[6]

Section 1: **Business Foresight in the Era of AI**

Chapter 1
Approach to Leading into the Future

Think back to your business a decade—or even two decades—ago. What are the shifts that, no matter how rigorous your strategic planning was, your organization simply didn't anticipate? Was it how marketing became algorithmic or how consumer trust began flowing toward digitally native brands? Was it about the geopolitical impact on your supply chains? Is it possible that as we enter the AI Era, other forces might blindside your business tomorrow?

Foresight isn't about prediction; it's about preparation. It's the structured capability to explore not just what is likely, but what is plausible, possible, and, as importantly, what is preferable.

As it becomes ever more difficult for business leaders to distinguish noise from relevant signals in everything related to AI, this chapter attempts to provide a framework to make sense of the relevant forces shaping business strategy in the Era of AI. Our ambition is to provide a better compass to navigate the uncertainty concerning AI and to help leaders better discover attractive opportunities for their businesses. Admittedly a second goal, only relevant for a subset of readers, is to channel the fear of missing out (FOMO) into energy dedicated to putting in place practical plans, to create tangible options that help the organization articulate critical choices under relevant, plausible, and challenging future scenarios. We start this exercise with two important premises:

- There are important uncertainties concerning the development of AI, the players, and the economics behind them. We need a useful framework to analyze and manage these uncertainties.
- Because of the speed of change, mobilizing organizations and resources would be just as important at setting the direction. The famous dilemma between strategy and execution has not yet gone away: excellent strategies poorly implemented continue to have little value.

Starting with "What Ifs"

Let's clear up an important point. As the title of the book states, the focus is about rethinking business strategy in the Era of AI; that is very different from defining a successful AI strategy. For companies that are not AI specialists, AI may be the enabler not the core. Yet these companies can be "AI first" in the same way that a couple of decades ago, Facebook—today Meta—chose to leapfrog web and be "mobile first." This choice yielded unprecedented market share

https://doi.org/10.1515/9783111583495-001

gains. At the core, Facebook was still about being a business-to-consumer (B2C) social and advertising platform. I believe in the value of anticipation and managing with foresight. We will use foresight approaches to help leaders pick up signals of change and act upon them before events trigger fundamental challenges or opportunities for business. Foresight can also be used to develop agility. Shell, the company that first embedded scenario planning into corporate strategic planning, famously used scenario planning to prepare for and adapt to oil shocks in 1972. As Arie de Geus, director of planning at Shell, puts it: "Managing internal change by foresight, rather than by crisis, is only possible if the change in the environment is seen on time."[7]

I believe that an undervalued skill of business leaders is their ability to look into the future to identify options and create tangible opportunities based on those options. There is a method to this madness. In using foresight approaches we leverage heavily into the thinking laid out by Peter Schwartz in *The Art of the Long View* (1991); his invitation to identify key drivers of uncertainty and map them into plausible, relevant, and challenging scenarios under which to design strategies that future proof the company under any scenario is the essence of Chapter 2. Let's unpack the meaning of plausible when doing foresight. Sometimes it is simpler to visualize it in the Futures Cone (as depicted in Figure 1) based on concepts of future thinking developed in the 1990s and 2000s.[8]

Imagine you have a weather radar. The possible weather conditions would include all types of weather ever experienced, from calm sunshine days to tornadoes and blizzards. Technically, all weather conditions are viable. They could conceivably happen without regard for likelihood; the only constraints would be the law of physics. The plausible futures are like forecasted storms based on current atmospheric models that could realistically happen. The constraints for plausible futures are defined by the current environment. Plausible futures are distinct from probable futures; the latter are the ones most models agree on, like "80 percent chance of rain tomorrow."

As a leader in your business, you will be most familiar with the projected future and the probable future. When thinking about the Era of AI, we will encourage you to think about plausible futures for your business. In other words, focus on what is scientifically and technically possible within the economic and sociopolitical environment you operate in.

Thinking about plausible scenarios is a fundamental change in mindset from what prevails in business today where we only approach uncertainty with a probability toolset and where the underlying assumption is that every move, every decision behaves according to a normal distribution. Probabilities are helpful for risk management, for prioritization of projects and resources, and for making tradeoffs. Probabilities have important limitations when managing strategic un-

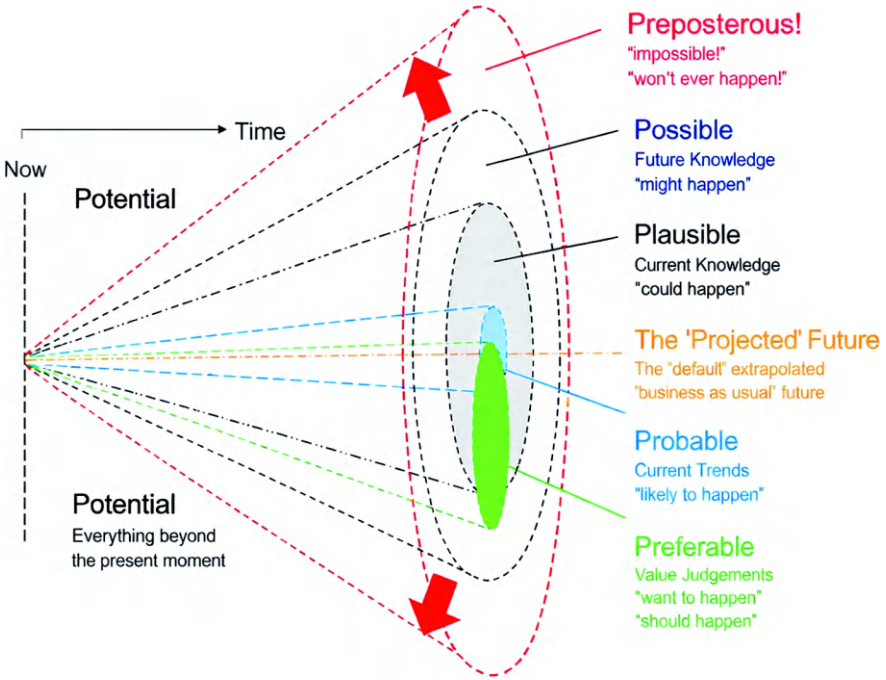


Figure 1: Futures cone.
Source: https://thevoroscope.com/2017/02/24/the-futures-cone-use-and-history/

certainty. For example, probabilities were critical during COVID-19 to predict infection rates, critical for resource planning. However, probabilities would not have uncovered supply chain disruptions or work-from-home acceleration. We would encourage you to go from what is probable, to leave room for “what if” type of questions—to focus on what is plausible.

When using foresight for strategy definition in addition to looking for plausible futures, we will look at those that are relevant and challenging. We will look for the “what ifs” that would substantially change the course of who you target as client, and how you deliver your product or service. If the identified uncertainties do not challenge your current business, they are of no use in this particular exercise. The often-cited taxonomy of knowledge of “known-knowns, known-unknowns and unknown-unknowns” is very helpful to manage and communicate uncertainty. As Donald Rumsfeld, U.S. Secretary of Defense, put it back in 2002:

> There are known-knowns—these are things we know that we know. There are known-unknowns—there are things that we know we do not know. But there are also unknown-unknowns—there are things we don't know we don't know.

Keeping this taxonomy in mind will give the reader the mindset to humbly start this exercise with the acknowledgment that we have gaps in our knowledge.

The first step that helps in turning noise to signal is identifying the known-knowns. In foresight methodology, they are called pre-determined variables. "[These] are trends or dynamics whose trajectory is relatively fixed based on known conditions and systemic inertia, allowing them to be confidently integrated into future scenarios."[9] They set the stage in which our businesses will have to compete.

The second step requires identifying known unknowns—key uncertainties: "Uncertainties are elements of the future that cannot be reliably predicted due to the complexity of influencing factors and the potential for emergent, unanticipated events."[10]

The third step, discussed in Chapter 3, creates future scenarios based on the most relevant interdependence between these uncertainties. When we depict future scenarios, we look to find relevant unknowns that would truly challenge existing business models and industry paradigms. Future scenarios are meant to help leaders challenge and test their current strategies (see Figure 2 with the step-by-step approach). It is about daring to ask the basic question.

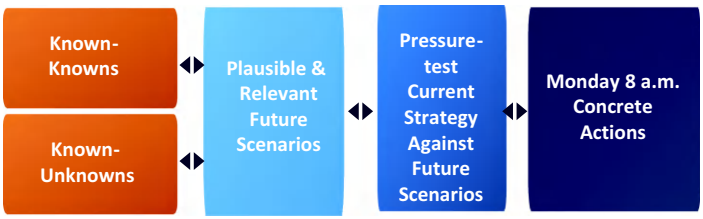


Figure 2: Foresight approach.

Let's go back to the Shell example.[11]

In 1985, Shell used scenario planning to prepare for another possible disruption to the industry. Group Planning brought together the top managers and presented them with the following scenario: "it's April 1986 and the price of oil has fallen from $30 at the end of 1985 to $16 today. What do you think your government will do? What do you think your competition will do? And what, if anything, will you do?" The price of oil at the time of the meeting was $28 and rising, de Geus recalls, and "$16 was the end of the world to oil people." However, the object

of the exercise was not to debate if the scenario would happen, but only what Shell would do if it did.

Shell assigned a large group of engineers to design a cheaper oil platform that would benefit from falling oil prices. By April 1986, oil prices dropped to $10 a barrel. While other companies announced layoffs and slowed investments, Shell announced two new oil fields in the North Sea and gave its operating companies more room to adapt to the changing conditions. Managers quickly adapted because they had already developed action plans to survive the steep price drops. They could react swiftly because they already had a clear "memory of the future."

Bias for Action

As Shell's history in implementing scenario planning has demonstrated, what matters is not so much being right in defining future scenarios; it is as important to convince managers who need to act on these alternative scenarios to do so. They will need to overcome their current mental models to better prepare the company for a successful future. An effective way of doing so is by creating "memories of the future," by engaging the organization in fleshing out relevant plausible and challenging scenarios so that they can quickly act on new signals.

Leaders with the ambition to create sustainable value for their business have a bias for action. Conceptual frameworks are critical to make sense of challenging and complex environments, but leaders also need a practical approach to drive necessary changes. I will build on Bill Sharpe's practical approach to foresight[12] and innovation, looking to design strategies for three potential horizons. Horizon one is oriented to short-term strategies to make current businesses thrive. This is the horizon where current managers shine. Horizon two is focused on managing for the medium term and expanding into adjacent business. This is the horizon in which entrepreneurs in organizations excel. Horizon three attempts to set a strategies that acknowledges key paradigm shifts in the industry. Typically, visionaries within companies drive this third scenario. The power of the three-horizon framework is that in the Era of AI the speed of change leads us to manage a portfolio of projects that cater to the different time horizons. As described in Figure 3, a portfolio approach to strategy can help address the need to deliver value across all three time horizons while better adapting to changes.

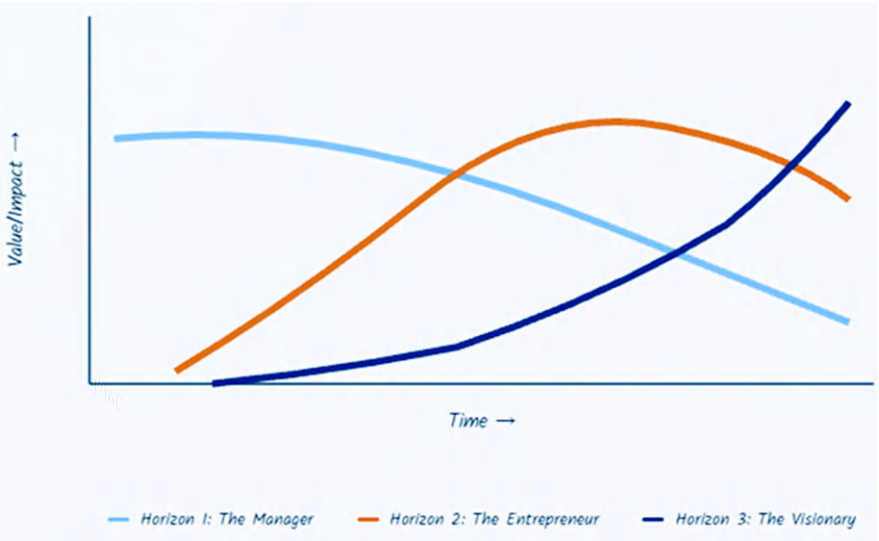


Figure 3: Visualization of impact of the three horizons.[13]

Chapter 2
Identifying Known-Knowns

Known-Knowns

Media is populated with headlines about the future of AI and how it will change society, geopolitics, or life of the average citizen. All these considerations are important, however when we narrow down on business strategy, there seem to be four important trends that are likely to continue, thereby setting the context for our businesses to thrive. We know the rough trajectory of these pre-determined trends, yet the precise speed and timing are difficult to predict. While there may be ample debate in the media on some of these topics, I will assume that these are known-knowns that are relevant and present, at least in the coming 10 years. These trends are:

- AI sovereignty will shape a nation's geopolitical power and influence. Governments are likely to be heavily involved in promoting the development of infrastructure and standards, but also their regional stars. Nations will want to control key elements of the AI ecosystem to guarantee AI sovereignty.
- AI Infrastructure will continue to improve. Technological progress will not hinder adoption.
- The cost of capital is likely to not be zero, and thus reasonable economic rewards within a reasonable timeframe need to be commensurate with risk. We will see less of the "build it, and they will come' mentality.
- While the development of AI comes with important global risks, governance guardrails are expected to prevent lethal use in the time horizon of this analysis.

AI: The Queen in the Geopolitical Chess Board

Historical Context

As AI leadership shapes the frontier of economic and political power, nations will continue to search for a combination of control and development of key enablers. It is difficult to imagine the next decade without the active engagement of governments and supranational entities in the creation of infrastructure, in setting AI-related standards, and directly or indirectly controlling critical access and uses. These elements are common in the development of general-purpose technologies, and AI is no different.

https://doi.org/10.1515/9783111583495-002

Electricity is a prime historical example; let´s take Europe. The process of electrification began in the late 19th and early 20th centuries, with a mix of municipalities and private companies building the foundations—from coal-fired and hydroelectric power plants to transmission and distribution infrastructure including high-voltage lines, poles, and transformers. Building this infrastructure demanded significant capital investment and quickly revealed path dependencies—where early decisions shaped long-term industrial development and economic activity. In countries like France and the UK, the state moved to nationalize electricity generation and distribution, creating Électricité de France (EDF) and the British Electricity Authority respectively. These moves were driven by multiple imperatives: the need to harmonize technical standards around voltage and frequency, broaden public access but also reconstruct economies after World War II.

Over time, these national utilities evolved into listed companies with mixed ownership, and the electricity sector transitioned from regulated monopolies to more competitive markets. Public authorities introduced price controls and mandated non-discriminatory access to the grid to ensure fair and standardized pricing across regions.

Crucially, electricity's integration laid the groundwork for deeper pan-European cooperation. The European Coal and Steel Community—a precursor to the EU—relied on shared infrastructure, and the creation of the European Network of Transmission System Operators for Electricity (ENTSO-E) enabled cross-border electricity flows and an integrated European energy market. None of this would have been possible without harmonized technical standards—particularly voltage and frequency alignment—across national grids.

In this context, centralized control and technical harmonization were not just industrial enablers—they were instruments of geopolitical power, both in post-war Europe and in its former colonies. In the early 20th century, electricity was not just a utility; it was a lever of influence and strategic advantage.

European governments are not the only ones to regulate and standardize the use of general-purpose technologies in their initial stages. Historically, the United States has followed a similar path, particularly with the development and expansion of the Internet, where public initiatives have played a crucial role in creating the initial infrastructure and regulatory frameworks. Aside from the fact that the Internet was born from a public need to create secure communication infrastructure during the Cold War, the initial R&D funding was exclusively public through the Defense Advanced Research Projects Agency (DARPA), which laid the groundwork for networking protocols (Internet Protocol (IP)/Transmission Control Protocol (TCP)). The initial infrastructure was built with the help of public participation, and the National Science Foundation (NSF) created a high-speed backbone

network for academic research institutions that was eventually extended to broader usage. After the initial public push, the government drove deregulation to drive infrastructure growth and greater competition, fostering the emergence of Internet Service Providers (ISPs). Later, the government also set rules for fair pricing and usage with net neutrality regulations to ensure a fair distribution of Internet traffic. This is like the European initiative of having national supervisory agencies control fair pricing and usage of electricity infrastructure. The Internet also required important public involvement in setting standards to facilitate broader cooperation and integration with alignment through organizations like ICANN, the Internet Corporation for Assigned Names and Numbers (ICANN) and the Internet Engineering Task Force (IETF), ensuring global interoperability and collaboration. For the United States, the Internet became a source of geopolitical power in its early days, for example, with tight export controls around encryption or programs like PRISM for Internet surveillance, using legal frameworks such as the Patriot Act 2001.

AI Sovereignty: Global Moves

In the next decade, active public involvement in setting AI sovereignty will lead to important infrastructure investments and controls. Early evidence includes Japan's Project Rapidus to gain technological sovereignty in chip manufacturing and make AI an economic engine for growth. It is expected to deploy more than ~ $65B by 2030, with an estimated multiplying effect of more than 20 times in its economy. To date, the economics of this project seem questionable, but other geopolitical considerations have made this project strategic for the country.

In China, we are also observing a strong focus on developing AI capabilities with a view to accentuating its global leadership. In 2017, China put forward a "New Generation Artificial Intelligence Development Plan" with a view to harness AI to further support its global ambitions by 2030. It has a multi-pronged approach, including the development of infrastructure and manufacturing capabilities, to ensure innovation and leadership in foundation models to develop competitive talent with robust AI skills. At the time of this writing they already had important advances under their belt; for example, DeepSeek, a Chinese AI start-up based on open-source AI technology, has delivered a surprising generative model with similar precision to the best in class currently in the market yet 90–95 percent more affordable training. Training models is currently the most expensive activity in building AI systems. This model seems to have been trained for millions rather than tens of millions or billions, as is common with models currently being developed in the United States.

Geopolitical power stemming from the early control of AI is evident in the political arena. For example, with the CHIPS and Science Act 2022 the United States is looking to incentivize semiconductor manufacturing domestically and with allies; in parallel the U.S. government is establishing export controls of what could be currently considered "strategic infrastructure": certain graphics processing units (GPUs) are forbidden for exports to China.

In 2024, the Indian government approved the India AI Mission with a budget outlay of approximately $1.25 billion with the intent to provide and control AI infrastructure of more than 10,000 GPUs for Indian businesses and to develop indigenous foundational models. Similarly, the United Arab Emirates's sovereign wealth fund ADQ created G42 to deliver the "Intelligence Grid," a scalable, interconnected infrastructure designed to provide on-demand AI services and capabilities for businesses, governments, and individuals; in 2022, it expanded the fund with $10 billion commitment. The European Union (EU) launched in February 2025 the initiative Invest AI, looking to invest €200B (~$227B) in developing AI, of which €20B (~$23B) was for building AI giga factories. Arguably this came as a response to the U.S. Stargate AI project looking to dedicate $500B to AI.

Governments are expected to continue to drive standards in general-purpose technologies like AI; the EU AI Act 2024 has rigorously classified the different uses and established a set of protocols to manage them, with strong governance over forbidden and high-risk applications.

According to the 2025 AI Index Report by the Stanford Institute for Human-Centered Artificial Intelligence (AI):

> In 2024, U.S. federal agencies introduced 59 AI-related regulations—more than double the number in 2023—and issued by twice as many agencies. Globally, legislative mentions of AI rose 21.3% across 75 countries since 2023, marking a ninefold increase since 2016.

As you have seen, there are diverse approaches to countries and regions developing their AI sovereignty. Which country will derive greater power and influence? The country investing more in innovation and infrastructure or the country diffusing the existing technology? The jury is still out.

In the early days of AI development, businesses expect their government to continue supporting the AI arms race, both directly and indirectly. It can come in the form of regulation, availability, and support for key infrastructure, and with special protection for strategic national or regional assets. The geopolitical dimension of AI can also be potentially used for the benefit of key players. It is worth noting that both Microsoft and Amazon have announced a European Sovereign Cloud service with autonomous governance from U.S. Headquarters, abiding by EU privacy and other relevant laws and managed by EU nationals. This strategy will potentially retain European clients, but as importantly it can deter European

competitors from entering the space if they do not have government support to reach critical scale.

As general-purpose technologies mature, their relevance in the geopolitical spheres decreases; but, for now, stay tuned to your national and international geopolitics around AI! It may open important opportunities.

Infrastructure for AI Will Continue to Improve

At a basic level, the infrastructure for AI deployment includes four key components: semiconductors driving the computing power behind AI; the cloud creating a scalable environment where AI runs; and foundation models, the pre-trained AI models where businesses can create applications and; the orchestration platform allowing applications and businesses to better leverage existing models.[14]

Let's imagine constructing a building. Semiconductors are like the cement and steel of our building. These raw materials are like the essential hardware that determines the robustness and performance of the building. The cloud is like the cranes and logistics; it is the machinery and systems that allow for construction to be efficient, fast, and at scale. The foundation models are like prefabricated modules, blueprints that ensure we can build anything with quality, with standards, and enough flexibility. With these prefabricated modules we can build anything: apps, tools, and services. In this metaphor our building is static, and the nature of AI infrastructure is dynamic; it constantly adapts and reconfigures. It is a smart building where steel and cement can be programmed and adaptable, where cranes – become ever more flexible based on the evolving needs during construction, the prefabricated modules, also learn and get reconfigured based on new information.

Technical Advancements in Compute and Cloud

Gordon Moore—co-founder of Intel—predicted that the compute power of chips would double every two years. The miniaturization of transistors and chips has made this law true from the 1970s to nowadays. This law embraced by the semiconductor industry has led to the wide availability of computers, servers, and mobile phones delivering better computing power for different applications.

Today, the basic infrastructure for the future development of AI relies on the efficiency improvement of semiconductors from chip designers such as Nvidia, Intel, AMD, and Qualcomm, as well as efficient manufacturing capabilities from the likes of TSMC, Intel, and Samsung. The semiconductor industry has managed to keep up with Moore's Law with a combination of improvements across the

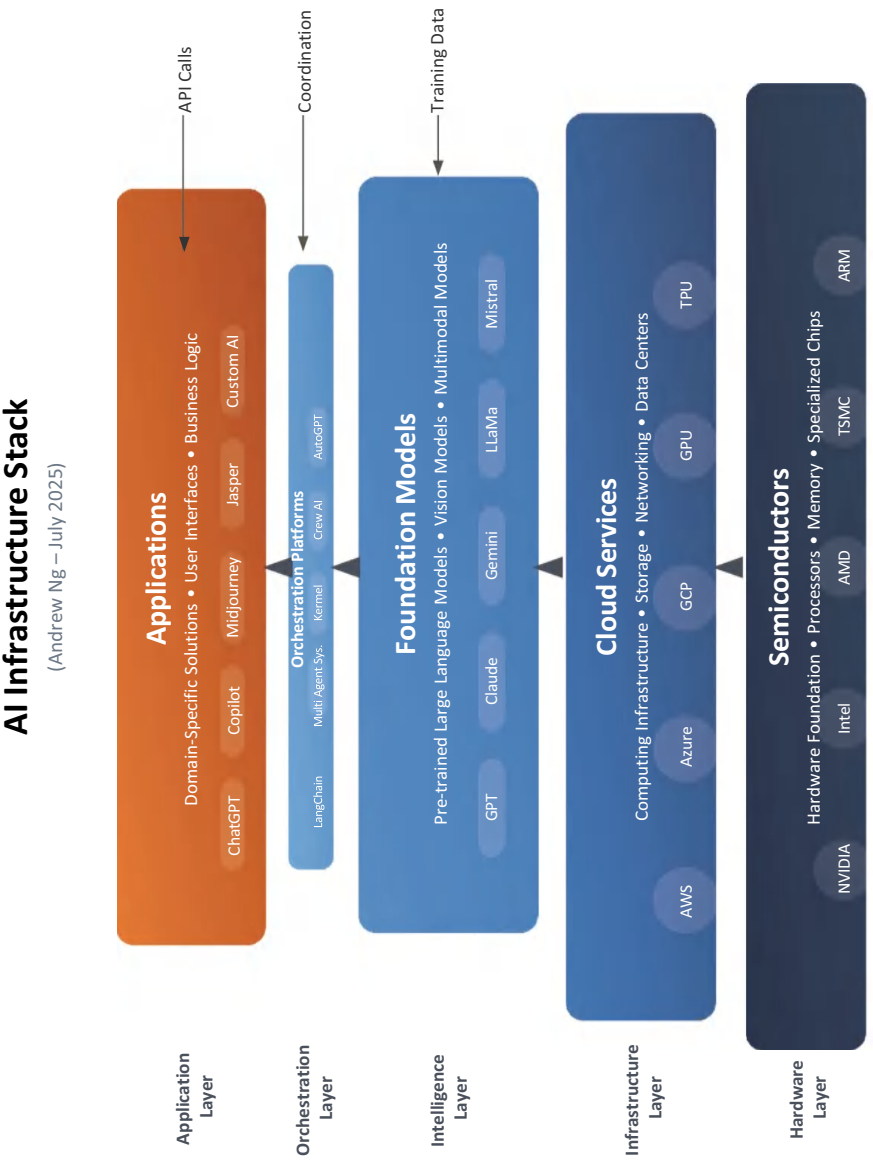


Figure 4: AI infrastructure stack (Andrew Ng).

value chain, including techniques for improved manufacturing that allow incorporation of even more transistors, such as extreme ultraviolet lithography, novel transistor structures, or even trying different materials other than silicon to improve electrical properties.

In addition to pure manufacturing improvements, there are other design improvements creating new more efficient supercomputer architectures: such as using specialized processors such as graphics processing units (GPUs) or high bandwidth memory (HBM) for inference or, improving the software that makes data flows more efficiently between component parts. AI is a catalyst for high-performance computing. Recent research by Epoch estimates that machine learning hardware performance has grown over the period 2008–2024 at an annual rate of approximately 43 percent, doubling every 1.9 years, as seen in Figure 5. We are in the age of "more than Moore", computing progress continues albeit not only from minaiturization.

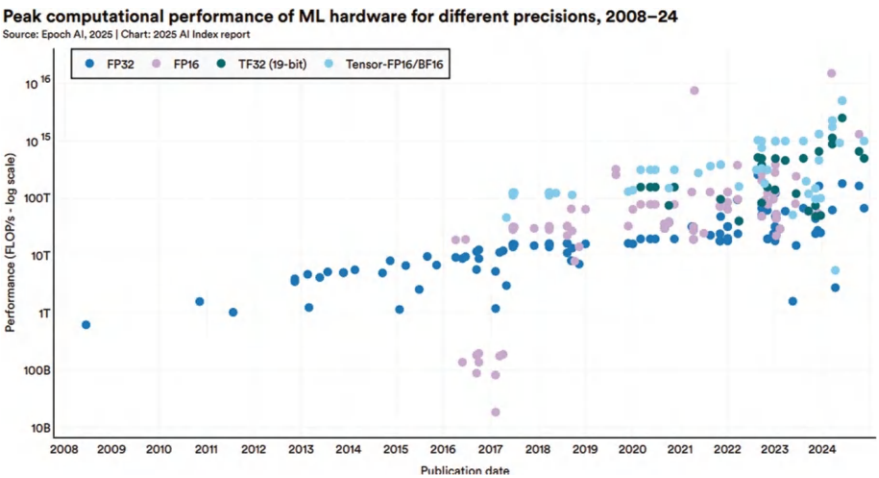


Figure 5: Peak computational performance in machine learning measured in 16 bits floating-point operations per second (2008–2024).
Source: Stanford HAI, *The 2025 AI Index Report*.[15]

Manufacturing semiconductors is a high fixed costs business—in the billions of dollars—with remarkably high barriers to entry. As such global scale is paramount to drive attractive economics. The semiconductor ecosystem has been investing over the past decade about half of its EBIT (Earnings Before Interests and Taxes) to foster innovation and growth.[16] Expectation for 2024 is approximately $540B investment. In addition to traditional semiconductor players looking to improve the compute infrastructure, we are also seeing hyperscalers like Amazon Web Services

(AWS) developing their own chips—working with TSMC to manufacture them. Interestingly, AWS is also making acquisitions to ensure scalable data center operations with more than $11B commitment.[17] This echoes the early e-commerce days when Amazon was not shy about investing in warehouses and eventually in aircrafts to ensure sustained global growth of e-commerce.

Jensen Huang, CEO of Nvidia, has a unique perspective on the potential growth of data centers given Nvidia´s market share in datacenters; he predicted in 2024 that the investments in data centers would double in the coming five years.

The availability of infrastructure by 2035 or 2040 seems less of a technical challenge and more of a financial challenge given the economics behind running these models. Energy consumption of running the models remains an unsolved problem.[18] Electric Power Research Institute (EPRI) estimates that a query in ChatGPT is 10 times more costly than a query in Google.[19] Data centers consume 2 percent of global electricity in 2022, according to the International Energy Agency, and are expected to continue growing. In the United States alone, in 2024, 4 percent of electricity was consumed by data centers (See Figure 6).

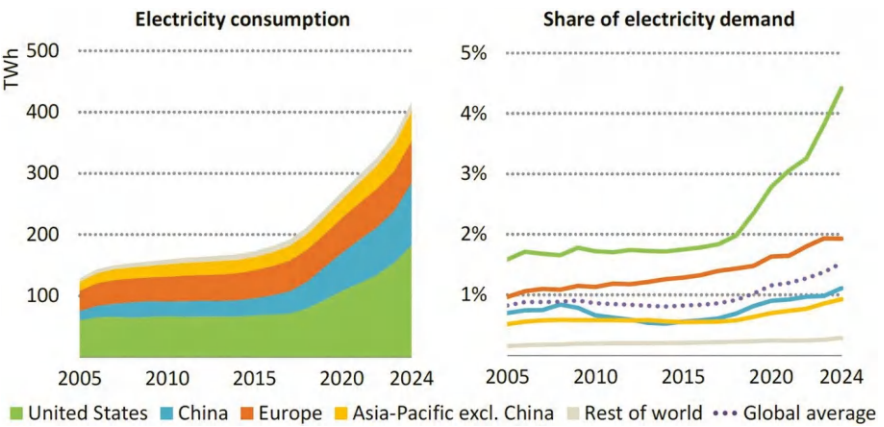


Figure 6: Data center electricity consumption by region: 2005–2024, per IEA.
Source: International Energy Agency (IEA), 2025. *Energy and AI – Analysis*. Paris: IEA. Available at: https://www.iea.org/reports/energy-and-ai

This is widely acknowledged as an important challenge. Currently important investments and R&D are dedicated to solving the power problem, from improving the design of chips (see Figure 7) to using cheaper sources of energy like solar or nuclear, to training models in distributed data centers with strong fiber connections between them. Tokens are pieces of a word; de facto they are like the unit of consumption of foundation models. Tokens are to foundation models what giga-

bytes are to data plans or kilowatts to energy. Taking Nvidia as an example, the energy required per large language model (LLM) tokens has been declining exponentially, especially in the early days (2014–2018).

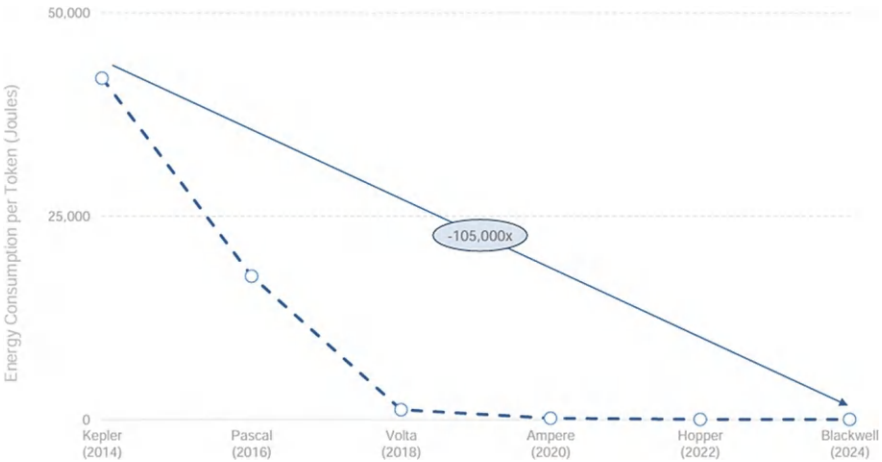


Figure 7: Energy required per LLM token (joules); Nvidia GPUS: 2014–2024, per Nvidia.
Source: Nvidia Investors Presentation, Company Overview, Feb 2025; https://investor.nvidia.com/events-and-presentations/presentations/default.aspx
Note: Kepler was released in 2012. NVIDIA materials mark the performance threshold shown above for Kepler as of 2014.

So, while the infrastructure is becoming more efficient, the wider use and the improvement opportunities continue to drive very heavy investments. Training foundation models is also a high fixed cost business with similar barriers to entry and need for global scale than we have seen in semiconductors. In Figure 8 we can see the orders of magnitude of training foundation models (in USD). From 2016 to 2024, training costs have increased approximately 2,400 times. According to Dario Amodei, co-founder and CEO of Anthropic, creator of Claude, the training of $10B models could start in 2025.[20]

Innovation in AI model design is also driving more efficient use of computing power. DeepSeek's V3 reasoning model was trained on a fraction of the cost of OpenAI's by more effectively calling only on parts of the model during training and inference and reducing memory usage. Inference is the process of querying a model. Investment and innovation are also driving significant improvements in performance of foundation models, as illustrated in Figure 9; though it seems to be asymptotic. Please beware that there is ample debate about the right definition

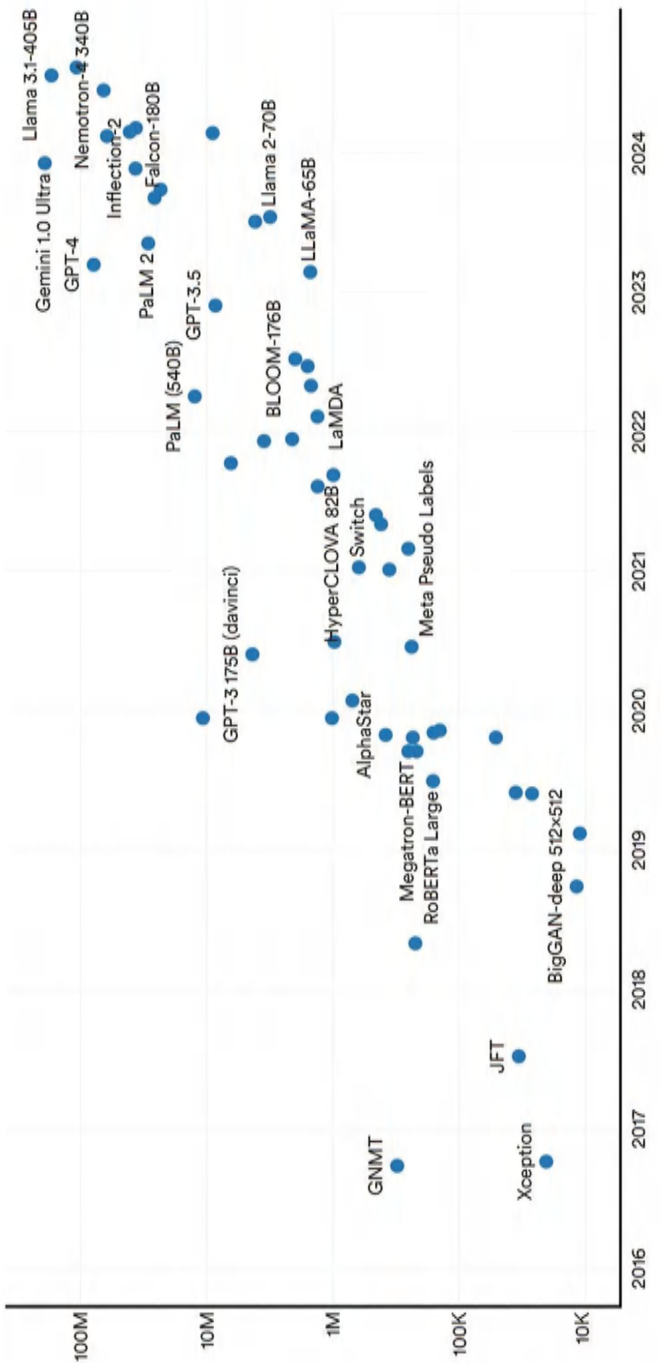


Figure 8: Estimated training cost of frontier AI models (2016–2024).
Source: Stanford HAI, *The 2025 AI Index Report*[21]; *In Good Company* (2024) podcast, June.
Note: Costs are estimates expressed in USD (Log Scale). Excludes most Chinese models due to lack of reliable cost data.

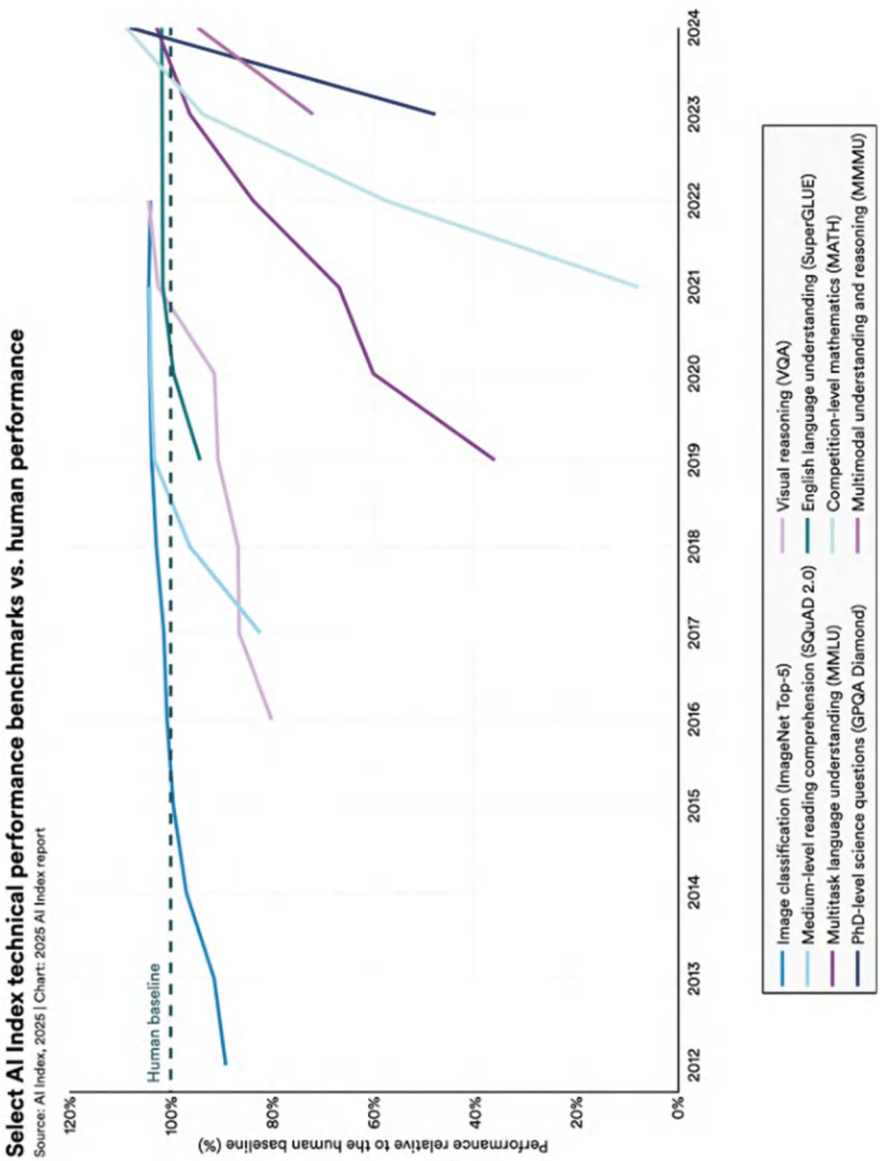


Figure 9: Performance of AI models versus human performance.
Source*:* Stanford HAI, *The 2025 AI Index Report.*[22]

of a human baseline, and it varies by topic. Let's assume, for now, that, relative to a constant baseline, performance has dramatically improved in the past decade.

Two potential issues may slow the development of infrastructure at the pace we have seen in the last five years. On the one hand is the availability of data to continue training models and, on the other, it is the demonstration that companies behind foundation models have an attractive economic model that justifies further refinement of the models.

On data availability, there is an increasing concern that training models' needs are growing faster than the ability of the Internet to generate publicly available human text. After all, there is only one Internet! According to preliminary research, we could run out of data by 2026–2032.[23] The alternative foundation models are evaluating is the use of synthetic data—that is, generated from existing data. There are concerns, however, that models will not work as well and will not improve at the pace we have seen thus far.

The other potential obstacle to infrastructure improvement may be the economics of the foundation model players. The semiconductor subsector has proven its economic viability even if certain players are experiencing economic challenges. The cloud subsector has also demonstrated the economic viability of their model despite their highly commoditized services. The historically robust economics of the semiconductor sector—with returns on equity (ROE) ranging from 17 percent to 22 percent depending on cycle and profile over the decade 2014–2024—and equally robust economics of leading cloud providers with historical ROEs more than 500 basis points greater than what we see in the semiconductor industry, leads us to believe that these two layers of the AI infrastructure will continue supporting the adoption of AI. However, builders of foundation models have yet to make a profit, especially the stand-alone players. In the case of Big Tech like Meta, Google, or Apple, there is no visibility yet of the stand-alone expenses or profitability of their models. For now, AI models are embedded as additional features within existing products. Big Tech only reports AI-related capex. We will review the apparent AI strategies of Big Tech in Chapter 3. The size of investments required in training and the fast convergence in performance paired with the incredibly competitive space raises concerns about future profitability of players solely focused on foundation models. We need to believe that companies like Anthropic, OpenAI, or Mistral can move down the experience curve fast enough and/or develop new monetization strategies.

Even if improvements in model performance were to slow down or the cost of using the models were to increase, none of these two facts may materially affect AI adoption in the coming 5 to 10 years. At this stage, there is barely one percent of companies reporting that their AI adoption is fully mature.[24] Therefore, let us consider developments in infrastructure as a known-known that will support AI adoption in your sector as it identifies the best use cases.

Capital Will Not be Free

In the 2010—2020-decade, we have seen reference interest rates at central banks in Europe and the United States at zero or even below zero in the case of Europe. This trend is part of a secular trend of declining interest rates in these geographies.

The capital bonanza to reactivate the economy after the global financial crisis of 2008 led to investors taking unparalleled risks with distorted investment decisions and corporate capital allocation. In this environment, unprofitable business models could exist for a long time, shaking up entire industries. In the Digital Era (2000–2020), new competitors would emerge out of the blue, and investors would tolerate subpar returns on investment under the promise of future returns. "Build it and they will come" was the mantra in Silicon Valley. The number of zombie firms that cannot cover interest payments with their Earnings Before Interest, Taxes, and Depreciation (EBITDA) ranges from 5 to 15 percent depending on the country.

There are reasons to believe we may be back in an era of normalization of interest rates. Interest rates are a key variable in investing. If investors get a 3 percent plus inflation on safe investments, this is the minimum they will ask to deploy capital. I do not pretend to make here a prediction of interest rates in our time horizon of focus—the next 5 to 10 years—but would like to unpack some of the healthy and unhealthy reasons why the very low interest rates seen pre-COVID-19 may not hold and any AI investment will demand, eventually, a solid economic rationale.

The healthy reason driving up interest rates is a greater demand of capital. There are three structural forces driving the capital revolution.[25] The first reason is AI. AI is a capital-intensive sector, as discussed earlier. In 2025, U.S. Big Techs alone are investing $300B. McKinsey, a consulting firm, estimates incremental capital outlays in data centers of approximately $7 trillion with the vast majority being driven by AI. Additionally, demand for greater electrification of economies with the goal to lower CO_2 emissions will necessarily include major investments in power plants of all sorts (from nuclear to renewable energies) in power transmission networks and storage. According to Boston Consulting Group, a consulting firm, global investment in energy transition in 2023 reached $1.8 trillion; it is expected to reach $19 trillion by 2030.[26] Finally, the other structural change driving the capital revolution is driven by geopolitical tensions and greater defense investments. In 2023, defense investments in developed economies reached $1.5 trillion of 2.5 percent of their gross domestic product (GDP) with the United States at 3.5 percent of GDP and European Union (EU) at 1.6 percent. Getting developed economies to the U.S. level of 3.5 percent of GDP will require an additional $620B. Ideally these massive investments will generate an increase in productivity and, therefore, more robust global

growth. To be clear, productivity growth has been lagging in all developed economies; so, it should not be taken as a given.

There are other, less healthy reasons why interest rates may be "normalized" like excessive fiscal deficits, supply shocks from tariffs, or geopolitical pressures on supply chains: inconsistent monetary policies that create inflation without growth. Some of these phenomena are also at play now. Looking at governments' level of indebtedness, debt to GDP ratios have declined in the Organization for Economic Co-operation and Development (OECD) from 125 percent of GDP in 2020 but the average remains north of 100 percent. If we look at fiscal deficits, OECD countries continue to have large fiscal deficits—average of 4.6 percent of GDP—far from the pre-pandemic 3 percent.

It is unclear if inflation will continue the post-pandemic downward trend given the geopolitical tensions on tariffs and supply chains. Small interest rates adjustments are likely, but for now, neither the U.S. Federal Reserve (Fed) nor the European Central Bank (ECB) is signaling radical declines in interest rates; at the time of writing Euribor is approximately 2 percent (for 12 months' rate) and the U.S. federal funds rate is about 4 percent.

Finally, looking at private capital markets, a critical source of capital for emerging AI infrastructure and applications, we see dry powder—uninvested raised capital—of private capital with more than $4 trillion by the end of 2023. This fact indicates that capital is not willing to reduce return expectations despite uncertainty and higher cost of capital relative to pre-COVID-19 years. Venture capitalists (VCs) continue to target net internal rate of return (IRR) in orders of 20–25 percent and private equity funds at the 15–20 percent level. While AI deals have been consistently growing, investors continue to look for solid financial performance with attractive unit economics and investments aligned with attractive profit prospects. Recent equity deals, akin to vendor financing, further validates that capital will not be free; arguably AI players are using the most expensive source of capital to fund their operations.

The good news is that incumbent businesses are less likely to be at risk of facing irrational new entrants with no pressure to generate reasonable returns on their capital. New AI-related businesses and projects will need to present a compelling business case.

AI Safety Will Prevail

Can AI end up destroying humanity? It looks like it is in the realm of the possible. AI can end up being used for evil purposes, models or agents may make harming decisions for us and we can no longer unplug them.

This debate is amplified by the possibility of artificial general intelligence (AGI) or artificial super intelligence (ASI).[27] More nuanced definitions will be addressed in the following chapters, but these concepts assume that artificial intelligence will eventually surpass human cognitive capabilities for a subset of, or all of its dimensions. The mere possibility of this happening leads to great philosophical discussions.[28] My working hypothesis is that human preservation instincts will prevail. It will take a combination of political will, proper standards and metrics, economic incentives, and effective enforcement mechanisms; the same set of measures we have already been using in the face of other major deadly risks to humanity.

Historically, technological innovations, from the harnessing of fire to the development of nuclear energy, have introduced clear benefits but also important risks. However, humanity has consistently identified ways to manage and mitigate these risks. Governance bodies and collaboration on self-imposed guidelines have so far been reasonably effective, for now, in controlling lethal damage to humanity. Two good examples with different governance approaches can help illustrate this point: nuclear weapons and climate abuse.

Modern nuclear weapons are a thousand times more powerful than the ones dropped in Hiroshima and Nagasaki. There is general agreement that a nuclear war could eradicate life on earth within minutes. The Nuclear Non-Proliferation Treaty (NPT) signed by more than 191 countries is the centerpiece of global efforts to prevent the spread of nuclear weapons and promote cooperation for peaceful uses of nuclear energy. The International Atomic Energy Agency (IAEA) plays a pivotal role in verifying the fulfillment of obligations under the treaty. The direct mechanisms for enforcement involve inspections and regular audits to ensure that nuclear materials and facilities are not diverted for weapon use, non-compliance reporting to the United Nations Security Council, and withdrawal of technical assistance. Admittedly, the IAEA can be compared to a dog with a strong bark but no bite. Action is only taken with the approval of the UN Security Council. The UN Security Council can use Chapter VII powers to impose sanctions, mandate corrective actions, or in extreme situations authorize force. While there has been important progress in controlling nuclear materials with effective interventions—South Africa and Iran in the 1990s, Libya in 2003—the truth of the matter is that North Korea exited the

NPT, and Israel, India, and Pakistan never signed it. Having these countries included in the NPT would be highly preferable, however, even then, there is no guarantee of no nuclear conflict. Additionally, geopolitical diplomacy and alliances are also at play. This imperfect system seems to have worked over the past 80+ years. This is no guarantee whatsoever that it will last another 80 years, or even one more year. It is just evidence of our self-preservation instincts. Let's hope they will continue to prevail over current geopolitical tensions or sheer evil intentions.

Humans' basic survival benefits from a stable climate; survival needs include clean water and air, food, and shelter. Greenhouse gases—carbon dioxide (CO_2) and methane (CH_4)—trap the heat in the atmosphere to keep the earth warm. Without it we may be at average temperatures of −18°C. While we humans badly need the greenhouse effect for survival, too much of it can cause overheating and provoke serious threats to humankind.

What have we done to prevent this lethal risk? At the risk of oversimplifying allow me to describe the global system in place to reduce the negative impact of climate change. There are also interesting insights for AI governance.

Let me use a sailing metaphor. We have created a navigation chart: the Intergovernmental Panel on Climate Change (IPCC), created in 1988 with 195 member countries, provides science-based, comprehensive assessments of climate change, its impact, and potential future risks. The creation of a common set of metrics, data sets, and scenarios can thereby guide global efforts to better understand the magnitude of climate-related challenges. Now that we have a navigation chart, we have constructed the hull for navigating through the ocean of climate risk. Kyoto (1997–-2012) and Paris (2015 to present) have created legally binding emission targets to reduce warming to no more than 2ªC. In theory, the Paris Agreement covers 95 percent of CO_2 emissions. However, U.S. exiting for the second time the Paris agreement and China not delivering on its commitments make this hull a patchy one. The Conference of the Parties (COP) within the United Nations Framework Convention on Climate Change (UNFCCC) provides a forum to drive global and regional alliances, to share best practices, and reinforce commitments. They are like ports where boats can be fixed and re-fueled. The fuel for our boat consists of a number of global and local financing mechanisms to put in place mitigation and adaptation strategies to reduce CO_2 emissions.

While none of the systems described have been effective in radically eliminating the root causes of these catastrophic risks, they are nonetheless precedents demonstrating humanity's capacity for collaboration in managing lethal threats.

Indeed, we are dealing with immensely powerful technologies, and we need to keep our eyes open to unknown-unknowns. Significant milestones have shown that global cooperation to control undesired effects from AI is possible. Just to name a few, The Bletchley Declaration (2023): signed by 28 countries, including

the U.S., China, the European Union (EU), and India, established international AI safety principles and cooperation frameworks. The OECD has put forward AI Principles adopted by 46 countries, promoting trustworthy AI development and emphasizing fairness, accountability, and transparency. They have influenced regulatory frameworks on AI safety, such as the EU AI Act and U.S. executive orders. Additionally, the United Nations AI Advisory Body (2023) was launched to coordinate global AI policy efforts, aiming to develop shared governance mechanisms and assess the risks tied to military and economic AI applications.

In addition to the efforts of supranational entities to collaborate and set guard rails, current players are looking to build and evaluate models according to critical safety standards. For example, Anthropic is developing a hierarchy of safety standards as part of its Responsible Scaling Policy. Similarly, Google DeepMind in a recent research paper[29] proposes specific approaches for preventing harmful situations from misuse in the hands of adversaries or to address problems stemming from misalignment or model mistakes. Furthermore, significant resources are going into AI safety research conducted by scientists at universities but also in non-for-profit organizations like the Center for AI Safety (CAIS) in California that goes as far as proposing an approach for national AI safety. VC capital from the likes of Sequoia Capital or Andreessen Horowitz is flowing into AI safety start-ups like Safe Superintelligence created by former OpenAI chief scientist Ilya Sutskever recently raising $1B in cash. To date it looks like there is encouraging progress in setting a global framework for cooperation and an initial set of standards.

If nuclear materials and climate change are to be of some guidance; it looks like we still need an organization with the ability to sanction and effectively intervene when needed. The EU Act and U.S. executive orders incorporate sanctioning mechanisms, but for the time being no single body has the power to evaluate and intervene on a global scale.

In sum, there seem to be important forces at work to control the undesired effects of AI. While the perils do exist, for the purpose of rethinking strategy in the Era of AI, we include into the category of known-knowns the control of lethal uses of AI for the coming 5 to 10 years. We want to be on the look-out for formal treaties against the weaponization of AI—much in the same way we have the Non-Proliferation Treaty for peaceful uses of nuclear materials. Additionally, we would like to have a globally recognized body for verification and control that can do inspections and audit models against the generally accepted guardrails. The most important feature to look for is the emergence of kill switches allowing the model to "unplug" to protect all of us from any lethal consequences from AI models.

As described earlier, the known-knowns still have some uncertainty around them, but as a business leader you want to focus your strategic mind and resources on the variables that can open up important opportunities or risks to your business, and where you have meaningful degree of influence. We will review them in the coming chapter, describing potential polarities impacting your business.

Chapter 3
Known-Unknowns: Navigating Polarities

Imagine that you are going down a river on a raft. You see the shores; you know your destination. If you have done a little bit of preparation, you know the known-knowns, the pre-determined variables we have just reviewed. For example, if you checked the weather forecast, you know you can expect rain, but it is unclear whether it will occur at noon or at 1 p.m. If you have checked the corresponding hydrographic authority, you may also know that at this time of year, you may encounter a few forks in the river. Some are gentle streams, whereas others hide extreme rapids or waterfalls. You cannot see what lies ahead of the fork itself, and once you choose a path, going back is impossible.

This analogy helps illustrate not only the challenges to perceptual blindness, but also the path dependency of certain decisions. It also prepares the discussion about what you want to carry on your raft knowing it can only take a certain weight; the mere acknowledgment of the fork in the river may drive decisions about what raft to take or with whom you will ride.

In foresight, we refer to these forks in the river as key uncertainties or known-unknowns; they encourage business leaders to contemplate non-linear futures over the next decade or so. If we go back to the Futures Cone, referred to in Chapter 1, we are looking for the elements that create relevant, plausible and challenging future scenarios. Awareness of these critical uncertainties cognitively primes leaders, equipping them to better separate noise from signal.

Timely capture of signals can activate a concrete course of action leading to the preferred future.

Trust in AI

Trust continues to be a prime catalyst for technology adoption. Public trust is fragile, and even if companies go out of their way to be transparent and have reliable third parties certify their products and services as robust and safe, it is difficult to control. For the purpose of this thinking riddle, let's agree that a combination of nations, regulators, and market participants will place the necessary guardrails in place to avoid lethal risks; however, there is an ample spectrum between lethal risk and totally safe AI use.

In the crucial initial stages of AI adoption, earning and maintaining public trust is not just important, it's vital. Once lost, trust is extremely hard to rebuild, and public perception can quickly outpace technical reality. This can stall AI de-

https://doi.org/10.1515/9783111583495-003

ployments and undermine the full potential of these technologies both at nation level and at company level. Building trust in AI is paramount to reap all its benefits. It is noteworthy that while the United States has led the innovation and investment efforts on AI thus far, it is the Chinese who seem to be trusting AI more and taking better advantage of its benefits.[30] The Edelman Trust Barometer reports attitudes toward AI comparing China and the United States, and the Chinese score on trust in AI doubles the U.S. scores. Chinese scores are also ahead in comfort with business use of AI and in willingness to embrace AI. Interestingly, comparing U.S. and Chinese companies adopting AI, it would appear that Chinese companies see a much greater valuation advantage relative to their respective local market than U.S. companies.[31] Additionally, AI systems thrive on high-quality data, but the supply of such data isn't growing at the same rate as model ambitions—after all, there's only one Internet. Data is fast becoming a strategically scarce resource, and a trust deficit will only tighten that bottleneck.

Let's take an example. Nuclear energy is considered a solid solution to mitigate the impact of climate change with $C0_2$ emissions through its lifecycle[32] comparable to solar and wind. Yet the public mistrust after accidents like the Windscale fire in England, Chernobyl in the former U.S.S.R., and Fukushima in Japan, have severely hampered adoption. Increasing pressure on governments has led to dramatic decisions, such as the one taken by Germany in 2011 after Fukushima, to decommission all nuclear power plants in the country. Similarly, public trust was also key in the early days of e-commerce demanded extra efforts to control fraud or ensure security of credit card payments to give consumers the confidence that what they were paying for was going to show up at their doorstep.

Business leaders will need to keep an eye on signposts that undermine or reinforce public trust in AI. This may affect different use cases or sectors differently. At a general level, research conducted on 30 models of Gen AI, including ChatGPT 4.0, Gemini, Llama 3, or Claude 3.7 showed that models lie under pressure; no models stayed honest more than 50 percent of the time. Large models tend to be more factually correct, but they do not always deter intentional falsehoods.[33]

There is still a striking silence around AI potential long-term effects on the human brain and individual well-being. We've already witnessed the psychological fallout of social media overuse in young people, and the cognitive strain caused by our culture of perpetual connectivity. Now, with the rise of AI assistants and companions, it's highly plausible we'll see profound neural and psychological shifts. These solutions may rewire our brains in ways that compromise our ability to distinguish truth from illusion. In this context, trust in AI isn't just a technical or ethical issue, it can be a mental health imperative. Building that trust will require intentional strategies designed not just for transparency or fairness, but for protect-

ing the integrity of human cognition itself. Leveraging existing foundation models for your business may be OK for low-risk activities, but you may need to consider adding additional layers of verification for mission critical activities. There may be a real need to explain and communicate transparently the magic behind the models. The fact that former OpenAI Chief Technology Officer Mira Murati created the start-up Thinking Machines to make AI systems more widely understood and trusted is an interesting sign of the importance of creating trusting tools.

Human Augmentation Versus Automation

Solow's Growth Model

Economist Robert Solow developed, in 1957, a foundational framework that has guided generations of economists to understand the disparities in economic growth across nations. At its core, the model emphasized technological progress as the primary growth engine. Although capital investments and human capital played crucial roles, they were respectively constrained by depreciation and diminishing returns. Only technological innovation can expand a country's growth frontier. However, for this innovation to translate into sustainable progress, it must be deeply embedded in the economy and supported by the right levels of human and physical capital.

About three decades later, in the 1980s, the model earned Robert Solow the Nobel Prize. Paradoxically, despite the relentless march of technological advancement and substantial global investments in information technology, productivity gains remained elusive. There was little evidence of increased output per hour worked—a phenomenon now known as the Productivity Paradox. As Solow wryly observed in a book review for *The New York Times*, "You can see the computer age everywhere but in the productivity statistics."[34]

It was not until the 1990s that enterprise software seamlessly, integrated into business processes, unlocked meaningful productivity improvements. Simultaneously, economists began recognizing that productivity measurements had some important pitfalls. Despite their undeniable value, variables such as more engaged customers or higher-quality company data are not adequately measured in company reports. Even today, debates persist on how to quantify these intangible assets, which, although critical to a firm's success, remain financially invisible. It is worth noting that this debate concerning augmentation versus substitution of human capabilities by key players in the technology sector has also shaped the product development approach of Big Tech. The idea of augmentation was beautifully described in 1990 by Steve Jobs in a documentary:[35]

> I remember reading an article when I was about 12 years old, I think it might have been in Scientific American, where they measured the efficiency of locomotion for all these species on planet earth, how many kilocalories did they expand to get from point A to point B, and the condor came in the top of the list, surpassed everything else, and humans came in about a third of the way down the list, which was not such a great showing for the crown of creation.
>
> But somebody there had the imagination to test the efficiency of a human riding a bicycle . . . and a human riding a bicycle blew away the condor all the way off the top of the list. And it made a really big impression on me that we humans are tool builders, and that we can fashion tools that amplify these inherent abilities that we have to spectacular magnitudes. And so, for me a computer has always been a bicycle of the mind, something that takes us far beyond our inherent abilities. I think we're just at the early stages of this tool, very early stages, and we've come only a very short distance, and it's still in its formation, but already we've seen enormous changes. I think that's nothing compared to what's coming in the next one hundred years.

Computers are like the bicycle of the mind. From 2007 to 2022, average annual productivity (measured as real GDP per hour worked) grew by 0.9 percent in the EU-27 and only slightly higher in the United States.[36] Looking ahead, we must ask ourselves: Will we see the AI Era everywhere except in productivity statistics? A key uncertainty looms ahead: will we once again witness a Productivity Paradox leading to vast investments in AI infrastructure and models that fail to produce visible economic gains?

AI Impact on the Workforce

Another crucial question is AI's impact on the workforce. If companies fixate on making AI mimic human capabilities, they may fall into what Stanford economist and productivity expert Erik Brynjolfsson calls the Turing Trap, a scenario in which AI is used to replicate human activities, looking principally at cost cutting rather than value creation. This could mirror the effects predicted by the famous 2013 study by Frey and Osborne (Oxford),[37] which projected that 47 percent of U.S. jobs were at risk of automation. This report ignited widespread debate, prompting researchers to further refine their analyses years later.

During the digital revolution of early 2000, most significant disruptions occurred in white-collar roles involving administrative and repetitive data processing such as accounting and claims administration. If the Turing Trap takes hold, AI will likely continue transforming these professions and make more white-collar jobs redundant. However, unlike previous waves of automation, blue-collar jobs may face equal disruption as next-generation robotics and autonomous systems become more capable with AI.

However, there is an alternative path in which AI enhances human capabilities rather than replacing them. AI has the potential to augment the workforce across industries, from white to blue collar workers and from doctors to call center agents. In an interview with the *Financial Times*, Brynjolfsson shared striking evidence from his research on the impact of AI on customer service:[38]

> There was an average of about a 14% productivity gain within just a few months. The least experienced agents saw a staggering 34% increase in productivity. Other studies on roles such as software development, management consulting, and writing tasks reveal a consistent pattern: (1) double-digit productivity gains, (2) greater benefits for less-experienced workers, and (3) rapid adoption within months.

At the same time, we also see evidence of the opposite trend—AI tools favoring top performers–especially in more complex activities such as research or management.[39]

Augmentation—Automation Paradox

Critically, automation and augmentation need not be mutually exclusive; instead, they may feed on each other. Research by Raisch et al. highlights that organizations achieve better results when combining both approaches. Consider Unilever, where AI-powered candidate screening accelerates hiring cycles while keeping final decisions human-led, resulting in a 16 percent boost in diversity hiring and £1 million annual savings.

This interplay between automation and augmentation forms what we can call the Augmentation–Automation Paradox: one process fuels the other. Automation frees up the time for augmentation, whereas augmentation refines automation. The typical pattern is augmentation, followed by automation.

In 2024, the World Economic Forum (WEF) projected that 85 million jobs would be displaced, while 97 million new roles would emerge. By 2025, the estimate had shifted to 92 million jobs expected to be lost (approximately. 8 percent of total formal employment), but 170 million jobs expected to be created by 2030. This tug-of-war between AI optimists (entrepreneurs, venture capitalists, and Big Tech) and AI pessimists (governments, regulators, and legacy industries) creates ample noise for business leaders. However, a focus on skills may turn this revolution into a more controlled evolutionary process. The question is no longer just which jobs will be automated but what new skills will be required. A recent analysis researchers Mäkelä and Stephany[40] looking at skills demanded in the United States, Australia, and New Zealand reports that growth in AI roles—those that include AI skills—is associated with growth in non-AI roles with complemen-

tary skills like teamwork, analytical thinking or adaptability. There is also a decline in non-AI roles with substitutable skills; but the net effect is a net positive with greater demand of non-AI roles with complementary skills more than offsetting the decline of the latter.

New Skills for Augmentation

There is ever-evolving research on critical skills for the future. The World Economic Forum (WEF) estimates that, on average, approximately 40 percent of skills on the average job will change from 2025 to 2030. While this may sound daunting to some of you, in a period of 5 years it requires changing 8 percent of the skills/year that are required for a given role. This feels like eminently doable in the right environment. Not surprisingly, resilience, agility, flexibility, and lifelong learning continue to be among the most sought-after skills. That is especially true if we think about organizations beyond 500 employees—the minimum threshold used in the WEF study.

As expected, all skills related to creating and using better AI continue to experience high growth according to the WEF *Future of Jobs Report 2025*, the loosely defined big data and AI, networks, and cybersecurity type of skills together with what they define as technological literacy are the fastest growing skills.

Metacognitive skills, such as self-awareness, planning with appropriate goal setting, adapting, and adjusting learning strategies are important skills. In this category, we can also include project business management skills, data skills, and obviously specific AI skills from designing to training models, creating AI products and agents, and translating business needs into AI-based solutions. Some of these skills have the unique characteristic of unlocking other valuable skills and opening new career paths. We called them hub skills. More on this topic in Section 2.

Human—AI Interaction

In this debate of whether AI will drive mostly automation versus augmentation for your workforce, it is worth a brief mention of the fact that this may happen with or without human agency. Human–AI interaction can take different forms: it can come about with individuals having maximum agency and deciding when the AI copilot or AI tool is used; however, it can also occur as ambient AI with much lower human agency. We already have many examples today, from thermostats in smart homes that automatically regulate the temperature, to AI assis-

tants at work that remind us of meeting conflicts, or even remote patient monitoring devices that help with preventive care of diabetes, hypertension, and weight management. Augmentation of humans may take unexpected forms: there is a disturbing new suite of AI audio companions[41] where the AI can emotionally connect with individuals, provide companionship, coaching, and other therapeutic services. It looks like the world described in the movie *Her* with Joaquin Phoenix (directed by Spike Jonze) filmed in 2013 is becoming a reality in 2025.

The balance between automation and augmentation and its form factors will differ by business and sector and is a clear fork in the river, a clear polarity for challenging future scenarios.

Role of Big Tech

Industry structure has always been a critical component of strategy. Big Tech can determine a critical polarity in the Era of AI: there are two potential outcomes of this known-unknown. AI infrastructure and primary services are predominantly managed by a limited number of entities that establish terms and pricing, similar to patterns observed in the Digital Era for certain services such as advertising. An alternative approach would involve more balanced competition, providing established companies and start-ups with a broader range of partnership options. It is important for businesses to recognize potential dependency on specific providers and take steps to control their long-term value generation strategies.

Let's review the context for this known-unknown and try to identify the signs to pay attention to.

Winner-Take-All Dynamics of the Digital Era

The Digital Era (2000–2020) brought about a unique and unprecedented scale of business, with selected players reaching billions of users. A combination of exceptionally large upfront investments paired with the low marginal cost of serving the next user drove certain players to uncontested dominant positions with the resulting attractive economics that come from large market share in scale-driven businesses. Additionally, platform business models attracting buyers and sellers compounded their scale benefits with powerful network effects creating exponential growth and, yet again, unparalleled economics. "Winner takes all" has dominated the competitive landscape of the Digital Era, with search and social media as the poster children.

This period lifted FAMGA companies (Facebook, Apple, Microsoft, Google, and Amazon.) to unprecedented scale and value. Microsoft and Apple were already present in 2004, but the others were mere start-ups, and, today, they collectively represent approximately 20 percent of the S&P 500's value.

Regardless of who the large-scale players in AI end up being, the more fundamental question for the average business is whether Big Tech will dominate the AI Era, or whether it will create room for more balanced competition. This polarity determines the speed of AI adoption, the economics behind AI, and potential choke points for incumbents and new businesses.

Think about digital advertising; in many countries it is in the hands of two players (Google and Meta). Their algorithms and pricing drive the economics of a multitude of businesses relying on them to attract customers. When Google changes its algorithm to improve user experience, to avoid spam or low-quality content, many businesses are heavily impacted without much room to maneuver. For example, Reddit's CEO blamed the recent Google algorithm change for falling short of its target for unique daily users causing a 15 percent drop in share price.[42] Think about e-commerce, if Amazon decides to no longer sell goods from a particular manufacturer, it may make a difference between life and death for that company. Reliance on Big Tech sends waves through multiple companies.

"Winner Takes Most" Dynamics of the AI Era

As we look at economic models behind AI service providers it looks like "winner-take-all" economics will continue to apply. So, do businesses need to operate with the assumption that their AI needs will be covered by a Big Tech that sets the rules of engagement and prices? Will we be in the hands of an oligopoly yet again? If so, who are the likely players?

While the prevailing assumption is that the core layers of the AI infrastructure—illustrated in Figure 4—will remain firmly controlled by Big Tech, forming oligopolies with outsized market power, an alternative and entirely plausible trajectory is emerging. Some layers of the stack may evolve into commodities, much like traditional infrastructure such as electricity distribution or mobile network towers—essential, but conferring little strategic power. At the same time, growing evidence suggests that smaller players can carve out highly differentiated positions within the ecosystem, building compelling businesses with attractive economics rather than being crowded out by scale alone. There are economic and political drivers that may challenge the current situation; business leaders will need to look for the signposts in this particular polarity.

The AI Era presents some meaningful differences from the "winner-take-all" economics we observed during the rise of digital platforms. While large capital commitments remain critical barriers to entry, the nature of competition has evolved. Building and training foundation models, along with constructing world-class data centers, demands investments in the hundreds of millions—if not billions—of dollars. In 2024 alone, the U.S. Big Tech giants (including Nvidia) poured more than $210 billion into AI infrastructure and research and development (R&D).[43]

Current infrastructure players are pulling ahead through unique scale advantages and self-reinforcing network effects. Models improve not just by accessing more data, but by harnessing better quality interactions—feedback loops where users train the system while the system adapts to users. This dynamic increases user lock-in and reinforces platform dominance. Moreover, the larger the scale, the more likely these players are to secure priority access to scarce graphics processing unit (GPU) capacity, further entrenching their leadership.

So yes, the combination of massive upfront investments and strong feedback loops still drives "winner-take-all" dynamics. But unlike the "as-a-service" or the platform business models of the Digital Age, the economics of the AI age come with a caveat: serving the incremental user is no longer "free"—or even cheap. The cost to serve the next user is material, especially compared to the near-zero marginal costs of past digital platforms. Each query consumes not just compute cycles and electricity but also requires substantial memory, particularly for long context windows—and low-latency delivery. That means dedicated GPU slices for each user, dramatically altering the scale economics. While these marginal costs are expected to decline over time, they remain non-trivial today—so much so that some providers discourage users from thanking their AI agents to avoid unnecessary token usage and cost.

In this context, Big Tech companies are no longer incentivized to chase scale at all costs. There's a built-in brake to the classic race to the bottom pricing dynamics. There is a chance we do not see again the repetition of Microsoft's historic strategy of giving Internet Explorer away for free, which eventually crushed Netscape's early lead. This marks a departure from the Digital Era. Netscape, the pioneering browser company, lost its edge not because of product flaws, but because it couldn't compete with Microsoft's distribution power and predatory pricing. In contrast, in AI, cost structure itself offers a form of strategic restraint, preserving space for differentiated players and limiting runaway consolidation—at least for now.

Additionally, there are alternatives to Big Tech models; there is evidence that open models—like Llama from Meta, DeepSeek, or Hugging Face – and highly specialized models – like Harvey for legal uses or Claude Code for coding, are

gaining traction. The winners will be those who can either operate at scale with minimal marginal costs to serve the next user, or those who control the key chokepoints—be that the hardware, the operating systems, or specialized content and/or proprietary data.

What Strategy is Big Tech pursuing Around AI?

Big Tech firms have taken different approaches to AI; for now they seem to make moves consistent with their core strategies. Let's briefly review them to identify particular Big Techs relevant for your business.

Starting with Apple, it seems to focus on providing value with its closed ecosystem and eventually drive more device sales by leveraging popular "good enough" open-source models for on-device use. It currently has basic on-device large language model (LLM) capabilities and a proprietary cloud infrastructure. The chip architecture allowing good on-device performance could be an interesting differentiator if it were truly integrated with the rest of its AI infrastructure. This assertion would lead to the development of its own AI models or to an acquisition—both would require a fundamentally different capital expenditure (capex) strategy from the current one. Apple seems to have opted for minimal AI-driven investment. While Apple is not immediately threatened by AI, early incursions into AI-driven hardware by other Big Techs—such as Meta's partnership with Ray-Ban glasses and OpenAI's acqui-hire of Jony Ive and other former Apple product designers—may erode Apple's dominance in premium hardware. Disappointing 2025 announcements around Siri's AI overhaul and growing open access to foundation models raise questions about the sustainability of its differentiation. Apple's tight control over hardware and operating systems still gives it strategic leverage, but the AI revolution has introduced an element of platform vulnerability not seen since the rise of Android.

Meta, in contrast, seems to be investing heavily, with capex commitments of $60–65B in 2025 and significantly larger for 2026, mostly dedicated to developing its open-source foundation models—called Llama. For Meta, AI should continue to drive engagement in their social platforms with enhanced content, as well as drive better monetization from advertising. As Meta is one of the main global advertising channels, this approach makes sense. It is fascinating to see that they may be willing to take their current outcomes-based pricing to another level with the ultimate business agent; in Zuckerberg's own words from a 2025 interview:[44] "Here's the business outcome that I want, here's what I'm willing to pay, I'm going to connect you to my bank account, I will pay you for as many business outcomes as you can achieve." Meta is also looking to leverage momen-

tum with WhatsApp in Asia to bet on business messaging with AI. Last but not least, Meta is building its own AI model. Meta AI has approximately one billion users across different Meta platforms (Facebook, Instagram, Threads, and others). However, there is the perception that Meta is playing catch-up—that their AI models are only "good enough." As a result, Meta has restructured their AI lab and has initiated a very visible war for talent, hiring talent from OpenAI, Microsoft, and Apple, with the specific goal to reach artificial general intelligence. Meta's ownership of content platforms and social data could give it a long-term advantage in model alignment and behavioral prediction, especially if LLM agents become integrated into social or advertising ecosystems.

Turning now to Google, it is the player under the greatest threat from generative AI given the substitution effect of their core search product. The perception thus far is that they are late to the game. Ironically, Google's DeepMind team generated the famous 2017 paper "Attention is all you need,"[45] which introduced the transformer architecture that underpins the generative AI revolution. Despite this seminal insight, Google seems to be on the defensive, committing $75B in 2025 on AI investments across its full stack: from developing its own chips to growing its cloud infrastructure and improving its multimodal foundation model, Gemini. In addition to its infrastructure investments, Google is developing new applications such as "AI overviews" to enhance its core search offering and protect advertising revenues. Google's ability to control both AI infrastructure (tensor processing units (TPUs) and massive data (YouTube, Gmail, Search queries) still gives it leverage. Looking at foundations models, Gemini seems to have finally delivered a competitive product against existing benchmarks. If search moves from showing a list of options to providing a single generated answer, monetization strategies must be re-evaluated; even a shift by OpenAI or Perplexity to advertising could significantly impact Google's core advertising revenue.

Amazon's AI strategy focuses on its core cloud infrastructure—Amazon Web Services—with expected investments of $100B in 2025, including development of its own chips for improved AI performance. It is currently partnered with Anthropic to access competitive foundation models. The combination of hardware (chips), infrastructure (Amazon Web Services (AWS)), and strategic model partnerships positions Amazon as a critical hyperscaler only contended by Microsoft. It may not control the final user, but it controls the essential infrastructure, and its model optionality can be a powerful differentiator compared to Google but not relative to Microsoft.

Likewise, Microsoft is committing $80B to grow its cloud infrastructure (Azure) and strengthen its partnership with OpenAI to deliver Microsoft's Copilot and other enterprise AI tools. Their bet seems to be on developing agentic AI products deeply integrated into the operating system and Microsoft suite, offering

business users continuous productivity gains. After years of exclusive partnership with OpenAI, Microsoft is becoming model agnostic, and business partners will be able to access via Azure a suite of models depending on performance and needs including xAI, Claude, or Mistral.

These strategies will evolve over the coming decade. History shows that the list of top companies by valuation will certainly only include a small subset of these names by 2040. From 2000 to 2020, only Microsoft and Walmart remained in the top ten most valuable global companies. New players will emerge. Nvidia, for example, has exceeded a $4Trillion valuation on the back of being the world's dominant AI hardware provider.

We cannot complete the review of Big Tech without including OpenAI—it does not have the scale of the players above but is a clear candidate to become a Big Tech player in the Era of AI. By July 2025, OpenAI had 700 million weekly active users; Meta, by the end of 2024, had less than 200 million weekly active users of its AI assistant—out of more than 3 billion total Meta users. With this momentum, OpenAI expects to raise $40B at a $300B valuation to expand research and build its own data centers. With expected 2025 revenues of $3.5B and a net loss of $5B, it's still orders of magnitude smaller than the other players, but a strong contender for the next top ten list. Control of proprietary data and usage at scale —not just model performance—will determine who turns research into strategic advantage.

The AI Era brings a different angle to the competitive landscape of Big Tech. While in the Digital Era (2000–2020) Big Tech had very distinct brand franchises and territories the AI Era tends to blur the boundaries and get these Big Tech to compete on a more direct basis. Google owned the franchise of search, Microsoft the franchise of business applications, Apple the franchise of hardware, Meta the franchise of social media, and Amazon the one from e-commerce. The AI Era is exposing much more evidently to the common battleground to provide AI infrastructure. Globally, Amazon Web Services (AWS), Azure from Microsoft, and Google Cloud hold more than 65 percent of the market share and continue to grow. In the distance, Alibaba and Baidu are also relevant players in selected geographies. There is a non-zero probability that current AI infrastructure (cloud services and foundational models) gets commoditized. Sheer differentiation becomes ever more difficult, leaving the decision to use one cloud provider versus another one to cost of using the different services of compute, storage, networking, database services, or analytics. Interestingly differentiation comes in the form of discounts or price per second. Cloud services have become more standardized, and they are increasingly interchangeable. Furthermore, the foundation models follow the same path. As recognized by Satya Nadella himself—CEO of Microsoft—in 2024, "As AI becomes more capable and agentic, models themselves become more of a

commodity, all value gets created by how you steer, ground, and finetune these models with your business data and workflow."[46] Microsoft has opened the door to accessing different foundational models (ChatGPT, xAI, or Mistral). So, access to proprietary models through a particular cloud does not seem to be a differentiating element for cloud providers. The unbundling process is likely to continue. At one point it seemed that cloud providers would be kingmakers for foundation model companies like Anthropic or Open-AI.

In addition, in the same way that electricity or the Internet became necessary goods and governments had to ensure access to electrical power and/or the Internet it is plausible that access to AI infrastructure becomes a necessary good—at least for businesses—and that basic access to cloud and foundational models becomes a commodity. Even if it takes decades to reach that point, the jury is still out on the ability to differentiate the AI infrastructure.

This fork on the road calls for a novel approach to dealing with your own AI capabilities as a business: either focused on complementing and enhancing Big Tech of your choice or, on looking for the smaller tech that best suits your business needs. Depending on the sector and country in which the business operates, this may not be a choice at all. Thus, the importance of strategically preparing for such a scenario.

Artificial General Intelligence

The degree of advancement of AI capabilities is another important known-unknown. A number of tech leaders believe that artificial general intelligence (AGI) could be a reality between 2025 and 2030. For example, Google DeepMind considers AGI achievable by 2030.[47] Other tech leaders and researchers tend to think there is a lot of hype, and we are further away than we think. A very vocal voice on the latter current is Yann LeCun, the ex-chief AI scientist at Meta. In the spirit of identifying key uncertainties that can be plausible in the time horizon of our analysis and challenging your existing strategies, as a leader you ought to consider the possibility of truly reaching AGI.

Since there is ample debate on the topic, let us pause to clarify definitions and align on expectations. Morris et al. (2023)[48] provide a good framework to define AGI, with "Exceptional AGI (Level 4) . . . defined as an AI system that matches or exceeds that of the 99th percentile of skilled adults on a wide range of non-physical tasks."

Some opinion leaders even venture the possibility of reaching artificial general super intelligence (AGSI) in the near future. Let us have DeepSeek define artificial super intelligence:

> Artificial Superintelligence (ASI) refers to a form of artificial intelligence that surpasses human intelligence in all domains, including creativity, problem-solving, emotional intelligence, and general wisdom. Unlike Artificial General Intelligence (AGI), which matches human-level intelligence across a broad range of tasks, ASI would exceed human capabilities in every way. It would not only perform tasks faster and more efficiently but also innovate, reason, and make decisions at a level beyond human comprehension. ASI would be capable of recursive self-improvement, meaning it could enhance its own algorithms, architecture, and capabilities without human intervention, leading to exponential growth in intelligence.

A more intuitive definition, by Ben Thompson, a technology analyst, is that AGI can correctly answer questions in a variety of fields, and ASI can produce the questions in the first place.[49]

Without getting into too many semantic nuances, it seems like a version of artificial general intelligence (AGI) is plausible within the time of the exercise proposed in this book. As we saw in Chapter 2, there are already models that, in addition to performing at par or beyond humans in specific tests demonstrating deep knowledge of certain fields of law, or coding or specific medical fields, are also capable of reasoning and, lately, of showing empathy and compassion. It is therefore plausible that, before 2040, AI can evolve into something akin to artificial general intelligence (AGI). For practical purposes, let's assume that the big polarity in your strategy is accepting AGI. If you prepare for AGI, I will hypothesize, your business is closer to adapt to ASI if it were to occur.

Chapter 4
Plausible Challenging Future Scenarios

So, as a leader, what do you do with all these known-knowns and known-unknowns? This section aims to equip you with the necessary skills to effectively distinguish valuable information from irrelevant data, to distinguish signals from noise, and to prepare you to pursue emerging AI-driven opportunities ahead of your competitors.

To be clear, the reason to identify future scenarios is not to forecast the future. Honestly, I would have a hard time giving a hardwired set of strategies for the Era of AI when it is hard to keep up with the pace of innovation. The goal of defining future scenarios is to identify early signals of opportunity and risk, and to rethink strategy ahead of competitors. Future thinking gives the confidence to leaders to "dive the deep" in uncertain times and overinvest when everyone else is selling or holding. The ability to anticipate also allows to build ahead of time certain organization capabilities that take time.

How do businesses take advantage of future thinking? By anticipating critical changes and creating optionality. In other words, by dedicating a subset of resources to investing in new potential business opportunities with outsized upside and limited downside. Optionality also serves as an insurance policy to protect the business from major risks. For some readers, scenario planning to improve risk management may be the more interesting approach. We will review in Section 2 the specific actions from scenario planning for your consideration.

For now, it is all about focusing on depicting a future that is relevant and plausible and that is designed with the main purpose of challenging existing business practices and models. It is plausible because there is already early evidence of the potential trend/business opportunity, something we call "pockets of the future." It is something that is happening in your industry, something that would have been lost in the excess of information we all receive, but because you have been cognitively primed by the scenarios you and your team have developed, this "something" becomes a valuable sign and helps trigger an action.

Being aware of the relevant "pockets of the future" for your business is like working at a company with an especially important advertising budget needing to improve its effectiveness and spotting the impact of search in the early 2000s. It is like starting to work with Google when it was private (its Initial Public Offering was in 2004) and reaping the full benefit of relevance-based ad placement before any of your competitors. Partnering with them might not only have yielded great Return on Investment (ROI) in ad investments, but, as importantly, the ability to build differential digital marketing capabilities ahead of competition. In

https://doi.org/10.1515/9783111583495-004

other words, spotting Google's potential in advertising early on may well have been the source of unparalleled growth opportunities.

Another example. This idea of spotting "pockets of the future" is like being a player in the automotive sector and looking for energy storage options in the early 2010s and identifying CATL (Contemporary Amperex Technology Co., Limited) from China. At the time, China was clearly stating in its five-year plans the priority of developing its battery innovation and rare earth supply security; several Chinese players emerged. CATL started in 2011, riding the wave of the push to electric vehicles (EV) and is now a company with more than ~$50B in revenues and almost 40 percent market share of EV batteries. For those who were looking, important strategic players were available and could have transformed their business in the following decade. Identifying "pockets of the future" can move the needle in one's business strategy but it requires a framework for separating noise from signal.

The mindset is to look for plausible scenarios, not probable. I insist that this is a quite different exercise from traditional corporate budgeting or strategic planning, where management identifies the most critical variables and does a sensitivity analysis on them generating the classic best, base, and worst-case scenarios. With this traditional approach, management focuses on a linear future; the approach is not designed to imagine alternative relevant futures. Often, this failure of imagination sequesters management teams and boards.

The possible combinations of the four known-unknowns described in the previous chapter yields 16 different scenarios.[50] The selection of four scenarios that follows has been selected because they can be truly challenging; they may not be probable, but they are certainly plausible. The selection has focused less on society and global order and more on the impact on businesses' value creation and competitive moats. To take this exercise from conceptual description to concrete changes we pick different sectors to illustrate the potential impact. It is not about predicting the evolution of a particular sector—it is about making tangible a potential reality. Companies cited exist at the time of writing and can represent a "pocket of the future."

The reader needs to go into this exercise looking for what components of these scenarios may be relevant to the business at hand and identify specific implications for their organization. See below – Figure 10 – a graphic representation of the polarities used to define provocative future scenarios.

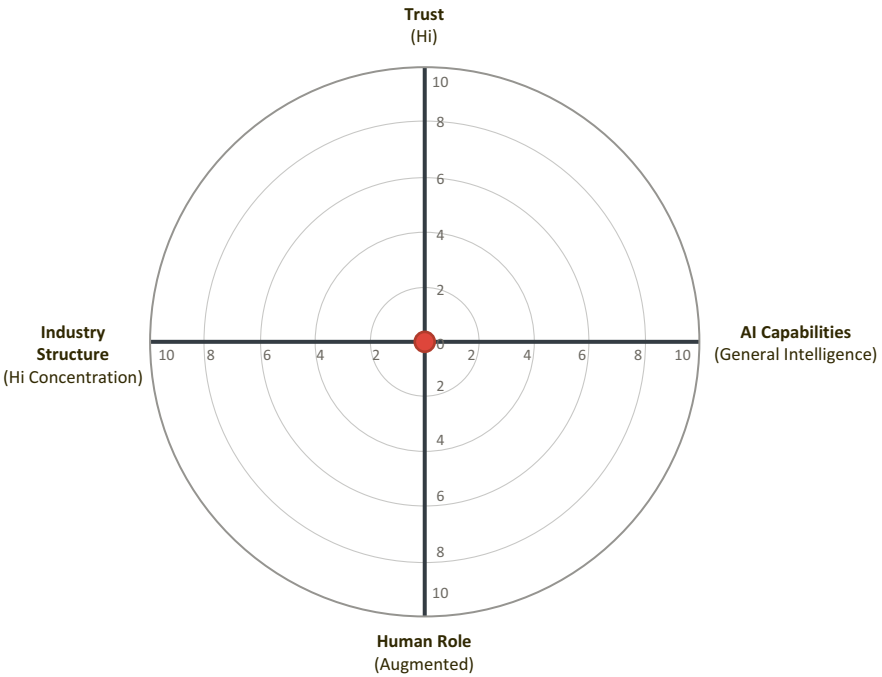


Figure 10: Key variables defining future plausible and challenging scenarios. Each access is graded from 0 to 10; there is no intention to establish a direct or indirect relationship between variables in opposite vectors.

The AI-Run Ecosystem

The Scenario

In this scenario, depicted in Figure 11, we see an AI system having automated a meaningful pool of jobs; artificial general intelligence is a reality, and it is trustworthy, partly due to checks and balances enforced by a distributed ecosystem of small tech players operating on an AI infrastructure that has become a commodity. Let's see what it looks like to operate in this scenario.

Let me share a story of what the world would look like under this scenario. In this case, we will use healthcare services as an example. We insist this is not a forecast of what the sector may look like, only an exercise with plausible changes that look to challenge current strategies. While the main company is imaginary, the other names mentioned in the story already exist and provide a "close enough" service to what we see in the story. They are the "pockets of the future"

that could lead us to believe the scenario is plausible. They are also the type of companies on which one may reflect in 10 years and think: "I wish I had partnered with these companies."

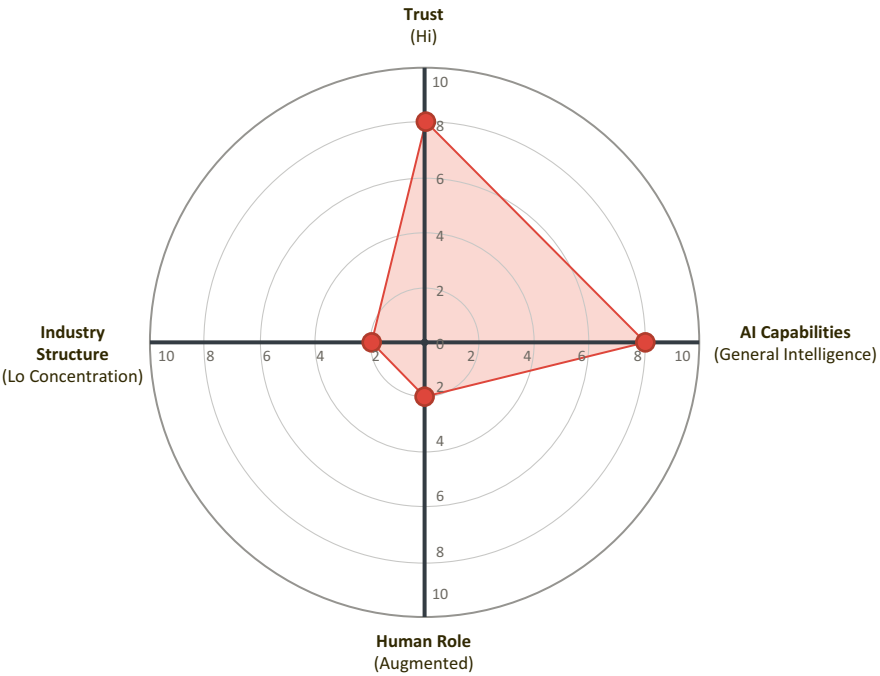


Figure 11: Visual summary of domain-specific AI-run Company Scenario.

The Story

It is 2030, Karine wakes up feeling unwell. Recently, she has been experiencing unexplained fatigue, intermittent numbness in her left hand, and occasional blurred vision—symptoms that are not only concerning her health but also threatening her ability to care for her children and maintain her income. She is 42, based in Raleigh, North Carolina, United States, divorced, and looking after two young children.

Detection and diagnosis: Rather than calling her primary care physician as she would have done in the past, Karine opens her K-Health symptom assessment tool. She answers a series of increasingly specific questions generated by its diagnostic AI. The system identifies potential neurological concerns and immediately connects her with nQ Medical, a specialized neurological screening service using passive

data from smartphones. Within hours, Karine receives a notification that a preliminary analysis suggests she should undergo further testing. Her healthcare portfolio manager from HealthCo (which evolved from insurance provider to healthcare orchestrator) automatically receives the alert and schedules appointments with specialized diagnostic services like Mindstrong Health. It provides a cognitive assessment through their specialized neurological function monitoring app. None of these services employ neurologists directly—instead, they each focus on perfecting specific types of neural assessment with specialized AI. The data is automatically assembled by her HealthCo portfolio manager and sent to Path AI, which specializes solely in neurological diagnosis patterns. Path AI delivers a preliminary diagnosis of multiple sclerosis, which is verified by a consulting neurologist at Summus Global who reviews and approves thousands of similar AI-generated diagnoses daily. For now, all the diagnostics have been performed with no human in the loop.

Treatment plan: Karine is referred to a single multiple sclerosis specialist who will not only receive the diagnosis but a suggested treatment from a constellation of specialists including medication from an online pharmacy, a neurological rehab plan from Neuro-Rehab, and a diet recommendation. The specialist focuses on explaining the disease and suggested treatment and, very importantly, focuses on empathizing with the patient. The specialist is there for emotional support. Several independent third parties have validated the treatment proposed. From then on, Karine has her special AI agent—that looks and sounds like a human—from Hippocratic AI that will be checking in that she is following the treatment. She visits her local hospital for her annual, in–depth check–up. The hospital relies on Sully.ai, an autonomous operating system that supports its operations through a suite of AI agents: an AI Receptionist that manages scheduling, an AI Triage Nurse that coordinates referrals and lab work, and an AI Scribe that transforms doctor–patient conversations into structured clinical notes.

Karine's healthcare expenses work differently than traditional insurance: Rather than paying co-pays to a single provider, micro-transactions flow to a constellation of specialized services. Her HealthCo portfolio automatically negotiates rates with each service based on quality performance metrics. For example, the price of the services is adjusted based on the number of highly reputed research sources consulted, or based on the number of external data used to triangulate the diagnostic, or based on the ability to explain how the model has been adjusted to the fact that Karine is a women in her 40s and, therefore, research without an explicit inclusion of her gender/age in the sample will be underweighted. She receives real-time cost–benefit analysis for each specialized intervention. When a new specialized treatment option emerges from Healx (rare disease treatment discovery platform), her HealthCo manager automatically evaluates its efficacy and cost compared to her current treatment constellation.

Food for Thought

In this scenario, Karine has taken full benefit of best diagnostic and treatment available globally regardless of her location; given she must look after two children it was especially convenient to do most of the diagnostic and treatment from home. She was also able to meet with a specialist, especially in the early phase of the diagnostic, to address the anxiety and need to understand the implications of her disease. This was only possible because the regulations and open-source models have made the relevant data and AI solutions highly trustworthy and interoperable. She also realizes the potential limitations of a system relying on coordinating multiple providers and ensuring consistency among them; the insurance portfolio manager played a key role.

The speed and performance-based economics of the interactions point to an efficient system. The ecosystem is managed by machines, not augmented, run by machines. AI agents, schedule, optimize, and decide; employees follow along. Incidentally, players in the ecosystem do not sell to other businesses or to individuals; they sell to agents. Brands do not exist.

How to uncover opportunities from this scenario for your business? In each organization with sufficient size, it is possible to find people who can think about the business in three different horizons we described in Chapter 1. Horizon one is predominant; most managers are looking to fulfill their budgets and look for incremental improvements in their top line or in the way they manage their cost base. Horizon two is three to five years down the road and is where most entrepreneurs thrive. Look for these profiles in your organization and provide them with an entrepreneurial space where they can pilot services, alliances, and so forth. In horizon three, 10 years down the road, the business model has been transformed. This is where visionaries operate. If the three profiles exist in your organization and have been identified, there is some clear tension among the three and a real battle for allocation of resources. Constructively managing this tension to create a portfolio of real options catering to different time horizons can be an attractive way to futureproof your company in the Era of AI.

The competitive moat in the AI-run Ecosystem Scenario relies on several variables. An incumbent in this space could only benefit from this scenario if it had a well-defined data architecture with clear roles and permissions and had interoperable systems with accepted protocols. This would be the focus of managers looking at horizon one. They may also start using agents to materially improve the diagnostic and treatment services and proactively reduce errors in recommendations before they happen.

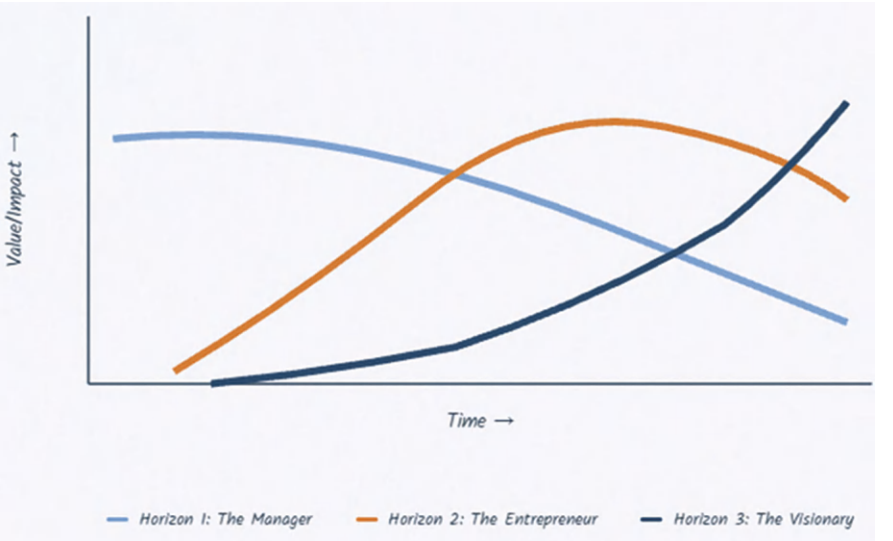


Figure 12: The three horizons.[51]

Another competitive moat, looking at horizon two – illustrated in Figure 12–, the entrepreneurs in this scenario, would be looking to test performance-based pricing in some of the Services. For example, an insurance player may lower prices to patients who rigorously follow preventive treatments; or HealthCo may pay a premium for diagnostic services that reveal all their sources and their ability to contextualize them based on the patient's personal profile. Other entrepreneurs in this ecosystem may want to be trust builders and create a business by validating diagnostics or giving second opinions on treatments.

The visionary looking at horizon three may try to run certain practices with maximum reliance on AI systems, removing unnecessary steps and people in the process. Innovation in this era can benefit from much faster prototyping and, in a year or two, could have a robust multiagent system pilot to deal with the majority of Karine's needs in both diagnostic and treatment phases.

In resource-scarce organizations it is often a matter of channeling existing innovative initiatives with a three-horizon lens. The limited range of options should be appropriately focused on the key opportunity or threat.

! ***Your prompts:***

- *Are you in a business susceptible to play in an AI-run Ecosystem Scenario with a constellation of micro-service companies?*
- *In this case, are you likely to be the orchestrator or the specialized provider? Who would your final customer be?*
- *Who would you want to partner with? Are there less visible partners that your team may have overlooked?*
- *What new capabilities will you need to build?*
- *Do you have the right data? Can you generate powerful and differentiated insights? Have they been validated by a relevant third party?*
- *Do you need to adapt your monetizing strategies? What is the expected impact on your current profit and loss (P&L)?*
- *What are the most relevant and plausible risks? How will you prepare the organization and the leadership team for this evolution?*

Digital Feudalism

The Scenario

It is a scenario where Big Tech has strong monopoly power: your data, your customers, your infrastructure all accessed on someone else's terms. You do not own the rails, Big Tech does. As described in Figure 13, The tools and systems have been deployed to mostly automate large pools of the workforce, but AI capabilities have not yet reached the level of general intelligence, largely due to lack of trust; commercial solutions and regulators of all nature have not managed to overcome the public mistrust in AI.

As in the previous scenario, I will take an illustrative sector—here we will use financial services as an example. The description of the scenario is not intended to give a forecast of the sector in the age of AI; it is only an exercise to shed some light into plausible and relevant challenges introduced by the Era of AI. The company names mentioned in the story already exist and may be delivering services like what is portrayed in the story, or close enough. They are used as "pockets of the future".

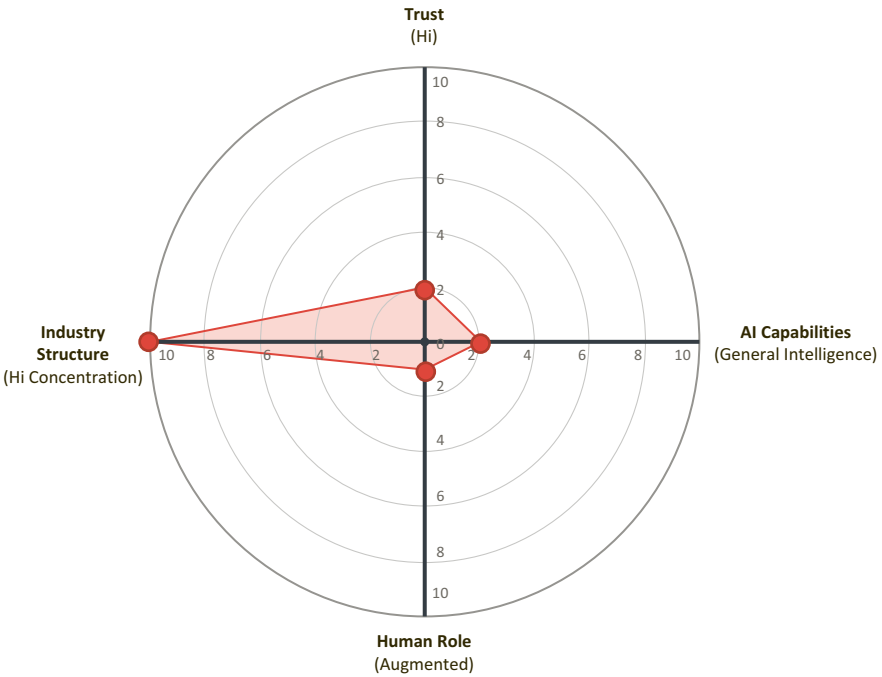


Figure 13: Visual summary of Digital Feudalism Scenario.

The Story

It is 2037. Pho Lin is sitting asking its AI assistant, Wei, to apply for a small loan for his business as a hail service driver in Hanoi, Vietnam. The answer comes back in three seconds. It has been denied, but he needs to wait three days until Wei comes back with the audited explanation as to why he did not get the loan. It looks like his hail cab activity is about to enter the low season, and the number of new licenses has grown exponentially in the last 12 months. The AI agent has placed Pho in the high probability of default segment for the next six months. Pho feels frustrated because he cannot appeal the decision; most banks use the same AI-based risk models provided by one of the large global AI hyperscalers, Alibaba in our story.

The hyperscaler also controls the layer of deep expert credit scoring models coming from two China players and one Indian player compliant with central bank regulation for each country in the region. The wealth and quality of data used to train these expert models have made them the unbeatable source to effectively assess critical variables like probability of default (PD), loss given default (LGD) and the updated true value of collateral assets. It is by licensing the suite of expert mo-

dels to banks that the hyperscaler makes attractive margins. AI infrastructure has been deemed an essential need by the government of Vietnam, and its usage needs to be affordable; AI infrastructure is profitable if used by at least one billion users.

The bank no longer owns the interface; they are behind a "curtain" controlled by one of the large consumer tech companies, incidentally also the hyperscaler, Alibaba. Alibaba has the final interface with consumers, but data privacy concerns and mistrust on both banks and AI systems have led to very stringent data usage requiring regular audits and controls that often slow transactions. Guardtime—a European company with the highest standards in privacy and security—has become the trusted third-party validator. Banking and other financial services are provided by an ecosystem of retail banks from the region who participate in an automated auction with algorithms that price loans and mortgages.

Pho receives notice after a few days from the company insuring his car that there will be an increase in the premium; it most certainly may have been triggered by the loan denial. He still has the option to revert to a parallel financing platform of peer-to-peer lending set up by players like Grab, the financial group developed out of the hailing service. It has higher costs, but it can meet its needs. Pho will use one of his three digital identities, to avoid "algorithmic contamination" across the different systems.

In 2037, retail banks have fundamentally redesigned their structures and activities. From customer service to internal audit and compliance departments, it all gets handled by significantly smaller teams from what they used to have in the early 2020s.

Their business model has evolved: from expert capital allocators across economic stakeholders to managers of trust and customer experience. Banks have become clients not owners of the banking intelligence stack. Some banks chose to overlay their own expert models to finetune pricing and credit risk management practices of extract value of the decimal basis point, others chose to build special capabilities dedicated to managing perceptions and trust by turning layers of compliance and controls into becoming explaining agents, others chose to double down on the personal relationships with clients.

In this world, the hyperscaler provides everything AI. It turns out that the Big Tech player has been heavily supported by regional political support to guarantee AI sovereignty. The truth of the matter is that this support has commoditized critical services like access to cutting edge foundational models; Big Tech has been driven to develop several layers of intelligence adapted to specific sectors. Hyperscalers are deriving unprecedented economics from high-volume–low margin cloud provision services to the more value-added high margin business of customer acquisition services. Their unique data set deepens their competitive advantage and makes their dominant position unchallenged. Hyperscaler has had a

few lawsuits, but it uses the fragmented anti-trust ecosystem in Asia to defend its business. Even if it has had to pay some penalties, it is an insignificant amount relative to overall profits; sanctions are a cost for doing business.

Food for Thought

This is a scenario where Big Tech is a real gatekeeper of both AI infrastructure but also in access to users. Automation has also led to important price competition in sectors like banking and yet there is fundamental public mistrust on both banks and AI services and recommendations. This has created a layer of controls and compliance demands that have hindered the pace of adoption of new AI functionalities, and, despite all the automation, it has risen operating costs significantly. In this scenario, banks have experienced remarkably high price pressure given commoditization of models and important level of automation. Scale dynamics have exacerbated the power of incumbents since AI agents allow for micro-targeting at massive scale with minimal marginal cost. Pressure on margins is brutal. Service differentiation becomes impossible when competitors must go through Big Tech to both acquire customers and service their AI models.

In horizon one, managers are hostage to Big Tech. They are focusing on better unpacking customer acquisition costs and on the effectiveness of the AI risk models compliant with banking regulations for optimizing the two largest cost components in their P&L. The business-to-business (B2B) business is still more relationship based and has, for now, escaped Big Tech.

In horizon two, community or sector focus can be an option to complement the extremely competitive core business. In our story, alternative financing platforms for businesses are looking to solve the financing needs of the employees of these businesses. Giving a loan in the low season to a business that has a good track record in the high season is less risky for these alternative platforms. For the entrepreneurs in this scenario, as in other scenarios with low trust, the world of compliance and verification opens a wide range of opportunities. Additionally, building complements to the Big Tech and having immediate access to their scale can be a profitable venture.

In horizon three, power has shifted from hyperscaler to alternative financing platforms, with intimate knowledge of users and a trusted relationship that has allowed for unprecedented share of the wallet of a financial service's needs—from micro loans to mortgages—to new investing and saving solutions based on new digital assets.

! ***Your prompts:***

- *Is your business likely to play in a Digital Feudalism Scenario?*
- *What is the Big Tech that may hold your business captive?*
- *How will you protect your business from Big Tech squeezing your margins?*
- *Across your business model, including the way you create value, the way you offer it, and the way you monetize it, what will be the basis for your true differentiation?*
- *Is it worth creating new options to complement your core business outside of Big Tech? How will you fund the required investments? Are there new partnerships you need to explore?*
- *What AI capabilities and non-AI capabilities can make a difference?*
- *How will you prepare the leadership team for this evolution?*
- *How will the shape of your P&L change? How will you offset the increase in costs of operating with Big Tech?*
- *What are the most relevant and plausible risks?*

The Augmented Enterprise

The Scenario

In this scenario, the augmentation of individuals prevails, and Big Tech players have worked with key organizations to ensure high levels of trust and reliability (see Figure 14). AGI is a reality, especially for certain verticals like high-value manufacturing, or energy distribution that has gone all-in in adopting AGI systems. Large purpose-built technologies dominate certain sectors; they provide centralized augmented intelligence further augmenting humans. All caveats from previous scenarios apply in that this is not a forecast, just a thought experiment. Humans are in the loop; however, companies have become over-reliant on these truly augmented individuals.

As in previous scenarios, I will take an illustrative sector—here we will use an example from the industrial sector. The description of the scenario is not intended to give a forecast of the sector in the age of AI. The main company—DesalinationCo—is made-up; the other names mentioned in the story already exist and may be delivering services like what is portrayed in the story, or close enough. They are used as "pockets of the future" to give some realism to the future scenario.

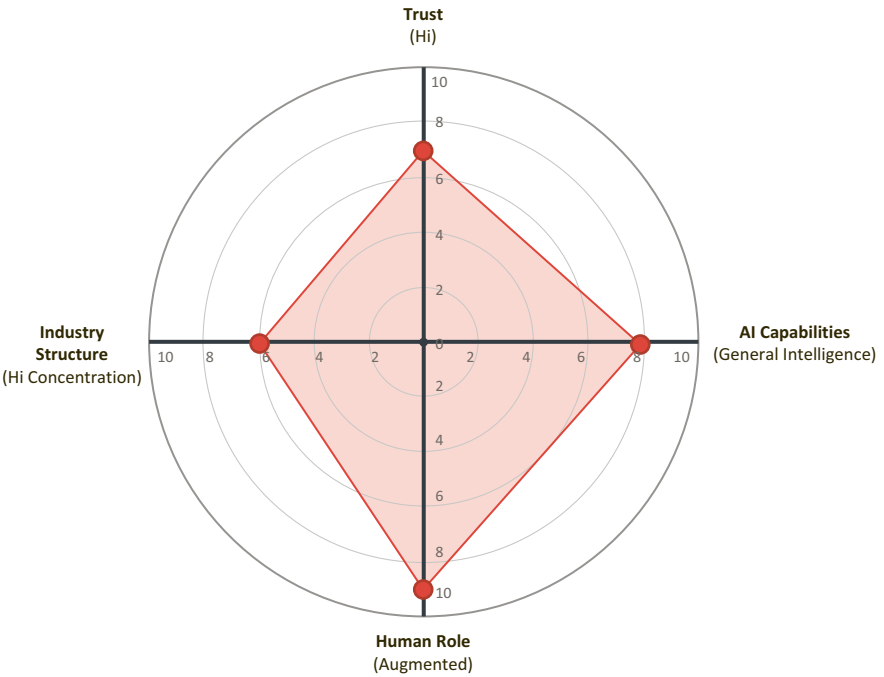


Figure 14: Visual summary of Augmented Enterprise Scenario.

The Story

Razan is entering the factory of DesalinationCo (imaginary company) on yet another day of work of this warm day of April 2032 in Kuwait. "Good morning, Razan," Pi's voice greets her; it is the automated co-pilot she was given when she joined. "I've analyzed the production anomalies from overnight and have three optimizations ready for your approval." Razan listens to the proposals and responds: "The third option looks most promising, let's implement it across plants 4 through 9."

With her approval, Pi instantly reconfigures dozens of robotic systems and smart materials membranes. What would have once required weeks of engineering meetings, testing protocols, and gradual deployment now happens in seconds. The network of plants responds like a living organism, its components communicating with one another through AQUAVISTA™ solution from Veolia—now incorporating AGI. Razan and her copilot are connected to AQUAVISTA™—like other plant managers across the global network of plants from DesalinationCo.

Later that day, Razan joins an operations planning meeting where her colleagues—both physically present and telepresent—merge their augmented insights through AQUAVISTA™ shared cognitive space that not only facilitates a form of collaborative thinking and problem solving that transcends traditional communication. Peter from South Africa is concerned, "I'm sensing potential resistance from the team using new graphene-enhanced membranes. Our lab operations analysis suggests possible contamination build-up under hypersaline conditions exceeding the limit we established." AQUAVISTA™ proposes a new membrane configuration, and the team agrees because it can guarantee 98 percent efficiency at the same energy consumption. There is complete trust in the AGI system; it has helped DesalinationCo reduce energy requirements for desalination by 78 percent over five years while increasing freshwater production through its plant network fivefold. The company has broken the paradigm where energy costs made desalinization cost prohibitive. The addressable market has multiplied thanks to AGI operating improvements.

Food for Thought

There are several competitive moats emerging from the Augmented Enterprise Scenario worth considering for existing industrial players: the specialized AI system developed by a pure player that became the Big Tech in that sector owning and managing AQUAVISTA™. In horizon one, the AI system developed a unique set of expert models for optimal running of desalinization plants that only a player with the scale and experience of Veolia could build. As a result, a mayor growth inhibitor for plants—energy—has been reduced to unthinkable limits, making the addressable market much more attractive. It is inspiring to think of AGI dedicated to solving a paradigm shifting issue for an industry.

Each plant can be managed by a small number of AGI augmented specialists supported by well-trained technicians who perform physical activities that cannot be automated. They are also supported by specialized co-pilots and a productivity platform such as Augmentir, which enhances their effectiveness from the outset and enables rapid dissemination of new best practices across DesalinationCo. All employees are connected through a network and consistently engaged in ongoing learning and optimization.

In horizon two, the entrepreneurs in the organization have developed a new business model of water-as-a-service model where base line pricing is determined as a function of quality, energy usage, and reliability. Not meeting the thresholds on any of these variables comes with a penalty accepted by DesalinationCo, a unique differentiating factor for municipalities buying from some of Desalinization-

Co competitors. This new business model keeps DesalinationCo on its toes, driving continuous improvement and sustained investment in refined use cases.

In horizon three, the visionaries at DesalinationCo, proposed a pilot for a potential adjacent business to process the brine left over—highly concentrated salt water—for key mineral extraction like lithium or magnesium. While this model may pose significant platform dependency risks on AQUAVISTA™, DesalinationCo has structured a contract based on shared value creation that keeps the balance of power. The potential of the adjacent business can, in time, balance the sources of profit and thus the reliance on specialized Big Tech.

Your prompts:

- *Is your business likely to play in the Augmented Enterprise Scenario?*
- *How will you protect your business from the Specialized Big Tech squeezing your margins? How can you balance their power?*
- *What AI capabilities and non-AI capabilities can make a difference?*
- *What data do you have that they do not? What insights do you have that they do not?*
- *What may be the new shape of your P&L?*
- *Do you have the right talent currently in the organization? Are they susceptible to being "augmented"? What skills and profiles do you need to look for? How do you retain the uniquely augmented individuals who run your business?*
- *How can you take full advantage of the amplified intelligence type solutions to scale your business faster? How will you manage the transition with your stakeholders—investors, clients, employees?*
- *Are your shareholders ready for a business model change? What if you moved to outcome-based pricing?*

The Personalized Tech-Mosaic

The Scenario

Trust in AI has failed; institutions need to operate in a highly fragmented environment that inhibits augmentation and productivity gains – see Figure 15. This is a scenario where small, highly specialized AI applications operate in an environment of great distrust and privacy concerns. AI solutions run on the edge, separated from institutional clouds. These applications truly augment individuals in a subset of functions. No AGI has truly emerged.

In this scenario we will use the sector of higher education. The main company—UniX—is a fictional university; the other names mentioned in the story are used as "pockets of the future".

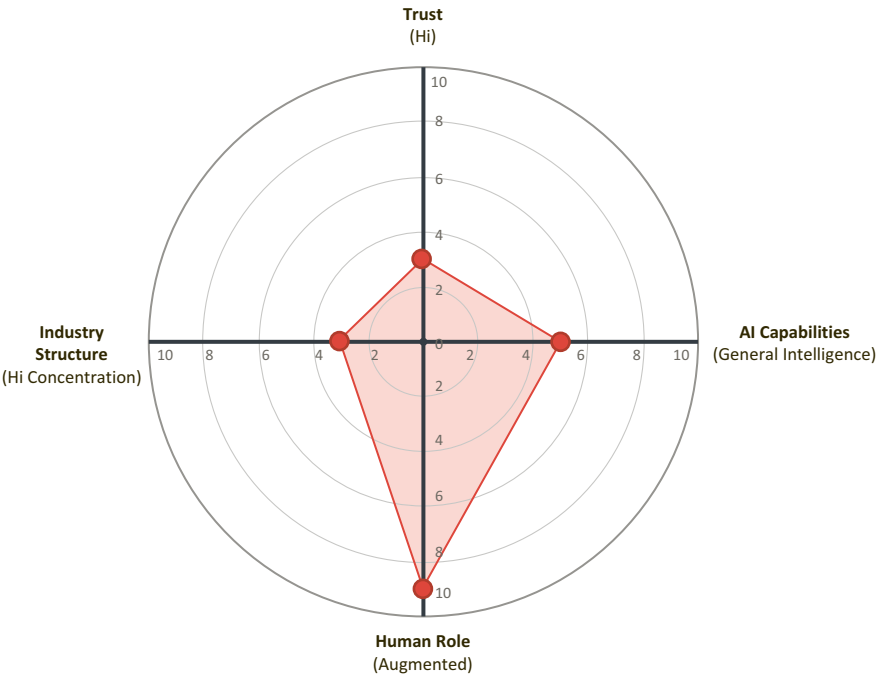


Figure 15: Visual summary of Personalized Tech-mosaic Scenario.

The Story

Professor Aswath Garcia is starting his Foundations of Finance class at public UniX in Marseille, France, in 2033. He opens his copilot with all his teaching materials, his research and databases compiled over the years with things like historical cost of capital by sectors, or equity risk premiums. He is teaching a set of augmented students with access to AI-powered tutors—like Khanmigo. Students got access to their personal AI tutor when they started college a couple of years ago. They have had the chance to fine-tune their personal tutor with the particular courses they have taken, so that the tutor can help them build interdisciplinary connections. In addition, the personal tutors hold the memory of all the exercises and projects done from the start of their studies at university; tutors have a deep understanding of the knowledge gaps and strengths, of the particular level of metacognitive skills, their students' preferred way to practice and receive feedback, and other information important for supporting students in their learning. They are typically used in airplane mode, and only a selected dataset is shared with faculty and UniX administrators. There are a wide variety of tutors used by

students, and UniX encourages students to bring their own tutor (BYOT). Students, if they choose, get to pay to keep their tutor after graduation, as a lifelong learning companion.

Incidentally, this university has recently broken with the ages and stages education paradigm, and, as a result, Aswath is teaching students in their twenties together with a few students in their forties and fifties who wanted to get a certificate in Foundations of Finance. Aswath gets from his co-pilot a set of questions to start the class discussion. They are based on what the co-pilot has identified as most problematic topics for this particular group of students. Aswath co-pilot has a way to know what this class is struggling with based on submitted exercises and questions. He proceeds to clarify how betas are used in cost of capital calculations and other concepts relevant for today's class on business valuation.

At UniX, both faculty and students know that learning is an emotional and social activity, and they use this fact to deliver more effective learning. Students had all of Aswath's materials before starting the class. Despite their augmentation, they had to put in at least a couple of hours of work to solve the case that Aswath had given them as homework. The case was on the cover page of the *Financial Times* earlier that week: IndustrialCo has made a bid to buy TechCo. Students need to prepare the relevant analysis to share with TechCo's board of directors about what to do. There will be an open debate, and Aswath will look into how students' reason, how they problem-solve, and how they react in real time to critical decisions. It turns out to be a fun exercise. Students learned from Aswath's depth of knowledge, his challenging approaches as well as from his charming personality. They also learned from the well-designed learning materials and from the personalized and empathetic AI tutor who forced them to practice before the class with a focus on the areas each student needed to emphasize to not to be left behind. Students also learned from playing in class, from practicing with a diverse group of students of different experiences and ages.

When the year is about to finish, UniX's administrators go through a painful process of gathering all relevant information for evaluation of students, faculty, and program. The information is in a variety of systems that cannot integrate easily due to lack of transparency and explainability of the AI models used. The task is even more challenging because of UniX's stand on privacy; protecting all data from students and faculty makes the process very cumbersome. The final evaluation system will integrate knowledge and skill-based evaluations from class and non-class activities, but also qualitative assessments about the preparedness of students to be better citizens. The Ministry of Education performs additional validation on students' identity and skills, which requires another set of complex processes to triangulate existing data before issuing the official degree. The Ministry gives special awards and economic incentives to faculty with demonstrated track

record for effective teaching; that is faculty who have achieved greatest improvements on "before and after" levels of knowledge and skills. Faculty with the awards for effective teaching are highly sought after by universities.

The business model of UniX, a public university, is heavily reliant on donations from alumni to support research as well as the ongoing improvements from the different AI tutors and co-pilots. If Alumni get beyond a threshold salary three years after graduation they commit to a gift back to the university so that others can enjoy the same good fortune. It is a small amount and only for another three years. By 2033, there are early signs in the EU that augmented education and research is worth funding with individual contributions.

Food for Thought

In the Tech-mosaic Scenario AI solutions can improve the effectiveness of different players in the higher education ecosystem. They are individual solutions that do not integrate well within the ecosystem, largely because of lack of economic incentives. Higher education continues to be resource constrained. There is, However, a myriad of easy to build AI solutions built on robust fundamentals of learning science like personalizing deliberate practice, leveraging social activities, or encouraging diversity of thinking to foster better problem solving.

Horizon-one-focused-managers are prototyping integrated evaluation solutions based on AI; they are looking for an improvement over traditional standardized exams that address integrity issues raised by the widespread use of generative AI tools. They are connecting with regional and global efforts and creating a highly reputed community of practice that regenerates assessments and evaluations recommendations. The main motivation here is risk management; it is to protect the value of their degrees.

Horizon-two-focused-entrepreneurs are looking for a suite of AI agents that enhance the entire learning journey of students and become genuine learning companions for the different areas of study at UniX. They have managed to get the endorsement of faculty who are keen on getting the Ministry of Education awards for Effective Teaching. UniX is driven by its mission to equip students with the skills needed to succeed in the job market and, more broadly, to help them become responsible and engaged citizens.

Visionaries concerned about horizon three are prototyping end-to-end AI systems that give individuals the agency to track their lifelong-learning journey results. The solution goes beyond a digital skills wallet; it includes personal, profes-

sional and learning information and is supported by a personal AI coach who gently nudges individuals to take the next step. Learners are willing to pay a small monthly fee for the learning coach in the same way they pay for high-value sports apps, or for movie streaming services. UniX is providing this service, partially financed with an EU grant aimed at fostering a common and transparent skill wallet for workers across the EU. UniX gets additional licensing fees from employers looking to tap into students with a combination of skills set and knowledge. The students allow UniX to share specific information under special agreement with the university because it enhances their employability. It is always shared at the discretion of the student who can, at any time, change the settings about what information is shared with what type of employer.

Your prompts:

- *Is the Tech-mosaic Scenario relevant for your business? Are you seeing a myriad of truly helpful solutions for your stakeholders that are not interoperable?*
- *What options should you be creating for each horizon?*
- *What pilots would you like to run? What do you need to learn from them?*
- *What is the best model for your organization to manage experimentation at the same time as you continue the maniacal drive for results?*
- *What AI capabilities and non-AI capabilities can make a difference?*
- *How will you manage the transition with your stakeholders—investors, clients, employees?*
- *Can you fund the level of investment required? How will you measure ROI?*

Section 2: **Monday 8 a.m.**

This quote from Bill Gates in his book *The Road Ahead* (1995) about the emerging Internet times continues to be of relevance in preparing for the AI Era: "We always overestimate the change that will occur in the next two years and underestimate the change that will occur in the next ten. Don't let yourself be lulled into inaction."[52]

It is not too late to prepare for what your organization can be in the next five to ten years. Let's build in this section of the book your strategy co-pilot with a set of clear guidelines.

Upon revisiting the assessment of known-unknowns, you have thoroughly examined potential polarities that may pose challenges to your organization. These pivotal decision points can also represent rare and significant opportunities for growth and advancement.

It is Monday at 8 a.m. when your bold ideas confront the hard edge of reality: legacy systems, slow approvals, fragmented data, confused teams. It's not that you lack insight: the future isn't a threat; it is terrain to be shaped. How do we go from vision to results?

https://doi.org/10.1515/9783111583495-005

Chapter 5
Dare Re-Coding Your Business Model

Business Model Innovation

Context

Future scenarios have helped uncover challenging yet plausible scenarios. In this process, you and your teams may have identified opportunities to serve new customers and new geographic markets. You may picture your business developing novel solutions and competing head-to-head with competitors who you barely encountered before.

As I stated in the preface of the book, I place a strong focus on superior and sustainable value creation. Questioning traditional views regarding market boundaries, core customers, competitors, and cost structures allows for the paradigm shifts needed to pursue the opportunities of the AI Era. To fully capitalize on these emerging opportunities, it helps to adopt a flexible and wholistic approach to your business model design. By business model design I mean how your business creates value, how it is offered, and how that value is captured.

Traditionally, when doing strategy exercises in the face of change we would review the industry structure following the five forces defined by Professor Michel Porter from Harvard Business School back in the 1980s. A review of the competitive intensity of existing players and of new entrants, an assessment of the power of suppliers and the power of clients together with the analysis of potential substitutes was the starting point of any strategic review. Industry structure and the assessments of the firms' unique assets and capabilities helped define a set of competitive advantages. Yet the speed of change, the less adversarial nature of both customers and suppliers, and the influence of data in shaping and evolving industry structures has made this framework insufficient for defining business strategies. Competitive advantages were no longer based on stable market structures; they also shortened their lifecycle and became transient—as described by Professor Rita McGrath from Columbia Business School. Designing business strategies had to include dynamic business model design. The Digital Era (2000–2020) further challenged some of the core tenets of classical strategy. It made evident that building unique differentiation was no longer mutually exclusive with the idea of striving for cost leadership. As a matter of fact, companies like Spotify with its personalized play list or like McGraw Hill Education or Pearson with their adaptive learning tools are good examples

https://doi.org/10.1515/9783111583495-006

of how the Digital Era models allowed personalization, therefore differentiation, at low cost.

I would argue that business model innovation compounded the positive effects of technological innovation in communications, in data storage, and in compute we saw in the Digital Era.

As we think about this holistic approach to business model innovation it helps to break it apart into its component parts described in Figure 16. Researchers tend to point to value proposition as the starting point of business model innovation, yet I would contest that, in certain situations, business model innovation has been triggered by changes in the ability to create value; in other words, supply side forces have led to value proposition changes. The Digital Era has some interesting examples.

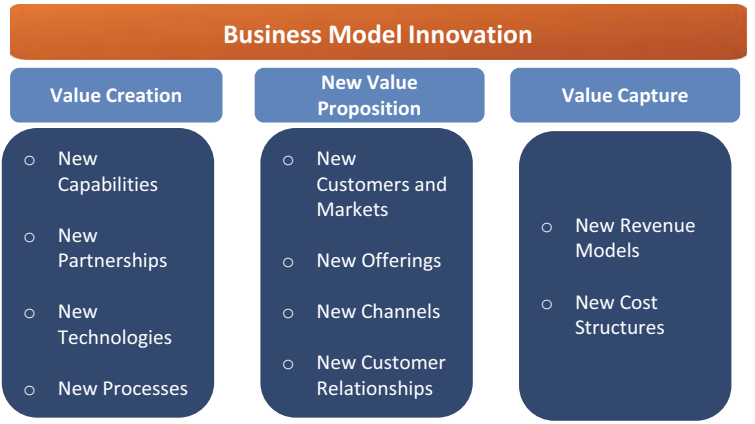


Figure 16: Business model innovation.[53]

Business Model Innovation from the Digital Era (2000–2020)

Business model innovation from the Digital Era generated unparalleled economics and can bring interesting inspiration for the AI Era.

"As a Service" Subscription Business Model

"Software as a Service" brought continuous innovation to unexpected levels. It emerged as a need to constantly adapt to new customer demands and provide

better technical solutions at speed and at competitive costs. It made traditional software sales in the form of packaged versions with a static set of features inadequate for market demands. The traditional model missed the complex, multifaceted approaches to value creation that actual Software as a Service (SaaS) models rely on today.

As per Figure 16, the strategic tenets of an "as a Service" (aaS) business model are predicated on both an interesting value proposition innovation and on value creation innovation. The first is about customer focus and continuous innovation as a source of value creation; it constantly offers performance improvements and new functionalities based on customer feedback. It drives differentiation relative to competition and ensures sustained relevance and customer stickiness. The latter is particularly relevant because a key value driver of "aaS" businesses is the Lifetime Value of customers. This business model relies heavily on driving better Lifetime Value of customers for every cohort of new customers: to achieve this, businesses work on optimizing customer acquisition costs (CAC) but equally important on ensuring customers stay for long and that they raise their average spent by buying additional features or services.

The second key strategic element of the "aaS" business model is its value creation innovation. That is the new processes that allow for scalability by design: it feeds itself on real time data with APIs, on modular architecture, and ensures operating costs that provide maximum operational leverage as the business grows. It is this very aspect of the business that drives low marginal costs to serve the next client. The users only consume the outcome; they have no visibility on the complexity of managing a scalable infrastructure.

In general, this business model has delivered unique economics with attractive recurring revenue based on monthly/annual fees, with gross margins above 70 percent, ROE in the mid to high 20 percent, and valuations averaging five times revenues.

Software as a Service (SaaS) gave way to a larger market of subscription based or "as a Service" suite of businesses that removed friction for users and enabled smoother delivery for manufacturers. Examples of this category of business model include Infrastructure as a Service (IaaS) (for example, Microsoft Azure or AWS), Data as a Service (DaaS) (Snowflake or Kpler), Analytics as a Service (AaaS) (Qlik), or Logistics as a Service (LaaS) (Ninja Van). Other examples include classic business-to-consumer (B2C) subscription services like Netflix or Amazon Prime. However, be mindful about the fact that there are many other businesses that are using the "as-a-service" label as a marketing gimmick; if there is no subscription fee that is competitive, that relies on low marginal costs to serve the incremental client, and that continuously improves the service, then it probably does not deserve the name.

In the Era of AI, we are already seeing this type of business model innovation promoting greater access to AI in a more scalable and flexible fashion. Let me describe a few examples that can either be an inspiration for your own business, or, as importantly, a better way to use AI-related services. We see at the infrastructure level there are new models emerging allowing for alternative uses of the cloud: Training as a Service (TaaS) allows for renting cloud compute, tools, and pipelines to train machine learning models at scale. Examples of this include players like Lightning AI, Paperspace Gradient, hyperscalers are also providing this type of service. In this category and given that inference costs (from querying models) are becoming an important cost item companies have started to offer Inference as a Service (IaaS)—outsourcing model inference to optimize for cost, latency, or energy efficiency. Examples include Runpod, Banana.dev. If we move up the AI tech stack into the model layer, we also see examples of Model as a Service (MaaS), pre-trained machine learning models delivered through APIs or platforms, rented on demand. Most current foundation model players offer a version of MaaS, providing access to their models via application programming interfaces (APIs) and charging a fee per token of queries (input) and answers (output). We of course continue to see the concept of Data as a Service (DaaS) that is becoming particularly relevant as models need to be fed with almost infinite data to train and fine tune their models. We see companies like Scale provide labeled training data for machine learning, or Mostly.ai provide synthetic data generation, or Quandl providing financial data sets on demand. We are even seeing Algorithm as a Service type business models that provide access to proprietary machine learning algorithms for optimization, search, or analytics without exposing the full model. In the application layer, it looks like the concept of Agent as a Service, could become a popular business model. The companies following this model allow for deployable autonomous AI agents that perform tasks like research, sales outreach, or code generation.

While “aaS” models from the Digital Era can be a source of inspiration for the AI Era, they also have their own challenges, especially as they incorporate AI-based services. The cost of serving the marginal customer is not as small as it used to be. Managing the inference costs—costs to query a model—is demanding a change in pricing schemes, from fixed pricing to a base price complemented by usage-based pricing for certain features and/or outcomes-based pricing. So far, most of the AI “aaS” models oriented to consumers seem to be free of advertising, but OpenAI has already opened the gate for advertising as a source of revenue with the potential to offset marginal costs to serve the next customer.

Platform and Ecosystem Business Models

Platform business models of the Digital Era are a genuine example of business model innovation in all three building blocks: new value proposition, new value creation infrastructure, and new value capture and monetization. Digital platform services emerge from the creation of a marketplace that attracts buyers and sellers where transactions are made in an effective manner: with low transaction costs and an important level of trust and reliability. Imagine a traditional farmers' market where the owner arranges the stalls and ensures there are clear rules followed by everyone so that buyers and sellers are attracted to go to that market. The owner does not make the products; they only guarantee the conditions for the trade. Well-known platform models include e-commerce players like Mercado Libre in Latin America, Alibaba in China, and Amazon, or, more specialized platforms, like Booking aggregating global hotel offerings for the traveler. These models are essentially open; that is, the criteria to participate in the platform is not very restrictive.

In contrast the ecosystem business model looks to attract buyers and sellers but instead of looking like the farmer's market it looks more like a "members-only mall," where all of its component parts have been carefully curated to provide a smooth experience; retailers, products, entertainment, and food options have been selected and prepared to harmoniously integrate with one another to provide a cohesive and unique experience. In here, the criteria for being part of the ecosystem is much more restrictive. Examples of ecosystem model include the Apple Store where all apps are scrutinized for certain experience and privacy standards. The main difference between these two business models is their focus: platforms' main focus is breadth of choice and transactions; the ecosystem's main focus is to provide a seamless experience. A sub-segment within the platform and ecosystem business models are the ones developed by the sharing economy allowing individuals or companies to tap into partial use of assets—think of car riding models like Cabify or Uber, as well as apartment sharing like Airbnb. The success of this model is largely based on commoditizing trust while providing smooth customer experience.

The business model innovation brought about by platform and ecosystem business models was a powerful new value proposition where goods, contacts, and reservations were only one click away. They were possible because of new value creation capabilities and technologies heavily reliant on partnerships and the purposeful activation of network effects. Platforms and ecosystems business models both rely on a more open market relative to what was possible in the pre-digital age—pre-2000—and in different degrees they benefit from a significantly more open sourcing of products and services; they both are dependent on powerful network effects. That is, the value of the business increases as the number of

participants—users and suppliers—grows. Generating network effects is there by design; these businesses look for new high-quality nodes that create non-linear increases in collective value. The magnitude and distribution of that value depend on network structure, incentives, and engagement dynamics.

Platforms brought about an interesting business model innovation in how they monetized value created. I hope this dimension is of inspiration to new business models of the AI Era: they introduced outcomes-based pricing in a fairly novel way. When these business models emerged, advertisers could finally know what the 50 percent of their advertising budget was being wasted. I am referring to the popular saying, started by John Wanamaker, an American department store merchant, who complained about half of its advertising budget being wasted, but the problem was to figure out which half. Google and Meta have developed a pricing system for advertisers based on expected outcomes. They price for impressions, for clicks, or, even better, for conversion, providing much better visibility on the return on investment of advertising dollars. The economics of these businesses tend to be extremely attractive as well, with gross margins that can reach 80–90 percent, operating incomes of 40 percent and ROE North of 30 percent. The dispersion around these metrics is high, but it is possibly the best business model ever created: network effects are a unique source of scale and growth, and for platforms that are asset light, the margins are unprecedented.

Platforms can be a great source of inspiration in terms of business model design for the AI Era, but as in the previous case of "aaS," these platforms are also experiencing their own challenges to adapt their business model to the AI Era. They may well become marketplaces for AI agents buying and selling from other AI agents. Advertising supported business models will need to evolve; current players are caught in Clay Christensen's innovator's dilemma.[54] As a reminder, Clayton Christensen, from Harvard Business School, explained why successful companies often fail when disruptive innovations emerge. Established firms in any sector, including the technology sector, tend to focus on sustaining innovations and improving existing products for their most profitable customers, because that's what their business models and incentives reward. But this focus blinds them to disruptive technologies that initially seem inferior or cheaper or to be serving niche markets. Over time, these disruptive innovations improve and capture mainstream demand, displacing incumbents who ignored them. The dilemma is that by doing what seems rational, listening to current customers and maximizing profits, incumbents unintentionally set themselves up to be disrupted.

Let me illustrate this dilemma by looking at the poster child of platform businesses, Google. Google has been the great sponsor of the Internet for the past 25 years, with great incentives to grow its usage and to provide an interface that

helped users navigate chaos. However, the market is evolving from a search interface to an answer interface that probably addresses the needs of more than 90 percent of the users more than 90 percent of the time. So now traffic stays on top of the page and does not have a need to scroll down web pages and click on the blue hyperlinks below the answer, and Google is no longer capable of driving traffic to its publishers, and its monetization schema of auction-based pricing to appear on top of the page is much less compelling. The incentives to protect current revenue streams challenge the ability to change—the typical innovator's dilemma. There is a growing need to develop a new business model for the AI Era —one in which AI companies, such as answer engines, share a portion of their revenue with the content creators whose work powers their systems. Some suggest a new business model based on pay-per-crawl (paying for accessing content on the web); others are looking to redefine the browser experience with a richer, more contextualized search that also allows agents to go check things (other prices, definitions, calendar availability) without losing the original tab. Whether monetization will be as attractive with these new models as it was with advertisers, it remains to be seen. This is all being figured out as I write.

Value Accelerators from the Digital Era Relevant for the AI Era

As some readers may come from non-digital businesses, for them I would like to emphasize the value of open systems and the deliberate pursuit of network effects; these two attributes that emerged in the Digital Era are powerful value accelerators that can be even more relevant in the AI Era.

Traditionally, organizations would keep a more internal focus to product innovation based on their own product and R&D teams and the ability of marketing teams to uncover unmet customer needs or underserved customer segments. Additionally, these organizations would tend to review their value chain with a focus on control and efficiency. This approach would limit the possibilities of incorporating external sources of inspiration for product and service design as well as for delivery options. One of the fundamental shifts of the Digital Era was how business innovation was done. The focus moved to continuous value co-creation; business models in the Digital Era create value through an open network of users, developers, and suppliers that benefit from participating in a more open ecosystem. These models do a good job at leveraging the collective intelligence of relevant stakeholders. Leveraging the power of "crowd sourcing" is a common theme of many of these business models. It is based on the principle that "the crowd" with diverse expertise will generate better solutions for complex problems than a single individual or a more homogenous team within an organization.[55] There is, by now, ample evidence that the benefits of collective intelligence could more

than offset the cost of coordination and alignment. The software industry has turned this system of collective intelligence into a key feature of its business model. While the level of openness in engaging with developers varies across software companies, the concept of incentivizing others to make your products better accelerates innovation and drives product-led growth generating unexpected yet invaluable use cases. The incentives include anything from a percentage of the revenue generated to more intangible benefits like building career capital and visibility for particular communities of the likes of GitHub.

A form of crowdsourcing has been widespread practice in health research but was made evident at the public level with the COVID-19 pandemic. The creation of vaccines was accelerated by regular sharing of insights from experiments and trials done in different parts of the world, eliminating duplicate work at a time when stopping COVID-19 expansion at a global level was urgent. The initial phases of discovery and trial were done in months when it typically takes years. Researchers sharing the genome of SARS-CoV-2 within weeks of the virus being identified (January 2020) and its diffusion on platforms like GISAID and GenBank together with pre-existing mRNA technologies allowed the exploratory phase of a typical vaccine to go from the standard one to five years to barely a few months in the case of COVID-19. In the pre-clinical phase, sharing clinical data on protein folding simulations, accelerated by multiple open-source collaborative platforms from academia and research organizations, allowed for faster identification of vaccines by players like Pfizer or Moderna. It is true that the collaborative approach and "crowdsourcing" of health researchers and academia results in potentially better solutions at an accelerated pace relative to "individual" research. However, I must mention that in the case of COVID-19, the secondary effects of such a fast process are not fully known yet. The thought of augmenting medical scientific discoveries powered by AI with a more open ecosystem of researchers gives plausibility to some of the aspirations to eradicate some of the lethal illness affecting humanity. I was struck by the impact of Every Cure, using the power of pattern recognition of AI to uncover new disease-drug combinations allowing to repurpose approved drugs to treat diseases once considered untreatable.[56] The Every Cure AI platform pulls its information from scientific publications across the world paired with curated biobanks with medical records from millions of people or repositories of chemicals and their safety profiles.

Back in the world of business, the research conducted by Henry Chesbrough on open innovation ventures that one of the reasons for the failure of IBM's Watson—their AI-based commercial solution—was the fact that IBM built a closed system all under its control. IBM did not include Software Development Kits (SDKs), no application programing interfaces (APIs), or custom integrations with other systems to engage with developers and users beyond IBM. IBM

had developed an early AI rule-based solution to solving chess problems; its Deep Blue computer won against chess master Garry Kasparov in 1997. Its follow-on version, Watson struggled to deliver on the economic promise, at least in healthcare. Watson Healthcare was sold for €1B after more than a decade of unsuccessfully competing in the AI space.

So, as future scenarios prompt ideas for business model innovation you want to stay mindful of the ecosystem your business is part of and look into how to tap into the collective intelligence, how to align incentives for others to improve your products. Remember Karine in the scenario of the AI-run ecosystem? HealthCo simply couldn't exist without its ecosystem. To support Karine through the diagnosis and treatment of her multiple sclerosis, HealthCo had to rely on a network of health specialists—and it needed to create strong incentives to keep them engaged and aligned around her care.

Network effects are created when new users increase the value of the product or service to other users. Network effects may be direct, when an increase in users directly attracts additional users, as in the case of WhatsApp. They may also be indirect, when growth in one user group enhances value for another, as illustrated by BlaBlaCar, a European long-distance car-sharing company, where an increase in drivers raises value for riders. Every new user added to a network can exponentially increase its value—a hallmark of platform business models. Today's leading companies intentionally design for these network effects, with social media standing out as the poster child: a platform where each new connection amplifies reach, engagement, and influence.

If we continue with the AI-run Company Scenario, HealthCo was able to benefit from network effects by building interoperability with different specialists by design, but also by using the value of incremental new data to pull players that can benefit from it (independent validators of diagnostics).

Dynamic Value Chains

Let me highlight a typical business model dilemma that has challenged companies over the years. In classical strategy, decisions concerning how to serve a target set of customers involved making choices about how to best deliver the key steps along the value chain—from sourcing to final delivery, even to post-sale support. The decisions of how far a company should drive vertical integration was a common conundrum driving decisions on where to allocate scarce resources and proprietary capabilities versus where and when to turn into outsourcing or partnering. Yet the definition of value creation as part of the business model ought to be dynamic and should change as the competitive dynamics change and industries mature.

Let's take a relevant example in the AI Era that is agonizing over this very set of decisions: Intel—the global semiconductor manufacturer that dominated the market for several decades and has lost its leadership to Nvidia in the AI Era.

Some historical context makes this story even more dramatic. Let us get back to the 1950s in the United States when William Shockley improved the original transistors—granting him the Nobel Prize—and eventually created the semiconductor in Mountain View, later called Silicon Valley. After that, more than 60 companies were created in the span of 20 years to manufacture semiconductors. One of them was created by scientists Gordon Moore and Robert Noyce, engineer Andy Grove, and financier Arthur Rock—one of the early VC investors in Mountain View. It is no other than Intel Corporation. The development work focused on Integrated Circuits (ICs)—-a device that integrated semiconductor devices such as transistors and diodes and passive devices such as resistors and bonded them on a printed circuit board (PCB). Robert Noyce invented and patented a new integrated circuit using silicon around the 1960s. In addition to the design of integrated printed circuit boards, Intel's strategy was focused on developing unique chip foundry capacity to create the first memory microprocessor destined to mainframe clients.

In an article in 1965, Gordon Moore had forecast that the number of components that could be packed into an integrated circuit would double every year (later changed to every two years) for at least a decade with no increase in power consumption. His revised forecast implied a compound annual growth (CAGR) of 41%. While Moore did not use empirical evidence in forecasting the historical trend would continue, his prediction held since 1975 and has since become known as a "law."[57]

From memory chips Intel invented the first microprocessor in 1977. Back in 1985 it had already set the foundation for its most successful processor x86, core to the later version of the Pentium in the 1990s. The unique chip and an aggressive strategy with Original Equipment Manufacturers (OEMs) led to the ultimate differentiation strategy with "Intel Inside" as a synonym for a PC with unique processing power. Intel created a unique moat based on its differentiation strategy with strong R&D investments doubling the power of chips every two years, continuing to upgrade and extend its manufacturing capabilities. The combination of both capabilities permitted speed to market moving faster than competition to next generation processors from the 8086 to the 186 286, 386, and 486. The moat was reinforced by endless suits against other chip manufacturers trying to copy or reverse engineering its patented chip architecture. To be clear, a differentiation strategy does not mean that costs do not matter. Microprocessors were only viable if waste was reduced to a minimum, and unit costs remained competi-

tive. Andy Grove was lead-engineer in charge of driving and optimizing manufacturing capacity, and he eventually became CEO.

As of the time of this writing, the world's fastest AI chip holds four trillion transistors. So, while the complexity of chips has grown exponentially, the fact that they need to be produced at nanoscale has driven up significantly the cost of building those factories. According to Arthur Rock, there is a less well-known second Moore's Law: the cost of manufacturing a chip plant will double every four years. Currently building a chip foundry may range from $10B to $20B. This elevated level of capex has become a real barrier to entry for new players.

Though a classic business model challenge, reassessing dynamic value chains is very relevant in preparing for the AI Era.

Value Chain Choices

The choice of strategy should determine the level of vertical integration, that is, the level of direct participation and control of the different steps of the value chain. The general principle was to control the parts of the chain of greater value addition. This criterion evolves as industries mature. For example, early mobile operators had to build towers to ensure proper cell-phone coverage, but as the industry matured and the operation and maintenance of towers became a commodity, telecom operators disposed of those assets. There is a need to constantly re-assess what parts of the value chain create a sustainable competitive advantage. Why did your business come to have the value chain you have today? Economists will highlight three main reasons:

- Controlling costs by avoiding intermediaries and the need to negotiate, monitor, and enforce contracts with suppliers. This is the line of thinking developed by economist and Nobel Laureate Oliver Williamson with its Transaction Cost Economics (TCE). For Intel, this integration should have resulted in lower transaction costs, thus greater efficiency; and it was one of the factors leading to its overwhelming success during the 1990s when it focused on microchips for PCs and servers.
- Achieving greater market power: the vertically integrated company has more control over its time-to-market and pricing power. In the language of economist Michael Porter, integration of suppliers and distributors generates higher barriers to entry and, therefore, a more sustainable competitive advantage. Intel achieved unprecedented market power in PCs and servers largely thanks to OEM agreements with the likes of IBM.

– Vertical integration can also be a way to hedge risks, as recently evidenced by COVID-19 or the war in Ukraine; companies that controlled their value chains were more flexible to address these shocks.

Did Intel forget to re-visit its direct participation in the different parts of the value chain with ever more complex chip designs and the fact that some of its competitors were offering manufacturing capacity to fabless players? Arguably, designing chips and manufacturing them have become two related but different businesses with fundamentally different cost structures and economics and a vastly different set of competitors.

On the one hand, some players in the semiconductor industry have focused on building foundries and second to none manufacturing capabilities; the poster child for this business model is TSMC—Taiwan Semiconductor Manufacturing Company created in 1987. Making the right equipment choices and looking to maximize scale efficiencies is key to making this part of the business profitable. Gross margins of this business are, at best, 15–20% of revenues. Cost efficiency is paramount, yet future sustainability is based on the ability to produce ever smaller chips. Now the race is for one or two nanometers chips.

On the other hand, there is a group of fabless players focused on the design of chips (like Nvidia and now AMD in the UK) for highly specialized uses. Gross margins in this group of players are more than double the margins of the previous group—50–60 percent of revenues—and their capex comes less in the form of factories and equipment and more in the form of R&D.

In this crossroads, Intel Corporation is one of the few companies in the industry following an integrated strategy—doing both design and manufacturing. On the one hand, it competes with the likes of Nvidia for the best-in-class design of chips; on the other, it competes with the best in class in manufacturing with the likes of TSMC. Pat Gelsinger, CEO of Intel up until 2024, initiated that year an effort to make transparent the profitability of the two main components of its value chain: the foundry business (a money losing business looking for customers to keep decent levels of utilization of costly manufacturing capacity) and the profitability of the fabless business. It was very poorly received by the market that seemed to have accepted that designing chips and manufacturing them are two different businesses, and unless there is a clear strategic rationale to be vertically integrated, all that truly matters is to either be credible in launching products in the "sexiest" segments (1 and 2 nanometers) or sustaining capacity utilization of manufacturing facilities over 97 percent. It remains unclear at this point whether Intel Corporation will be able to pull out its strategy of vertical integration in a value additive way for its shareholders. New CEO Lip-Bu Tan is now focused on cutting costs at Intel but has not yet defined how it will make the company more

competitive and whether it will address the core question around what parts of its value chain does it want to lean-in for this next chapter. To make things even more complex, this difficult business decision is influenced by other geopolitical forces pushing the United States to have chip manufacturing capabilities in their territory. The business model decisions around key capabilities, key activities and partnerships, and the resulting cost structure are transcendental, especially in established business with greater legacy and lower flexibility to pivot. Let's use the Intel story and an invitation to challenge the rationale for your current value chain in light of the relevant future scenarios for your business.

AI as a Catalyst for Business Model Innovation

As you explore future scenarios, AI can be a trigger to re-code your business model. The opportunity is not just to identify a few promising AI use cases. AI can act like a catalyst in a chemical reaction: sometimes subtly accelerating what already works, sometimes enabling entirely new processes, and occasionally triggering a full structural transformation. The key is being clear about which type of change you are aiming for, and being attentive to the conditions needed to enable it.

One clear shift in this era is the speed and cost of experimentation. Prototyping and testing business model elements—whether related to value creation, delivery, or monetization—has become more accessible than ever. This makes it possible to treat experimentation not as a side activity, but as a structured learning process: one that combines deliberate intent with openness to unexpected outcomes.

Most organizations are still in the early phases. According to MIT's Project Nanda Report (2025) many AI pilots have yet to generate clear returns—not because the tools are ineffective, but because process and workflow adaptation takes time. The familiar J-curve of technology adoption still applies! Productivity often dips before gains materialize.

This underlines the importance of linking experimentation to learning. Some parts of the business may benefit from careful piloting; others—once workflows and tools are aligned—may be ready to move to a straight redesign. In several cases, operational and back-office applications have shown the most immediate impact, even though many initial pilots focused on revenue growth.

As momentum builds, conditions may emerge for more radical redesigns. Feedback from customers, suppliers, and peers can clarify which parts of the business model are ready for change—and which are not. Timing matters. Being

alert to inflection points and external signals can help determine when to move from incremental adjustment to deeper transformation.

Rather than following a fixed sequence from experimentation to redesign to transformation, it may be more useful to think in terms of optionality. Building the organizational agility and the learning mindset can be a pre-condition to fully extracting the value of your business model options. See an illustrative example in Figure 17 on how to place the different AI bets across the business model.

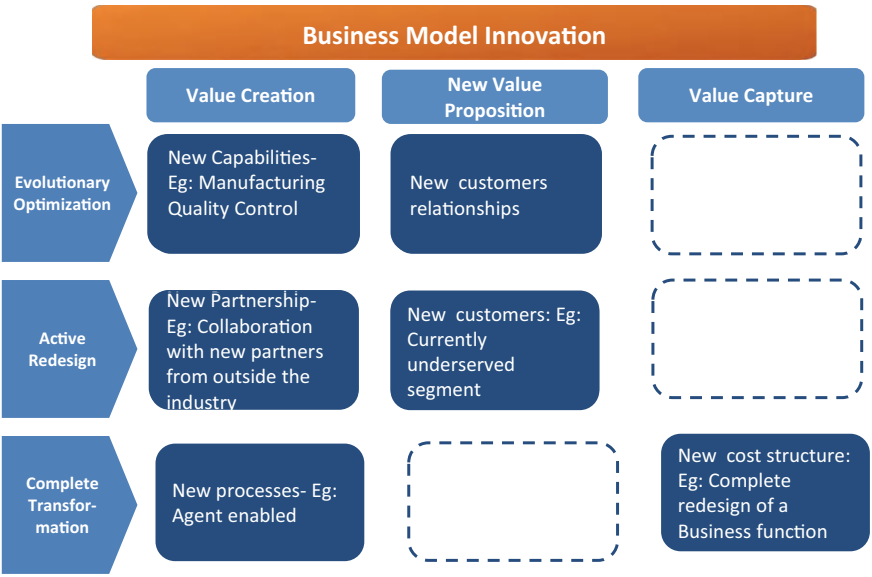


Figure 17: Business model innovation: Mapping your prototypes (illustrative example).

Can AI agents turn your services into more of a subscription that keeps customers locked-in like in the Tech-mosaic Scenario with Aswath? Can you create or magnify network effects in your business like in the healthcare example where Karine could leverage an orchestrator platform of healthcare providers harmonized by AI agents? Can AI give you more visibility around final outcomes and thus allow you to price for real value generated, such as in the Augmented Enterprise Scenario with Razan at the desalinization plant? Can new partnerships open new businesses that give you greater degrees of freedom if you need to operate in the Digital Feudalism Scenario? Creating the relevant suite of options is likely to require a combination of bottom-up suggestions, often incremental, as well as some top-down initiatives inspired by your company's ability to anticipate future scenarios.

Chapter 6
Re-Design for Optionality, Not Perfection

Optionality in Light of Growing Uncertainty

Paradoxically, in an era of rapid change, certain capabilities take time to build, and certain decisions should be made thoughtfully and slowly. Investing time and resources to build optionality is one of them. This is particularly relevant as you experiment with different variables of your business model and as you push its boundaries. Optionality is not only good in terms of managing your portfolio of products and services, but also generally a good approach to hedge against other sources of uncertainty beyond technology. Several indicators lead us to believe that regardless of the evolution of AI, global economic volatility is increasing. According to the Global Economic Policy Uncertainty Index, a widely used index for measuring uncertainty based on citations about economic policy, January 2025 has reached the highest level of uncertainty since 1997. When looking beyond economic policy, other more general indexes like the World Uncertainty Index (WUI) point to the same trend. Notwithstanding the current economic and geopolitical uncertainty, we have observed a step change in volatility since mid-2014 (see Figure 18).

In this context, technological developments in AI further exacerbate uncertainty. With greater uncertainty, options provide strategic flexibility. Deliberately creating options based on potential polarities can be very valuable for businesses. The concept of real options provides a framework for managing uncertainty, as it generates tangible opportunities in addition to offering risk-management tools.

Indeed, the idea of building a portfolio of options is not new in business; as a matter of fact it is core to most innovation projects and to the active management of a portfolio of products and services. Some of the projects are doomed to fail but the cost of trying is embedded in the overall approach. Somehow, when success happens, economic rewards can not only offset the cost of all the failed projects but can also generate unique insights into unmet needs of target customers, or the strengh of our own capabilities.

When confronted with a list of initiatives that open potential opportunities for the future of your business, how do you select and prioritize amongst them? In this context, real options give you the power to test promising directions without betting the farm. These aren't vague experiments; they're small, structured initiatives with defined potential outcomes and costs. The idea is simple: look for low capped downside and for high unlimited upside. Let's review a simple illustrative example. Imagine your company finds itself in the AI-run Ecosystem Scenario—the healthcare example—or in the domain-specific Augmented Enterprise

https://doi.org/10.1515/9783111583495-007



Figure 18: Monthly Global Economic Policy Uncertainty Index (September 2025).
Source*: Economic Policy Uncertainty.*

Scenario—the desalination plant example. Both scenarios will pose a real challenge to your business. Your team decides that a pilot with a small, specialized AI company integrating it as part of your offering is worth investing in. Illustratively, let's imagine the pilot requires $1M of one-time integration investment, and an incremental operating expense of $150,000/year in the next three years if you choose to extend the pilot that long. In three years, you expect to generate:

– $3M of incremental operating income with a 50 percent probability
– $1.5M with 30 percent probability and
– $5M with a 20 percent probability

You can break the commitment in the initial three years with no penalties; after that, if you choose to continue you can reinvest an additional $1M to deliver a better array of services. You are in fact buying a call option with a cost of $1M and a strike price of $150K every year and a pay-out of $3.3M. This is likely to be a positive investment under a wide range of discount rates.[58] Ideally you want to go into it with maximum certainty about the maximum downside (in this

case $1M + $150,000 for three years). The maximum loss is known. The fundamental consideration is whether the value of the learning process outweighs the $1,450,000 investment.

Your team also decides that a plausible future also includes one of your global suppliers developing a highly specialized suite of functions for your sector. If they do so and become the market standard your business may be hostage to their overall terms and pricing. The team feels it is also worth putting energy behind another pilot with them because there is a chance you can influence the functionality of the new product. The team proposes another pilot with the large, specialized tech player. This is more of a defensive move, and you are looking to better understand how Big Tech thinks and what their real strengths are. In this case you have fewer out-of-pocket expenses, but, in the end, you must dedicate approximately $1,450,000 of your team's time and resources to conducting the pilot that is also set to last three years. If things go well, you estimate you can save approximately $5,000,000 in costs by truly augmenting your team.

- You give this an 80 percent probability because of the longstanding relationship with this supplier and the fact that they have always looked to your company as a test bed for their innovation.
- You give a 20 percent probability that the pilot is a total waste of time and resources.

In this case the maximum loss is also $1,450,000; but it comes in the form of opportunity costs—no direct outlays. The pay-out is $4,000,000.

You would most certainly do the second pilot, and depending on bandwidth of teams, cash-flow considerations, and approach to strategic innovation you may want to also conduct the first pilot with the AI specialized firm. You have two real options and three years to figure out if your options are in the money, which is when exercising it will result in a positive contribution. We tend to be less demanding with the second type of options that require the same value invested but they are not out-of-pocket expenses.

To be clear, optionality is not an excuse to turn a bad investment with negative Net Present Value (NPV) into a good one. It does require real discipline to validate on a regular basis if both costs and value expected are delivered. In this example we have only used a decision tree type approach for simplicity; your company may want to leverage more targeted approaches like Monte Carlo simulations or dynamic programming to calculate the value of the real options. Let me note however, that the biggest obstacle to adopting real options in businesses is the complexity of the valuation models.

As a company preparing for the Era of AI, you need to have several initiatives—often under the category of innovation projects that give your company optionality.

Options Can Come in Different Forms

I like to think about options as shaping the future in two basic forms: as a way to generate new business, and as a way to hedge business risks. When looking for new businesses, the standard approach is to build a portfolio of innovation projects that can go after new profit pools. In this case, these innovation projects tend to have clearer ROI metrics to go after, and they can be evaluated with a portfolio approach as opposed to on a stand-alone basis. That is, the sum of the projects needs to generate positive NPV. It is like giving management teams in your company a sandbox for innovation: the investment and resources are finite; the results can come from one of the multiple projects within the sandbox. This portfolio approach is typical of antifragile companies. Nassim Taleb, author of *Incerto*, the five book series on uncertainty and randomness of which the *Black Swam* is one of the better-known books, described antifragile as:

> Some things benefit from shocks; they thrive and grow when exposed to volatility, randomness, disorder, and stressors and love adventure, risk, and uncertainty. Yet, in spite of the ubiquity of the phenomenon, there is no word for the exact opposite of fragile. Let us call it antifragile.[59]

Antifragility is when something gains from disorder. It is not only about being resilient and robust to avoid breaking; antifragile businesses become stronger with stressors, with uncertain volatile markets. Antifragile companies sometimes operate with multiple projects that are doomed to fail or barely succeed yet they look to have at least one project that delivers 10x. They have an investment strategy akin to venture capital funds. The AI Era may open a window of opportunity to create an antifragile business: as you extract opportunities for your business from your future scenarios, be on the look-out for those opportunities that benefit from randomness and volatility. In the same way, disorder on the Internet benefited Google's search business with its overarching goal of organizing information, maybe the disorder created by a multitude of black-box unexplainable models may create an opportunity for companies that are able to explain and create a layer of trust in the AI applications ecosystem.

Options can also come in the form of risk hedges: the second example above (the pilot with the specialized Big Tech) is mostly justified as a risk hedge to learn more about the relevant Big Tech that can hold the future of your business hostage. In the period 2016 to 2020 in the United States, it was evident that the trade tensions between China and the United States were raising and that may have an impact on companies doing business in, or with China. Apple had built its core manufacturing base in China relying heavily on Foxconn to supply its iPhones. As the tensions arose, Apple expanded capacity into India—supposedly to serve the

Indian market demand for iPhones and to Vietnam to manufacture some of the accessories. These shifts in the supply chain take time because they require important capital investment but, as importantly, they require vast amounts of workforce training to be able to deliver the right quality standards. Under the tariff war started in early 2025 by the United States, Apple decided to move part of the iPhone manufacturing capacity to India to also serve the U.S. market. In later months, Apple made an announcement that it will serve the totality of U.S. market demand for iPhones from India. This is a simple yet powerful example of building optionality ahead of the need. As a leader in your business, you do not want to be caught with no options when certain important business risks arise, especially if they are "written on the wall." Arguably, the Indian operation may have taken a while to be as efficient as the Chinese operation—assuming it ever reached the same level—but this sacrifice in efficiency seemed worth it when an increase in tariffs by 145 percent for goods coming from China into the United States was announced in April 2025. Even if Apple still has numerous components manufactured in China, the investments in India and Vietnam's manufacturing capability now seem quite a relief for Apple.

Managing the Tension of Investing in the Core Versus Creating Options Beyond the Core

As we have reviewed, optionality is not only giving your business strategic flexibility and the opportunity to learn first and then decide; it is a way to protect your firm from unforeseen situations. No different from getting insurance that covers the firm from bad or catastrophic situations. It has a cost, but if the situation arises, the organization can go back to serving its clients with minimal disruption.

Yet while the concept is simple, investing in creating optionality at an institutional level encounters some important challenges. Most resource-conscious companies struggle with the actual execution of this idea. Undoubtedly, options have a cost. At an institutional level, creating optionality can be challenging because it can be perceived as a distraction of resources and of management attention; there is a constant reminder of the opportunity cost of not investing in the core where outcomes are predictable and the return on investment more certain. In addition, when looking to implement strategic optionality the other big challenge is that it is hard to measure and monitor.

Classic strategy recommends a strong focus on the core business first. It feels like in the new AI Era developing a true competitive advantage in your core product, core service, and core set of clients continues to be the best option for suc-

cess. A strong core with robust strategic and economic fundamentals is a never-ending effort that requires constant focus and continuous investments.

Once you have a strong, well-optimized core business, classic strategy recommends growing into adjacent businesses. This move is known among economists as related diversification. The new adjacent business leverages existing capabilities or assets like similar customer segments, or channels, or suppliers. The ideal new adjacent business reinforces the success of the core business. The more related to the core business the higher the odds of success and the higher the odds of generating attractive and sustainable value. Typical adjacency moves include selling complements: growing from women apparel to accessories. It can also include tapping into new customer segments with the same winning formula: like Nike going from running shoes and apparel to basketball shoes and apparel, and successively covering different sports categories while leveraging existing suppliers, channel relationships, and related capabilities from brand building and marketing to retailing. In this case, related diversification delivers greater economies of scale and builds a deeper competitive moat. This is an example of ideal adjacent business. Adjacent businesses may also reside in the existing value chain. For example, using existing capacity in your value chain to third parties may not only lower your own costs but also be a profit generator. Amazon built its most successful business to date—Amazon Web Services—out of this very idea. Back in the 2000s with a growing online retail business, Amazon found its IT teams building the same infrastructure repeatedly (servers, storage, databases). It moved to build standardized, reusable, cost-effective infrastructure. They realized that at the time of explosion of online business, others may have similar needs, and they started building it for others. Amazon also realized the opportunity for business model innovation and instead of selling a one-time solution, they moved to subscription services—Infrastructure as a Service (IaaS)—allowing them to shift the paradigm of infrastructure usage. This move not only served Amazon by setting it on the path to lead the global cloud market but also triggered a fundamentally new operating model. Amazon Web Services has generated over an extended period of time the majority of Amazon's operating income. Incidentally, the cloud business model was key to the booming Digital Era. It gave access to cloud services on a variable cost basis, thus lowering the barriers to access compute power and storage to multiple start-ups and small and medium business. This phenomenon turned an unfathomable high fixed cost investment into purely variable costs.

Amazon had a strong set of capabilities and assets to build this adjacent business, from data center expertise to software skills, to maniacal customer orientation. This was a solid foundation for the new adjacent business. I suppose, at that time, Amazon may have envisioned multiple futures: one where they lead the

new cloud market with potential to grow exponentially; another future where cloud grows slowly and Amazon's profits from this business to cover their internal infrastructure; and a final scenario where cloud does not take off and all benefits are exclusively internal; their online retail business becomes more efficient and scalable. It is a classic approach of adjacency expansion with good odds of success and high option value. With 20–20 hindsight it is easy to see that AWS was a good decision, yet, back in 2006, it may not have been that obvious for a B2C online retail business to move into a B2B cloud infrastructure subscription model.

From a strategic perspective, building optionality through related diversification seems an attractive approach for the Era of AI. We see most Big Techs following a similar approach. Let's investigate Meta. When thinking about AI priorities, they start by improving their core business: improving advertising effectiveness. Meta aims to be a one-stop shop for businesses where Meta can understand well the desired outcomes and they will create the most effective combination of content generation, the targeting, the channels, and lastly, the monitoring and measurement of the process so that ultimately desired outcomes are delivered—the creation of a super advertising agent for businesses. The second priority is also to reinforce its core business and use AI to drive up engagement by becoming much better at generating and recommending content.

Even if the concept of building a portfolio of options is accepted, and the tension between investing in the core with more certain outcomes versus investing in optionality is managed appropriately in your business, there is always a difficult balance between the two. Bain & Company, a consulting firm, goes as far as saying that given the increasingly turbulent times, creating a whole second engine of growth separate from your core is highly recommended. According to their research based on an analysis of two decades of sustainable growth companies,[60] 60 percent of them are deriving important valuation benefits from what they label as Engine 2. This second business often comes from an adjacency: about one-third of the time, Engine 2 is the next generation of business model, often competing with the Engine 1 model, and two thirds of the time, this Engine 2 is actually a distant adjacency made available because of changes in technology. Ecolab, a company in industrial cleaning based in Minnesota (United States), realized that its clients were increasingly concerned about clean water. Today, they have a water treatment and purification business started with acquisitions—a common way to build Engine 2 growth—that in 2025 represented more than 70 percent of total revenues and turning into its core business today. The company has leveraged their expertise in clean and disinfected areas into sophisticated water treatment solutions for a variety of industries: from food and beverage to manufacturing micro-electronic products. They may well be on their way to

launch a new engine, maybe Engine 3 of this coming decade, with their digital services around water management systems for data centers. Because water cooling systems are 20 times more effective than air cooling systems, and cooling is critical to lower energy consumption, their water management know-how is uncovering a new $5B opportunity in serving the AI infrastructure market.

It is noteworthy to emphasize that none of the Engine 2 business out of the Bain & Company study emerged from corporate venturing. They estimated that the frequency with which a moderately sized Engine 2 emerged from corporate incubators was slim to none.

Much Beyond the Core

Is unrelated diversification not relevant or interesting in the Era of AI? Companies with unrelated businesses have traditionally traded at a discount. There is extensive research showing that in a conglomerate with unrelated businesses, the parts of the business, when valued individually, are often worth more than the whole company. Economists have used different arguments to justify the discount of diversified firms—from inefficient capital allocation where a conglomerate should be deploying capital into businesses with better returns, instead of spreading a scarce resource across many different businesses with suboptimal resources. Think of a farmer with many fields; distributing scarce water resources equally across all fields may yield suboptimal crops. Along similar lines, other economists would argue that public markets do a better job at diversifying capital that businesses do, so the excess cash flow should be returned to shareholders to freely diversify based on their particular risk/return profiles. Another reason for discounting unrelated diversification is the loss of strategic focus by spreading organizations too thinly across different businesses which may lead to losing deep domain expertise and its competitive edge.

And yet, there have been visible exceptions to the idea of conglomerates being less valuable than the sum of their parts: one is General Electric during the leadership of Jack Welch (1981–2001), and the other is Berkshire Hathaway (BH) under the leadership of Warren Buffet. They have some common key success factors that may be relevant for unrelated diversification in the AI Era. General Electric maintained a diverse portfolio that included industrial operations—ranging from aircraft engine manufacturing to locomotive and healthcare equipment production—as well as financial services spanning consumer and real estate finance. GE managed the portfolio with ruthless capital allocation discipline; it only owned and deployed capital in companies that had the number one or number two position in their markets. If businesses were not able to extract the full value

of leadership economics, they would be sold. In addition to this ruthless capital discipline the diversified portfolio shared some critical transversal capabilities: lean manufacturing capabilities paired with a unique leadership grooming that was shared across the portfolio. The businesses were closely managed and monitored from the center. During the Jack Welch tenure GE managed to avoid the pitfalls of conglomerates, but it subsequently declined, in part because of the increasing complexity and the relaxation of the key principles that made it so successful. In the case of Berkshire Hathaway, its portfolio includes wholly owned companies from insurance to retail to transportation. Berkshire Hathaway is also known for the investments in companies like Apple or Coke where they own minority stakes. BH shares the same capital allocation discipline and thrust for high performance as GE had, however the units operate with total autonomy.

For unrelated diversification to be relevant in the AI Era, we would have to believe that in addition to the capital allocation discipline, diversity of business can generate unique value via the use and application of AI models. For example, we would have to believe that by accessing unique proprietary data from unrelated business—say logistics and healthcare management—the corporate center can develop better demand predictive models, better staffing models, and so forth. Or that sheer scale can give differential terms when accessing critical AI infrastructure or type of data.

Managing the Portfolio of Options

As you shape your innovation portfolio for the AI Era, it's essential to strike a deliberate balance across the Three Horizons framework outlined in Chapter 4. Horizon one should secure steady growth from your core business, while horizon two projects pave the way for near-term expansions, and horizon three bets explore bold, future-defining opportunities. AI will often show up subtly—perhaps as incremental gains in client solutions or operational efficiencies—yet its strategic impact can be profound. Not every initiative needs to scream "AI," but every option must reflect thoughtful innovation. The key question is: Are you placing enough smart bets? Have new products, services, or markets contributed meaningfully to growth in the past five years? Or are most of your innovative moves simply reactions to external pressures rather than proactive, future-shaping plays? The entrepreneurs in your business, more focused on horizon two, should be driving a portfolio of innovative projects that pursue deeper innovation and chase adjacent businesses or potentially the next Engine 2 for the AI Era. The visionaries in your organization ought to be looking into the future scenarios and into existing innovative projects to dare re-coding your business model. Sometimes it may

help to just draw a simple map like the one in Figure 19 to check if your portfolio of options is exceedingly simple or too complex. Bear in mind that in this illustrative example, we are looking at strategic projects; however, with a digital operating model, companies have the opportunity to run thousands of experiments and accelerate learning around key business hypothesis. At this stage we are looking to the portfolio of strategic options. Are there any blatant gaps? Are the metrics of success clear and are there the right processes to control and monitor proper execution?

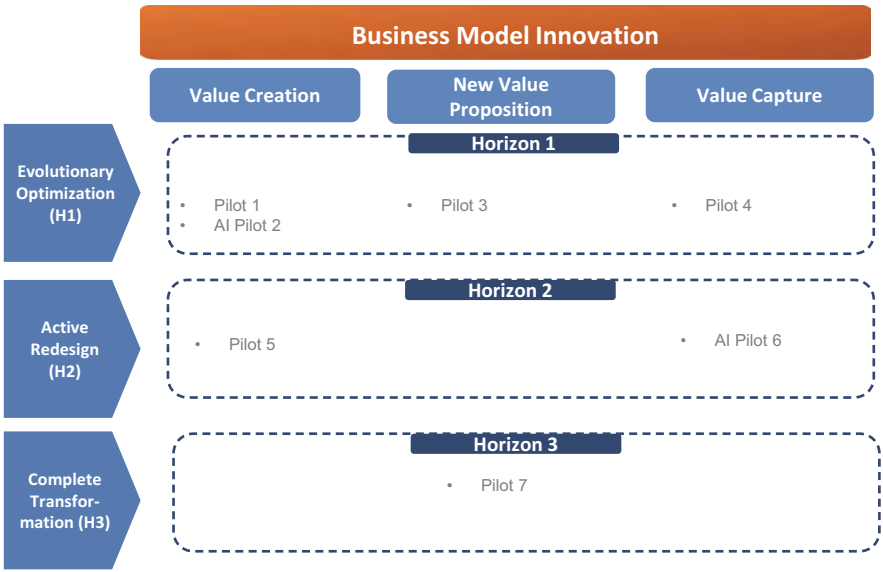


Figure 19: Managing your portfolio of options.

Where your company sits on the spectrum of bold innovator versus cautious operator will shape how you approach AI. On one end, the risk is *FOMO*-fueled chaos—chasing every shiny pilot without discipline. Guardrails like strategic alignment, links to plausible future scenarios, clear execution milestones, and solid business cases are essential to tame this urge. At the other end, if your organization tends to shy away from innovation, AI can serve as a powerful spark to unlock creativity and experimentation. That may mean lowering barriers: easing rigid business-case requirements, carving out resource pools, or giving employees access to AI tools—even outside traditional innovation or product teams. In both cases, the challenge is the same: channel the energy of AI into purposeful, value-creating optionality for the future of your business rather than scattered activity.

Chapter 7
Re-Wire your Operating Model for the AI Factory

The Burning Platform

As a leader driving a business, I am sure you regularly revisit your structure and your people; these two elements of the operating model may even be the subject of discussions at the board of directors. You are likely to spend far less time re-examining the other two core elements of your operating model: the core processes and infrastructure behind it all. The board of directors would be lucky if you bring these topics to their attention once every two years.[61] Yet, there is no future scenario in the Era of AI where robust data and technology architecture are not a pre-requisite to being successful in this era. There is still an unusual number of firms—big and small—in any geography—including the United States and Europe—that are dragging their feet in investing behind data and technology. The challenges in putting these topics on the executive agenda range from lack of financial resources to lack of expertise, to sheer paralysis generated by the speed of technological change. Make no mistake, without a truly digital operating model you will not be able to take full advantage of the unique opportunities AI is unlocking for your current and future competitors.

Let's just take a bit of perspective on the urgent point of having appropriate data architecture. Data creation according to the World Economic Forum went from 2 billion gigabytes of data in 2,000 to 600X expansion, reaching 1,200 billion gigabytes in 2010.[62] Several driving forces led to this exponential creation of data —from greater user-generated content on the Internet (blogs, new aps), social media, and e-commerce, to the explosion of mobile devices—smartphone taking the number of subscribers from 740 million worldwide at the start of the 2000s to more than 8 billion in 2020. There were more cellphones than people. The data creation engines have evolved, and AI has compounded this trend. In 2025, the expectation is to reach 181,000 billion gigabytes of data (or 181 zettabytes); that is more than 20,000X what we had in 2010. Over the past years, most forecasts about data creation have fallen short. This one may too.

There is no doubt your business will need to rely on the right data for your products and services. There are several helpful frameworks around data maturity of organizations that can help frame this internal reflection. If most of your data uses are for descriptive or diagnostic uses, your organization is likely to be at basic levels of maturity; if, on the contrary, most of your uses are predictive or prescriptive (as in with AI agents) your organization will be better prepared for

https://doi.org/10.1515/9783111583495-008

the changes that are coming. Let me try to provide you with an intuitive description, for non-tech leaders, of the key building blocks of your data architecture.

Your company will use two sets of data: structured data and unstructured data. The first category includes everything about your products, your employees, your vendors, and your clients. You need to make sure it is clean, updated, well labeled, and properly integrated with your business processes. Unstructured data is an untapped resource; it can include anything from conversations with clients at the call center, to your competitors' postings on social media. Industry estimates that more than 80 percent of data used in businesses is unstructured data. On the architecture front: you need robust data ingestion and storage mechanisms in a lake house, that is, efficient storage with a structure adapted to your specific business processes. You may need to deal with the combination of siloed data storage that is not interoperable with more up-to-date data systems. Interoperability requires common definitions and clear protocols. Easier said than done! You need an effective transformation layer that turns the collected data—structured and unstructured—into relevant insights or information for decision making. This layer can comprise analytics, machine learning, AI agents, and other tools that enable users to access relevant data through queries and visualizations. Ideally, you would like to move toward a scenario where end users can independently find the relevant data and models. This ideal state is not conceivable without a proper governance system that not only guarantees the quality of the data but as importantly the permissions and access controls. The governance model needs to provide full trust and explainability. See a visualization of a simplified data architecture diagram in Figure 20.

Invisible Yet Invaluable Asset

No matter where you are on the maturity curve, building the right data infrastructure can feel overwhelming. But don't think of it only as a long march toward the AI Era. The payoff starts much sooner: with the right data foundations, you unlock important value long before the AI future fully arrives. Data infrastructure is one of these assets that is mostly invisible yet invaluable. This is an intangible asset like others in your organization, the effectiveness of R&D investments, the quality of your brand, or the quality of your human capital. Accountants do not do a good job of reflecting these unique assets in the financial statements. They are often treated as operating expenditure when in fact these types of capabilities are building the long-term future of the company and should be capitalized and reflected on the balance sheet as part of the intangible assets.

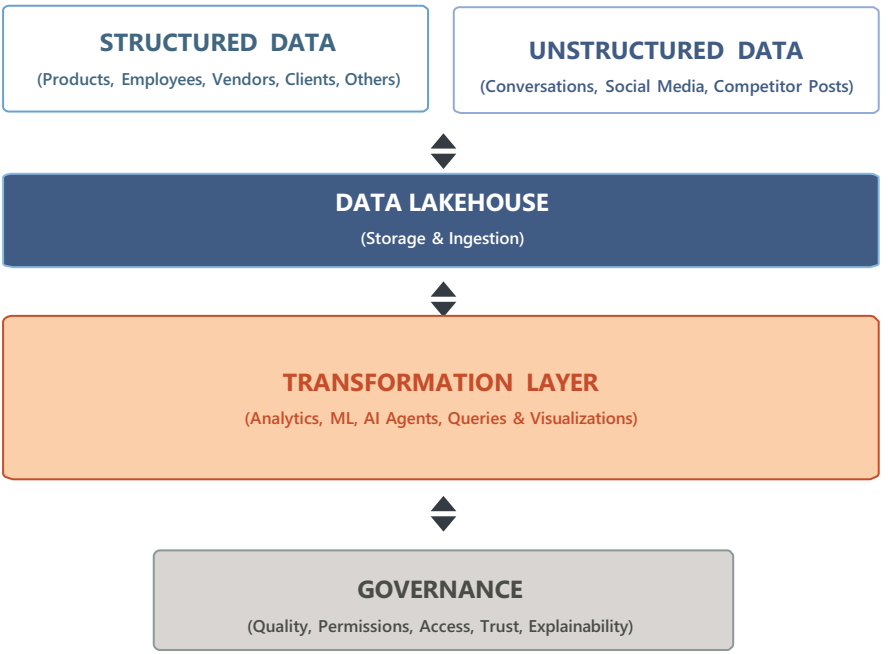


Figure 20: Simplified data architecture.

These strategic capabilities do show in an intrinsic valuation type of exercise. Let's review briefly how your company may see the value of investing in AI-ready data infrastructure. If we look at the intrinsic enterprise value of your business, we will discount expected cashflows for the years to come.

$$Value\ of\ Business = \frac{E(Cash\ Flow_1)}{(1+r)^1} + \frac{E(Cash\ Flow_2)}{(1+r)^2} + \ldots + \frac{E(Cash\ Flow_{n+1})}{(r-g_n)(1+r)^n}$$

The effect of strong intangible assets must come from either better expected cashflows or lower discount rate. So, on the free cash flows one would have to believe that good data architecture allows for faster possibility to go after new markets (greater volume) in a differentiated way (greater pricing power), driving greater revenue growth and/or greater ability to scale efficiently and thus obtain greater operating margins. Alternatively, these intangible assets can show up in the form of a lower discount rate because your business can easily access equity capital and thus lower the cost of equity; or that (given a level of risk) your company can borrow at below market rates.

If you are not convinced of the value you can create in investing behind a robust, integrated, and flexible data architecture we can go back in history and to 2002 when Jeff Bezos famously addressed its team with a controversial memo as the company was drowning in complexity and costs and its developers were slowed down by data architecture bottle necks and silos. The memo included the following points:

- All teams must expose their data and functionality through service interfaces (Application Programming Interfaces).
- Teams must communicate with each other exclusively through these interfaces.
- No other form of inter-process communication was allowed—no direct linking, no direct reads of another team's data store.
- All service interfaces had to be designed to be externally exposable from day one.
- Anyone who didn't follow this approach would be fired.[63]

According to some internal employees this was the memo that turned Amazon from a retailer to a software company, fundamentally changing its operating model. This happened in 2002. In 2025, no matter whether you are in retail or in software, as a leader of your organization you need to have put some serious thought into your data and your technology stack. As a litmus test, ask yourself about how much time your executive board spends on these issues? How well do they understand the point of arrival and the necessary steps to get there? What is the role of the IT and data team in your organization in shaping the direction of your business?

Building your AI Factory (Not to be Confused with Data Centers)

Digital and non-digital business have been affected in their core processes, people, and organizational structures by the data revolution. The transformation is not easy. Like we saw earlier in this section, even Amazon found itself replicating IT and data infrastructure for each business unit. IT tends to be at the back of the business and mostly focused on handling thousands of requests catering to users operating in siloed units.

Yet in the Era of AI, not having properly built the AI factory with the flexibility for different users to access data from a common repository and giving a more horizontal perspective to every business unit can be a major inhibitor of AI adoption. Let us think about Netflix, for example; they have an AI-centric opera-

ting model where they collect ample data, train their models, and make decisions based on the output of those models across every part of the business. They use data to recommend movies, but they also use reviews received from customers with movie studios to better negotiate financial terms. Netflix also uses data and AI to decide what content they should produce directly. Their data is sourced and organized for transversal use across the company.

According to Iansiti and Lakhani,[64] there are four components of the AI factory essential for the new era.

- A *data pipeline* that collects, cleans, and organizes data in a systematic and scalable way and with the appropriate governance. Companies need to put effort into using high-quality data for their models and applications to avoid the risk of the popular saying "garbage in, garbage out"; keeping the data pipeline updated and precise on a timely basis is increasingly challenging given the size if the data, the complexity and different formats it comes in, and the speed at which it needs to be updated.
- *Algorithm and model development* that help predict and act on the next step in a business process. Data needs to be turned into insights and needs to support business decisions. This transformation layer needs to be closely aligned with the competitive differentiation that the company pursues. Not all algorithms are created equal. Going back to our future scenarios, the AI-run-company HealthCo will look to create unique algorithms to improve diagnostics or the reliability of different providers; in the case of the Digital Feudalism, banks will look to focus on personalization algorithms and recommendation engines.
- *Infrastructure* that enables workflows, computing, and analytics. Applications and algorithms need to run with appropriate response time and the appropriate security controls.
- A *connected experimentation platform* where developments are tested prior to moving into full production. With the right AI factory, you can run extensive randomized control trials (RCTs) and pressure test changes in your value proposition in your capabilities before a full roll-out. This experimentation platform can help test the benefits of product changes, of marketing campaigns, of price changes, and so forth.

This AI factory not only allows for more integrated business processes across organizational units with more connected decisions; it also allows for greater agility and relevance to respond to market needs. You want to aim for a centralized and consistent enterprise architecture with maximum flexibility, interoperability, and autonomy for end-users in the organization. The flexibility objective will drive you to create an orchestration layer—as described in the AI infrastructure stack

by Andrew Ng in Section One of the book—that gives you the possibility to incorporate as many AI domains and function-specific applications and switch to the most appropriate foundation model.

Moderna, the biotech company based in Cambridge, Massachusetts, USA, was able to go from sequencing the COVID-19 DNA to producing the vaccine in one month. This is because their operating model focuses on speed to market. It is the poster child of an AI-centric operating model. Moderna, initially ModeRNA, started off in 2010 with the idea to build an *mRNA technology platform*. The innovative idea: messenger RNA could be engineered to instruct a patient's own cells to produce proteins to prevent, treat, or cure disease. This platform approach allows for many potential applications (vaccines, therapeutics) rather than a single product. This single unified data platform—from research, clinical development, manufacturing to commercialization gives them unprecedented time to market in addition to other R&D benefits across products. Moderna is developing different therapeutic categories like oncology or rare diseases that can replicate the agility demonstrated during the COVID-19 pandemic in other treatments. Moderna has built an AI-centric operating model at the same time as it invested heavily in manufacturing facilities. Unlike many biotech companies that tend to rely on third-party manufacturing, Moderna has created a competitive moat around speed to market, and both its digital and physical infrastructure is built to reinforce its moat.

Companies in other sectors are also developing AI architectures in what are essentially physical businesses. For example, Siemens in Germany has also developed a strong AI architecture integrating cloud services, data architecture, and Internet of Things type technologies. Their Digital Industries unit has developed additional capabilities and technologies to develop a new value proposition to their clients including digitalizing the entire value chain for manufacturing and process industries, enabling enhanced innovation and efficiency. The business has become an important source of growth within industrial business where it currently belongs and, as importantly, it has the potential to contribute to meaningful margin expansion for the group.

I would also like to elaborate on the concept of the experimentation platform. In the book from this series, *Artificial Intelligence: Shaping the future of innovation*, Alfred Essa[65] describes the benefits of experimentation, but more particularly the ability to build quick prototypes. With today's AI tools, prototypes can be accelerated to a fraction of the time and effort. He argues that a rough prototype yields better learnings than a polished one. The latter does not easily open up for improvements, while a simple demo invites for feedback and creativity. In my own direct experience, after a year of multiple teams experimenting with AI tools and agents, the average collaborator has been able to create prototypes that are

then shared and improved across teams. This experimentation mindset across the organization accelerated the identification of impactful use cases and its scalability. Interestingly, the role of IT and data teams also changes. They evolve from developers of technical solutions to scalers and orchestrators of user-driven, even user-developed prototypes. This low-risk experimentation and prototyping changes organizational dynamics from planning to playing. It inverts the burden of the proof: from top-down transformation to bottom-up driven transformation where there is no need to wait for the business cases, often hard to make, or for a pipeline of IT projects to clear. This approach of more open experimentation and prototyping can help firms evolve from AI-assisted to AI-integrated capabilities across key teams.

Break Silos, Focus on Decisions

In strategy, there is a generally held belief—tested by research—that high-quality fast decisions drive better returns and greater competitive advantage. The Era of AI tests the quality and speed of decisions because relevant information can be made available to decision makers much faster and, conversely, the windows of opportunity to make decisions become narrower. As you upgrade your operating model, be strategic about your decision-making processes.

Decision Quality and Speed

Studies by several consulting firms demonstrate that quality and speed of decisions are at the source of business outperformance.[66,67] The ability to fully execute these decisions is also critical. Academic research has corroborated that in highly dynamic environments speed of decision making is a true competitive advantage.[68] Yet there are some important nuances: a recent study by Shepherd et al.[69] highlights the importance of layering on this analysis, the availability of resources, the opportunities for growth in that market, and the competitive intensity. In an environment where leaders feel there is no way back to a given decision, the quality of decisions should override speed. You, as a leader, will know in which environment you operate in.

Types of Decisions

At any rate, the clarity of the decision types and who makes them based on what information seems to be a fundamental strategic principle to upgrade your operating model for the Era of AI.

When it comes to decision-making, you can lean on any of the frameworks from consulting firms—or take a page from Jeff Bezos' playbook with his distinction between *Type One* and *Type Two* decisions. *Type One* decisions are "one-way doors": once you step through, there's no easy way back. Think of reshaping your technology stack, entering a new geography, or making an acquisition. These require deliberation, robust testing of assumptions, and involvement from top leadership. *Type Two* decisions, by contrast, are "two-way doors": reversible, adaptable, and best handled quickly by small teams with sound judgment and proximity to the facts. In this case, missteps aren't failures—they're valuable lessons that generate sharper insights.

In the AI Era, however, there's a new dimension we will call *Type Three* decisions. These are the decisions we might delegate to intelligent AI agents. Here, the challenge shifts to designing systems that learn correctly and make reliable calls. We already see this in action—from dynamic pricing in airlines and ride-hailing, to credit card approvals, to ad placement algorithms, and resource allocation across data centers. The real leadership task is not whether to delegate, but how to ensure that AI-driven choices align with strategy, ethics, and long-term value creation.

Using the stated principles, you may want to consider starting with mission-critical business decisions and working backwards to make higher quality and faster decisions where it matters most.

In sum, upgrading your digital operating model will better position your enterprise to take full advantage of the AI futures most relevant to your organization.

Chapter 8
Learn, Learn, Re-Learn

Setting your GPS

Your organization needs to create its own GPS (Global Positioning System), its own mechanism to adapt to changing market conditions. A learning organization can be your best GPS in the Era of AI. As we have seen in the chapter on future scenarios, your competitors may become collaborators in the AI-run Ecosystem Scenario, your supplier may become your fiercest competitor in the Digital Feudalism Scenario, and your unique enterprise capabilities may now depend on a handful of augmented individuals, as described in the Augmented Enterprise Scenario. In a context where business definition is fluid, the ability to learn and adapt is a first principle of rethinking strategy in the Era of AI. I like how Satya Nadella, CEO of Microsoft, described the transformation he was going to lead when he took over from Steve Ballmer in 2014. At that time, Microsoft dominated the PC market, yet it had missed important market opportunities—namely mobile—and was about to leave the global cloud market to Amazon Web Services. As he puts it in his book *Hit Refresh*, a fundamental change was needed. He wanted people to move from a "know-it all mindset to a learn-it all mindset."[70]

For this cultural change, he leaned heavily on the insights from Carol Dweck from Stanford University about her work on learning mindsets. Based on her research of children in school environments she had identified two types of mindsets among students. The fixed mindset was typical of students who thought intelligence was static. They tended to avoid challenges, and they were frustrated when they failed. They would play by the rules and look to perform as best they could according to these rules. This first group would take no risks if it meant getting lower grades. Among this group, we would find the usual "I am no good at math, or no good at French." They genuinely believed this was a permanent situation. There was a second group of students who were naturally more curious and more optimistic about their capabilities to learn. They knew that if they put in the effort, they could get good grades. In this group we would find students who may say, "I am not *yet* good at Math, or not *yet* good at French." This second group has a growth mindset. In her findings,[71] Carol Dweck proved that the second group would do better at Math over the course of the two-year analysis. While the research was only partially confirmed in subsequent studies among other students, the idea of a growth mindset to change learning behaviors in corporate environments has been a good catalyst for change. Nadella was highly inspired by this approach and invested in embedding the "not yet" culture. Microsoft's stunning

https://doi.org/10.1515/9783111583495-009

rise—nearly a tenfold increase in market value over the past decade—can't be chalked up to culture alone. Bold strategic bets lit the way: doubling down on Azure, opening products to iOS and other ecosystems, acquiring LinkedIn, GitHub, and Activision, and backing OpenAI before anyone else paved the way to this extraordinary lift in market value. The lesson for the AI Era is clear: strategic decisions matter, but their impact multiplies when reinforced by the right organizational mindsets and behaviors. There are numerous approaches to how to get an organization to adapt to change: the digital transformation in the 2000s, the great financial crisis, and lately the COVID-19 pandemic have all provided no shortage of recipes. But I would argue that a crisis is not needed for an organization to make a case for change.

Wiring into the culture that the only thing that does not change is the need for change seems like a helpful pre-requisite for this era as well. Great performance had been achieved with systems designed for "continuous improvement": whether agile methodologies or kanban systems, the idea of bringing in the insight from the team to improve the next phase of production is of utmost importance in an environment where the AI arms race delivers relevant innovation every three to six months. So, designing feedback loops and being curious regarding input from your workforce, your clients, and your suppliers continues to be relevant for the AI Era.

The need to create learning organizations is not new. The different schools of thought combined a balance of individual learning with institutional learning. A learning organization can be your best GPS in the Era of AI.

Demanded Skills

When discussing navigating polarities in the age of AI we referred to the Augmentation–Automation Paradox, highlighting that they do not need to be mutually exclusive. Initial augmentation can lead to automation, and this can lead to further augmentation. You can probably see in your organization, teams whose roles will be more prone to substitution and others who can benefit from the Augmentation–Automation Paradox. So, what are the augmentation skills that will set your talent apart?

The general guidance given by the World Economic Forum (WEF) is a good backdrop of what the market is demanding. Even if it focuses on general skills, it is one of the few studies that combines qualitative insights from experts and employers with quantitative data from large market platforms like LinkedIn or Burning Glass Institute. There is a set of skills that continues to be on the list of the top ten most demanded skills for the last ten years. They include complex

problem solving, analytical and critical thinking, and creativity. Interestingly "active learning and learning strategies" are emerging skills in high demand. So far, the WEF study does not report specific AI–human collaboration skills, but, interestingly, research from Mäkelä and Stephany (2025)[72] highlights that AI increases the demand for human skills such as teamwork, resilience, and analytical thinking in AI roles. To be clear, AI roles are not necessarily in the technology sector. These are roles that demand some sort of AI-related skill; so they truly cut across all sectors and functions. It is worth noting there is also a growing demand for ethical reasoning.

When it comes to AI-specific skills, the landscape is still taking shape. There's no universally agreed language yet, and robust data is scarce. One of the clearest signals comes from Stanford's *Artificial Intelligence Index Report 2024*,[73] which shows that generative AI skills are mentioned in 60 percent of AI-related job postings in the United States—ahead of large language models, ChatGPT, and prompt engineering. But what exactly counts as "generative AI skills"? Are employers asking for deep expertise in building and deploying LLMs, or do they mostly mean the ability to prompt effectively? The lack of precision here mirrors the early days when "Microsoft Office" was first recognized as a formal skill category; today, "ChatGPT skills" seems to be reaching that same standard. However, in addition to the specific technical skills around AI like Machine Learning, Natural Language Processing, Model Deployment, Computer Vision, Data Engineering, and many others, there are a set of AI augmented skills that are becoming equally relevant.

To understand what those skills might encompass, it's worth looking at how people actually use ChatGPT. While it is early days, there are some intriguing pointers on how ChatGPT may be augmenting skills at work. A study[74] from OpenAI, Duke, and Harvard tracked more than 700 million weekly users sending more than 2.5 billion messages daily (as of July 2025). Surprisingly, about 70 percent of those queries had nothing to do with work. Among general use cases, practical guidance, writing, and information seeking account for nearly 80 percent of the uses. Among work use, however, ChatGPT's role is sharper: about 40 percent of use is writing-related, from refining emails and reports to translating or reformatting text—essentially, re-creating user-generated content. When looking at intent, workers lean on ChatGPT as an advisor and research assistant, especially among more educated users. The common thread? Across functions and industries, its biggest value at work lies in decision support—helping people make better, more informed choices mostly on routine tasks. No strong evidence yet from uses of ChatGPT as agents, for now only assistants.

AI-Augmented Skills

I would hypothesize there are five broad categories of AI-augmented skills that will be in high demand in the AI Era.

- *High-quality decision making* is a skill that is eminently available. As we saw in the future scenarios, AI can augment decision making from running a plant to improving medical diagnostics. The skill here is not "letting AI decide," but interpreting and challenging AI outputs. Workers must learn probabilistic thinking, bias detection, and ethical reasoning. This is already being done; we see Banco Bandesco in Brazil, training their employees in interpreting and validating AI credit scoring. Ping An, an insurance company in China, uses AI to perform medical diagnostic, but medical doctors always have the final word. For these types of skills, AI is like a co-pilot; in turbulent times, the pilot decides whether to trust the instruments or her instincts.
- *Self-learning*: Learning how to learn is becoming one of the most valuable skills. Active learning strategies are in high demand, and it's no coincidence that about 10 percent of ChatGPT's global use already happens in tutoring and teaching. That surge comes largely from schools and universities—but the real frontier is the workplace. With roughly 40 percent of the tasks in most jobs changing every three to five years, reskilling is no longer a side project; it's survival. That's where AI-assisted learning at work comes in. Think of it as building a personal coach. Personally, I've found it powerful to create my own AI tutor—even though I can't code—on topics I want to master. With any generative AI subscription, anyone can do this. I tell my tutor the depth I want (basic, intermediate, advanced, expert), and I ask it to use metaphors from areas I know well—so the new knowledge "clicks" faster, like fitting a new puzzle piece into a familiar picture. To check if I'm truly learning, I ask the tutor to test me with questions, following Bloom's taxonomy:[75] from recalling facts, to applying concepts, to creating something new. When I make mistakes, the AI doesn't just say "wrong"—it gives me detailed personalized feedback so I can practice and improve. In short, AI turns learning at work into a much more engaging and effective activity. Instead of being passengers on the road of reskilling, we now have the tools to sit in the driver's seat—with AI as our coach and sparring partner. Companies encourage self-learning skills in many ways; some leverage AI tools in their learning platforms like Infosys in India, which runs AI-assisted re-skilling at scale for 300,000 employees.
- *Effective communication*: It is no surprise that writing is the most popular work use of Chat GPT at this stage. There is high value in communicating effectively at work and the ability to strike the right tone, to critically express

facts and insights, and to do so with cross-cultural empathy. If we manage to keep the authenticity, AI can be a great communication assistant. Increasingly, it can do this in a multimodal way including voice and video. For example, Mercado Libre, the e-commerce company in Latam, uses AI chatbots for multilingual customer engagement but trains employees on "human-in-the-loop" corrections to keep the conversations as authentic as possible.

- *Creativity and innovation skills:* as we saw in the previous section, in the Era of AI there is value in fast experimentation, sometimes in the form of prototypes. AI-assisted co-creation can become a highly valuable skill akin to good old problem-solving skills but just with new tools at our disposal. A recent study at MIT,[76] with a randomized controlled trial with 444 professionals, found that ChatGPT "restructures tasks toward idea-generation and editing and away from rough-drafting." So human ideation together with an AI-trained co-pilot can help generate new products, new service features, or even the kind of business model innovation your business may need.
- *Project management and collaboration:* in an environment where jobs get redefined and where creating the AI factory demands transversal collaboration, AI-augmented skills in these two areas are paramount. There will be a need to revisit workflows; we will likely see the following steps: AI generates–human curates–AI refines–human delivers. So tasks and roles decomposition and redesign are likely to be critical skills beyond traditional re-engineering teams. We can easily imagine how professional services like consultants and lawyers will reshape their service processes along the lines described. The AI—human integration and orchestration is likely to become a highly sought after skill.

What are the skills that AI systems are unlikely to develop in the next decade? Yann LeCun, chief scientist officer at Meta, and Turing Award winner, argues that AI models cannot understand our physical environment; according to him a cat has a much better understanding of space and laws of gravity than any AI model developed so far. AI has made important progress in cognitive empathy; it can detect human feelings based on verbal and body language cues, and it can also deliver an expected response. AI may be able to do this much better than any of your colleagues. However, AI still struggles to understand complex social and emotional situations where judgement is required. Finally, human evaluation skills to assess the output of AI are likely to be needed specially for type one business decisions. Even with AI handling quality control and validation, irreversible mission-critical decisions will likely stay in human hands.

Hub Skills

The value of a skill is context specific; it is a function of the set of specific skills it complements. This is what the research by Stephany et al.[77] reveals. Programming in Python may be more valuable for a non-tech job like an admin role, than for a tech job. By "valuable" we mean commanding a higher wage premium relative to base wage for the role. The authors develop the interesting concept of hub skills, as the skills that are capable of unlocking value for many roles. Hub skills can be like API connectors; they can unlock access to more learning and faster adaptation. These hub skills are especially important in an era when execution is increasingly handed over to AI tools and humans need to do the higher order of thinking, planning, and setting direction.[78] These hub skills can also prepare the workforce for a wider range of roles and job opportunities. Currently, these hub skills include AI-specific skills like machine learning and natural language processing, but also less technical skills like prompting or project management.

Valuable skills do not necessarily mean resilience to automation. I would argue the more the skills are a connector across different functions the more protected they are from automation. For instance, within our executive education programs, we observe strong demand for advanced finance competencies that extend beyond auditing, financial reporting, and investment evaluation. When these skills are applied in sales or technology roles, they offer significant value and enduring relevance. In your organization, there may be a special type of hub skill that may be worth identifying.

Change in Job Skills DNA

It is important, however, to recognize that while the concept of augmentation through complementary skills is very attractive and less traumatic for many roles, it is key to acknowledge that certain roles will require total redefinition. For example, let's look at the role of actuaries. Actuaries do a lot of modeling; machine learning does a lot of Math and modeling, too. As of 2022, only 2 percent of actuaries are required to have machine learning skills. Yet, when they are required to have those skills, the job's entire skill DNA changes because the application of machine learning to the role shifts the focus of the work entirely.[79]

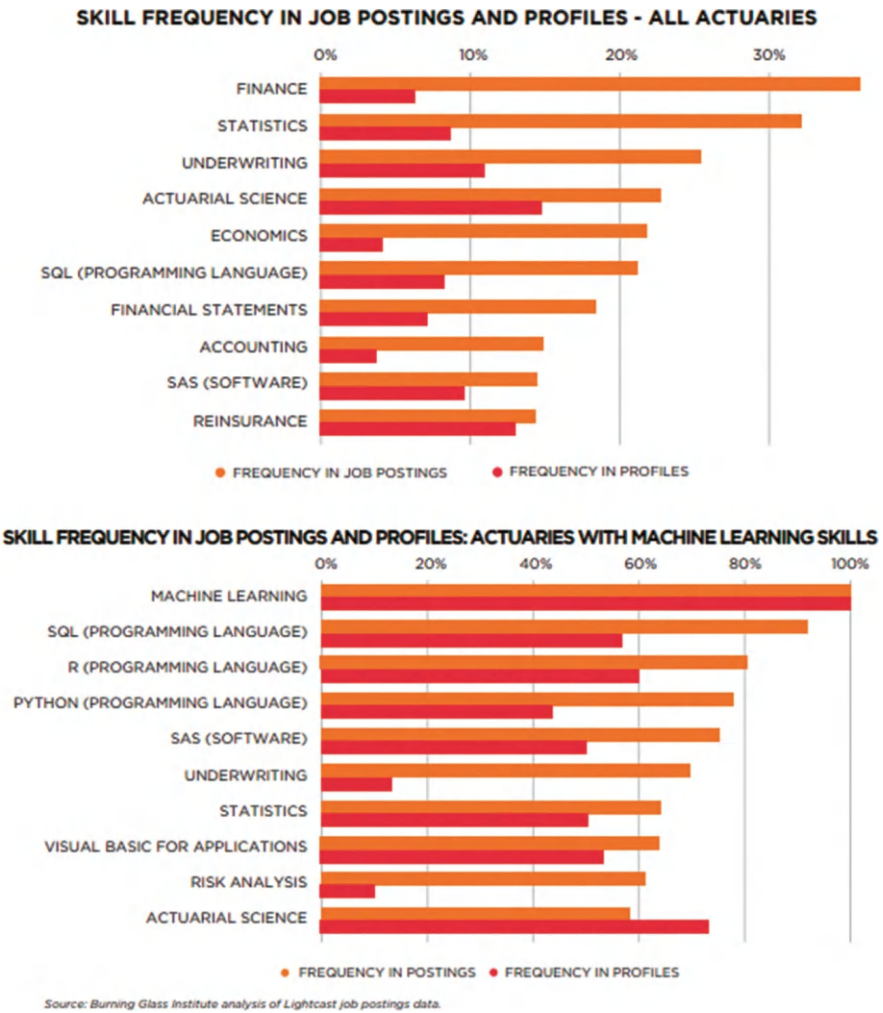


Figure 21: Skill frequency: Example of actuaries.

New Roles

Additionally, there will be new specific AI roles created. From pure transformational roles like embedding AI champions in smaller teams that foster learning and experimentation to more execution-oriented roles like agent creator, agent orchestrator, model evaluator, or AI researcher who designs the next hypothesis to be tested by your own models.

How to Learn: Developing and Training your Workforce

In addition to traditional upskilling and reskilling efforts, the Era of AI is likely to demand new skills. In the previous section we referred to hub skills—skills that unlock value across a variety of roles.

As in the previous eras, progress from technological advances depends partly on technical and data architectures but, as importantly, on people and processes. To avoid "seeing AI everywhere but in the productivity and value creation statistics,"[80] it is paramount to building learning and adaptation mechanisms. Be mindful of the J-curve: the productivity improvements do not start when you give your workforce AI tools. The J-curve highlights that there are productivity losses before the productivity gains are seen. Investing in people, in teaching them new tools, and new ways of working is a necessary first step, often leading to productivity losses. There are great conceptual frameworks out there on how to build learning organizations; many are failing to be adopted because of the challenges around implementation and disconnect from the short-term economic pressures of companies.

Connecting Learning to Value Creation

Institutional learning needs to always be connected to a driver of value creation in the business: whether it is better customer service, better product differentiation, greater efficiency, better decision making, or better risk management. I am surprised at how many corporate training courses bear little to no connection with adding value in the business.

The 70/20/10 rule for corporate learning will still be relevant in the Era of AI. It states that 70 percent of the learning happens in the flow of work, connected to the value drivers. There is another 20 percent of the learning that comes from social learning, from collaboration or coaching at work. Developing granular use cases shared with closely related colleagues is a great recipe for adaptation and ongoing improvement in preparation for deeper AI-driven changes. Encouraging your teams to use no-code AI tools to create prototypes that either make their jobs simpler or deliver value in novel ways is not only a major source of learning, but it can be a source of innovation too. The number of agents created or shared, as well as their actual use, can be an early key performance indicator (KPI) of effective AI adoption. Additionally, principles that have proven successful in upskilling and re-skilling the workforce during the digital age will continue to apply; the execution has just been made easier. Moving "one skill at a time" makes it easier for most organizations to digest. Here is when training becomes a neces-

sary companion. Formal training, the 10 percent in the 70/20/10 rule, provides a base, and lowers the fears and obstacles to exploring and experimenting with new skills. One of the most overlooked benefits of formal training is its power to shape culture and reinforce a shared vision across the organization. When training is tightly aligned with strategic priorities, it becomes a vehicle for clarity and cohesion—a source of real value. Arguably, a value multiplier.

Listening Mode

In the initial stages of experimentation and learning, being watchful about the lessons learned in your organization, identifying the roadblocks, and addressing them will be a critical step in scaling the truly relevant learning. This somewhat reactive approach works well for engagement at scale on new AI skills. Eventually, certain roles will require a fundamental redefinition, and your own workforce equipped with their AI companions may tell you which ones they are. Reinventing roles for the age of AI can be scary; however, if your business avoids inaction, much can be achieved in a couple of years.

In summary, building an organization in constant learning and adaptation is likely to be one of the critical pillars of a rethinking strategy in the AI Era.

Engage and Mobilize the Team: There is Value in the Process

There's no question: top-down direction matters. Clear operating principles can give your organization focus, alignment, and speed. Sometimes, leadership must step in—to define what truly hard problems look like, to decide which options are worth serious investment, and to name the future-critical skills that simply aren't negotiable. However, if you put in place your unique listening and learning mechanisms the output may surprise you. The combination of top-down and bottom-up approaches is likely to take your business further faster. Defining and pressure-testing current strategy can be a learning exercise and a process by which your enterprise learns how to learn.

Typical strategy exercises focus on systematic reviews of current performance as well as the current market, competitors, and products. They often lack creativity and a wider look at the relevant context. There tends to be little room for the "what ifs." Taking the time with your team to develop relevant plausible and challenging scenarios will cognitively prime your leadership team and build greater agility when signs leading to potential polarities emerge.

The process of building the scenarios is more akin to a social process looking to engage with different stakeholders in the organization, including the front line, the R&D team if you have one, the IT and data teams, the manufacturing or delivery teams, and other less obvious participants in standard strategic exercises. As an example, the early scenario planning team at Shell took the time to travel to the Middle East back in the 1970s and to interview analysts and key opinion leaders to deeply understand the macroeconomics and geopolitics of Iran and Saudi Arabia and other oil producers in the region in order to define the relevant plausible scenarios that would challenge Shell in the late 1970s.

Who would your team need to interview to get a challenging perspective of what the AI Era will bring to your business? How will you make sure your organization is actively engaging in its own learning and transformation?

Conclusion

I hope this brief journey has helped you imagine non-linear futures for your business—futures that stretch beyond straight-line forecasts and fragmented information. By looking at the *known-knowns* and the *known-unknowns*, you've explored how to avoid the false confidence that comes from piecing together isolated data points.

The future will not politely wait for your next planning cycle. It is already unfolding. While predicting it with certainty is impossible, preparing for it is non-negotiable. As you engage your team in this work, aim to strike a balance between *opportunity and risk*. Be wary of the common reflex: *"Yes, that risk exists—but not for us."* That denial is often the seed of disruption. You now have practical tools: scenario building and identifying "pockets of the future"; the three horizons model; shifting lenses between the visionary, the entrepreneur, and the manager; and planting "strategic seeds" in the form of real options. Used well, these tools won't just help you anticipate change—they'll help you shape it in the Era of AI.

Most of all, I hope you've seen that AI is not just a tool. It is a force that can profoundly transform both your business and your life. Practicing foresight and managing the known-unknowns gives you more than decision-making power; it provides perspective and agency in times of turbulence.

Finally, a note of humility. Writing this book has been a reminder that facts age quickly—by the time these pages reach you, some will already be outdated. But the real value lies not in the facts, but in the *approach*: a way of thinking about your own known-unknowns. That is timeless.

I will continue sharing and updating these approaches—mostly on LinkedIn—not to chase the latest headlines, but to refine the methods and mindsets that will help you and your team navigate uncertainty with clarity, courage, and imagination.

Prompts from your strategy co-pilot

!

It is three years from now. It is Sunday morning. You are peacefully reading the newspaper. You get a notification from your strategy co-pilot; it is to check-in on your strategy for the Era of AI. Here are some of the prompts you may want to feed it with.

News from the future

- What known-knowns went unexpectedly fast or surprisingly slow in AI? Was it the AI infrastructure? Was it the geopolitical context? Was it the consumer behavior?
- What events that were once "known-unknowns" have now materialized—and with what consequences?
- Which of your three to five core future scenarios we identified still feel plausible? Which now seem either naïve or insufficiently bold?

https://doi.org/10.1515/9783111583495-010

- What pockets of the future have materialized faster than expected?
- Are there new plausible, relevant, and challenging scenarios to your business we should explore?

Business model bets

- Where has AI fundamentally improved your value proposition, beyond productivity gains?
- What is the real engine of value capture from our AI initiatives? Are we merely saving costs, or architecting new capabilities?
- How successful have we been in uncovering new monetization models?
- Have we achieved a balanced portfolio of options? Are we more at the optimization end of the spectrum or at the transformation end?
- Do we have encouraging prospects for a potential Engine 2?

Operating model bets: People and processes

- Is our "AI factory" truly operating as a system, or as a collection of pilots?
- What decisions are still made by humans that AI should own? Which are owned by AI that should be owned by humans?
- What governance is in place to ensure AI agents are aligned, supervised, and resilient to system shocks?
- What is the sentiment in the organization: are they learning alongside AI or are they fearing obsolescence?
- What is the balance of AI-native, AI-augmented, and AI-immune skills? Where are we falling short?
- Are learning incentives aligned with long-term value creation?

List of Figures

Figure 1 Futures cone — **5**
Figure 2 Foresight approach — **6**
Figure 3 Visualization of impact of the three horizons — **8**
Figure 4 AI infrastructure stack (Andrew Ng) — **14**
Figure 5 Peak computational performance in machine learning measured in 16 bits floating-point operations per second (2008–2024) — **15**
Figure 6 Data center electricity consumption by region: 2005–2024, per IEA — **16**
Figure 7 Energy required per LLM token (joules); Nvidia GPUS: 2014–2024, per Nvidia — **17**
Figure 8 Estimated training cost of frontier AI models (2016–2024) — **18**
Figure 9 Performance of AI models versus human performance — **19**
Figure 10 Key variables defining future plausible and challenging scenarios — **43**
Figure 11 Visual summary of domain-specific AI-run Company Scenario — **44**
Figure 12 The three horizons — **47**
Figure 13 Visual summary of Digital Feudalism Scenario — **49**
Figure 14 Visual summary of Augmented Enterprise Scenario — **53**
Figure 15 Visual summary of Personalized Tech-mosaic Scenario — **56**
Figure 16 Business model innovation — **66**
Figure 17 Business model innovation: Mapping your prototypes (illustrative example) — **78**
Figure 18 Monthly Global Economic Policy Uncertainty Index (September 2025) — **80**
Figure 19 Managing your portfolio of options — **88**
Figure 20 Simplified data architecture — **91**
Figure 21 Skill frequency: Example of actuaries — **103**

https://doi.org/10.1515/9783111583495-011

Notes

The author used artificial intelligence tools, including Chat GPT, Claude, Gemini, Mistral and others to experience first hand the power and functions of the tools, to improve clarity of writting, as well as to support the identification of relevant research. Responsibility for final content, critical interpretation and accuracy remains with the author.

1 Kieren, Pascal, Mertes, Lukas, and Weber, Martin (2024) "Information partitioning, learning and beliefs," September.
2 Note: Beware, GPT can also stand for generative pre-trained transformer, as used in the name ChatGPT.
3 Friede, G., Busch, T., and Bassen, A. (2015) "ESG and financial performance: Aggregated evidence from more than 2000 empirical studies," *Journal of Sustainable Finance & Investment*, 5(4).
4 Serafeim, G. (2020) "Social impact efforts that create real value," *Harvard Business Review*
5 Bain & Company (2024) "Financial or stakeholder value for shareholders? Both are best (Snap Chart)." Available at: https://www.bain.com/insights/financial-or-stakeholder-value-for-shareholders-both-are-best-snap-chart
6 https://es.linkedin.com/in/teresa-martin-retortillo
7 de Geus, Arie P. (1990). *Stockton Lecture – London Business School*, 3 May 1990. (Transcript on ariedegeus.com)
8 The Futures Cone was copyrighted by Joseh Voros.
9 Schwartz, P. (1991) *The Art of the Long View*. New York: Doubleday.
10 Sharpe, B. and Hodgson, A. (2006) *The Three Horizons: The Patterning of Hope*. Axminster: Triarchy Press.
11 Lannon, Colleen (n.d.) "Scenario-based planning: Managing by foresight." Available at: https://thesystemsthinker.com/scenario-based-planning-managing-by-foresight/
12 Sharpe, B. and Hodgson, A. (2006) *The Three Horizons: The Patterning of Hope*. Axminster: Triarchy Press.
13 Sharpe, B. and Hodgson, A. (2006) *The Three Horizons: The Patterning of Hope*. Axminster: Triarchy Press.
14 Ng, A. (2025) "Building faster with AI," Y Combinator AI Start-up School Keynote, July. Available at: https://www.youtube.com/watch?v=RNJCfif1dPY (accessed September 21, 2025).
15 Maslej, N., Fattorini, L., Perrault, R., Gil, Y., Parli, V., Kariuki, N., Capstick, E., Reuel, A., Brynjolfsson, E., Etchemendy, J., Ligett, K., Lyons, T., Manyika, J., Niebles, J.C., Shoham, Y., Wald, R., Walsh, T., Hamrah, A., Santarlasci, L., Betts Lotufo, J., Rome, A., Shi, A. and Oak, S., 2025. *The AI Index 2025 Annual Report*. Stanford, CA: AI Index Steering Committee, Institute for Human-Centered AI, Stanford University.
16 See ASML Investor Presentation, November 2024.
17 Interview with AWS CEO Matt Garman, *The Verge*, 13 January, 2025.
18 See ASML Investor Presentation, November 2024, with the outlook for the semiconductor industry to which it supplies its equipment.
19 Electric Power Research Institute (EPRI), 2024. *Powering intelligence: analysing artificial intelligence and data center energy consumption*. Palo Alto, CA: Electric Power Research Institute
20 Amodei, D. (2023) "Anthropic CEO on Claude, AGI & the future of AI & humanity | Lex Fridman Podcast #452," *Lex Fridman Podcast*, June 21. Available at: https://www.youtube.com/watch?v=XXXX

https://doi.org/10.1515/9783111583495-012

21 Maslej, N., Fattorini, L., Perrault, R., Gil, Y., Parli, V., Kariuki, N., Capstick, E., Reuel, A., Brynjolfsson, E., Etchemendy, J., Ligett, K., Lyons, T., Manyika, J., Niebles, J.C., Shoham, Y., Wald, R., Walsh, T., Hamrah, A., Santarlasci, L., Betts Lotufo, J., Rome, A., Shi, A. and Oak, S., 2025. *The AI Index 2025 Annual Report*. Stanford, CA: AI Index Steering Committee, Institute for Human-Centered AI, Stanford University.

22 Maslej, N., Fattorini, L., Perrault, R., Gil, Y., Parli, V., Kariuki, N., Capstick, E., Reuel, A., Brynjolfsson, E., Etchemendy, J., Ligett, K., Lyons, T., Manyika, J., Niebles, J.C., Shoham, Y., Wald, R., Walsh, T., Hamrah, A., Santarlasci, L., Betts Lotufo, J., Rome, A., Shi, A. and Oak, S., 2025. *The AI Index 2025 Annual Report*. Stanford, CA: AI Index Steering Committee, Institute for Human-Centered AI, Stanford University.

23 Villalobos, P., Ho, et al. 2024. *Position: Will we run out of data? Limits of LLM scaling based on human-generated data*. Proceedings of the 41st International Conference on Machine Learning, Proceedings of Machine Learning Research, vol. 235, pp. 49523–49544. Available at: https://proceedings.mlr.press/v235/villalobos24a.html

24 Singla, A., Sukharevsky, A., Yee, L., Chui, M. & Hall, B., 2025. *The state of AI: How organizations are rewiring to capture value.* March. McKinsey & Company, QuantumBlack, Available at: https://www.mckinsey.com/~/media/mckinsey/business%20functions/quantumblack/our%20insights/the-state-of-ai/2025/the-state-of-ai-how-organizations-are-rewiring-to-capture-value_final.pdf

25 As described by Harvard Business School professor Luis Viceira at a conference in 2025.

26 Fitz, R., Berns, M., Seshadri, P., Sastry, A., Nielsen, J., Gottesdiener, B. & Winnike, B., 2023. *Bridging the $18 trillion gap in net zero capital*. November. Boston Consulting Group (BCG) Center for Energy Impact. Available at: https://web-assets.bcg.com/3a/48/8261667a46ceba6573f51b3af0b7/bcg-bridging-the-18-trillion-gap-in-net-zero-capital-nov-2023.pdf

27 See Chapter 2 for a deeper explanation of AGI and ASI.

28 Interesting debate by two AI world known masterminds on the existential risks of AI. Yudkowsky, E. & Wolfram, S., 2024. 'Eliezer Yudkowsky and Stephen Wolfram on AI X-risk'. *Machine Learning Street Talk* (podcast), 11 November. Available at: https://podcasts.apple.com/ca/podcast/eliezer-yudkowsky-and-stephen-wolfram-on-ai-x-risk/id1510472996?i=1000676537725.

29 Shah, R., Irpan, A., Turner, A. M., Wang, A., Conmy, A., Lindner, D., Brown-Shah, R., Irpan, A., Turner, A. M., Wang, A., Conmy, A., Lindner, D., Brown-Cohen, J., Ho, L., Nanda, N., Popa, R. A., Jain, R., Greig, R., Albanie, S., Emmons, S., Farquhar, S., Krier, S., Rajamanoharan, S., Bridgers, S., Ijitoye, T., Everitt, T., Krakovna, V., Varma, V., Mikulik, V., Kenton, Z., Orr, D., Legg, S., Goodman, N., Dafoe, A., Flynn, F. & Dragan, A., 2025. An approach to technical AGI safety and security. arXiv preprint arXiv:2504.01849. Available at: https://arxiv.org/abs/2504.01849

30 *The Economist* (2025) "China and America are racing to develop the best AI. But who is ahead in using it?," 3 April.

31 *The Economist* (2025) "China and America are racing to develop the best AI. But who is ahead in using it?," 3 April.

32 That includes uranium mining or construction of reactors.

33 Katzke, C. & Hendrycks, D., 2025. *AI Safety Newsletter #49: Superintelligence Strategy*. 6 March. Centre for AI Safety: Substack. Available at: https://newsletter.safe.ai/p/ai-safety-newsletter-49-superintelligence

34 Solow, R.M., 1987. *We'd Better Watch Out*. New York Times Book Review, 12 July,

35 Jobs, S., 1990. Interview by WBGH Staff, Boston, 14 May. Available at: https://allaboutstevejobs.com/verbatim/quotes

36 CaixaBank Research, 2024. *Productivity growth in Europe: low, uneven and slowing*. IM06 No. 24-09 Dossier, June. Available at: https://www.caixabankresearch.com/sites/default/files/content/file/2024/06/21/34454/im06_24-09-dossier-3-en.pdf

37 Frey, C.B. & Osborne, M.A., 2013. *The Future of Employment: How Susceptible Are Jobs to Computerisation?* Oxford Martin School, University of Oxford, Oxford. Available at: https://www.oxfordmartin.ox.ac.uk/publications/the-future-of-employment/

38 Kelly, J., 2024. *Erik Brynjolfsson: 'This could be the best decade in history — or the worst'*. *Financial Times*, 30 January. Available at: https://www.ft.com/content/b71759fe-397b-4688-bc81-b082edb25f31

39 *The Economist* (2025) "How AI will divide the best from the rest," February 13. Available at: https://www.economist.com/finance-and-economics/2025/02/13/how-ai-will-divide-the-best-from-the-rest (accessed September 21, 2025).

40 Mäkelä, E. & Stephany, F., 2025. *Complement or Substitute? How AI Increases the Demand for Human Skills*. SSRN Working Paper. Available at: https://ssrn.com/abstract=5153230

41 Auren, Sesame, or Pi among others.

42 Niemeyer, K., 2025. *Reddit Stock Fell 15% After the CEO Said a Google Tweak Hurt Its User Growth*. *Business Insider*, 12 February. Available at: https://www.businessinsider.com/reddit-stock-drops-google-algorithm-tweak-hurt-its-user-growth-2025-2

43 Bond (2025) "Trends: Artificial Intelligence," May.

44 Thompson, B., 2025. *An Interview with Meta CEO Mark Zuckerberg About AI and the Evolution of Social Media*. Stratechery, 1 May. Available at: https://stratechery.com/2025/an-interview-with-meta-ceo-mark-zuckerberg-about-ai-and-the-evolution-of-social-media/

45 Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A.N., Kaiser, Ł. & Polosukhin, I., 2017. *Attention is all you need*. In: Advances in Neural Information Processing Systems 30 (NeurIPS 2017). Available at: https://arxiv.org/abs/1706.03762

46 Nadella, S., 2024. *LinkedIn post*, 22 June. Available at: https://www.linkedin.com/posts/satyanadella_as-ai-becomes-more-capable-and-agentic-models-activity-7241464002021617664-G2rN/

47 Shah, R. et el. 2025. *An approach to technical AGI safety and security*. arXiv preprint arXiv:2504.01849. Available at: https://arxiv.org/abs/2504.01849

48 Morris, M.R. et al. (2023) "Levels of AGI: Operationalizing progress on the path to AGI," Proceedings of the International Conference on Machine Learning.

49 Thompson, B., 2025. *Deep Research and Knowledge Value*. Stratechery, 10 February. Available at: https://stratechery.com/2025/deep-research-and-knowledge-value/

50 Assumes there are mainly four variables and each with only two possible outcomes.

51 Sharpe, B. and Hodgson, A. (2006) *The Three Horizons: The Patterning of Hope*. Axminster: Triarchy Press.

52 Gates, B., 1995. *The Road Ahead*. New York: Viking Penguin.

53 Adapted from Bocken et al. (2018), Osterwalder & Pigneur (2010) and Richardson (2008).

54 Christensen, C.M., 1997. *The Innovator's Dilemma: When New Technologies Cause Great Firms to Fail*. Boston, MA: Harvard Business School Press

55 McAfee, A. and Brynjolfsson, E. (2017) *Machine, Platform, Crowd*. New York: W.W. Norton.

56 Khullar, D. (2025) "Can A.I. find cures for untreatable diseases—using drugs we already have?," *The New Yorker*, July.

57 Sanjiv, 2024. *Intel Corporation: A mighty history but an uncertain future*. Long-Term Investing (Substack), 2 July. Available at: https://sanj2f3.substack.com/p/intel-inc-intc-a-mighty-history-but-an-uncertain-future

58 Note: One of the challenges of real options is the complexity of estimating volatility of outcomes. In the case described looking at expected pay-off is a solid starting point.
59 Taleb, N.N., 2012. *Antifragile: Things That Gain from Disorder*. New York: Random House.
60 Bain & Company, 2020. *The Engine 2 Imperative: New Business Innovation and Profitable Growth under Turbulence*. 4 December. Boston: Bain & Company. Available at: https://www.bain.com/insights/engine-2-imperative-new-business-innovation-and-profitable-growth-under-turbulence/
61 Note: I am referring to non-tech companies.
62 World Economic Forum (2020) "Here's how technology has changed—and changed us—over the past 20 years," November. Available at: https://www.weforum.org/agenda/2020/11/heres-how-technology-has-changed-and-changed-us-over-the-past-20-years//
63 Source: Verma, K., 2019. *A copy (for posterity) of Steve Yegge's internal memo in Google about what platforms are and how Amazon learnt to build them*. Gist (Public), 26 Dec. Memo reported by a former employee at Amazon that later moved to Google and shared the highlights of the memo employees had received.
64 Iansiti, M. and Lakhani, K. (2020) *Competing in the Age of AI: Strategy and Leadership when Algorithms and Networks Run the World*. Boston, Harvard Business Review Press.
65 Essa, A., 2025. *Artificial Intelligence: Shaping the Future of Innovation*. Berlin: De Gruyter.
66 Blenko, M., Bain, C., and Mankins, M. (2010) "Decide and deliver: Five steps to breakthrough performance in your organization," Harvard Business Review Press. Available at: https://www.bain.com/insights/the-five-steps-to-better-decisions
67 McKinsey & Company (2019) "Decision making in the age of urgency." Available at: https://www.mckinsey.com/capabilities/people-and-organizational-performance/our-insights/decision-making-in-the-age-of-urgency
68 Eisenhardt, K.M. (1989) "Making fast strategic decisions in high-velocity environments," *Academy of Management Journal*, 32(3), 543–576.
69 Shepherd, G., Rudd, J.M., & Greenley, G.E., 2021. Deciding fast: Examining the relationship between strategic decision speed and decision quality across multiple environmental contexts. *European Management Review*, 18(2), pp.119–140. https://doi.org/10.1111/emre.12429
70 Nadella, S., 2017. *Hit Refresh: The Quest to Rediscover Microsoft's Soul and Imagine a Better Future for Everyone*. New York: Harper Business.
71 Dweck, C.S., 2006. *Mindset: The New Psychology of Success*. New York: Random House.
72 Mäkelä, E. & Stephany, F., 2025. *Complement or Substitute? How AI Increases the Demand for Human Skills*. SSRN Working Paper. Available at: https://ssrn.com/abstract=5153230
73 AI Index Steering Committee (2024) *AI Index 2024 Annual Report*. Stanford, CA: Institute for Human-Centered AI, Stanford University.
74 Chatterji, A. et al. (2025) "How people use ChatGPT." OpenAI, Duke University, Harvard University, September.
75 Benjamin Bloom was an educational psychologist whose work in the 1950s produced the popular Bloom's taxonomy used in many educational settings to structure a course. The taxonomy involves five learning stages (1. Remember, 2 Understand and Apply, 3. Analyze, 4. Evaluate, and 5. Create). The taxonomy allows for the creation of clear learning goals and assessments on a given course.
76 Noy, S. and Zhang, W. (2023) "Experimental evidence on the productivity effects of generative artificial intelligence." Working Paper. Cambridge, MA: Massachusetts Institute of Technology.
77 Stephany, F., Teutloff, O., and Lehdonvirta, V. (2022) "What is the price of a skill? Revealing the complementary value of skills." Oxford Internet Institute, University of Oxford; Humboldt In-

stitute for Internet and Society; University of Copenhagen. Available at: https://mpra.ub.uni-muenchen.de/114874/

78 Korinek, A. (2023) "Language models and cognitive automation for economic research." NBER Working Paper No. 30957. Cambridge, MA: National Bureau of Economic Research.

79 Burning Glass Institute, 2022. *How Skills Are Disrupting Work: The Transformational Power of Fast-Growing, In-Demand Skills*. Cambridge, MA: The Burning Glass Institute. Available at: https://static1.squarespace.com/static/6197797102be715f55c0e0a1/t/6388b6daaae0b3075d6c7658/1669904091972/SkillsDisruption_Final_2022.pdf

80 Refers to Robert Solow's quote of "seeing the computer age everywhere except in the productivity statistics." Solow, R.M., 1987. *We'd Better Watch Out*. New York Times Book Review, 12 July,

Selected Bibliography

Ansón, J., Lanzerotti, L., Mankins, M.C. & Schor, M., 2020. *The Engine 2 Imperative: New Business Innovation and Profitable Growth under Turbulence*. 4 December. Boston: Bain & Company. Available at: https://www.bain.com/insights/engine-2-imperative-new-business-innovation-and-profitable-growth-under-turbulence/

Burning Glass Institute, 2022. *How Skills Are Disrupting Work: The Transformational Power of Fast-Growing, In-Demand Skills*. Cambridge, MA: The Burning Glass Institute. Available at: https://static1.squarespace.com/static/6197797102be715f55c0e0a1/t/6388b6daaae0b3075d6c7658/1669904091972/SkillsDisruption_Final_2022.pdf

Chatterji, A., Cunningham, T., Deming, D., Hitzig, Z., Ong, C., Shan, C. & Wadman, K., 2025. *How People Use ChatGPT*. NBER Working Paper No. w34255, September. Available at: https://www.nber.org/papers/w34255

Christensen, C.M., 1997. *The Innovator's Dilemma: When New Technologies Cause Great Firms to Fail*. Boston, MA: Harvard Business School Press

Dweck, C.S., 2006. *Mindset: The New Psychology of Success*. New York: Random House.

Eisenhardt, K.M. (1989) "Making fast strategic decisions in high-velocity environments," *Academy of Management Journal*, 32(3), 543–576.

Electric Power Research Institute (EPRI), 2024. *Powering intelligence: analysing artificial intelligence and data center energy consumption*. Palo Alto, CA: Electric Power Research Institute

Essa, A., 2025. *Artificial Intelligence: Shaping the Future of Innovation*. Berlin: De Gruyter.

Frey, C.B. & Osborne, M.A., 2013. *The Future of Employment: How Susceptible Are Jobs to Computerisation?* Oxford Martin School, University of Oxford, Oxford. Available at: https://www.oxfordmartin.ox.ac.uk/publications/the-future-of-employment/

Friede, G., Busch, T., and Bassen, A. (2015) "ESG and financial performance: Aggregated evidence from more than 2000 empirical studies," *Journal of Sustainable Finance and Investment*, 5(4), 210–233.

Fitz, R., Berns, M., Seshadri, P., Sastry, A., Nielsen, J., Gottesdiener, B. & Winnike, B., 2023. *Bridging the $18 trillion gap in net zero capital*. November. Boston Consulting Group (BCG) Center for Energy Impact. Available at: https://web-assets.bcg.com/3a/48/8261667a46ceba6573f51b3af0b7/bcg-bridging-the-18-trillion-gap-in-net-zero-capital-nov-2023.pdf

Gates, B., 1995. *The Road Ahead*. New York: Viking Penguin.

Iansiti, M. and Lakhani, K. (2020) *Competing in the Age of AI: Strategy and Leadership when Algorithms and Networks run the World*. Harvard Business Review Press

Jobs, S., 1990. Interview by WBGH Staff, Boston, 14 May. Available at: https://allaboutstevejobs.com/verbatim/quotes

Katzke, C. & Hendrycks, D., 2025. *AI Safety Newsletter #49: Superintelligence Strategy*. 6 March. Centre for AI Safety: Substack. Available at: https://newsletter.safe.ai/p/ai-safety-newsletter-49-superintelligence

Kelly, J., 2024. *Erik Brynjolfsson: 'This could be the best decade in history — or the worst'*. *Financial Times*, 30 January. Available at: https://www.ft.com/content/b71759fe-397b-4688-bc81-b082edb25f31

Kieren, Pascal, Mertes, Lukas, and Weber, Martin (2024) "Information partitioning, learning and beliefs," September.

Korinek, A. (2023) "Language models and cognitive automation for economic research." NBER Working Paper No. 30957. Cambridge, MA: National Bureau of Economic Research.

https://doi.org/10.1515/9783111583495-013

Mäkelä, E. & Stephany, F., 2025. *Complement or Substitute? How AI Increases the Demand for Human Skills*. SSRN Working Paper. Available at: https://ssrn.com/abstract=5153230

Maslej, N., Fattorini, L., Perrault, R., Gil, Y., Parli, V., Kariuki, N., Capstick, E., Reuel, A., Brynjolfsson, E., Etchemendy, J., Ligett, K., Lyons, T., Manyika, J., Niebles, J.C., Shoham, Y., Wald, R., Walsh, T., Hamrah, A., Santarlasci, L., Betts Lotufo, J., Rome, A., Shi, A. and Oak, S., 2025. *The AI Index 2025 Annual Report*. Stanford, CA: AI Index Steering Committee, Institute for Human-Centered AI, Stanford University.

McAfee, A. and Brynjolfsson, E. (2017) *Machine, Platform, Crowd*. New York: W.W. Norton.

Morris, M.R., Sohl-Dickstein, J., Fiedel, N., Warkentin, T., Dafoe, A., Faust, A., Farabet, C., and Legg, S. (2023) "Levels of AGI: Operationalizing progress on the path to AGI," Proceedings of the International Conference on Machine Learning.

Nadella, S., 2017. *Hit Refresh: The Quest to Rediscover Microsoft's Soul and Imagine a Better Future for Everyone*. New York: Harper Business.

Noy, S. and Zhang, W. (2023) "Experimental evidence on the productivity effects of generative artificial intelligence." Working Paper. Cambridge, MA: Massachusetts Institute of Technology.

Sanjiv, 2024. *Intel Corporation: A mighty history but an uncertain future*. Long-Term Investing (Substack), 2 July. Available at: https://sanj2f3.substack.com/p/intel-inc-intc-a-mighty-history-but-an-uncertain-future

Schwartz, P. (1991) *The Art of the Long View*. New York: Doubleday.

Serafeim, G. (2020) "Social impact efforts that create real value," *Harvard Business Review*.

Singla, A., Sukharevsky, A., Yee, L., Chui, M. & Hall, B., 2025. *The state of AI: How organizations are rewiring to capture value.* March. McKinsey & Company, QuantumBlack, Available at: https://www.mckinsey.com/~/media/mckinsey/business%20functions/quantumblack/our%20insights/the-state-of-ai/2025/the-state-of-ai-how-organizations-are-rewiring-to-capture-value_final.pdf

Shah, R., Irpan, A., Turner, A. M., Wang, A., Conmy, A., Lindner, D., Brown-Cohen, J., Ho,L., Nanda, N., Popa, R. A., Jain, R., Greig, R., Albanie, S., Emmons, S., Farquhar, S., Krier, S., Rajamanoharan, S., Bridgers, S., Ijitoye, T., Everitt, T., Krakovna, V., Varma, V., Mikulik, V., Kenton, Z., Orr, D., Legg, S., Goodman, N., Dafoe, A., Flynn, F. & Dragan, A., 2025. *An approach to technical AGI safety and security*. arXiv preprint arXiv:2504.01849. Available at: https://arxiv.org/abs/2504.01849

Sharpe, B. and Hodgson, A. (2006) *The Three Horizons: The Patterning of Hope*. Axminster: Triarchy Press.

Shepherd, G., Rudd, J.M., & Greenley, G.E., 2021. Deciding fast: Examining the relationship between strategic decision speed and decision quality across multiple environmental contexts. *European Management Review*, 18(2), pp.119–140. https://doi.org/10.1111/emre.12429.

Solow, R.M., 1987. *We'd Better Watch Out*. New York Times Book Review, 12 July,

Stephany, F., Teutloff, O., and Lehdonvirta, V. (2022) "What is the price of a skill? Revealing the complementary value of skills." Oxford Internet Institute, University of Oxford; Humboldt Institute for Internet and Society; University of Copenhagen. Available at: https://mpra.ub.uni-muenchen.de/114874/

Taleb, N.N., 2012. *Antifragile: Things That Gain from Disorder*. New York: Random House.

Thompson, B., 2025. *An Interview with Meta CEO Mark Zuckerberg About AI and the Evolution of Social Media*. Stratechery, 1 May. Available at: https://stratechery.com/2025/an-interview-with-meta-ceo-mark-zuckerberg-about-ai-and-the-evolution-of-social-media/

Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A.N., Kaiser, Ł. & Polosukhin, I., 2017. *Attention is all you need*. In: Advances in Neural Information Processing Systems 30 (NeurIPS 2017). Available at: https://arxiv.org/abs/1706.03762

Villalobos, P., Ho, et al. 2024. *Position: Will we run out of data? Limits of LLM scaling based on human-generated data*. Proceedings of the 41st International Conference on Machine Learning, Proceedings of Machine Learning Research, vol. 235, pp. 49523–49544. Available at: https://proceedings.mlr.press/v235/villalobos24a.html

Yudkowsky, E. & Wolfram, S., 2024. 'Eliezer Yudkowsky and Stephen Wolfram on AI X-risk'. *Machine Learning Street Talk* (podcast), 11 November. Available at: https://podcasts.apple.com/ca/podcast/eliezer-yudkowsky-and-stephen-wolfram-on-ai-x-risk/id1510472996?i=1000676537725

More in this Series

Alfred Essa, *Artificial Intelligence: Shaping the Future of Innovation*, De Gruyter, 2025.

Index

AI factory operating model
– burning platform (*See* burning platform)
– centralized and consistent enterprise 93–94
– components 93
– COVID-19 DNA 94
– decision quality and speed 95
– decision types 96
– digital and non-digital business 92
– Era of AI 92–93
– Moderna 94
– mRNA technology platform 94
– Netflix 92–93
– user-developed prototypes 95
AI infrastructure
– Big Tech 20
– cloud subsector 20
– components 13
– data center electricity consumption 16
– efficient manufacturing capabilities 13
– future development 13
– *vs.* human performance 19
– innovation 17
– machine learning hardware performance 15
– potential issues 20
– stack 14
– tokens 16
– training foundation models 17
– transistors and chips 13, 15
AI-Run Company
– diagnostic and treatment 46
– Horizon 46–47
– regulations and open-source models 46
– resource-scarce organizations 47
– scenario 43–44
– speed and performance-based economics 46
– the story 44–45
AI sovereignty 11–13
Amazon Web Services (AWS) 15–16, 37, 38–39, 84, 97
Apple Store 69
application programming interfaces (APIs) 68
AQUAVISTA™ solution 52–54
artificial general intelligence (AGI) 23, 39–40
artificial general super intelligence (AGSI) 39
Artificial Intelligence Index Report 2024 99
Artificial Intelligence: Shaping the future of innovation (Essa) 94
artificial super intelligence (ASI) 23
The Art of the Long View (Schwartz) 4
"as a Service" (aaS) 66–68
Augmented Enterprise Scenario 97

Bain & Company 85, 86
Berkshire Hathaway 87
Big Tech 20, 29
– advertising 33
– Digital Era 33–34
– industry structure 33
– strategy 35–39
– "winner-take-all" economics 34–35
The Bletchley Declaration (2023) 24
bring their own tutor (BYOT) 57
British Electricity Authority 10
burning platform, AI factory
– data architecture 89–90
– data infrastructure 90
– governance model 90
– intangible assets 91
– interoperability 90
– intrinsic enterprise value 91
– operating model 89
– simplified data architecture 91
– structured and unstructured data 90
business model innovation
– "as a Service" (aaS) 66–68
– customers and suppliers 65
– design 65
– Digital Era (2000–2020) 66, 71–73
– dynamic value chains 73–77
– future scenarios 65
– platform and ecosystem 69–71
– strategy exercises 65
– superior and sustainable value creation 65
business-to-business (B2B) 51
business-to-consumer (B2C) 4

https://doi.org/10.1515/9783111583495-014

capital-intensive sector 21
CATL (Contemporary Amperex Technology Co., Limited) 42
Center for AI Safety (CAIS) 25
ChatGPT skills 99
CO_2 emissions 24
Conference of the Parties (COP) 24
COVID-19 pandemic 72
creativity and innovation skills 101
crowd sourcing 71–72
customer acquisition costs (CAC) 67

Data as a Service (DaaS) 68
DeepSeek's V3 reasoning model 17
Defense Advanced Research Projects Agency (DARPA) 10
DesalinationCo 52–54
Digital Industries 94
dynamic programming 81
dynamic value chains
– AI, catalyst 77–78
– definition 73–74
– economists 75–76
– Intel Corporation 74
– microprocessor 74
– vertical integration 75

Earnings Before Interests and Taxes (EBIT) 15
Earnings Before Interest, Taxes, and Depreciation (EBITDA) 21
Edelman Trust Barometer 28
effective communication 100–101
Électricité de France (EDF) 10
Electric Power Research Institute (EPRI) 16
electric vehicles (EV) 42
Essa, Alfred 94
European Central Bank (ECB) 22
European Coal and Steel Community 10
European Network of Transmission System Operators for Electricity (ENTSO-E) 10

fear of missing out (FOMO) 3
financial crisis 98
Financial Times 31, 57
future/future scenarios
– Augmented Enterprise Scenario 52–55
– business strategy 3
– business-to-consumer (B2C) 4
– conceptual frameworks 7
– cone 5
– definition 41
– digital feudalism 48–51
– Era of AI 3, 41
– fear of missing out (FOMO) 3
– foresight approach 4, 6
– group planning 6–7
– horizon 7–8
– "memory of the future" 7
– optionality 41
– Personalized Tech-mosaic Scenario 55–59
– plausible and challenging 4, 42–43
– "pockets of the future" 41–42
– pre-determined variables 6
– premises 3
– probabilities 4–5
– weather radar 4

General Electric (GE) 86
geopolitical power
– AI sovereignty 11–13
– and economic 9
– electricity 10
– general-purpose technologies 9, 10
Global Economic Policy Uncertainty Index 79
Global Positioning System (GPS)
– AI-augmented skills 100–101
– AI-run Company Scenario 97
– corporate environments 97
– defining and pressure-testing 105
– demanded skills 98–99
– financial crisis 98
– hub skills 102
– learning organization 97, 104
– learning to value creation 104–105
– new roles 103
– rethinking strategy 97
– school environments 97
– skill frequency 102–103
– social process 106
Google's DeepMind 37
graphics processing unit (GPU) 12, 15, 35
greenhouse gases 24

gross domestic product (GDP) 21–22

HealthCo 45
high bandwidth memory (HBM) 15
high-quality decision making 100
human augmentation *vs.* automation
– AI impact on the workforce 30–31
– human–AI interaction 32–33
– new skills 32
– paradox 31–32
– Solow's growth model 29–30
hyperscaler 49–50

Inference as a Service (IaaS) 68
Infrastructure as a Service (IaaS) 84
integrated circuits (ICs) 74
Intel Corporation 74, 76
"Intelligence Grid" 12
Intergovernmental Panel on Climate Change (IPCC) 24
International Atomic Energy Agency (IAEA) 23
International Energy Agency 16
Internet Corporation for Assigned Names and Numbers (ICANN) 11
Internet Engineering Task Force (IETF) 11
Internet Service Providers (ISPs) 11

key performance indicator (KPI) 104
K-Health symptom assessment tool 44
known-knowns
– AI safety 23–26
– capital 21–22
– geopolitical power (*See* geopolitical power)
– pre-determined trends 9
– trends 9
known-unknowns
– pre-determined variables 27
– trust in AI 27–29

large language model (LLM) 17, 36
Lifetime Value of customers 67

Mäkelä, E. 31, 99
messenger RNA 94
microprocessor 74
Microsoft Office 37, 99
Model as a Service (MaaS) 68
Moderna 94
Monte Carlo simulations 81
Moore's Law 13
Morris, M. R. 39
multi-pronged approach 11

National Science Foundation (NSF) 10–11
Netflix 92–93
net present value (NPV) 81
New Generation Artificial Intelligence Development Plan 11
The New York Times 29
nuclear energy 28
Nuclear Non-Proliferation Treaty (NPT) 23

optionality
– AI-run Company Scenario 79
– antifragility 82
– Augmented Enterprise Scenario 79–80
– building 85
– capital investment 83
– classic strategy recommends 83–84
– conglomerate 86
– customer segments 84
– economic and geopolitical uncertainty 79
– innovation portfolio 87–88
– innovation projects 82
– institutional level 83
– leadership economics 86–87
– Net Present Value (NPV) 81
– plausible future 81
– products and services, portfolio of 79
– resource-conscious companies 83
– risk-management tools 79
– selling complements 84
– sustainable growth companies 85
– U.S. market 83
– variety of industries 85–86
Organisation for Economic Co-operation and Development (OECD) 22
Original Equipment Manufacturers (OEMs) 74

Patriot Act 2001 11
pay-per-crawl 71
"pockets of the future" 43–44, 107
pre-determined variables 6
printed circuit board (PCB) 74

Productivity Paradox 29
project management and collaboration 101

Responsible Scaling Policy 25
return on investment (ROI) 41
returns on equity (ROE) 20

Schwartz, Peter 4
self-learning 100
semiconductors 13
Siemens 94
social media platforms 69
Software as a Service (SaaS) models 66–67
Software Development Kits (SDKs) 72
Solow's growth model 29–30
Stephany, F. 31, 99, 102
strategic infrastructure 12

Taiwan Semiconductor Manufacturing Company (TSMC) 76
tensor processing units (TPUs) 37
Training as a Service (TaaS) 68
Transaction Cost Economics (TCE) 75

United Nations AI Advisory Body (2023) 25
United Nations Framework Convention on Climate Change (UNFCCC) 24
UniX 57
UN Security Council 23
user-developed prototypes 95

water-as-a-service model 54
World Economic Forum (WEF) 31–32, 89, 98–99
World Uncertainty Index (WUI) 79